

---

# Bayesian Coresets: Revisiting the Nonconvex Optimization Perspective

---

Jacky Y. Zhang  
UIUC

Rajiv Khanna  
UC Berkeley

Anastasios Kyrillidis  
Rice University

Oluwasanmi Koyejo  
UIUC

## Abstract

Bayesian coresets have emerged as a promising approach for implementing scalable Bayesian inference. The Bayesian coreset problem involves selecting a (weighted) subset of the data samples, such that the posterior inference using the selected subset closely approximates the posterior inference using the full dataset. This manuscript revisits Bayesian coresets through the lens of sparsity constrained optimization. Leveraging recent advances in accelerated optimization methods, we propose and analyze a novel algorithm for coreset selection. We provide explicit convergence rate guarantees and present an empirical evaluation on a variety of benchmark datasets to highlight our proposed algorithm’s superior performance compared to state-of-the-art on speed and accuracy.

## 1 Introduction

Bayesian coresets have emerged as a promising approach for scalable Bayesian inference (Huggins et al., 2016; Campbell & Broderick, 2018, 2019; Campbell & Beronov, 2019). The key idea is to select a (weighted) subset of the data such that the posterior inference using the selected subset closely approximates the posterior inference using the full dataset. This creates a trade-off, where using Bayesian coresets as opposed to the full dataset exchanges approximation accuracy for computational speedups. We study Bayesian coresets as they are easy to implement, effective in practice, and come with useful theoretical guarantees that relate the coreset size with the approximation quality.

The main technical challenge in the Bayesian coreset

problem lies in handling the combinatorial constraints – we desire to select a few data points out of many as the coreset. In terms of optimization, previous approaches mainly rely on two ideas: *convexification* and *greedy methods*. In convexification (Campbell & Broderick, 2019), the sparsity constraint – i.e., selection of  $k$  data samples – is relaxed into a convex  $\ell_1$ -norm constraint. This allows them to use out-of-the-box solvers such as Frank-Wolfe (FW) type-of methods (Frank & Wolfe, 1956; Jaggi, 2013). An alternative approach is by using greedy methods (Campbell & Broderick, 2018), which constructs a sparse weight vector based on local decisions to greedily optimize the approximation problem (Tropp & Gilbert, 2007; Needell & Tropp, 2009). The resulting method, greedy iterative geodesic ascent (GIGA), achieves linear convergence with no hyper-parameter tuning and optimal scaling (Campbell & Broderick, 2018). More recently, sparse variational inference (SparseVI) is considered for Bayesian coreset construction. SparseVI also employs a greedy algorithm to minimize a KL divergence objective. The method achieves state-of-the-art accuracy, but at a cost of higher computational requirements. Therefore, existing work illustrates the trade-off between accuracy and efficiency, opening a gap for improvements.

We revisit Bayesian coresets through the lens of sparsity constrained optimization. Sparsity, a kind of nonconvexity, appears in a variety of applications in machine learning and statistics. For instance, compressed sensing (Donoho et al., 2006; Candes, 2008) is an example where sparsity is used as a complexity measure for signal representation. Leveraging and building upon recent advances in non-convex optimization, we solve the Bayesian coreset problem based on hard thresholding algorithms (Blumensath & Davies, 2009) that directly work on the non-convex sparsity constraint. Hard-thresholding schemes are highly flexible, and easily accommodate variations such as subspace exploration (Dai & Milenkovic, 2009), de-bias steps (Needell & Tropp, 2009), adaptive step size selections (Kyrillidis & Cevher, 2011), as well as different types of sparsity constraints, such as group sparsity (Baldassarre et al., 2016), sparsity within groups (Kyrillidis et al.,

2015), and generic structured sparsity (Baraniuk et al., 2010). The thresholding step involves a projection onto the  $k$ -sparsity constraint set to determine the selected sample set in each iteration. While we achieve state-of-the-art accuracy using direct application of this algorithm, re-building the set in every iteration makes it slower than previous works. To fix this, we employ line search for step size selection and momentum based techniques (Khanna & Kyrillidis, 2018) to accelerate the algorithm, also achieving state-of-the-art speed.

**Contributions.** In this paper, we adapt accelerated iterative hard thresholding schemes to the Bayesian coreset problem. Despite directly attacking the non-convex optimization problem, we provide strong convergence guarantees. To summarize our contributions:

- We revisit the Bayesian coreset problem via a non-convex (sparse) optimization lens, and provide an IHT-based algorithm that combines hard thresholding and momentum steps;
- We analyze its convergence based on standard assumptions;
- We provide extensive empirical evaluation to show superior performance of the proposed method vis-à-vis state-of-the-art algorithms in terms of approximation accuracy as well as speed.

## 2 Problem Formulation

Given  $n$  observations, one can compute the log-likelihood  $\mathcal{L}_i(\theta)$  of each of the observations, parameterized by  $\theta$ . Assuming observations are conditionally independent given  $\theta$ , one can represent the likelihood of all the observations as the sum of individual log-likelihoods, i.e.,  $\mathcal{L}(\theta) = \sum_{i=1}^n \mathcal{L}_i(\theta)$ . With prior density  $\pi_0(\theta)$ , the posterior density can be derived as:

$$\pi(\theta) := \frac{1}{Z} \cdot e^{\mathcal{L}(\theta)} \cdot \pi_0(\theta),$$

where  $Z = \int e^{\mathcal{L}(\theta)} \pi_0(\theta) d\theta$  is a normalization factor.

However, for most applications, exact posterior estimation is intractable; i.e.,  $\pi$  is too hard to evaluate exactly. Practitioners use algorithms for approximate inference that may approximate the  $\pi$  in a closed-form (e.g., using variational inference), or allow for sampling from the posterior without providing a closed-form expression (e.g., MCMC methods). Such algorithms often scale at least linearly with the size of the dataset  $n$ , which makes them prohibitively expensive for large datasets. As such, designing algorithms to speed up inference is an area of active research.

One solution to the scalability problem is to use coresets. Coresets approximate the empirical log-likelihood  $\mathcal{L} =$

$\sum_{i=1}^n \mathcal{L}_i$  using a *weighted sum of a subset of all the log-likelihoods*  $\mathcal{L}_i$ . In other words, we use  $\mathcal{L}_w = \sum_{i=1}^n w_i \mathcal{L}_i$  to approximate the true  $\mathcal{L}$ , where  $w \in \mathbb{R}_+^n$  is a non-negative sparse vector. It will be useful to view that  $\mathcal{L}, \mathcal{L}_i$  and  $\mathcal{L}_w$  are functions in a Hilbert space, and we will use  $L^2$ -norm to denote the 2-norm defined in function space, differentiating with the  $\ell_2$ -norm defined in Euclidean space. We enforce the sparsity constraint as  $\|w\|_0 \leq k$ , for  $k < n$ ; here  $\|\cdot\|_0$  denotes the pseudo-norm that counts the number of non-zero entries.

When  $k < n$ , posterior estimation (e.g., using MCMC or variational inference) is less expensive on the coreset as opposed to the entire dataset. However, sparsifying  $w$  involves dropping some samples, which in turn implies deviating from the best performance possible from using the full dataset. The Bayesian coreset problem is formulated to minimize this loss in performance.

**The Bayesian Coreset Problem.** *The Bayesian coreset problem is to control the deviation of coreset log-likelihood from true log-likelihood via sparsity:*

$$\begin{aligned} \arg \min_{w \in \mathbb{R}^n} \quad & f(w) := \text{DIST}(\mathcal{L}, \mathcal{L}_w) \\ \text{s.t.} \quad & \|w\|_0 \leq k, w_i \geq 0, \forall i. \end{aligned} \quad (1)$$

Key components are (i) the weights  $w \in \mathbb{R}_+^n$  over  $n$  data points, (ii) the function  $f(\cdot)$  that controls the deviation between the full-dataset log-likelihood  $\mathcal{L}$  and the coreset log-likelihood  $\mathcal{L}_w$  using the distance functional  $\text{DIST}(\cdot, \cdot)$ , and (iii) the non-convex sparsity constraint that restricts the number of nonzeros in  $w$ , thus constraining the number of active data points in the coreset. Examples of  $\text{DIST}(\cdot, \cdot)$  include the weighted  $L^2$ -norm (Campbell & Broderick, 2019) and the KL-divergence (Campbell & Beronov, 2019). In this manuscript, we consider the  $L^2(\hat{\pi})$ -norm as the distance metric in the embedding Hilbert space, i.e.,

$$\begin{aligned} \text{DIST}(\mathcal{L}, \mathcal{L}_w)^2 &= \|\mathcal{L} - \mathcal{L}_w\|_{\hat{\pi}, 2}^2 \\ &= \mathbb{E}_{\theta \sim \hat{\pi}} [(\mathcal{L}(\theta) - \mathcal{L}_w(\theta))^2], \end{aligned} \quad (2)$$

where  $\hat{\pi}$  is a weighting distribution that has the same support as true posterior  $\pi$ . Ideally,  $\hat{\pi}$  is the true posterior, which is obviously unknown. However, one can employ Laplace approximation to derive an inexpensive and reasonable approximation for  $\hat{\pi}$  (Campbell & Broderick, 2019).

To account for the shift invariance, we write  $g_i = \mathcal{L}_i - \mathbb{E}_{\theta \sim \hat{\pi}} \mathcal{L}_i(\theta)$ , so the equivalent optimization problem is now: minimize  $\|\sum_{i=1}^n g_i - \sum_{i=1}^n w_i g_i\|_{\hat{\pi}, 2}^2$ . Further, noting that the  $L^2(\hat{\pi})$ -norm is in the form of expectation (equation (2)), it can be approximated by a finite-dimensional  $\ell_2$ -norm which replaces the function with a vector of sampled evaluations  $\theta \sim \hat{\pi}$ . Thus,

given  $S$  samples  $\{\theta_j\}_{j=1}^S, \theta_j \sim \hat{\pi}$ , and using

$$\hat{g}_i = \frac{1}{\sqrt{S}} \cdot [\mathcal{L}_i(\theta_1) - \bar{\mathcal{L}}_i, \dots, \mathcal{L}_i(\theta_S) - \bar{\mathcal{L}}_i]^\top \in \mathbb{R}^S$$

as projections from function space to standard Euclidean space, where  $\bar{\mathcal{L}}_i = \frac{1}{S} \sum_{j=1}^S \mathcal{L}_i(\theta_j)$ , the Bayesian coreset problem (1) becomes a *finite-dimensional sparse regression problem*:

$$\begin{aligned} \arg \min_{w \in \mathbb{R}^n} f(w) &:= \left\| \sum_{i=1}^n \hat{g}_i - \sum_{i=1}^n w_i \hat{g}_i \right\|_2^2 \\ \text{s.t.} \quad &\|w\|_0 \leq k, \quad w_i \geq 0, \forall i. \end{aligned} \quad (3)$$

The resulting sparse regression problem is non-convex due to the combinatorial nature of the constraints. Previous methods that use this  $\ell_2$ -norm formulation (Campbell & Broderick, 2019, 2018) offers less satisfactory approximation accuracy compared to the state-of-the-art sparse variational inference method (Campbell & Beronov, 2019). However, the high computational cost of the latter method makes it impractical for real-world large datasets. Nonetheless, as we will show, our approach for solving equation (3) using a variant of iterative hard thresholding, achieves better accuracy and speed.

### 3 Our approach

---

#### Algorithm 1 Vanilla IHT

---

**input** Objective  $f : \mathbb{R}^n \rightarrow \mathbb{R}$ ; sparsity  $k$ ; step size  $\mu$   
 1: Initialize  $w$   
 2: **repeat**  
 3:  $w \leftarrow \Pi_{\mathcal{C}_k \cap \mathbb{R}_+^n}(w - \mu \nabla f(w))$   
 4: **until** Stop criteria met  
 5: **return**  $w$

---

For clarity of exposition, we gradually build up our approach for solving the optimization problem (3). The fundamental ingredient of our approach is the vanilla Iterative Hard Thresholding (IHT) method presented in Algorithm 1. We develop our approach by augmenting IHT with momentum updates, step size selection for line search and active subspace expansion techniques to accelerate and automate the algorithm (Algorithms 2 & 3). Details follow.

#### 3.1 Iterative Hard Thresholding (IHT)

The classical IHT (Blumensath & Davies, 2009) is a projected gradient descent method that performs a gradient descent step and then projects the iterate onto the non-convex  $k$ -sparsity constraint set. We denote the orthogonal projection of a given  $z \in \mathbb{R}^n$  to a space  $\mathcal{C} \subseteq \mathbb{R}^n$  as:  $\Pi_{\mathcal{C}}(z) := \arg \min_{w \in \mathcal{C}} \|w - z\|_2$ . Define the sparsity restricted space as:  $\mathcal{C}_k = \{w \in \mathbb{R}^n :$

$|\text{supp}(w)| \leq k\}$ , where  $\text{supp}(w) = \{i | w_i \neq 0\}$  denotes the support set of  $w$ . Here, we describe the plain sparsity case, but one can consider different realizations of  $\mathcal{C}_k$  as in (Baldassarre et al., 2016; Kyrillidis et al., 2015; Baraniuk et al., 2010). The projection step in the classical IHT, *i.e.*,  $\Pi_{\mathcal{C}_k}$ , can be computed easily by selecting the top- $k$  elements in  $O(n \log k)$  time; but projection can be more challenging for more complex constraint sets, *e.g.*, if the variable is a distribution on a lattice (Zhang et al., 2019).

For our problem, we require that the projected sparse vector only has non-negative values. For vector variate functions, the projection step in Algorithm 1, *i.e.*,  $\Pi_{\mathcal{C}_k \cap \mathbb{R}_+^n}(w)$  is also straightforward; it can be done optimally in  $O(n \log k)$  time by simply picking the top  $k$  largest *non-negative* elements. More discussions about the projections are presented in section B in appendix.

#### 3.2 Accelerated IHT

For clarity, we rewrite the problem in equation (3) as:

$$w^* = \arg \min_{w \in \mathcal{C}_k \cap \mathbb{R}_+^n} f(w) := \|y - \Phi w\|_2^2,$$

where  $y = \sum_{i=1}^n \hat{g}_i$  and  $\Phi = [\hat{g}_1, \dots, \hat{g}_n]$ . In this case,  $\nabla f(w) \equiv -2\Phi^\top(y - \Phi w)$ .

**Step size selection in IHT:** Classical results on the performance of IHT algorithms come with rigorous convergence guarantees (under regularity conditions) (Blumensath & Davies, 2009; Foucart, 2011). However, these results require step size assumptions that either do not work in practice, or rely on strong assumptions. For example, in (Blumensath & Davies, 2009; Foucart, 2011) strong isometry constant bounds are assumed to allow step size  $\mu = 1$  for all the iterations, and thus remove the requirement of hyper-parameter tuning. Moreover, the authors in (Blumensath & Davies, 2010) present toy examples by carefully selecting  $\Phi$  so that the vanilla IHT algorithm diverges without appropriate step size selection. In this work, given the quadratic objective  $f(w)$ , we perform exact line search to obtain the best step size per iteration (Blumensath & Davies, 2010; Kyrillidis & Cevher, 2011):  $\mu_t := \|\tilde{\nabla}_t\|_2^2 / 2 \|\Phi \tilde{\nabla}_t\|_2^2$ ; details in Algorithm 2.

**Memory in vanilla IHT:** Based upon the same ideas as step size selection, we propose to include adaptive momentum acceleration; we select the momentum term as the minimizer of the objective:  $\tau_{t+1} = \arg \min_{\tau} f(w_{t+1} + \tau(w_{t+1} - w_t)) = \frac{\langle y - \Phi w_{t+1}, \Phi(w_{t+1} - w_t) \rangle}{2 \|\Phi(w_{t+1} - w_t)\|_2^2}$ , which also comes out as a closed-form solution. The step  $z_{t+1} = w_{t+1} + \tau_{t+1}(w_{t+1} - w_t)$  at the end of the algorithm captures memory in the algorithm based on the results on acceleration by Nesterov (1983) for convex optimization.

---

**Algorithm 2** Automated Accelerated IHT
 

---

**input** Objective  $f(w) = \|y - \Phi w\|_2^2$ ; sparsity  $k$   
 1:  $t = 0, z_0 = 0, w_0 = 0$   
 2: **repeat**  
 3:  $\mathcal{Z} = \text{supp}(z_t)$   
 4:  $\mathcal{S} = \text{supp}(\Pi_{\mathcal{C}_k \setminus \mathcal{Z}}(\nabla f(z_t))) \cup \mathcal{Z}$  where  $|\mathcal{S}| \leq 3k$   
 5:  $\tilde{\nabla}_t = \nabla f(z_t)|_{\mathcal{S}}$   
 6:  $\mu_t = \arg \min_{\mu} f(z_t - \mu \tilde{\nabla}_t) = \frac{\|\tilde{\nabla}_t\|_2^2}{2\|\Phi \tilde{\nabla}_t\|_2^2}$   
 7:  $w_{t+1} = \Pi_{\mathcal{C}_k \cap \mathbb{R}_+^n}(z_t - \mu_t \nabla f(z_t))$   
 8:  $\tau_{t+1} = \arg \min_{\tau} f(w_{t+1} + \tau(w_{t+1} - w_t))$   
      $= \frac{\langle y - \Phi w_{t+1}, \Phi(w_{t+1} - w_t) \rangle}{2\|\Phi(w_{t+1} - w_t)\|_2^2}$   
 9:  $z_{t+1} = w_{t+1} + \tau_{t+1}(w_{t+1} - w_t)$   
 10:  $t = t + 1$   
 11: **until** Stop criteria met  
 12: **return**  $w_t$

---

**Automated Accelerated IHT for coreset selection:** Combining the ideas above leads to Automated Accelerated IHT, as presented in Algorithm 2. The algorithm alternates between the projection step (steps 6 and 7) after the gradient updates, and the momentum acceleration step (step 8). It thus maintains two sets of iterates that alternatively update each other in each iteration at only a constant factor increase in per iteration complexity. The iterate  $w_t$  at iteration  $t$  is the most recent estimate of the optimizer, while the iterate  $z_t$  models the effect of momentum or “memory” in the iterates. We have shown exact line search that solves one dimensional problems to automate the step size selection ( $\mu$ ) and the momentum parameter ( $\tau$ ) for acceleration. In practice, these parameters can also be selected using a backtracking line search.

**Using de-bias steps in Automated Accelerated IHT:** Based on pursuit methods for sparse optimization (Needell & Tropp, 2009; Dai & Milenkovic, 2009; Kyrillidis & Cevher, 2014), we propose a modification that improves upon Algorithm 2 both in speed and accuracy in empirical evaluation. The modified algorithm is presented in Algorithm 3 in section A in appendix due to space limitations. The key differences of Algorithm 3 from Algorithm 2 are that, with additional de-bias steps, one performs another gradient step and a line search in the sparsified space in each iteration for further error reduction. We omit these steps in the algorithmic description to maintain clarity, since these steps do not provide much intellectual merit to the existing algorithm, but help boost the practical performance of Automated Accelerated IHT.

### 3.3 Theoretical Analysis

In this subsection, we study the convergence properties of our main algorithm Automated Accelerated IHT in

Algorithm 2. We make a standard assumption about the objective – the Restricted Isometry Property or RIP (Assumption 1), which is a standard assumption made for analysis of IHT and its variants. In RIP,  $\alpha_k$  reflects the convexity and  $\beta_k$  reflects the smoothness of the objective in some sense (Khanna & Kyrillidis, 2018; Kyrillidis & Cevher, 2014). We note that the assumption may not be necessary but is sufficient to show convergence theoretically. For example, if the number of samples required to exactly construct  $\hat{g}$  is less than the coreset size ( $a_k = 0$  in RIP), so that the system becomes under-determined, then a local minimum can also be global achieving zero error without assuming that the RIP holds. On the other hand, when the number of samples goes to infinity, RIP is saying that the restricted eigenvalues of covariance matrix,  $\text{cov}[\mathcal{L}_i(\theta), \mathcal{L}_j(\theta)]$  where  $\theta \sim \hat{\pi}$ , are upper bounded and lower bounded away from 0. It is an active area of research in random matrix theory to quantify RIP constants e.g. see (Baraniuk et al., 2008).

**Assumption 1** (Restricted Isometry Property (RIP)). *The matrix  $\Phi$  in the objective function satisfies the RIP property, i.e., for  $\forall w \in \mathcal{C}_k$*

$$\alpha_k \|w\|_2^2 \leq \|\Phi w\|_2^2 \leq \beta_k \|w\|_2^2.$$

RIP generalizes to restricted strong convexity and smoothness (Chen & Sanghavi, 2010); thus our results could potentially be extended to general convex  $f(\cdot)$  functions. We present our main result next, and defer the details of the theory to section B in the appendix.

**Theorem 1.** *In the worst case scenario, with Assumption 1, the solutions path found by Automated Accelerated IHT satisfies the following iterative invariant.*

$$\begin{aligned} \|w_{t+1} - w^*\|_2 &\leq \rho |1 + \tau_t| \cdot \|w_t - w^*\|_2 \\ &\quad + \rho |\tau_t| \cdot \|w_{t-1} - w^*\|_2 + 2\beta_{3k} \sqrt{\beta_{2k}} \|\epsilon\|_2, \end{aligned}$$

where  $\rho = \left( 2 \max\left\{ \frac{\beta_{2k}}{\alpha_{3k}} - 1, 1 - \frac{\alpha_{2k}}{\beta_{3k}} \right\} + \frac{\beta_{4k} - \alpha_{4k}}{\alpha_{3k}} \right)$ , and  $\|\epsilon\|_2 = \|y - \Phi w^*\|_2$  is the optimal error.

The theorem provides an upper bound invariant among consecutive iterates of the algorithm. To have a better sense of convergence rate, we can derive linear convergence from our iterative invariant.

**Corollary 1.** *Given the iterative invariant as stated in Theorem 1, and assuming the optimal solution achieves  $\|\epsilon\|_2 = 0$ , the solution found by Algorithm 2 satisfies:*

$$f(w_{t+1}) - f(w^*) \leq \phi^t \left( \frac{\beta_{2k}}{\alpha_{2k}} f(w_1) + \frac{\rho \tau \beta_{2k}}{\phi \alpha_k} f(w_0) \right),$$

where  $\phi = (\rho(1 + \tau) + \sqrt{\rho^2(1 + \tau)^2 + 4\rho\tau})/2$  and  $\tau = \max_{i \in [t]} |\tau_i|$ . It is sufficient to show linear convergence to the global optimum, when  $\phi < 1$ , or equivalently  $\rho < 1/(1 + 2\tau)$ .

Thus, Algorithm 2 generates a sequence of iterates that decrease the quadratic objective in equation (3) at a geometric rate. The quadratic objective can upper bound the sum of forward KL and reverse KL divergences between the constructed coresets posterior and the true posterior under certain conditions, as shown in Proposition 2 by Campbell & Beronov (2019), which further justifies our approach of using this objective.

Our theory and algorithm differ from the work by Khanna & Kyrillidis (2018) in several ways. The non-negative constraint is unique to the Bayesian coreset problem, and extending the analysis from the original IHT to our setting is non-trivial (see Section B in appendix). Further, the new analysis we present does not work with the restricted gradient used by Khanna & Kyrillidis (2018), which is why we choose to use the full gradient instead (line 7 in Algorithm 2). We also observe empirically in our experiments that using the full gradient performs better for the coreset problem. Further, we also automated the step-size selection, the momentum selection, and the de-bias step selection to minimize the need of tuning. Recall that vanilla IHT (Algorithm 1) is much slower than the greedy approach by Campbell & Broderick (2018), and so the enhancements we propose are crucial to ensure that the overall algorithm is both faster as well as better performing than the state-of-the-art.

## 4 Related Work

Other scalable approaches for Bayesian inference include subsampling and streaming methods for variational Bayes (Hoffman et al., 2013; Broderick et al., 2013), subsampling methods for MCMC (Welling & Teh, 2011; Ahn et al., 2012; Korattikara et al., 2014; Maclaurin & Adams, 2015), and consensus methods for MCMC (Srivastava et al., 2015; Rabinovich et al., 2015; Scott et al., 2016). These algorithms are motivated by empirical performance and come with few or no theoretical optimization-based guarantees on the inference quality, and often do not scale to larger datasets. Bayesian coresets could be used as part of these approaches, thus resulting into a universal tool for approximate MCMC and variational inference.

There have been few studies that study convergence properties of approximate inference algorithms. Campbell & Beronov (2019) presented a linear convergence rate, but the assumptions they make are non-standard as the rate of convergence depends on the how well individual samples correlate with the overall loss. Approximation guarantees in terms of KL-divergence are provided (Koyejo et al., 2014; Khanna et al., 2017) for structured sparse posterior inference using the greedy forward selection procedure. Locatello et al. (2017,

2018) study convergence rates for a boosting based algorithm for iteratively refined variational inference.

Thresholding based optimization algorithms have been attractive alternatives to relaxing the constraint to a convex one or to greedy selection. Bahmani et al. (2013) provide a gradient thresholding algorithm that generalizes pursuit approaches for compressed sensing to more general losses. Yuan et al. (2018) study convergence of gradient thresholding algorithms for general losses. Jain et al. (2014) consider several variants of thresholding based algorithms for high dimensional sparse estimation. Additional related works are discussed in Section D in the appendix.

## 5 Experiments

We empirically examine the performance of our algorithms to construct coresets for Bayesian posterior approximation. Three sets of experiments are presented: Gaussian posterior inference, Bayesian radial basis function regression, and Bayesian logistic and Poisson regression using real-world datasets.

Besides the Automated Accelerated IHT (Algorithm 2), we propose Automated Accelerated IHT - II (Algorithm 3 in appendix), that adds a de-bias step that further improves Algorithm 2 in practice. We refer to the appendix for detailed explanation and discussion of Algorithm 3 due to space limitation.

The proposed algorithms, Automated Accelerated IHT (A-IHT) and Automated Accelerated IHT II (A-IHT II), are compared with three baseline algorithms, *i.e.*, Random (Uniform), Greedy Iterative Geodesic Ascent (GIGA) (Campbell & Broderick, 2018) and Sparse Variational Inference (SparseVI) (Campbell & Beronov, 2019). We note that the Frank-Wolfe (FW) method proposed in (Campbell & Broderick, 2019) has been shown to be inferior to GIGA and SparseVI in the two corresponding articles, and thus we believe that comparing with GIGA and SparseVI is sufficient. It can be observed that A-IHT and A-IHT II obtain better coresets than both GIGA and SparseVI most of the time, within time comparable to GIGA, while SparseVI pays a significant computational cost to achieve better coresets than GIGA, with construction time typically  $\times 10^4$  more than IHT in our experiments. We calculate the Kullback-Leibler (KL) divergence between the constructed coresets posterior  $\pi_w$  and the true posterior  $\pi$ . We measure both the forward KL divergence  $D_{\text{KL}}(\pi \parallel \pi_w)$  and reverse KL divergence  $D_{\text{KL}}(\pi_w \parallel \pi)$ .

Both A-IHT and A-IHT II require minimal tuning, *i.e.*, only the stopping criterion is required:  $\|w_t - w_{t-1}\| \leq 10^{-5} \|w_t\|$ , or number of iterations  $> 300$  for both A-IHT and A-IHT II.




Figure 1: (a): Bayesian coresets for synthetic Gaussian posterior inference. (b): Experiments on Bayesian radial basis function regression, with the difference between true posterior and coreset posterior measured in both forward KL and reverse KL. For both (a) and (b),  $k$  is the sparsity setting, and the solid lines are the median KL divergence between the constructed coreset posterior and true posterior over 10 trials. The shaded area is the KL divergence between 25<sup>th</sup> and 75<sup>th</sup> percentiles.

## 5.1 Synthetic Gaussian posterior inference

We examine the algorithms in this synthetic experiment, where we have closed-form exact expressions. Specifically, we compare each of these algorithms in terms of optimization accuracy without errors from sampling. For the  $D$ -dimensional Gaussian distribution, we set the parameter  $\theta \sim \mathcal{N}(\mu_0, \Sigma_0)$  and draw  $N$  i.i.d. samples  $x_n \sim \mathcal{N}(\theta, \Sigma)$ , which results in a Gaussian posterior distribution with closed-form parameters, as shown in (Campbell & Beronov, 2019). We set the dimension  $D = 200$ , number of samples  $N = 600$ , and maximal sparsity  $k$  is set to be  $1, \dots, 300$ . The initial covariance matrix is set to be  $\Sigma_0 = \Sigma = I$ . The learning rate for SparseVI is  $\gamma_t = 1/t$ , and the number of weight update iterations for Sparse VI is 100, as suggested by their paper.

Comparison among all the 5 algorithms measuring the forward KL divergence between the true posterior and the coreset posterior is presented in Figure 1 (a), which shows that IHT outperforms SparseVI and GIGA, achieving nearly optimal results. We observe that SparseVI stops improving once it hits certain sparsity level, which we suspect is due to the limitations of its greedy nature. It can also be observed that A-IHT II converges faster than A-IHT. Additional results are put in the section E in appendix.

## 5.2 Bayesian Radial Basis Function Regression

In this subsection, we explore the performance of proposed methods versus the baselines in terms of the both

forward KL and reverse KL divergence. The SparseVI algorithm optimizes reverse KL; we show this does not always imply reduction in the forward KL. Indeed selecting more points to greedily optimizing the reverse KL can cause an increase in the forward KL!

We aim to infer the posterior for Bayesian radial basis function regression. Given the dataset<sup>1</sup>  $\{(x_n, y_n) \in \mathbb{R}^2 \times \mathbb{R}\}_{n=1}^N$ , where  $x_n$  is the latitude/longitude coordinates and  $y_n$  is house-sale log-price in the United Kingdom, the goal is to infer coefficients  $\alpha \in \mathbb{R}^D$  for  $D$  radial basis functions  $b_d(x) = \exp(-\frac{1}{2\sigma_d^2}(x - \mu_d)^2)$  for  $d \in [D]$ . The model is  $y_n = b_n^\top \alpha + \epsilon_n$ , where  $\epsilon_n \sim \mathcal{N}(0, \sigma^2)$  with  $\sigma^2$  be the variance of  $\{y_n\}$ , and  $b_n = [b_1(x_n), \dots, b_D(x_n)]^\top$ . We set prior  $\alpha \sim \mathcal{N}(\mu_0, \sigma_0^2 I)$ , where  $\mu_0, \sigma_0^2$  are empirical mean and second moment of the data. We subsampled the dataset uniformly at random to  $N = 1000$  records for the experiments, and generated 50 basis functions for each of the 6 scales  $\sigma_d \in \{0.2, 0.4, 0.8, 1.2, 1.6, 2.0\}$  by generating means  $\mu_d$  for each basis uniformly from data. Except for the 300 basis functions, an additional near-constant basis of scale 100, with mean corresponding to the mean latitude and longitude of the data, is added. Therefore,  $D = 301$  basis functions are considered. Each of the algorithms has access to the closed-form of posterior distribution and covariance (see (Campbell & Beronov, 2019) for detailed derivation).

<sup>1</sup>The task is to predict housing prices from the UK land registry data (<https://www.gov.uk/government/statistical-data-sets/price-paid-data-downloads>) using latitude/longitude coordinates from the Geonames postal code data (<http://download.geonames.org/export/zip/>) as features.

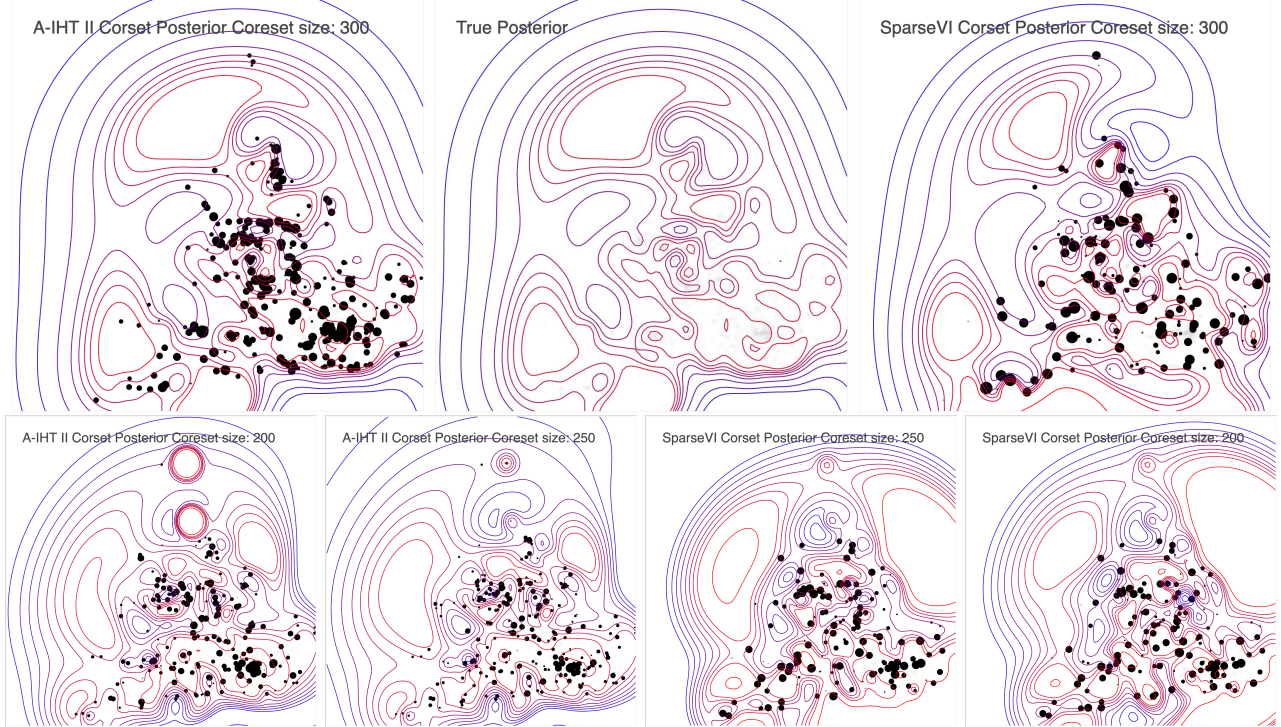


Figure 2: Experiments on Bayesian radial basis function regression, where coreset sparsity setting  $k = 200, 250, 300$ . Coreset points are presented as black dots, with their radius indicating assigned weights. When  $k = 300$ , posterior constructed by Accelerated IHT II (top left) shows almost exact contours as the true posterior (top middle), while posterior constructed by SparseVI (top right) shows deviated contours from the true posterior distribution.

Specific settings for the algorithms are as follows. For SparseVI, the exact covariance can be obtained, and the weight update step can be done without Monte Carlo estimation. For IHT and GIGA, we use true posterior for constructing the  $\ell_2$  loss function. The learning rate for SparseVI is set to be  $\gamma_t = 1/t$ , and iteration number  $T = 100$ , which is the setting SparseVI uses for the experiment (Campbell & Beronov, 2019).

IHT’s objective indicates both bounded forward KL and reverse KL. However, SparseVI, which optimizes the reverse KL, offers no guarantee for the forward KL. As shown in Figure 1 (b), SparseVI increasingly deviates from the true distribution in forward KL as more coreset points are selected. However, IHT methods offers consistently better coresets in both forward KL and reverse KL metric.

The reverse KL divergence alone is not enough to indicate good approximation, as shown in Figure 2. We plot the posterior contours for both the true posterior and coreset posterior at a random trial when sparsity level  $k = 200, 250, 300$ . The coreset posterior constructed by our Algorithm 3 recovers the true posterior almost exactly at  $k = 300$ , unlike SparseVI. The results for other trials are provided in section F in the appendix.

### 5.3 Bayesian logistic and Poisson regression

We consider how IHT performs when used in real applications where the closed-form expressions are unattainable. As the true posterior is unknown, a Laplace approximation is used for GIGA and IHT to derive the finite projection of the distribution, *i.e.*,  $\hat{g}_i$ . Further, Monte Carlo sampling is used to derive gradients of  $D_{\text{KL}}$  for SparseVI. We compare different algorithms estimating the posterior distribution for logistic regression and Poisson regression. The reverse KL and forward KL between the coreset posterior and true posterior are estimated using another Laplace approximation. The experiment was proposed by Campbell & Broderick (2019), and is used in (Campbell & Broderick, 2018) and (Campbell & Beronov, 2019). Due to space limitations, we refer to section G in the appendix for details of the experimental setup, and extensive additional results.

For logistic regression, given a dataset  $\{(x_n, y_n) \in \mathbb{R}^D \times \{1, -1\} \mid i \in [N]\}$ , we aim to infer  $\theta \in \mathbb{R}^{D+1}$  based on the model:

$$y_n \mid x_n, \theta \sim \text{Bern} \left( \frac{1}{1 + e^{-z_n^\top \theta}} \right),$$

where  $z_n = [x_n^\top, 1]^\top$ . We set  $N = 500$  by uniformly

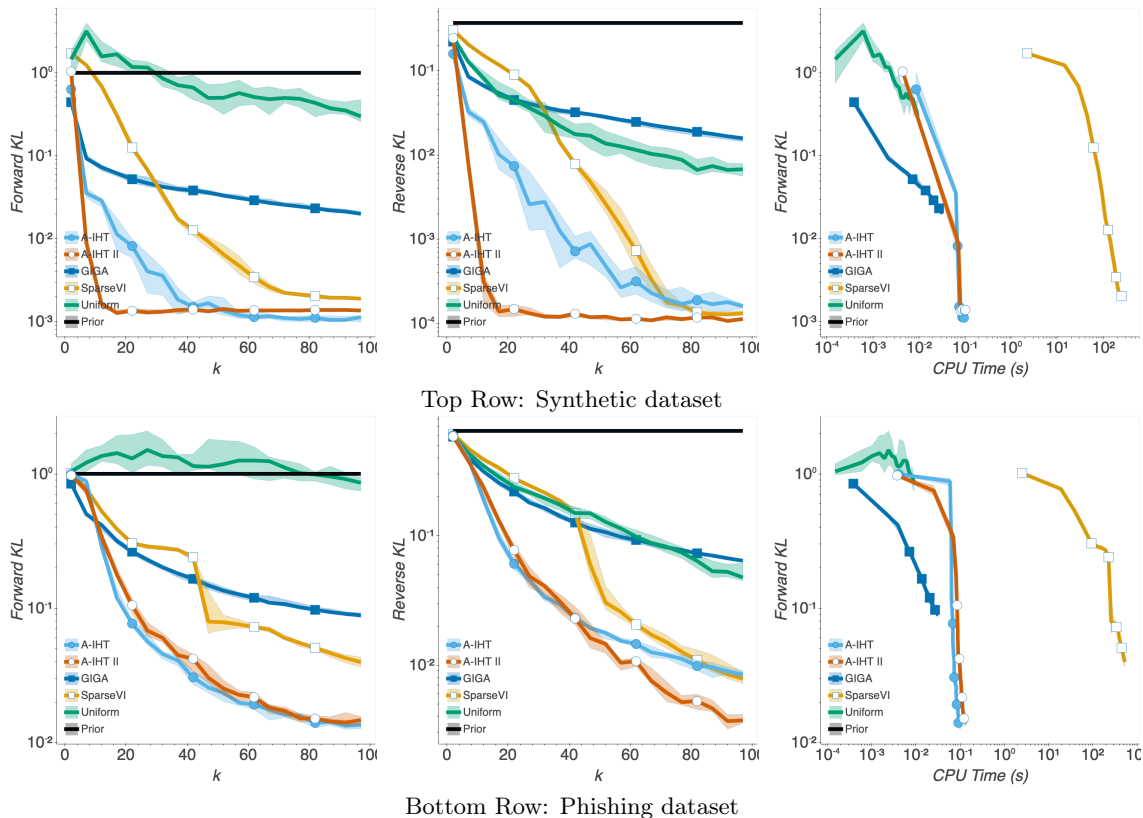


Figure 3: Bayesian coreset construction for logistic regression (LR) using the synthetic dataset (top row) and the phishing dataset (bottom row). All the algorithms are run 20 times, and the median as well as the interval of 35<sup>th</sup> and 65<sup>th</sup> percentile, indicated as the shaded area, are reported. Different maximal coreset size  $k$  is tested from 1 to 100. Forward KL (left column) and reverse KL (middle column) divergence between estimated true posterior and coreset posterior indicate the quality of the constructed coreset. The smaller the KL divergence, the better the coreset is. The running time for each algorithms is also recorded (right column).

sub-sampling from datasets due to the high computation cost of SparseVI. Three datasets are used for logistic regression. The synthetic dataset consists of  $x_n$  sampled i.i.d. from standard normal distribution  $\mathcal{N}(0, I)$ , and label  $y_n$  sampled from Bernoulli distribution conditioned on  $x_n$  and  $\theta = [3, 3, 0]^T$ . The phishing dataset<sup>2</sup> is preprocessed (Campbell & Beronov, 2019) via PCA to dimension of  $D = 10$  to mitigate high computation by SparseVI. The chemical reactivities dataset<sup>3</sup> has  $D = 10$ .

We present two sets of experiments, *i.e.*, logistic regression using the synthetic dataset and the phishing dataset, in Figure 3. One other set of experiments on logistic regression, and three sets of experiments on Poisson regression are deferred to section G in appendix.

It is observed that A-IHT and A-IHT II achieve state-of-

the-art performance. The IHT algorithms often obtain coresets with smaller KL between the coreset posterior and true posterior than GIGA and SparseVI, with computing time comparable to GIGA and significantly less than SparseVI. The experiments indicate that IHT outperforms the previous methods, improving the trade-off between accuracy and performance.

**Conclusion** In this paper, we consider the Bayesian coreset construction problem from a sparse optimization perspective, through which we propose a new algorithm that incorporates the paradigms of sparse as well as accelerated optimization. We provide theoretical analysis for our method, showing linear convergence under standard assumptions. Finally, numerical results demonstrate the improvement in both accuracy and efficiency when compared to the state of the art methods. Our viewpoint of using sparse optimization for Bayesian coresets can potentially help to consider more complex structured sparsity, which is left as future work.

<sup>2</sup><https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/binary.html>

<sup>3</sup><http://komarix.org/ac/ds>



## References

- Sungjin Ahn, Anoop Korattikara, and Max Welling. Bayesian posterior sampling via stochastic gradient fisher scoring. In *29th International Conference on Machine Learning, ICML 2012*, pp. 1591–1598, 2012.
- Sohail Bahmani, Bhiksha Raj, and Petros T. Boufounos. Greedy sparsity-constrained optimization. *J. Mach. Learn. Res.*, 14(1):807–841, March 2013. ISSN 1532-4435.
- Luca Baldassarre, Nirav Bhan, Volkan Cevher, Anastasios Kyrillidis, and Siddhartha Satpathi. Group-sparse model selection: Hardness and relaxations. *IEEE Transactions on Information Theory*, 62(11):6508–6534, 2016.
- Richard Baraniuk, Mark A. Davenport, Ronald A. DeVore, and Michael B. Wakin. A simple proof of the restricted isometry property for random matrices. *Constructive Approximation*, 28:253–263, 2008.
- Richard G Baraniuk, Volkan Cevher, Marco F Duarte, and Chinmay Hegde. Model-based compressive sensing. *IEEE Transactions on information theory*, 56(4):1982–2001, 2010.
- Amir Beck and Marc Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Imaging Sciences*, 2:183–202, 2009.
- T. Blumensath and M. Davies. Iterative hard thresholding for compressed sensing. *Applied and computational harmonic analysis*, 27(3):265–274, 2009.
- Thomas Blumensath. Accelerated iterative hard thresholding. *Signal Process.*, 92(3):752–756, March 2012. ISSN 0165-1684. doi: 10.1016/j.sigpro.2011.09.017. URL <https://doi.org/10.1016/j.sigpro.2011.09.017>.
- Thomas Blumensath and Mike E Davies. Normalized iterative hard thresholding: Guaranteed stability and performance. *IEEE Journal of selected topics in signal processing*, 4(2):298–309, 2010.
- Tamara Broderick, Nicholas Boyd, Andre Wibisono, Ashia C Wilson, and Michael I Jordan. Streaming variational bayes. In *Advances in Neural Information Processing Systems*, pp. 1727–1735, 2013.
- Trevor Campbell and Boyan Beronov. Sparse variational inference: Bayesian coresets from scratch. In *Advances in Neural Information Processing Systems*, pp. 11457–11468, 2019.
- Trevor Campbell and Tamara Broderick. Bayesian coreset construction via greedy iterative geodesic ascent. In *International Conference on Machine Learning*, pp. 697–705, 2018.
- Trevor Campbell and Tamara Broderick. Automated scalable bayesian inference via hilbert coresets. *The Journal of Machine Learning Research*, 20(1):551–588, 2019.
- Emmanuel J Candes. The restricted isometry property and its implications for compressed sensing. *Comptes rendus mathematique*, 346(9-10):589–592, 2008.
- Yuxin Chen and Sujay Sanghavi. A general framework for high-dimensional estimation in the presence of incoherence. In *2010 48th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pp. 1570–1576. IEEE, 2010.
- Wei Dai and Olghica Milenkovic. Subspace pursuit for compressive sensing signal reconstruction. *IEEE transactions on Information Theory*, 55(5):2230–2249, 2009.
- David L Donoho et al. Compressed sensing. *IEEE Transactions on information theory*, 52(4):1289–1306, 2006.
- Simon Foucart. Hard thresholding pursuit: an algorithm for compressive sensing. *SIAM Journal on Numerical Analysis*, 49(6):2543–2563, 2011.
- Marguerite Frank and Philip Wolfe. An algorithm for quadratic programming. *Naval research logistics quarterly*, 3(1-2):95–110, 1956.
- E. Ghadimi, H. R. Feyzmahdavian, and M. Johansson. Global convergence of the heavy-ball method for convex optimization. In *2015 European Control Conference (ECC)*, pp. 310–315, July 2015.
- Matthew D Hoffman, David M Blei, Chong Wang, and John Paisley. Stochastic variational inference. *The Journal of Machine Learning Research*, 14(1):1303–1347, 2013.
- Jonathan Huggins, Trevor Campbell, and Tamara Broderick. Coresets for scalable bayesian logistic regression. In *Advances in Neural Information Processing Systems*, pp. 4080–4088, 2016.
- Martin Jaggi. Revisiting frank-wolfe: Projection-free sparse convex optimization. In *Proceedings of the 30th international conference on machine learning*, pp. 427–435, 2013.
- Prateek Jain, Ambuj Tewari, and Purushottam Kar. On iterative hard thresholding methods for high-dimensional m-estimation. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 1, NIPS’14*, pp. 685–693, Cambridge, MA, USA, 2014. MIT Press.
- Prateek Jain, Nikhil Rao, and Inderjit S Dhillon. Structured sparse regression via greedy hard thresholding. In *Advances in Neural Information Processing Systems*, pp. 1516–1524, 2016.

- Rajiv Khanna and Anastasios Kyrillidis. Iht dies hard: Provable accelerated iterative hard thresholding. In *International Conference on Artificial Intelligence and Statistics*, pp. 188–198, 2018.
- Rajiv Khanna, Joydeep Ghosh, Russell Poldrack, and Oluwasanmi Koyejo. Information projection and approximate inference for structured sparse variables. In *Artificial Intelligence and Statistics*, pp. 1358–1366, 2017.
- Anoop Korattikara, Yutian Chen, and Max Welling. Austerity in mcmc land: Cutting the metropolis-hastings budget. In *International Conference on Machine Learning*, pp. 181–189, 2014.
- Oluwasanmi Koyejo, Rajiv Khanna, Joydeep Ghosh, and Russell A. Poldrack. On prior distributions and approximate inference for structured variables. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 1*, NIPS’14, pp. 676–684, Cambridge, MA, USA, 2014. MIT Press.
- Anastasios Kyrillidis and Volkan Cevher. Recipes on hard thresholding methods. In *2011 4th IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, pp. 353–356. IEEE, 2011.
- Anastasios Kyrillidis and Volkan Cevher. Matrix recipes for hard thresholding methods. *Journal of Mathematical Imaging and Vision*, 2(48):235–265, 2014.
- Anastasios Kyrillidis, Luca Baldassarre, Marwa El Halabi, Quoc Tran-Dinh, and Volkan Cevher. Structured sparsity: Discrete and convex approaches. In *Compressed Sensing and its Applications*, pp. 341–387. Springer, 2015.
- Xingguo Li, Tuo Zhao, Raman Arora, Han Liu, and Jarvis Haupt. Stochastic variance reduced optimization for nonconvex sparse learning. In Maria Florina Balcan and Kilian Q. Weinberger (eds.), *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 917–925, New York, New York, USA, 20–22 Jun 2016. PMLR.
- Francesco Locatello, Rajiv Khanna, Joydeep Ghosh, and Gunnar Rätsch. Boosting variational inference: an optimization perspective. In *AISTATS*, 2017.
- Francesco Locatello, Gideon Dresdner, Rajiv Khanna, Isabel Valera, and Gunnar Rätsch. Boosting black box variational inference. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NIPS’18, pp. 3405–3415, Red Hook, NY, USA, 2018. Curran Associates Inc.
- Yi-An Ma, Yuansi Chen, Chi Jin, Nicolas Flammarion, and Michael I. Jordan. Sampling can be faster than optimization. *Proceedings of the National Academy of Sciences*, 116(42):20881–20885, 2019. ISSN 0027-8424. doi: 10.1073/pnas.1820003116.
- Dougal Maclaurin and Ryan Prescott Adams. Firefly monte carlo: Exact mcmc with subsets of data. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015.
- Deanna Needell and Joel A Tropp. Cosamp: Iterative signal recovery from incomplete and inaccurate samples. *Applied and computational harmonic analysis*, 26(3):301–321, 2009.
- Yurii E Nesterov. A method for solving the convex programming problem with convergence rate  $o(1/k^2)$ . In *Dokl. akad. nauk Sssr*, volume 269, pp. 543–547, 1983.
- Nam Nguyen, Deanna Needell, and Tina Woolf. Linear convergence of stochastic iterative greedy algorithms with sparse constraints. *IEEE Transactions on Information Theory*, 63:6869–6895, 2014.
- Maxim Rabinovich, Elaine Angelino, and Michael I Jordan. Variational consensus monte carlo. In *Advances in Neural Information Processing Systems*, pp. 1207–1215, 2015.
- Steven L Scott, Alexander W Blocker, Fernando V Bonassi, Hugh A Chipman, Edward I George, and Robert E McCulloch. Bayes and big data: The consensus monte carlo algorithm. *International Journal of Management Science and Engineering Management*, 11(2):78–88, 2016.
- Shai Shalev-Shwartz, Nathan Srebro, and Tong Zhang. Trading accuracy for sparsity in optimization problems with sparsity constraints. *SIAM J. on Optimization*, 20(6):2807–2832, August 2010. ISSN 1052-6234. doi: 10.1137/090759574.
- Sanvesh Srivastava, Volkan Cevher, Quoc Dinh, and David Dunson. Wasp: Scalable bayes via barycenters of subset posteriors. In *Artificial Intelligence and Statistics*, pp. 912–920, 2015.
- Joel A Tropp and Anna C Gilbert. Signal recovery from random measurements via orthogonal matching pursuit. *IEEE Transactions on information theory*, 53(12):4655–4666, 2007.
- Max Welling and Yee W Teh. Bayesian learning via stochastic gradient langevin dynamics. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pp. 681–688, 2011.
- Xiao-Tong Yuan, Ping Li, and Tong Zhang. Gradient hard thresholding pursuit. *Journal of Machine Learning Research*, 18(166):1–43, 2018.
- Jacky Y Zhang, Rajiv Khanna, Anastasios Kyrillidis, and Oluwasanmi O Koyejo. Learning sparse distributions using iterative hard thresholding. In *Advances*

*in Neural Information Processing Systems 32*, pp.  
6757–6766. Curran Associates, Inc., 2019.

---

## Bayesian Coresets: Revisiting the Nonconvex Optimization Perspective Supplementary Materials

---

### A Automated Accelerated IHT with De-bias Step

---

**Algorithm 3** Automated Accelerated IHT - II
 

---

**input** Objective  $f(w) = \|y - \Phi w\|_2^2$ ; sparsity  $k$ 

- 1:  $t = 0, z_0 = 0, w_0 = 0$
  - 2: **repeat**
  - 3:  $\mathcal{Z} = \text{supp}(z_t)$
  - 4:  $\mathcal{S} = \text{supp}(\Pi_{\mathcal{C}_k \setminus \mathcal{Z}}(\nabla f(z_t))) \cup \mathcal{Z}$  where  $|\mathcal{S}| \leq 3k$  {active subspace expansion}
  - 5:  $\tilde{\nabla}^{(1)} = \nabla f(z_t)|_{\mathcal{S}}$
  - 6:  $\mu_t^{(1)} = \arg \min_{\mu} f(z_t - \mu \tilde{\nabla}^{(1)}) = \frac{\|\tilde{\nabla}^{(1)}\|_2^2}{2\|\Phi \tilde{\nabla}^{(1)}\|_2^2}$  {step size selection}
  - 7:  $x_t = \Pi_{\mathcal{C}_k \cap \mathbb{R}_+^n}(z_t - \mu_t^{(1)} \tilde{\nabla}^{(1)})$  {projected gradient descent}
  - 8:  $\tilde{\nabla}^{(2)} = \nabla f(x_t)|_{\text{supp}(x)}$
  - 9:  $\mu_t^{(2)} = \arg \min_{\mu} f(x_t - \mu \tilde{\nabla}^{(2)}) = \frac{\|\tilde{\nabla}^{(2)}\|_2^2}{2\|\Phi \tilde{\nabla}^{(2)}\|_2^2}$  {step size selection}
  - 10:  $w_{t+1} = \Pi_{\mathbb{R}_+^n}(x_t - \mu_t^{(2)} \tilde{\nabla}^{(2)})$  {de-bias step}
  - 11:  $\tau_{t+1} = \arg \min_{\tau} f(w_{t+1} + \tau(w_{t+1} - w_t)) = \frac{\langle y - \Phi w_{t+1}, \Phi(w_{t+1} - w_t) \rangle}{2\|\Phi(w_{t+1} - w_t)\|_2^2}$
  - 12:  $z_{t+1} = w_{t+1} + \tau_{t+1}(w_{t+1} - w_t)$  {momentum step}
  - 13:  $t = t + 1$
  - 14: **until** Stop criteria met
  - 15: **return**  $w_t$
- 

In the main text, we mention that Algorithm 2 can be boosted better in practice using de-bias steps. Here we present the algorithm with de-bias step, as shown in Algorithm 3.

Like Automated Accelerated IHT, Algorithm 3 also starts with active subspace expansion, i.e., line 3 & 4. As  $\mathcal{Z} = \text{supp}(z_t) = \text{supp}(w_{t-1}) \cup \text{supp}(w_t)$  is a  $2k$ -sparse index set, the expanded index set  $\mathcal{S}$  is a  $3k$ -sparse index set that is the union of the support of three elements, i.e.,

$$\mathcal{S} = \text{supp}(w_{t-1}) \cup \text{supp}(w_t) \cup \text{supp}(\Pi_{\mathcal{C}_k \setminus \mathcal{Z}}(\nabla f(z_t))).$$

We note that, with a little abuse of notation, we use  $\mathcal{Z}$  to denote both the support set  $\mathcal{Z} \subset [n]$ , and the subspace restricted by the support, i.e.,  $\{x \in \mathbb{R}^n \mid \text{supp}(x) \subseteq \mathcal{Z}\}$ .

The subspace corresponding to this index set  $\mathcal{S}$  is a subspace that the algorithm considers as potential to achieve low loss within. Therefore, in the next step we perform projected gradient descent in this expanded subspace. Note that we use  $\nabla f(\cdot)|_{\mathcal{S}}$  to denote a sparse subset  $\mathcal{S}$  of the gradient, i.e., setting the  $i^{\text{th}}$  entry of  $\nabla f(\cdot)$  to 0 if  $i \notin \mathcal{S}$ .

The projected gradient descent step consists of three sub-steps, i.e., step size selection (line 6), gradient descent (line 7), and projection to non-negative  $k$ -sparse restricted domain (line 7). The step size selection is performed by an exact line search to obtain good step size automatically. The projection step (line 7) is where we do ‘‘hard thresholding’’ to obtain a  $k$ -sparse solution  $x_t$ . As mentioned before, this projection step can be done optimally in the sense of  $\ell_2$ -norm by simply sorting and piking the  $k$ -largest non-negative elements.

Then, we come to the key difference between Algorithm 2 and Algorithm 3, i.e., the de-bias step at line 8, 9 & 10. With additional de-bias steps, we adjust the solution  $k$ -sparse solution  $x_t$  inside its own sparse space, i.e., the space corresponding to  $\text{supp}(x_t)$ , such that a better  $k$ -sparse solution is found. After computing the gradient (line 8), another exact line search is performed (line 9). By gradient descent and imposing the non-negativity constraint (line 10), we have the solution  $w_{t+1}$  for this iteration.

Lastly, the momentum step (line 11 & 12) is the same as Algorithm 2. We select the momentum term as the minimizer of the objective:  $\tau_{t+1} = \arg \min_{\tau} f(w_{t+1} + \tau(w_{t+1} - w_t))$ , and then apply the momentum to our solutions  $w_{t+1}$  and  $w_t$  as  $z_{t+1} = w_{t+1} + \tau_{t+1}(w_{t+1} - w_t)$  to capture memory in the algorithm. Momentum can offer faster convergence rate for convex optimization (Nesterov, 1983).

## B Theoretical Analysis

In this section, we provide a detailed theoretical analysis that is abstracted in the main paper due to space limitation. All of the proofs are defer to section C for clarity. To begin with, let us show that all of the projection operators used in our algorithms can be done optimally and efficiently.

Given an index set  $\mathcal{S} \subseteq [n]$ , the projection of  $w$  to the subspace with support  $\mathcal{S}$  is  $\Pi_{\mathcal{S}}(w)$ , which can be done optimally by setting  $w_{\mathcal{S}^c} = 0$ , where  $\mathcal{S}^c$  denotes the complement of  $\mathcal{S}$ . We note that, with a little abuse of notation, we use  $\mathcal{S}$  to denote both the support set  $\mathcal{S} \subset [n]$ , and the subspace restricted by the support, i.e.,  $\{x \in \mathbb{R}^n \mid \text{supp}(x) \subseteq \mathcal{S}\}$ . The projection to non-negative space, i.e.,  $\Pi_{\mathbb{R}_+^n}(w)$ , can also be done optimally and efficiently by setting the negative entries to zero. Moreover,  $\Pi_{\mathcal{C}_k}$  is shown to be optimal by simply picking the top  $k$  largest (in absolute value) entries. It is also the case for  $\Pi_{\mathcal{C}_k \cap \mathbb{R}_+^n}(w)$ , where it can be done by picking the top  $k$  largest non-negative entries. The optimality for the above projections is in terms of Euclidean distance.

Let us show the optimality for  $\Pi_{\mathcal{C}_k \cap \mathbb{R}_+^n}(w)$ . Given a  $k$ -sparse support  $\mathcal{S}$ , the optimal projection of  $w \in \mathbb{R}^n$  to its restricted sparsity space intersecting the non-negative orthant is  $w' = \Pi_{\mathcal{S} \cap \mathbb{R}_+^n}(w)$ . We can see that for entry  $i \in [n]$ ,  $w'_i = w_i$  if  $i \in \mathcal{S}$  and  $w_i \geq 0$ , and  $w'_i = 0$  otherwise. Therefore, the distance between  $w$  and its projection to  $\mathcal{S} \cap \mathbb{R}_+^n$  is  $\|w' - w\|_2^2 = \|w\|_2^2 - \sum_{i \in \mathcal{S}, w_i > 0} w_i^2$ . As  $\Pi_{\mathcal{C}_k \cap \mathbb{R}_+^n}(w) = \min_{\mathcal{S}: |\mathcal{S}| \leq k} \Pi_{\mathcal{S} \cap \mathbb{R}_+^n}(w)$ , we can see that it is the support with  $k$  largest  $w_i$  that has the least distance. Therefore, simply picking top  $k$  largest non-negative entries gives the optimal projection.

We give the convergence analysis for our main algorithm Automated Accelerated IHT in Algorithm 2. One standard assumption about the objective is required for the theory to begin, i.e., RIP property, which is a normal assumption in IHT context, reflecting convexity and smoothness of the objective in some sense (Khanna & Kyrillidis, 2018; Kyrillidis & Cevher, 2014). We note that the assumption is not necessary but is sufficient. For example, if the number of samples required to exactly construct  $\hat{g}$  is less than the coreset size ( $a_k = 0$  in RIP), so that the system becomes underdetermined, then local minima can be global one achieving zero-error without the RIP. On the other hand, when the number of samples goes to infinity, RIP ensures the eigenvalues of covariance matrix,  $\text{cov}[\mathcal{L}_i(\theta), \mathcal{L}_j(\theta)]$  where  $\theta \sim \hat{\pi}$ , are lower and upper bounded. It is an active area of research in random matrix theory to quantify RIP constants e.g. see (Baraniuk et al., 2008).

**Assumption 1** (Restricted Isometry Property). *Matrix  $\Phi$  in the objective function satisfies the RIP property, i.e., for  $\forall w \in \mathcal{C}_k$*

$$\alpha_k \|w\|_2^2 \leq \|\Phi w\|_2^2 \leq \beta_k \|w\|_2^2.$$

It is known that there are connections between RIP and restricted strong convexity and smoothness assumptions (Chen & Sanghavi, 2010); thus our results could potentially generalized for different convex  $f(\cdot)$  functions.

Leading to our main theorem, some useful technical properties are presented. An useful observation is that, for any set  $\mathcal{S} \subseteq [n]$ , the projection operator  $\Pi_{\mathcal{S}} : \mathbb{R}^n \rightarrow \mathbb{R}^n$  is in fact a linear operator in the form of a diagonal matrix

$$\Pi_{\mathcal{S}} = \{\text{diag}(\delta_i)\}_{i=1}^n,$$

where  $\delta_i$  is an indicator function:  $\delta_i = 1$  if  $i \in \mathcal{S}$ , and  $\delta_i = 0$  otherwise. This leads to our first lemma.

**Lemma 1.** *Supposing  $\Phi$  satisfies the RIP assumption, given a sparse set  $\mathcal{S} \subseteq [n]$  and  $|\mathcal{S}| \leq k$ , for  $\forall w \in \mathbb{R}^n$  it*

holds that

$$\alpha_k \|\Pi_{\mathcal{S}} w\|_2 \leq \|\Pi_{\mathcal{S}} \Phi^\top \Phi \Pi_{\mathcal{S}} w\|_2 \leq \beta_k \|\Pi_{\mathcal{S}} w\|_2.$$

Lemma 1 reveals a property of the eigenvalues of  $\Pi_{\mathcal{S}} \Phi^\top \Phi \Pi_{\mathcal{S}}$ , which leads to the following lemma that bounds an iterated projection using the RIP property.

**Lemma 2.** *Supposing  $\Phi$  satisfies the RIP assumption, given two sets  $\mathcal{S}_1, \mathcal{S}_2 \subseteq [n]$  and  $|\mathcal{S}_1 \cup \mathcal{S}_2| \leq k$ , for  $\forall w \in \mathbb{R}^n$  it holds that*

$$\|\Pi_{\mathcal{S}_1} \Phi^\top \Phi \Pi_{\mathcal{S}_1^c} \Pi_{\mathcal{S}_2} w\|_2 \leq \frac{\beta_k - \alpha_k}{2} \cdot \|\Pi_{\mathcal{S}_2} w\|_2.$$

Armed with above two lemmas, we are ready to prove convergence for Automated Accelerated IHT (Algorithm 2). A key observation is that solution  $w_{t+1}$  found by Algorithm 2 is derived by the following two steps:

$$\{w_t, w_{t-1}\} \xrightarrow[\text{line 9}]{\textcircled{1}} z_t \xrightarrow[\text{line 7}]{\textcircled{2}} w_{t+1}.$$

Procedure  $\textcircled{1}$  is a momentum step, with momentum size chosen automatically; procedure  $\textcircled{2}$  aims for exploration in an expanded subspace spanned by a  $3k$ -sparse subset  $\mathcal{S}$ , and projecting to  $k$ -sparse non-negative subspace.

We break down the proof into two parts. Denoting the optimal solution as

$$w^* = \arg \min_{w \in \mathcal{C}_k \cap \mathbb{R}_+^n} \|y - \Phi w\|_2^2,$$

we propose the following two lemmas for the two steps respectively.

**Lemma 3.** *For procedure  $\textcircled{1}$ , the following iterative invariant holds.*

$$\|z_t - w^*\|_2 \leq |1 + \tau_t| \cdot \|w_t - w^*\|_2 + |\tau_t| \cdot \|w_{t-1} - w^*\|_2.$$

For the second procedure, we consider the actual step size  $\mu_t$  automatically chosen by the algorithm. Noting that  $|\text{supp}(\tilde{\nabla}_t)| \leq 3k$ , according to RIP we can see that the step size  $\mu_t = \frac{\|\tilde{\nabla}_t\|_2^2}{2\|\Phi \tilde{\nabla}_t\|_2^2}$  is bounded as

$$\frac{1}{2\beta_{3k}} \leq \mu_t \leq \frac{1}{2\alpha_{3k}}.$$

Therefore, using the Lemma 1 and Lemma 2, one can prove the following lemma.

**Lemma 4.** *For procedure  $\textcircled{2}$ , the following iterative invariant holds.*

$$\|w_{t+1} - w^*\|_2 \leq \rho \|z_t - w^*\|_2 + 2\beta_{3k} \sqrt{\beta_{2k}} \|\epsilon\|_2,$$

where  $\rho = \left(2 \max\left\{\frac{\beta_{2k}}{\alpha_{3k}} - 1, 1 - \frac{\alpha_{2k}}{\beta_{3k}}\right\} + \frac{\beta_{4k} - \alpha_{4k}}{\alpha_{3k}}\right)$ , and  $\|\epsilon\|_2 = \|y - \Phi w^*\|_2$  is the optimal objective value.

Combining the above two lemmas leads to our main convergence analysis theorem.

**Theorem 1 (Restated).** *In the worst case scenario, with Assumption 1, the solutions path find by Automated Accelerated IHT (Algorithm 2) satisfy the following iterative invariant.*

$$\|w_{t+1} - w^*\|_2 \leq \rho |1 + \tau_t| \cdot \|w_t - w^*\|_2 + \rho |\tau_t| \cdot \|w_{t-1} - w^*\|_2 + 2\beta_{3k} \sqrt{\beta_{2k}} \|\epsilon\|_2,$$

where  $\rho = \left(2 \max\left\{\frac{\beta_{2k}}{\alpha_{3k}} - 1, 1 - \frac{\alpha_{2k}}{\beta_{3k}}\right\} + \frac{\beta_{4k} - \alpha_{4k}}{\alpha_{3k}}\right)$ , and  $\|\epsilon\|_2 = \|y - \Phi w^*\|_2$  is the optimal objective value.

The theorem provides an upper bound invariant among consecutive iterates of the algorithm. To have better sense of convergence rate, we assume the optimal solution achieves  $\|\epsilon\|_2 = 0$ . Theorem 1 then implies

$$\|w_{t+1} - w^*\|_2 \leq \rho(1 + |\tau_t|) \|w_t - w^*\|_2 + \rho |\tau_t| \cdot \|w_{t-1} - w^*\|_2.$$

Given the above homogeneous recurrence, we can solve for the following corollary that shows linear convergence of the proposed algorithm under given conditions.

**Corollary 1** (Restated). *Given the iterative invariant as stated in Theorem 1, and assuming the optimal solution achieves  $\|\epsilon\|_2 = 0$ , the solution found by Algorithm 2 satisfies:*

$$f(w_{t+1}) - f(w^*) \leq \phi^t \left( \frac{\beta_{2k}}{\alpha_{2k}} f(w_1) + \frac{\rho\tau\beta_{2k}}{\phi\alpha_k} f(w_0) \right),$$

where  $\phi = (\rho(1 + \tau) + \sqrt{\rho^2(1 + \tau)^2 + 4\rho\tau})/2$  and  $\tau = \max_{i \in [t]} |\tau_i|$ . It is sufficient to show linear convergence to the global optimum, when  $\phi < 1$ , or equivalently  $\rho < 1/(1 + 2\tau)$ .

## C Proofs

This section provides proofs for the theoretical results the presented in the previous section. For the sake of good readability, the lemma/theorem to be proven is also restated preceding its proof.

### C.1 Proof of Lemma 1

**Lemma 1** (Restated). *Supposing  $\Phi$  satisfies the RIP assumption, given a sparse set  $\mathcal{S} \subseteq [n]$  and  $|\mathcal{S}| \leq k$ , for  $\forall w \in \mathbb{R}^n$  it holds that*

$$\alpha_k \|\Pi_{\mathcal{S}} w\|_2 \leq \|\Pi_{\mathcal{S}} \Phi^\top \Phi \Pi_{\mathcal{S}} w\|_2 \leq \beta_k \|\Pi_{\mathcal{S}} w\|_2.$$

*Proof.* Recall that  $\Pi_{\mathcal{S}}$  is a linear operator that projects a vector  $w \in \mathbb{R}^n$  to sparse restricted set with support  $\mathcal{S}$  by simply setting  $w_i = 0$  for each  $i \notin \mathcal{S}$ . As a result, for a  $k$ -sparse set  $\mathcal{S}$ ,  $\Pi_{\mathcal{S}} w$  is a  $k$ -sparse vector. Given that  $\Phi \in \mathbb{R}^{m \times n}$  satisfies RIP property, for  $\forall w \in \mathbb{R}^n$ , it holds that

$$\alpha_k \|\Pi_{\mathcal{S}} w\|_2^2 \leq \|\Phi \Pi_{\mathcal{S}} w\|_2^2 \leq \beta_k \|\Pi_{\mathcal{S}} w\|_2^2. \quad (4)$$

Let us denote  $b = \Phi \Pi_{\mathcal{S}} w$ , and  $\langle \cdot, \cdot \rangle$  as standard Euclidean inner product. With regular linear algebra manipulation, the following stands:

$$\begin{aligned} \|\Pi_{\mathcal{S}} \Phi^\top b\|_2^2 &= \max_{x \in \mathbb{R}^n: \|x\|_2=1} (\langle \Pi_{\mathcal{S}} \Phi^\top b, x \rangle)^2 \\ &= \max_{x \in \mathbb{R}^n: \|x\|_2=1} (b^\top \Phi \Pi_{\mathcal{S}} x)^2 \\ &= \max_{x \in \mathbb{R}^n: \|x\|_2=1} (\langle b, \Phi \Pi_{\mathcal{S}} x \rangle)^2 \\ &= \max_{x \in \mathbb{R}^n: \|x\|_2=1} (\langle \Phi \Pi_{\mathcal{S}} w, \Phi \Pi_{\mathcal{S}} x \rangle)^2, \end{aligned} \quad (5)$$

where the second equality is due to the fact that  $\Pi_{\mathcal{S}}$  is symmetric, i.e.,  $(\Pi_{\mathcal{S}} \Phi^\top b)^\top = b^\top \Phi \Pi_{\mathcal{S}}$ .

Letting  $x^*$  be the solution of (5), we have the upper bound of (5):

$$(5) = (\langle \Phi \Pi_{\mathcal{S}} w, \Phi \Pi_{\mathcal{S}} x^* \rangle)^2 \leq \|\Phi \Pi_{\mathcal{S}} w\|_2^2 \cdot \|\Phi \Pi_{\mathcal{S}} x^*\|_2^2,$$

where the inequality is by Cauchy-Schwarz inequality applying on inner product.

On the other hand, the lower bound can be obtained by removing the maximizing operator and setting  $x = \Pi_{\mathcal{S}} w / \|\Pi_{\mathcal{S}} w\|_2$ , as follows. Denoting  $x' = \Pi_{\mathcal{S}} w / \|\Pi_{\mathcal{S}} w\|_2$ , we have,

$$(5) \geq (\langle \Phi \Pi_{\mathcal{S}} w, \Phi \Pi_{\mathcal{S}} x' \rangle)^2 = \|\Phi \Pi_{\mathcal{S}} w\|_2^2 \cdot \|\Phi \Pi_{\mathcal{S}} x'\|_2^2,$$

where the last equality is due to that  $\Pi_{\mathcal{S}} w$  and  $x'$  are parallel.

Applying (4) to the above upper bound and lower bound, it follows that

$$\begin{aligned} (5) &\leq \|\Phi \Pi_{\mathcal{S}} w\|_2^2 \cdot \|\Phi \Pi_{\mathcal{S}} x^*\|_2^2 \leq \beta_k \|\Pi_{\mathcal{S}} w\|_2^2 \cdot \beta_k \|\Pi_{\mathcal{S}} x^*\|_2^2, \\ (5) &\geq \|\Phi \Pi_{\mathcal{S}} w\|_2^2 \cdot \|\Phi \Pi_{\mathcal{S}} x'\|_2^2 \geq \alpha_k \|\Pi_{\mathcal{S}} w\|_2^2 \cdot \alpha_k \|\Pi_{\mathcal{S}} x'\|_2^2. \end{aligned} \quad (6)$$

Noting that  $x^*$  is a unit-length vector, and the projection  $\Pi_{\mathcal{S}}$  is done by setting elements to zero, we can see that  $\|\Pi x^*\|_2 \leq 1$ . As  $x' = \Pi_{\mathcal{S}} w / \|\Pi_{\mathcal{S}} w\|_2$  has already been a sparse vector in the restricted space by  $\mathcal{S}$ , we can see that  $\|\Pi_{\mathcal{S}} x'\|_2 = \|x'\|_2 = 1$ . Plugging them in (6), it holds that

$$\alpha_k^2 \|\Pi_{\mathcal{S}} w\|_2^2 = \alpha_k \|\Pi_{\mathcal{S}} w\|_2^2 \cdot \alpha_k \|\Pi_{\mathcal{S}} x'\|_2^2 \leq (5) \leq \beta_k \|\Pi_{\mathcal{S}} w\|_2^2 \cdot \beta_k \|\Pi_{\mathcal{S}} x^*\|_2^2 \leq \beta_k^2 \|\Pi_{\mathcal{S}} w\|_2^2.$$

Plugging that (5) =  $\|\Pi_{\mathcal{S}} \Phi^\top b\|_2^2 = \|\Pi_{\mathcal{S}} \Phi^\top \Phi \Pi_{\mathcal{S}} w\|_2^2$ , and taking the square root, we finally have

$$\alpha_k \|\Pi_{\mathcal{S}} w\|_2 \leq \|\Pi_{\mathcal{S}} \Phi^\top \Phi \Pi_{\mathcal{S}} w\|_2 \leq \beta_k \|\Pi_{\mathcal{S}} w\|_2.$$

□

## C.2 Proof of Lemma 2

**Lemma 2 (Restated).** *Supposing  $\Phi$  satisfies the RIP assumption, given two sets  $\mathcal{S}_1, \mathcal{S}_2 \subseteq [n]$  and  $|\mathcal{S}_1 \cup \mathcal{S}_2| \leq k$ , for  $\forall w \in \mathbb{R}^n$  it holds that*

$$\|\Pi_{\mathcal{S}_1} \Phi^\top \Phi \Pi_{\mathcal{S}_1^c} \Pi_{\mathcal{S}_2} w\|_2 \leq \frac{\beta_k - \alpha_k}{2} \cdot \|\Pi_{\mathcal{S}_2} w\|_2.$$

*Proof.* Similar to the proof of Lemma 1, we first write the norm in the form of an inner product. Given two sets  $\mathcal{S}_1, \mathcal{S}_2 \subseteq [n]$  and  $|\mathcal{S}_1 \cup \mathcal{S}_2| \leq k$ , for  $\forall w \in \mathbb{R}^n$ , with regular linear algebra manipulation, it holds that

$$\begin{aligned} & \|\Pi_{\mathcal{S}_1} \Phi^\top \Phi \Pi_{\mathcal{S}_1^c} \Pi_{\mathcal{S}_2} w\|_2 \\ &= \max_{b \in \mathbb{R}^n: \|b\|_2=1} |\langle b, \Pi_{\mathcal{S}_1} \Phi^\top \Phi \Pi_{\mathcal{S}_1^c} \Pi_{\mathcal{S}_2} w \rangle| \\ &= \max_{b \in \mathbb{R}^n: \|b\|_2=1} |\langle \Phi \Pi_{\mathcal{S}_1} b, \Phi \Pi_{\mathcal{S}_1^c} \Pi_{\mathcal{S}_2} w \rangle|, \end{aligned} \quad (7)$$

where the second equality is due to the fact that  $\Pi_{\mathcal{S}_1}$  is symmetric.

Define two unit-length vectors

$$X = \frac{\Pi_{\mathcal{S}_1^c} \Pi_{\mathcal{S}_2} w}{\|\Pi_{\mathcal{S}_1^c} \Pi_{\mathcal{S}_2} w\|_2}, \quad Y = \frac{\Pi_{\mathcal{S}_1} b}{\|\Pi_{\mathcal{S}_1} b\|_2},$$

and we can see that  $\langle X, Y \rangle = 0$ , as  $\mathcal{S}_1^c$  and  $\mathcal{S}_1$  are disjoint. As a result,  $\|X + Y\|_2^2 = \|X\|_2^2 + \|Y\|_2^2 = 2$ . Moreover, given that  $|\mathcal{S}_1 \cup \mathcal{S}_2| \leq k$ , we can see that  $X + Y$  is  $k$ -sparse. Applying the RIP property, the following holds:

$$2\alpha_k = \alpha_k \|X + Y\|_2^2 \leq \|\Phi X + \Phi Y\|_2^2 \leq \beta_k \|X + Y\|_2^2 = 2\beta_k.$$

Similarly,  $\|X - Y\|_2^2 = 2$  and  $X - Y$  is also  $k$ -sparse:

$$2\alpha_k \leq \|\Phi X - \Phi Y\|_2^2 \leq 2\beta_k.$$

Noting that

$$\langle \Phi X, \Phi Y \rangle = \frac{\|\Phi X + \Phi Y\|_2^2 - \|\Phi X - \Phi Y\|_2^2}{4},$$

we can see the following,

$$-\frac{\beta_k - \alpha_k}{2} \leq \langle \Phi X, \Phi Y \rangle \leq \frac{\beta_k - \alpha_k}{2}. \quad (8)$$

Recall that

$$(7) = \max_{\|b\|_2=1} |\langle \Phi X, \Phi Y \rangle| \cdot \|\Pi_{\mathcal{S}_1} b\|_2 \cdot \|\Pi_{\mathcal{S}_1^c} \Pi_{\mathcal{S}_2} w\|_2,$$

and apply (8) to the above, we conclude that

$$\begin{aligned} (7) &\leq \max_{\|b\|_2=1} \frac{\beta_k - \alpha_k}{2} \cdot \|\Pi_{\mathcal{S}_1} b\|_2 \cdot \|\Pi_{\mathcal{S}_1^c} \Pi_{\mathcal{S}_2} w\|_2 \\ &\leq \frac{\beta_k - \alpha_k}{2} \|\Pi_{\mathcal{S}_2} w\|_2. \end{aligned}$$

□



### C.3 Proof of Lemma 3

**Lemma 3 (Restated).** For procedure ①, the following iterative invariant holds.

$$\|z_t - w^*\|_2 \leq |1 + \tau_t| \cdot \|w_t - w^*\|_2 + |\tau_t| \cdot \|w_{t-1} - w^*\|_2.$$

*Proof.* According to line 9 in Algorithm 2, with some regular linear algebra manipulation, we can derive

$$\begin{aligned} \|z_t - w^*\|_2 &= \|w_t + \tau_t(w_t - w_{t-1}) - w^*\|_2 \\ &= \|(1 + \tau_t)(w_t - w^*) + \tau_t(w^* - w_{t-1})\|_2 \\ &\leq |1 + \tau_t| \|w_t - w^*\|_2 + |\tau_t| \|w_{t-1} - w^*\|_2, \end{aligned}$$

where the last inequality is done by triangle inequality.  $\square$

### C.4 Proof of Lemma 4

**Lemma 4 (Restated).** For procedure ②, the following iterative invariant holds.

$$\|w_{t+1} - w^*\|_2 \leq \rho \|z_t - w^*\|_2 + 2\beta_{3k} \sqrt{\beta_{2k}} \|\epsilon\|_2,$$

where  $\rho = \left(2 \max\left\{\frac{\beta_{2k}}{\alpha_{3k}} - 1, 1 - \frac{\alpha_{2k}}{\beta_{3k}}\right\} + \frac{\beta_{4k} - \alpha_{4k}}{\alpha_{3k}}\right)$ , and  $\|\epsilon\|_2 = \|y - \Phi w^*\|_2$  is the optimal objective value.

*Proof.* Denoting  $v = z_t - \mu_t \nabla f(z_t)$ , and set  $\mathcal{S}_* = \text{supp}(w_{t+1}) \cup \text{supp}(w^*)$ , we begin by the projection at line 7 in Algorithm 2. Applying the triangle inequality,

$$\|w_{t+1} - w^*\|_2 \leq \|w_{t+1} - \Pi_{\mathcal{S}_*} v\|_2 + \|\Pi_{\mathcal{S}_*} v - w^*\|_2. \quad (9)$$

As  $\mathcal{S}_* = \text{supp}(w_{t+1}) \cup \text{supp}(w^*)$ , we can observe that  $\langle w_{t+1}, \Pi_{\mathcal{S}_*} v \rangle = 0$  and  $\langle w^*, \Pi_{\mathcal{S}_*} v \rangle = 0$ . As a result,

$$\begin{aligned} \|w_{t+1} - \Pi_{\mathcal{S}_*} v\|_2^2 &= \|w_{t+1} - v + \Pi_{\mathcal{S}_*} v\|_2^2 \\ &= \|w_{t+1} - v\|_2^2 + \|\Pi_{\mathcal{S}_*} v\|_2^2 + 2\langle w_{t+1} - v, \Pi_{\mathcal{S}_*} v \rangle \\ &= \|w_{t+1} - v\|_2^2 + \|\Pi_{\mathcal{S}_*} v\|_2^2 + 2\langle -v, \Pi_{\mathcal{S}_*} v \rangle \\ &\leq \|w^* - v\|_2^2 + \|\Pi_{\mathcal{S}_*} v\|_2^2 + 2\langle -v, \Pi_{\mathcal{S}_*} v \rangle \\ &= \|w^* - v\|_2^2 + \|\Pi_{\mathcal{S}_*} v\|_2^2 + 2\langle w^* - v, \Pi_{\mathcal{S}_*} v \rangle \\ &= \|w^* - v + \Pi_{\mathcal{S}_*} v\|_2^2 \\ &= \|w^* - \Pi_{\mathcal{S}_*} v\|_2^2, \end{aligned}$$

where the inequality is due to the projection step  $w_{t+1} = \Pi_{\mathcal{C}_k \cap \mathbb{R}_+^n} v$  is done optimally, and  $w^* \in \mathcal{C}_k \cap \mathbb{R}_+^n$ . Plugging the above inequality into (9), it holds that

$$\|w_{t+1} - w^*\|_2 \leq 2\|\Pi_{\mathcal{S}_*} v - w^*\|_2. \quad (10)$$

Expanding  $v$  and denoting  $\epsilon = \Phi w^* - y$ , we have

$$\begin{aligned} v &= z_t - \mu_t (\nabla f(z_t)) \\ &= z_t - \mu_t (2\Phi^\top (\Phi z_t - y)) \\ &= z_t - \mu_t (2\Phi^\top \Phi (z_t - w^*) + 2\Phi^\top (\Phi w^* - y)) \\ &= z_t - 2\mu_t \Phi^\top \Phi (z_t - w^*) - 2\mu_t \Phi^\top \epsilon. \end{aligned}$$

Plugging the above into inequality (10), we can further expand

$$\begin{aligned} \|w_{t+1} - w^*\|_2 &\leq 2\|\Pi_{\mathcal{S}_*} (z_t - 2\mu_t \Phi^\top \Phi (z_t - w^*) - 2\mu_t \Phi^\top \epsilon) - w^*\|_2 \\ &= 2\|\Pi_{\mathcal{S}_*} (z_t - w^*) - 2\mu_t \Pi_{\mathcal{S}_*} \Phi^\top \Phi (z_t - w^*) - 2\mu_t \Pi_{\mathcal{S}_*} \Phi^\top \epsilon\|_2 \\ &\leq 2\|\Pi_{\mathcal{S}_*} (z_t - w^*) - 2\mu_t \Pi_{\mathcal{S}_*} \Phi^\top \Phi (z_t - w^*)\|_2 + 4\mu_t \|\Pi_{\mathcal{S}_*} \Phi^\top \epsilon\|_2 \\ &= 2\|\Pi_{\mathcal{S}_*} (z_t - w^*) - 2\mu_t \Pi_{\mathcal{S}_*} \Phi^\top \Phi I (z_t - w^*)\|_2 + 4\mu_t \|\Pi_{\mathcal{S}_*} \Phi^\top \epsilon\|_2. \end{aligned} \quad (11)$$

Expanding the identity matrix by  $I = \Pi_{\mathcal{S}_*} + \Pi_{\mathcal{S}_*^c}$ , we have

$$(11) \leq \underbrace{2\|(I - 2\mu_t \Pi_{\mathcal{S}_*} \Phi^\top \Phi \Pi_{\mathcal{S}_*}) \Pi_{\mathcal{S}_*} (z_t - w^*)\|_2}_A + \underbrace{4\mu_t \|\Pi_{\mathcal{S}_*} \Phi^\top \Phi \Pi_{\mathcal{S}_*^c} (z_t - w^*)\|_2}_B + \underbrace{4\mu_t \|\Pi_{\mathcal{S}_*} \Phi^\top \epsilon\|_2}_C.$$

Now we bound the three terms respectively.

Noting that  $|\mathcal{S}_*| \leq 2k$ , according to Lemma 1, in the subspace with support  $\mathcal{S}_*$ , i.e.,  $\{w \mid \text{supp}(w) = \mathcal{S}_*\}$ , the eigenvalues  $\alpha_{2k} \leq \lambda_{\mathcal{S}_*}(\Pi_{\mathcal{S}_*} \Phi^\top \Phi \Pi_{\mathcal{S}_*}) \leq \beta_{2k}$ . Therefore, eigenvalues

$$\lambda_{\mathcal{S}_*}(I - 2\mu_t \Pi_{\mathcal{S}_*} \Phi^\top \Phi \Pi_{\mathcal{S}_*}) \in [1 - 2\mu_t \beta_{2k}, 1 - 2\mu_t \alpha_{2k}],$$

which means

$$\begin{aligned} A &\leq 2 \max\{2\mu_t \beta_{2k} - 1, 1 - 2\mu_t \alpha_{2k}\} \|\Pi_{\mathcal{S}_*} (z_t - w^*)\|_2 \\ &\leq 2 \max\{\beta_{2k}/\alpha_{3k} - 1, 1 - \alpha_{2k}/\beta_{3k}\} \|z_t - w^*\|_2. \end{aligned}$$

For term B, demoting  $\mathcal{S}' = \text{supp}(z_t) \cup \text{supp}(w^*)$ , it can be observed that

$$B = 4\mu_t \|\Pi_{\mathcal{S}_*} \Phi^\top \Phi \Pi_{\mathcal{S}_*^c} \Pi_{\mathcal{S}'} (z_t - w^*)\|_2.$$

Noting that  $|\mathcal{S}' \cup \mathcal{S}_*| \leq 4k$ , by directly applying Lemma 2 we have

$$\begin{aligned} B &\leq 4\mu_t \frac{\beta_{4k} - \alpha_{4k}}{2} \|\Pi_{\mathcal{S}'} (z_t - w^*)\|_2 \\ &\leq \frac{\beta_{4k} - \alpha_{4k}}{\alpha_{3k}} \|z_t - w^*\|_2. \end{aligned}$$

To complete the proof, let us deal with the last piece. Similar to the techniques used in the proof on Lemma 1,

$$\begin{aligned} \|\Pi_{\mathcal{S}_*} \Phi^\top \epsilon\|_2 &= \max_{x \in \mathbb{R}^n: \|x\|_2=1} \langle \Pi_{\mathcal{S}_*} \Phi^\top \epsilon, x \rangle \\ &= \max_{x \in \mathbb{R}^n: \|x\|_2=1} \epsilon^\top \Phi \Pi_{\mathcal{S}_*} x \\ &= \max_{x \in \mathbb{R}^n: \|x\|_2=1} \langle \epsilon, \Phi \Pi_{\mathcal{S}_*} x \rangle \\ &\leq \max_{x \in \mathbb{R}^n: \|x\|_2=1} \|\epsilon\|_2 \cdot \|\Phi \Pi_{\mathcal{S}_*} x\|_2 \\ &\leq \sqrt{\beta_{2k}} \|\epsilon\|_2, \end{aligned}$$

where the last inequality is done by directly applying the definition of RIP. Therefore,

$$C \leq 4\mu_t \sqrt{\beta_{2k}} \|\epsilon\|_2 \leq 2\beta_{3k} \sqrt{\beta_{2k}} \|\epsilon\|_2.$$

Combining the 3 pieces together, we finally derive

$$\begin{aligned} \|w_{t+1} - w^*\|_2 &\leq 2 \max\left\{\frac{\beta_{2k}}{\alpha_{3k}} - 1, 1 - \frac{\alpha_{2k}}{\beta_{3k}}\right\} \|z_t - w^*\|_2 \\ &\quad + \frac{\beta_{4k} - \alpha_{4k}}{\alpha_{3k}} \|z_t - w^*\|_2 + 2\beta_{3k} \sqrt{\beta_{2k}} \|\epsilon\|_2. \end{aligned}$$

Rearranging the inequality completes the proof.  $\square$

### C.5 Proof of Theorem 1

**Theorem 1 (Restated).** *In the worst case scenario, with Assumption 1, the solutions path find by Automated Accelerated IHT (Algorithm 2) satisfy the following iterative invariant.*

$$\|w_{t+1} - w^*\|_2 \leq \rho|1 + \tau_t| \cdot \|w_t - w^*\|_2 + \rho|\tau_t| \cdot \|w_{t-1} - w^*\|_2 + 2\beta_{3k}\sqrt{\beta_{2k}}\|\epsilon\|_2,$$

where  $\rho = \left(2 \max\left\{\frac{\beta_{2k}}{\alpha_{3k}} - 1, 1 - \frac{\alpha_{2k}}{\beta_{3k}}\right\} + \frac{\beta_{4k} - \alpha_{4k}}{\alpha_{3k}}\right)$ , and  $\|\epsilon\|_2 = \|y - \Phi w^*\|_2$  is the optimal objective value.

*Proof.* Lemma 3 suggests

$$\|z_t - w^*\|_2 \leq |1 + \tau_t| \|w_t - w^*\|_2 + |\tau_t| \|w_{t-1} - w^*\|_2.$$

Combining with lemma 4, *i.e.*,

$$\|w_{t+1} - w^*\|_2 \leq \rho \|z_t - w^*\|_2 + 2\beta_{3k}\sqrt{\beta_{2k}}\|\epsilon\|_2,$$

where  $\rho = \left(2 \max\left\{\frac{\beta_{2k}}{\alpha_{3k}} - 1, 1 - \frac{\alpha_{2k}}{\beta_{3k}}\right\} + \frac{\beta_{4k} - \alpha_{4k}}{\alpha_{3k}}\right)$ , and  $\|\epsilon\|_2 = \|y - \Phi w^*\|_2$ , we have

$$\|x_t - w^*\|_2 \leq \rho|1 + \tau_t| \|w_t - w^*\|_2 + \rho|\tau_t| \|w_{t-1} - w^*\|_2 + 2\beta_{3k}\sqrt{\beta_{2k}}\|\epsilon\|_2,$$

which completes the proof.  $\square$

### C.6 Proof of Corollary 1

**Corollary 1 (Restated).** *Given the iterative invariant as stated in Theorem 1, and assuming the optimal solution achieves  $\|\epsilon\|_2 = 0$ , the solution found by Algorithm 2 satisfies:*

$$f(w_{t+1}) - f(w^*) \leq \phi^t \left( \frac{\beta_{2k}}{\alpha_{2k}} f(w_1) + \frac{\rho\tau\beta_{2k}}{\phi\alpha_k} f(w_0) \right),$$

where  $\phi = (\rho(1 + \tau) + \sqrt{\rho^2(1 + \tau)^2 + 4\rho\tau})/2$  and  $\tau = \max_{i \in [t]} |\tau_i|$ . It is sufficient to show linear convergence to the global optimum, when  $\phi < 1$ , or equivalently  $\rho < 1/(1 + 2\tau)$ .

*Proof.* Theorem 1 provides an upper bound invariant among consecutive iterates of the algorithm. To have better sense of convergence rate, we assume the optimal solution achieves  $\|\epsilon\|_2 = 0$ . Theorem 1 then implies

$$\begin{aligned} \|w_{t+1} - w^*\|_2 &\leq \rho(1 + |\tau_t|) \|w_t - w^*\|_2 + \rho|\tau_t| \cdot \|w_{t-1} - w^*\|_2 \\ &\leq \rho(1 + \tau) \|w_t - w^*\|_2 + \rho\tau \cdot \|w_{t-1} - w^*\|_2. \end{aligned}$$

Rearranging the inequality with some regular algebraic manipulations, we have

$$\begin{aligned} \|w_{t+1} - w^*\|_2 + \frac{\rho\tau}{\phi} \|w_t - w^*\|_2 &\leq \phi \left( \|w_t - w^*\|_2 + \frac{\rho\tau}{\phi} \|w_{t-1} - w^*\|_2 \right) \\ &\leq \phi^t \left( \|w_1 - w^*\|_2 + \frac{\rho\tau}{\phi} \|w_0 - w^*\|_2 \right), \end{aligned}$$

where  $\phi = \frac{\sqrt{\rho^2(1+\tau)^2 + 4\rho\tau} + \rho(1+\tau)}{2}$ .

Noting that all  $\rho, \tau, \phi$  are non-negative, we can relax the inequality a bit to be

$$\|w_{t+1} - w^*\|_2 \leq \phi^t \left( \|w_1 - w^*\|_2 + \frac{\rho\tau}{\phi} \|w_0 - w^*\|_2 \right). \quad (12)$$

It is sufficient for linear convergence when  $\phi < 1$ , i.e.,

$$\begin{aligned}
 & \frac{\sqrt{\rho^2(1+\tau)^2 + 4\rho\tau} + \rho(1+\tau)}{2} < 1 \\
 \iff & \sqrt{\rho^2(1+\tau)^2 + 4\rho\tau} < 2 - \rho(1+\tau) \\
 \iff & \begin{cases} \rho^2(1+\tau)^2 + 4\rho\tau < (2 - \rho(1+\tau))^2 \\ 0 < 2 - \rho(1+\tau) \end{cases} \\
 \iff & \begin{cases} \rho(1+2\tau) < 1 \\ \rho(1+\tau) < 2 \end{cases} \\
 \iff & \rho < 1/(1+2\tau)
 \end{aligned}$$

In our case, this also indicates the linear convergence of function values. Noting that  $(w_{t+1} - w^*)$  and  $(w_1 - w^*)$  are at most  $2k$ -sparse, and  $(w_0 - w^*) = -w^*$  is  $k$ -sparse, we have the following statements according to RIP property:

$$\begin{aligned}
 \|\Phi(w_{t+1} - w^*)\|_2^2 &\leq \beta_{2k} \|w_{t+1} - w^*\|_2^2 \\
 \|\Phi(w_1 - w^*)\|_2^2 &\geq \alpha_{2k} \|w_1 - w^*\|_2^2 \\
 \|\Phi(w_0 - w^*)\|_2^2 &\geq \alpha_k \|w_0 - w^*\|_2^2
 \end{aligned}$$

As we assume  $\|\epsilon\|_2 = \|y - \Phi w^*\|_2 = 0$ , i.e.,  $y = \Phi w^*$  and  $f(w^*)$ , we can see that

$$\begin{aligned}
 f(w_{t+1}) &= \|\Phi w_{t+1} - y\|_2^2 \leq \beta_{2k} \|w_{t+1} - w^*\|_2^2 \\
 f(w_1) &= \|\Phi w_1 - y\|_2^2 \geq \alpha_{2k} \|w_1 - w^*\|_2^2 \\
 f(w_0) &= \|\Phi w_0 - y\|_2^2 \geq \alpha_k \|w_0 - w^*\|_2^2
 \end{aligned}$$

Plugging these into (12) completes the proof. □

## D Additional Related Work

Thresholding based optimization algorithms have been attractive alternatives to relaxing the constraint to a convex one or to greedy selection. Bahmani et al. (2013) provide a gradient thresholding algorithm that generalizes pursuit approaches for compressed sensing to more general losses. Yuan et al. (2018) study convergence of gradient thresholding algorithms for general losses. Jain et al. (2014) consider several variants of thresholding based algorithms for high dimensional sparse estimation. Nguyen et al. (2014); Li et al. (2016) discuss convergence properties of thresholding algorithms for stochastic settings, while in (Jain et al., 2016) the algorithm is extended to structured sparsity. Greedy algorithms (Shalev-Shwartz et al., 2010) for cardinality constrained problems have similar convergence guarantees and smaller per iteration cost but tend to underperform when compared to threshold based algorithms (Khanna & Kyrillidis, 2018).

Acceleration using momentum term (Beck & Teboulle, 2009; Ghadimi et al., 2015) allows for faster convergence of first order methods without increasing the per iteration cost. In the context of accelerating sparsity constrained first order optimization, Khanna & Kyrillidis (2018); Blumensath (2012) use momentum terms in conjunction with thresholding and prove linear convergence of their method. We extend their work by also including additional constraints of non-negativity. More recently, there have also been works (Ma et al., 2019) that study acceleration in sampling methods such as MCMC that are relevant to Bayesian coresets.

## E Additional Results for Synthetic Gaussian Posterior Inference

Additional results for experiments in section 5.1 are provided in this section.

From an optimization perspective, one may be curious about the convergence speed of the two proposed algorithms, i.e., A-IHT and Accelerated A-IHT II (Algorithm 2 & 3). The convergence for the two algorithms compared to

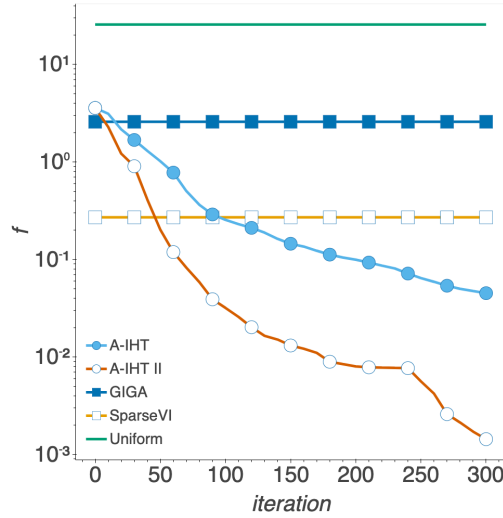


Figure 4: Convergence results for synthetic Gaussian posterior inference (subsection 5.1) when sparsity setting  $k = 200$  in the first trial. For GIGA, SparseVI and Uniform, each of the objective function values  $f$  is calculated by the final output of each algorithms.

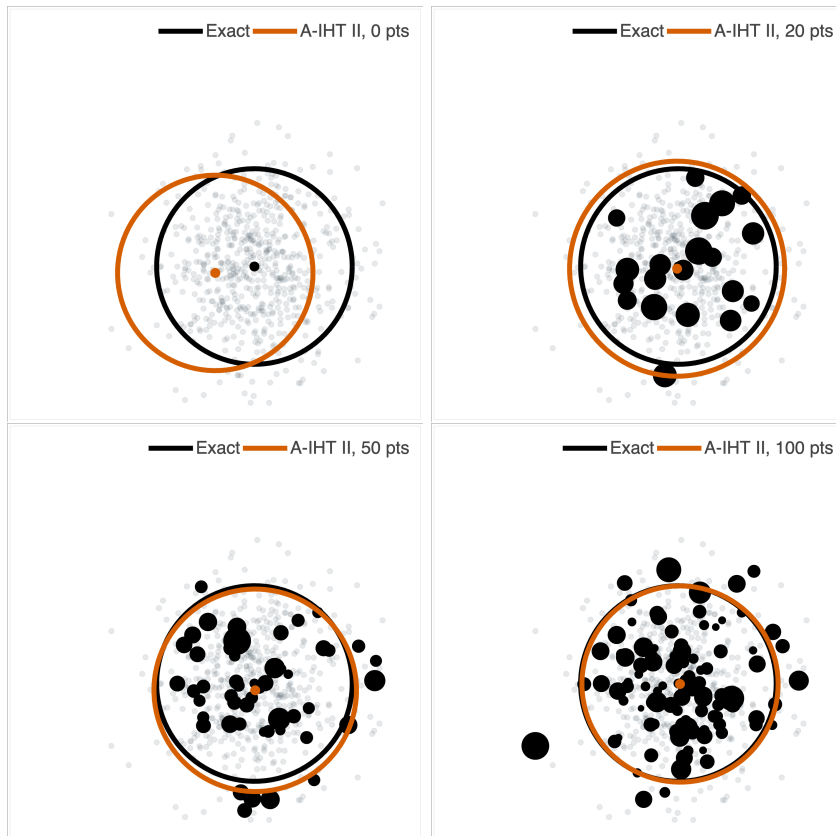


Figure 5: Illustration of true posterior and posterior constructed by A-IHT II after projecting to 2-dimensional plane for synthetic Gaussian posterior inference (Section 5.1). Results at different sparsity level are shown. The ellipses indicate  $2\sigma$ -prediction of the posterior distribution, and the black dots represent coreset points selected with their radius denoting the respective weights.

the solutions by baselines are presented in Figure 4. The x-axis is iteration number for A-IHT and A-IHT II, and

the y-axis is the objective function to be minimized, *i.e.*,

$$f(w) = \|y - \Phi w\|_2^2,$$

where  $y = \sum_{i=1}^n \hat{g}_i$  and  $\Phi = [\hat{g}_1, \dots, \hat{g}_n]$ .

The two IHT algorithms' fast convergence speed reflects what our theory suggests. They surpass GIGA within about 30 iterations, and surpass SparseVI within 50 iterations (A-IHT II) and within 100 iterations (A-IHT), respectively. Although we should note that the objective function which SparseVI minimizes is reverse KL divergence instead of  $l_2$  distance, the two IHT algorithms can achieve much better solutions when considering KL divergence as well, as shown in Figure 1. Moreover, the tendency of further decrease in objective value is still observed for the two IHT algorithms at 300<sup>th</sup> iteration.

Illustration of the coresets constructed by A-IHT II in the first trial after projecting to 2D is presented in Figure 5.

## F Additional Results for Radial Basis Regression

In this section, we provide additional experimental results of posterior contours for the radial basis regression experiment (section 5.2).

We plot the posterior contours for both the true posterior and coreset posterior when sparsity level  $k = 300$  in the first four random trials out of ten trials. The coreset posterior constructed by our Algorithm 3 recovers the true posterior almost exactly, unlike SparseVI. Results are shown in Figure 6.

## G Details and Extensive Results of the Bayesian logistic and Poisson regression Experiments

We consider how IHT performs when used in real applications where the closed-form expressions are unattainable. As the true posterior is unknown, a Laplace approximation is used for GIGA and IHT to derive the finite projection of the distribution, *i.e.*,  $\hat{g}_i$ . Further, Monte Carlo sampling is needed to derive gradients of  $D_{\text{KL}}$  for SparseVI. We compare different algorithms estimating the posterior distribution for logistic regression and Poisson regression. The reverse KL and forward KL between the coreset posterior and true posterior are estimated using another Laplace approximation. The experiment was proposed by Campbell & Broderick (2019), and is used in (Campbell & Broderick, 2018) (GIGA) and (Campbell & Beronov, 2019) (SparseVI). The experimental settings for each baseline algorithms are set following their original settings for this experiment.

For logistic regression, given a dataset  $\{(x_n, y_n) \in \mathbb{R}^D \times \{1, -1\} \mid i \in [N]\}$ , we aim to infer  $\theta \in \mathbb{R}^{D+1}$  based on the model:

$$y_n \mid x_n, \theta \sim \text{Bern} \left( \frac{1}{1 + e^{-z_n^\top \theta}} \right),$$

where  $z_n = [x_n^\top, 1]^\top$ . Three datasets are used for logistic regression. The *synthetic dataset for logistic regression* consists of data  $x_n$  sampled i.i.d. from standard normal distribution  $\mathcal{N}(0, I)$ , and label  $y_n$  sampled from Bernoulli distribution conditioned on  $x_n$  and  $\theta = [3, 3, 0]^\top$ . The original *phishing* dataset<sup>4</sup> consists of  $N = 11055$  data points with dimension  $D = 68$ . The *phishing* dataset used in this experiment is preprocessed (Campbell & Beronov, 2019) via principle component analysis to project each data points to dimension of  $D = 10$  to mitigate high computation by SparseVI. The original *chemical reactivities* dataset<sup>5</sup> has  $N = 26733$  data points with dimension  $D = 10$ . We uniformly sub-sample  $N = 500$  data points from each datasets for this experiment, due to the high computation cost of SparseVI.

For Poisson regression, given  $\{(x_n, y_n) \in \mathbb{R}^D \times \mathbb{N} \mid i \in [N]\}$ , we aim to infer  $\theta \in \mathbb{R}^{D+1}$  from model

$$y_n \mid x_n, \theta \sim \text{Pois} \left( \log \left( 1 + e^{-z_n^\top \theta} \right) \right),$$

where  $z_n = [x_n^\top, 1]^\top$ . Three other datasets are used for Poisson regression: the *synthetic dataset for Poisson regression* consists of data  $x_n$  sampled i.i.d. from a standard normal distribution  $\mathcal{N}(0, 1)$ , and target  $y_n$  sampled

<sup>4</sup><https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/binary.html>

<sup>5</sup><http://komarix.org/ac/ds>

from Poisson distribution conditioned on  $x_n$  and  $\theta = [1, 0]^\top$ . The *biketrips* dataset<sup>6</sup> consists of  $N = 17386$  data points with dimension  $D = 8$ . The *airportdelays* dataset<sup>7</sup> has  $N = 7580$  data points with dimension  $D = 15$ . Same as logistic regression, we uniformly sub-sample  $N = 500$  data points from each datasets for this experiment.

The comparison of the algorithms for Bayesian coresets construction for logistic regression are shown in Figure 7, and Bayesian coresets construction for Poisson regression are shown in Figure 8. The left column shows forward KL divergence given sparsity setting  $k$ , the middle column shows reverse KL divergence, and the right column presents the running time for coresets construction for each algorithm.

It is observed that A-IHT and A-IHT II achieve state-of-the-art performance. The IHT algorithms often obtain coresets with smaller KL than GIGA and SparseVI, with computing time comparable to GIGA, significantly less than SparseVI. The experiments indicate that IHT outperforms the previous methods, improving the trade-off between accuracy and performance.

---

<sup>6</sup><http://archive.ics.uci.edu/ml/datasets/Bike+Sharing+Dataset>

<sup>7</sup>The *airportdelays* dataset was constructed (Campbell & Broderick, 2019) by combining flight delay data (<http://stat-computing.org/dataexpo/2009/the-data.html>) and weather data (<https://www.wunderground.com/history/>).

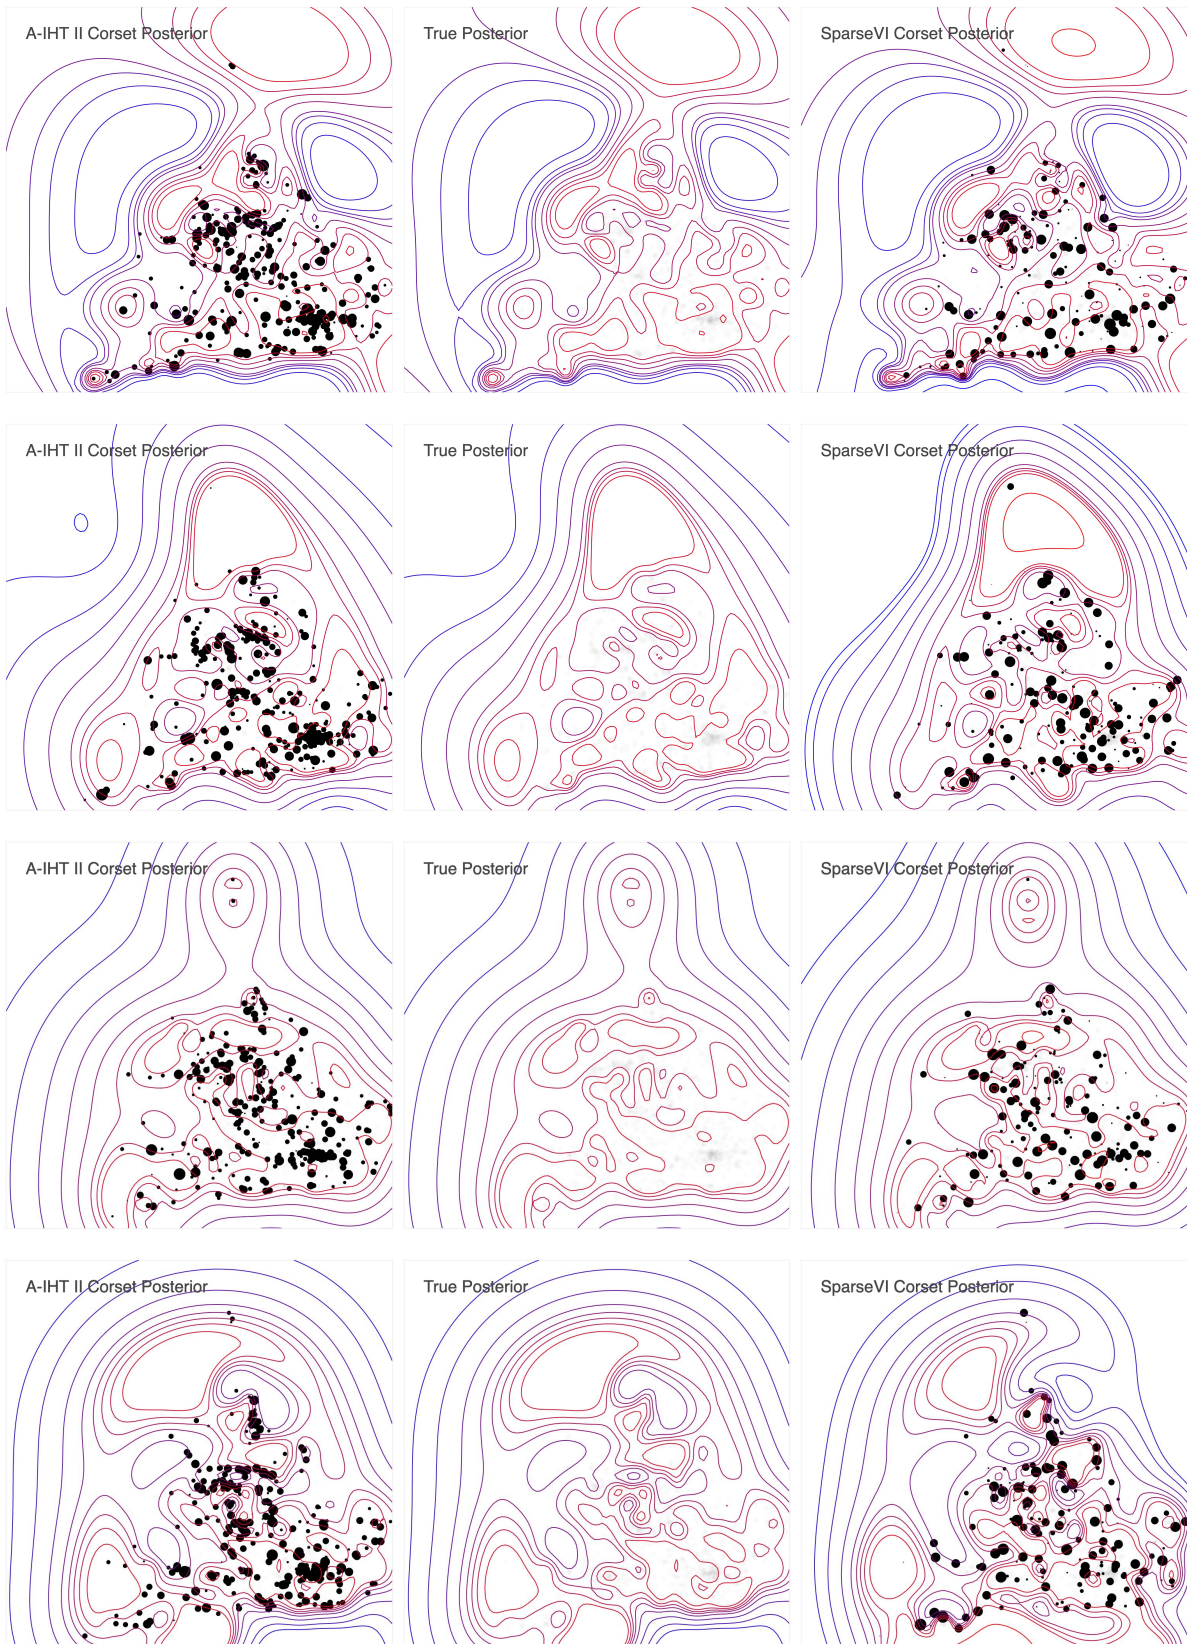




Figure 6: Experiments on Bayesian radial basis function regression in the first four random trials out of ten trails, where coreset sparsity setting  $k = 300$ . Coreset points are presented as black dots, with their radius indicating assigned weights. Posterior constructed by Accelerated IHT II (left) shows almost exact contours as the true posterior distribution (middle), while posterior constructed by SparseVI (right) shows deviated contours from the true posterior distribution.

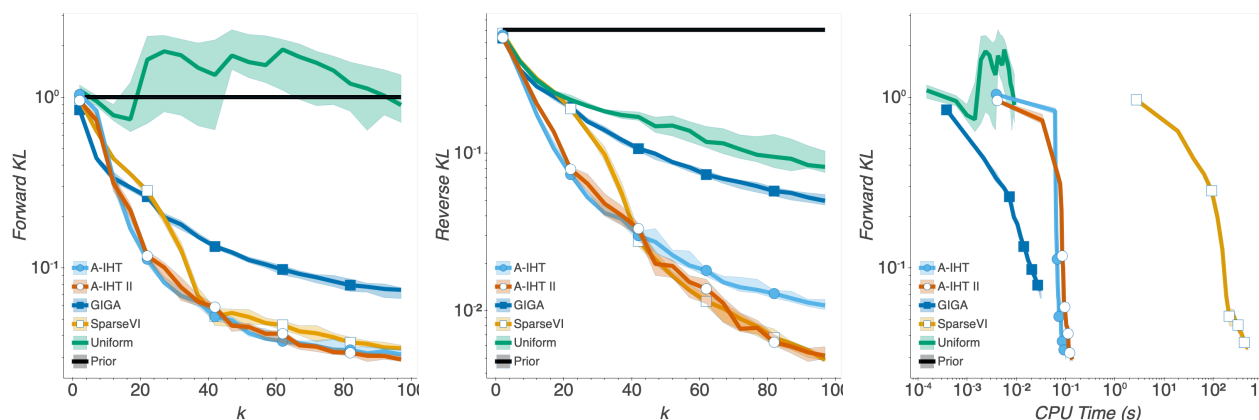




(a) synthetic dataset for logistic regression




(b) phishing dataset for logistic regression




(c) chemical reactivities dataset for logistic regression

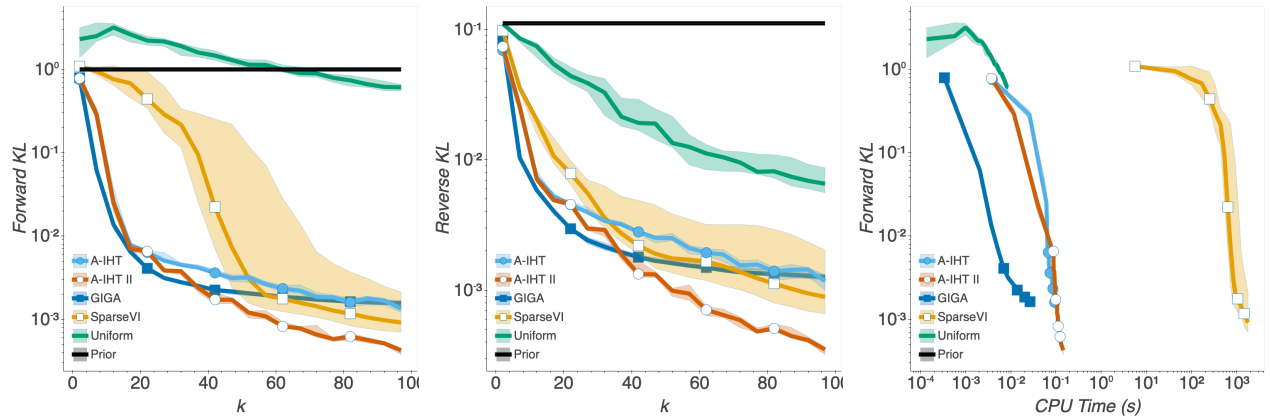
Figure 7: Bayesian coreset construction for logistic regression (LR) using the three different datasets. All the algorithms are run 20 times, and the median as well as the interval of 35<sup>th</sup> and 65<sup>th</sup> percentile, indicated as the shaded area, are reported. Different maximal coreset size  $k$  is tested from 1 to 100. Forward KL (left) and reverse KL (middle) divergence between estimated true posterior and coreset posterior indicate the quality of the constructed coreset. The smaller the KL divergence, the better the coreset is. The running time for each algorithms is also recorded (right).



(a) synthetic dataset for Poisson regression



(b) biketrips dataset for Poisson regression



(c) airportdelays dataset for Poisson regression

Figure 8: Bayesian coreset construction for Poisson regression (PR) using the three different datasets. All the algorithms are run 20 times, and the median as well as the interval of 35<sup>th</sup> and 65<sup>th</sup> percentile, indicated as the shaded area, are reported. Different maximal coreset size  $k$  is tested from 1 to 100. Forward KL (left) and reverse KL (middle) divergence between estimated true posterior and coreset posterior indicate the quality of the constructed coreset. The smaller the KL divergence, the better the coreset is. The running time for each algorithms is also recorded (right).