

M-IHS: An Accelerated Randomized Preconditioning Method Avoiding Costly Matrix Decompositions

Ibrahim K. Ozaslan^{a,*}, Mert Pilanci^{b,1}, Orhan Arikan^c

^a*Department of Electrical and Computer Engineering, University of Southern California, Los Angeles, 90089, CA, USA*

^b*Department of Electrical Engineering, Stanford University, Stanford, 94305, CA, USA*

^c*Department of Electrical and Electronics Engineering, Bilkent University, Ankara, 06800, Turkey*

Abstract

Momentum Iterative Hessian Sketch (M-IHS) techniques, a group of solvers for large scale regularized linear Least Squares (LS) problems, are proposed and analyzed in detail. The proposed M-IHS techniques are obtained by incorporating Polyak's heavy ball acceleration into the Iterative Hessian Sketch algorithm and they provide significant improvements over the randomized preconditioning techniques. By solving the linear systems arising in the sub-problems during the iterations approximately, the proposed techniques are capable of avoiding all matrix decompositions and inversions, which is one of the main advantages over the alternative solvers such as the Blendenpik and the LSRN. Similar to the Chebyshev semi-iterations, the M-IHS variants do not use any inner products and eliminate the corresponding synchronization steps in hierarchical or distributed memory systems, yet the M-IHS converges faster than the Chebyshev semi-iteration based solvers. Lower bounds on the required sketch size for various randomized distributions are established through the error analyses. Unlike the previously proposed approaches to produce a solution approximation, the proposed M-IHS techniques can use sketch sizes that are proportional to the statistical dimension which is always smaller than the rank of the coefficient matrix.

Keywords: Tikhonov regularization, ridge regression, random projection, randomized preconditioning, acceleration

2000 MSC: 65F08, 65F10, 65F20, 68W20

1. Introduction

We are presenting a group of solvers, named as Momentum Iterative Hessian Sketch, M-IHS, that is designed for solving large scale linear system of equations in the form of

$$Ax_0 + \omega = b, \tag{1}$$

*Corresponding author

Email address: ozaslan@usc.edu (Ibrahim K. Ozaslan)

¹Mert Pilanci's work was partially supported by the National Science Foundation under grant IIS-1838179.

where $A \in \mathbb{R}^{n \times d}$ is the given data or coefficient matrix, b is the given measurement vector contaminated by the noise or computation/discretization error ω , and x_0 is the vector desired to be recovered. Due to contaminated measurements, solutions can differ according to the constraints imposed on the problem. In this article, we are particularly interested in the ℓ_2 -norm regularized Least Squares (LS) solution:

$$x^* = \underset{x}{\operatorname{argmin}} \underbrace{\frac{1}{2} \|Ax - b\|_2^2 + \frac{\lambda}{2} \|x\|_2^2}_{f(x)}, \quad (2)$$

which is known as Tikhonov regularization or the ridge regression [1]. Problem (2) frequently arises in various large scale applications of science and engineering. For example, such regularized problems appear in the discretization of Fredholm integral equations of the first kind [2]. In those cases, the data matrix might be ill conditioned and the linear system can be either over-determined or square. When the system is under-determined, although sparse solutions are recently popularized by compressed sensing [3], the least norm solutions occupy a fundamental place in statistics applications such as the support vector machines [4]. Solutions to the problem in both regimes, i.e, $n \geq d$ and $n < d$, are often required as intermediate steps of rather complicated algorithms such as the interior point and the ADMM that are widely used in machine learning and image processing applications [5, 6, 7].

Throughout the manuscript, it is assumed that a proper estimate for the regularization parameter λ is available. In the absence of such an estimate, risk estimators such as the discrepancy principle, unbiased prediction risk estimate, Stein's unbiased risk estimate, and generalized cross validation can be directly used to obtain an estimate for the regularization parameter in the moderate size problems [8]. In large scale problems, these risk estimators can be adapted for the lower dimensional sub-problems that arise during the iteration of the first-order iterative solvers [9]. A hybrid scheme that adaptively selects the regularization parameter along with the iterations is also suitable for the proposed M-IHS solvers. Indeed, we have developed such a technique that can be found in Chapter 4 of [10], but we will present it in a separate manuscript due to page length constraints.

The regularized solution in (2) can be obtained by using *direct* methods such as the Cholesky decomposition for square A , or the QR decomposition for rectangular A . However, $O(nd \min(n, d))$ computational complexity of any full matrix decomposition becomes prohibitively large as the dimensions increase. For large scale problems, linear dependence on both dimensions might seem acceptable and can be realizable by using the first-order iterative solvers that are based on the Krylov subspaces [1]. These methods require only a few matrix-vector and vector-vector multiplications for each iteration, but the number of iterations that is needed to reach a high level of accuracy is highly sensitive to the condition number of the coefficient matrix. For the problem given in (2), the convergence rate of the first-order iterative solvers based on the Krylov subspace iterations including the conjugate gradient (CGLS), LSQR, LSMR, and Chebyshev Semi-iterative (CS) technique is

characterized by the following inequality:

$$\|x^i - x^*\|_2 \leq \left(\frac{\sqrt{\kappa(A^T A + \lambda I_d)} - 1}{\sqrt{\kappa(A^T A + \lambda I_d)} + 1} \right)^i \|x^1 - x^*\|_2,$$

where x^* is the optimal solution of (2), x^1 is the initial guess, x^i is the i -th iterate of the solver for $i < 1$, and the condition number $\kappa(\cdot)$ is defined as the ratio of the largest singular value to the smallest singular value of its argument [11]. Since for ill conditioned matrices $\kappa(A^T A + \lambda I_d)$ can be large, the rate of convergence may be extremely slow.

The computational complexity of the Krylov subspace-based iterative solvers is $O(nd)$ for each iteration, hence the total complexity is less than $O(nd \min(n, d))$ if the number of iterations can be significantly fewer than $\min(n, d)$. However, in the applications such as big data where A is very large dimensional, the computational complexity is not the only metric for the feasibility of the algorithms. For example, if the coefficient matrix is too large to fit in a single working memory and it could be merely stored in a number of distributed computational nodes, then at least two distributed computations of matrix-vector multiplications are required at each iteration of algorithms such as CGLS or LSQR [12, 13]. Therefore, in addition to the computational complexity, the number of iterations should be counted as an important metric to measure the overall complexity of an algorithm. One way to reduce the number of iterations in the iterative solvers is to use preconditioning to transform an ill conditioned problem to a well conditioned one [14]. In the deterministic settings, finding a low-cost and effective preconditioning matrix is still a challenging task unless the coefficient matrix has a particular structure [15].

In addition to the number of iterations, the number of inner products in each iteration also plays an important role in the overall complexity. Each inner product calculation constitutes a synchronization step in parallel computing and therefore is undesirable for distributed or hierarchical memory systems [14]. The CS technique can be preferred in this kind of applications, since it does not use any inner products, and therefore, eliminates some of the synchronization steps that are required by the techniques such as CG or GMRES. However, CS technique requires prior information about the ellipsoid that contains all the eigenvalues of A , which is typically not available in practice [16].

These aforementioned drawbacks of the direct and the iterative methods can be alleviated using the Random Projection (RP) techniques [17, 18]. These techniques are capable of both reducing the dimensions and bounding the number of iterations with statistical guarantees, while they are quite convenient for parallel and distributed computations [19, 20]. There are two main approaches for the applications of the RP-based techniques to the regularized LS problem in (2). In the first approach, referred to as the *classical sketching*, the coefficient matrix A and the measurement vector b are projected down onto a lower dimensional $(m \ll n)^2$ subspace by using a randomly constructed *sketching matrix* $S \in \mathbb{R}^{m \times n}$, to obtain

²Without loss of generality we can assume the linear system is over-determined; otherwise, we can take the Lagrange dual to obtain an over-determined problem as elaborated later in the text.

efficiently an ζ -optimal solution with high probability for the *cost approximation* [21, 22]:

$$\tilde{x} = \operatorname{argmin}_x \frac{1}{2} \|SAx - Sb\|_2^2 + \frac{\lambda}{2} \|x\|_2^2, \text{ such that } f(\tilde{x}) \leq (1 + \zeta)f(x^*). \quad (3)$$

For both the sparse and dense systems, in [23], the best known lower bounds on the sketch size for obtaining an ζ -optimal cost approximation have been derived showing that the sketch size can be chosen proportional to the *statistical dimension* $\mathbf{sd}_\lambda(A) = \operatorname{tr}(A(A^T A + \lambda I_d)^{-1} A^T)$. Although the cost approximation is sufficient for many machine learning problems, the *solution approximation* which aims to produce solutions that are close to the optimal solution is a more preferable metric for the problems arisen from, for example, discretization of Fredholm integrals [1, 8]. However, as shown in [24], the classical sketching is sub-optimal in terms of the minimum sketch size for obtaining a solution approximation.

In the second approach of *randomized preconditioning*, by iteratively solving a number of low dimensional sub-problems constituted by $(SA, \nabla f(x^i))$ pairs, algorithms with reasonable sketch sizes obtain an η -optimal solution approximation \hat{x} such that

$$\|\hat{x} - x^*\|_W \leq \eta \|x^*\|_W, \quad (4)$$

where W is a positive definite weight matrix. In [25], RP techniques have been proposed to construct a preconditioning matrix for CG-like algorithms by using the inverse of R factor in the QR decomposition of the *sketched matrix* SA . Later, implementation of similar ideas resulted in Blendenpik and LSRN which have been shown to be faster than some of the deterministic solvers of LAPACK [26, 27]. To solve the preconditioned problems, as opposed to Blendenpik which uses LSQR, LSRN uses the CS technique for parallelization purposes and deduce the prior information about the eigenvalues based on the results of the random matrix theory. The main drawback of LSRN and Blendenpik is that regardless of the desired accuracy η , one has to pay the whole cost, $O(md^2)$, of a full $m \times d$ dimensional matrix decomposition, which is the dominant term in the computational complexity of these algorithms. Iterative Hessian Sketch (IHS) proposed in [24] enables the use of the sketched Hessian as preconditioning matrix in the gradient descent method [28] thereby providing a reduction in the dominant complexity term $O(md^2)$ to $O(md)$. For the IHS, instead of computing a full decomposition or an inversion, a linear system can be approximately solved for a pre-determined tolerance. The ability to use this *inexact* approach becomes more important in the large scale inverse problems such as 3D imaging [29] where even a decomposition of $m \times d$ -dimensional sketched matrix is infeasible to compute. By using the preconditioning idea of the IHS in the CG technique, Accelerated IHS (A-IHS) has been proposed in [28]. Lastly, in [30], it has been showed that if the linear system is strongly over-determined, then the momentum parameters of the heavy ball method can be robustly estimated by using Marchenko Pastur Law (MPL) [31]. This analysis results in a prototype solver M-IHS that we study here in detail.

The statistical lower bounds obtained in the current literature suggest that the sketch size in randomized preconditioning algorithms can be chosen proportional to the rank of the

problem, which can be significantly larger than the statistical dimension. Some lower bounds on the sketch size that are proportional to the statistical dimension have been obtained in ridge regression [32, 33]. However, these analyses are limited to kernel ridge regression or under-determined regime where $n \leq d$ and do not consider the acceleration in the iteration of the methods.

2. Contributions

In this article, we propose a group of random projection based iterative solvers for large scale regularized LS problems. As shown by detailed analyses of their convergence behaviours, the proposed M-IHS variants can be used for any dimension regimes with significant computational savings if the statistical dimension of the problem is sufficiently smaller than at least one size of the coefficient matrix. Our guarantees, presented in Theorem 1 and Corollary 2, are based on the solution approximation metric given in (4) as opposed to the results obtained for cost approximation metric given in (3). In Lemma 3, we improved the known lower bounds on the sketch size of various randomized distribution for obtaining a pre-determined convergence rate with a constant probability. These guarantees can be readily extended to any other sketching types by using the *Approximate Matrix Multiplication* (AMM) property defined in [34]. When tighter bounds for the AMM property will be available in the future, the bounds derived in this work can be automatically improved as well. Additionally, we provide some approximate bounds for the sketch size and the rate of convergence in Corollary 4 which is remarkably tight as demonstrated through numerical experiments. Lastly, in Algorithm 6, we extend the idea of LSQR into the linear problems in the form of $A^T Ax = b$ which we need to solve during the iterations of all proposed *Inexact M-IHS* variants and of the Newton Sketch [7]. Similar to the stability advantage of the LSQR over the CGLS technique [35], the proposed method solves the system in the above form without squaring the condition number as opposed to the techniques such as the symmetric CG and the symmetric Lanczos techniques. In the following link, implementations of the proposed solvers in *MATLAB* together with the codes that generate the figures in the article can be found: <https://github.com/ibrahimkurban/M-IHS>.

3. The proposed M-IHS solvers for the regularized LS problems

The naive IHS algorithm approximates the Hessian in the Newton method to gain computational savings while solving the Newton sub-systems, and it iteratively minimizes the quadratic objective function given in (2) by performing the following updates [24]:

$$x^{i+1} = \operatorname{argmin}_{x \in \mathbb{R}^d} \|S_i A(x - x^i)\|_2^2 + \lambda \|x - x^i\|_2^2 + 2\langle \nabla f(x^i), x \rangle. \quad (5)$$

where the gradient is given by $\nabla f(x) = A^T(Ax - b) + \lambda x$. Here, we aim to gain further computational efficiency and improved convergence rate for the naive IHS technique. We realized these two goals by incorporating Polyak's momentum method [36] into the iterations

in (5) and propose the following Momentum-IHS (M-IHS) updates:

$$\begin{aligned}\Delta x^i &= \operatorname{argmin}_{x \in \mathbb{R}^d} \|SAx\|_2^2 + \lambda \|x\|_2^2 + 2 \langle \nabla f(x^i), x \rangle, \\ x^{i+1} &= x^i + \alpha_i \Delta x^i + \beta_i (x^i - x^{i-1}),\end{aligned}\tag{6}$$

where the same sketching matrix S is used for all iterations. For a properly chosen momentum parameters α_i and β_i , there is no need to change the sketching matrix in the above iterations unlike the IHS technique. Moreover, the optimal fixed momentum parameters α and β , that maximize the convergence rate, can be estimated by using the random matrix theory as completely independent of the spectral properties of the coefficient matrix A [30]. Here, the linear system is assumed to be strongly over-determined, i.e., $n \gg d$. By using the dual formulation, the theory can be straightforwardly extended to the strongly under-determined case of $d \gg n$ as follows. A dual of Problem (2) is given by

$$\nu^* = \operatorname{argmin}_{\nu \in \mathbb{R}^n} \underbrace{\frac{1}{2} \|A^T \nu\|_2^2 + \frac{\lambda}{2} \|\nu\|_2^2 - \langle b, \nu \rangle}_{g(\nu)},\tag{7}$$

with the following relation between the solutions of the primal and dual problems

$$\nu^* = (b - Ax^*)/\lambda \iff x^* = A^T \nu^*.\tag{8}$$

The corresponding M-IHS iterations for the dual problem are:

$$\begin{aligned}\Delta \nu^i &= \operatorname{argmin}_{\nu \in \mathbb{R}^n} \|SA^T \nu\|_2^2 + \lambda \|\nu\|_2^2 + 2 \langle \nabla g(\nu^i), \nu \rangle, \\ \nu^{i+1} &= \nu^i + \alpha \Delta \nu^i + \beta (\nu^i - \nu^{i-1}),\end{aligned}\tag{9}$$

where $\nabla g(\nu) = (AA^T + \lambda I)\nu - b$ and the solution of the primal problem can be obtained through the relation in (8). We refer to this algorithm as the **Dual M-IHS**. The convergence rate of the M-IHS and the **Dual M-IHS** solvers together with the optimal fixed momentum parameters are stated in the Theorem 1 below.

Theorem 1. *Let A and b be the given data in (1) with singular values σ_i in descending order $1 \leq i \leq \min(n, d)$, $x^* \in \mathbb{R}^d$ and $\nu^* \in \mathbb{R}^n$ are as in (2) and (7), respectively. Let $U_1 \in \mathbb{R}^{\max(n, d) \times \min(n, d)}$ consists of the first n rows of an orthogonal basis for $[A^T \sqrt{\lambda} I_d]^T$ if the problem is over-determined, and consists of the first d rows of an orthogonal basis for $[A \sqrt{\lambda} I_n]^T$ if the problem is under-determined. Let the sketching matrix $S \in \mathbb{R}^{m \times \max(n, d)}$ be drawn from a distribution \mathcal{D} such that*

$$\mathbb{P}_{S \sim \mathcal{D}} (\|U_1^T S^T S U_1 - U_1^T U_1\|_2 \geq \epsilon) < \delta, \quad \epsilon \in (0, 1).\tag{10}$$

Then, the M-IHS applied on (2) and the Dual M-IHS applied on (7) with the following

momentum parameters

$$\beta^* = \left(\frac{\epsilon}{1 + \sqrt{1 - \epsilon^2}} \right)^2, \quad \alpha^* = (1 - \beta^*) \sqrt{1 - \epsilon^2},$$

converge to the optimal solutions, x^* and ν^* , respectively, at the following rate with a probability of at least $(1 - \delta)$:

$$\begin{aligned} \|x^{i+1} - x^*\|_{D^{-1}} &\leq \frac{\epsilon}{1 + \sqrt{1 - \epsilon^2}} \|x^i - x^*\|_{D^{-1}}, \\ \|\nu^{i+1} - \nu^*\|_{D^{-1}} &\leq \frac{\epsilon}{1 + \sqrt{1 - \epsilon^2}} \|\nu^i - \nu^*\|_{D^{-1}}, \end{aligned}$$

where D^{-1} is the diagonal matrix whose diagonal entries are $\sqrt{\sigma_i^2 + \lambda}$.

Proof. In the following analysis we denote $A = U\Sigma V^T$ as the compact SVD with $U \in \mathbb{R}^{n \times r}$, $\Sigma = \mathbf{diag}(\sigma_1, \dots, \sigma_r) \in \mathbb{R}^{r \times r}$ and $V \in \mathbb{R}^{d \times r}$ where $r = \min(n, d)$. To prove the theorem for M-IHS and Dual M-IHS, we mainly combine the idea of *partly exact* sketching, that is proposed in [23], with the analysis used in [30]. We define the diagonal matrix $D := (\Sigma^2 + \lambda I_r)^{-1/2}$ and the *partly exact* sketching matrix S as:

$$\widehat{S} = \begin{bmatrix} S & \mathbf{0} \\ \mathbf{0} & I_r \end{bmatrix}, \quad S \in \mathbb{R}^{m \times \max(n, d)}.$$

The proof for M-IHS. Let

$$\widehat{A} = \begin{bmatrix} U\Sigma D \\ \sqrt{\lambda}VD \end{bmatrix} = \begin{bmatrix} U_1 \\ U_2 \end{bmatrix}, \quad \widehat{A}^T \widehat{A} = I_d, \quad \widehat{b} = \begin{bmatrix} b \\ \mathbf{0} \end{bmatrix},$$

so that U_1 consists of the first n rows of an orthogonal basis for $[A^T \ \sqrt{\lambda}I_d]^T$ as required by the condition in (10) of the theorem. To simplify the analysis, the following LS problem will be used:

$$y^* = \operatorname{argmin}_{y \in \mathbb{R}^d} \|\widehat{A}y - \widehat{b}\|_2^2 \quad (11)$$

which is equivalent to Problem (2) due to the one-to-one mapping $y^* = D^{-1}V^T x(\lambda)$. For Problem (11), the equivalent of M-IHS given in (6) is the following update:

$$\begin{aligned} \Delta y^i &= \operatorname{argmin}_y \|\widehat{S}\widehat{A}y\|_2^2 - 2\langle \widehat{A}^T(\widehat{b} - \widehat{A}y^i), y \rangle \\ y^{i+1} &= y^i + \alpha \Delta y^i + \beta(y^i - y^{i-1}) \end{aligned}$$

with sketched matrix

$$\widehat{S}\widehat{A} = \begin{bmatrix} SU\Sigma D \\ \sqrt{\lambda}VD \end{bmatrix} = \begin{bmatrix} SU_1 \\ U_2 \end{bmatrix}.$$

Thus, we can examine the following bipartite transformation to find out the convergence properties of the M-IHS:

$$\begin{bmatrix} y^{i+1} - y^* \\ y^i - y^* \end{bmatrix} = \underbrace{\begin{bmatrix} (1 + \beta)I_d - \alpha(\widehat{A}^T \widehat{S}^T \widehat{S} \widehat{A})^{-1} & -\beta I_d \\ I_d & \mathbf{0} \end{bmatrix}}_T \begin{bmatrix} y^i - y^* \\ y^{i-1} - y^* \end{bmatrix}.$$

The contraction ratio of the transformation, which is determined by the eigenvalues, can be found analytically by converting the matrix T into a block diagonal form through the similarity transformation $T = P^{-1} \mathbf{diag}(T_1, \dots, T_d)P$ where

$$T_i := \begin{bmatrix} 1 + \beta - \alpha\mu_i & \beta \\ 1 & 0 \end{bmatrix}, \quad P = \begin{bmatrix} \Psi & 0 \\ 0 & \Psi \end{bmatrix} \Pi, \quad \Pi_{i,j} = \begin{cases} 1 & i \text{ is odd } j = i, \\ 1 & i \text{ is even } j = r + i, \\ 0 & \text{otherwise,} \end{cases}$$

$\Psi \Xi \Psi^T$ is the spectral decomposition of $(\widehat{A}^T \widehat{S}^T \widehat{S} \widehat{A})^{-1}$, and μ_i is the i^{th} eigenvalue [30]. The characteristic polynomials of each block is

$$u^2 - (1 + \beta - \alpha\mu_i)u + \beta = 0, \quad \forall i \in [r]. \quad (12)$$

If the following condition holds

$$\beta \geq (1 - \sqrt{\alpha\mu_i})^2, \quad \forall i \in [r], \quad (13)$$

then both of the roots are imaginary and both have a magnitude $\sqrt{\beta}$ for all μ_i 's. In this case, all linear dynamical systems driven by the above characteristic polynomial will be in the underdamped regime and the contraction rate of transformation T , through all directions, not just one of them, will be exactly $\sqrt{\beta}$. If the condition in (13) is not satisfied for a μ_i with $i \in [r]$, then the linear dynamical system corresponding to μ_i will be in the overdamped regime and the contraction rate in the direction through the eigenvector corresponding to this overdamped system will be smaller compared to the others. As a result, the overall algorithm will be slowed down (see [37] for details). If the condition in (10) of Theorem 1 holds,

$$\|\widehat{A}^T \widehat{S}^T \widehat{S} \widehat{A} - I_r\|_2 = \|U_1^T S^T S U_1 + U_2^T U_2 - I_r\|_2 = \|U_1^T S^T S U_1 - U_1^T U_1\|_2 \leq \epsilon,$$

then, we have the following bounds:

$$\sup_{\|v\|_2=1} v^T \widehat{A}^T \widehat{S}^T \widehat{S} \widehat{A} v \leq 1 + \epsilon \quad \text{and} \quad \inf_{\|v\|_2=1} v^T \widehat{A}^T \widehat{S}^T \widehat{S} \widehat{A} v \geq 1 - \epsilon,$$

which are equivalent to:

$$\underset{i \in [r]}{\text{maximize}} \mu_i \leq \frac{1}{1 - \epsilon} \quad \text{and} \quad \underset{i \in [r]}{\text{minimize}} \mu_i \geq \frac{1}{1 + \epsilon}.$$

Consequently, the condition in (13) can be satisfied for all μ_i 's by the following choice of β that minimizes the convergence rate over step size α

$$\sqrt{\beta^*} = \underset{\alpha}{\text{minimize}} \left(\max \left\{ 1 - \frac{\sqrt{\alpha}}{\sqrt{1+\epsilon}}, \frac{\sqrt{\alpha}}{\sqrt{1-\epsilon}} - 1 \right\} \right) = \frac{\epsilon}{1 + \sqrt{1-\epsilon^2}},$$

where the minimum is achieved at $\alpha^* = \frac{4(1-\epsilon^2)}{(\sqrt{1+\epsilon} + \sqrt{1-\epsilon})^2} = (1 - \beta^*)\sqrt{1 - \epsilon^2}$ as claimed.

The proof for the Dual M-IHS. The proof of the under-determined case is parallel to the over-determined one except for the following modifications. Let

$$\widehat{A}^T = \begin{bmatrix} V\Sigma D \\ \sqrt{\lambda}UD \end{bmatrix} = \begin{bmatrix} U_1 \\ U_2 \end{bmatrix}, \quad \widehat{A}\widehat{A}^T = I_n \quad \text{and} \quad \widehat{S}\widehat{A}^T = \begin{bmatrix} SV\Sigma D \\ \sqrt{\lambda}UD \end{bmatrix} = \begin{bmatrix} SU_1 \\ U_2 \end{bmatrix}, \quad (14)$$

so that U_1 is the first d rows of an orthogonal basis for $[A \ \sqrt{\lambda}I_n]$ as required by the theorem. Similar to M-IHS case, the analysis can be simplified using the following reformulation of Problem (7)

$$w^* = \underset{w \in \mathbb{R}^n}{\text{argmin}} = \frac{1}{2} \|\widehat{A}^T w\|_2^2 - \langle DU^T b, w \rangle,$$

where the solutions of two problems are related through the one-to-one mapping $w^* = D^{-1}U^T \nu(\lambda)$. For this form, the equivalent of Dual M-IHS is

$$\begin{aligned} \Delta w^i &= \underset{w}{\text{argmin}} \|\widehat{S}\widehat{A}^T w\|_2^2 - 2\langle DU^T b - \widehat{A}\widehat{A}^T w^i, w \rangle, \\ w^{i+1} &= w^i + \alpha \Delta w^i + \beta (w^i - w^{i-1}). \end{aligned}$$

Therefore, we can analyze the following bipartite transformation to figure out the convergence properties of the Dual M-IHS

$$\begin{bmatrix} w^{i+1} - w^* \\ w^i - w^* \end{bmatrix} = \underbrace{\begin{bmatrix} (1 + \beta)I_n - \alpha(\widehat{A}\widehat{S}^T\widehat{S}\widehat{A}^T)^{-1} & -\beta I_n \\ I_n & \mathbf{0} \end{bmatrix}}_T \begin{bmatrix} w^i - w^* \\ w^{i-1} - w^* \end{bmatrix}.$$

The rest of the proof can be completed straightforwardly by following the same analysis steps as in the proof for the M-IHS case. \square

Note that, Theorem 1 remains valid in unregularized problems where A is full rank, if, instead of (10), the following condition is satisfied

$$\mathbb{P}_{S \sim \mathcal{D}} (\|U^T S^T S U - I_d\|_2 \geq \epsilon) < \delta, \quad \epsilon \in (0, 1).$$

When the necessary conditions are met, the number of iterations needed for both algorithms to reach a certain level of accuracy is stated in the following corollary.

Corollary 2. For some $\epsilon \in (0, 1/2)$ and arbitrary η , if the sketching matrix meets the condition in (10) and the fixed momentum parameters are chosen as in Theorem 1, then the number of iterations for the *M-IHS* and the *Dual M-IHS* to obtain an η -optimal solution approximation in ℓ_2 -norm is upper bounded by

$$N = \left\lceil \frac{\log(\eta) \log(C)}{\log(\epsilon) - \log(1 + \sqrt{1 - \epsilon^2})} \right\rceil$$

where $C = \sqrt{\kappa(A^T A + \lambda I_d)}$ for *M-IHS* and $C = \kappa(A) \sqrt{\kappa(AA^T + \lambda I_n)}$ for *Dual M-IHS*.

Corollary 2 is an immediate result of Theorem 1. To satisfy the condition in (10), a set of cases for the sketching matrix S are given in Lemma 3. We note that the statistical dimension of matrix A is defined as $\mathbf{sd}_\lambda(A) = \sum_{i=1}^{\min(n,d)} \sigma_i^2 / (\sigma_i^2 + \lambda)$.

Lemma 3. If the sketching matrix S is chosen in one of the following cases, the condition in (10) of Theorem 1 is satisfied.

- i. S is a sparse subspace embedding [20] with single nonzero element in each column, with a sketch size

$$m = \Omega(\mathbf{sd}_\lambda(A)^2 / (\epsilon^2 \delta))$$

where $\Omega(\cdot)$ notation is defined as $a(n) = \Omega(b(n))$, if there exists two integers k and n_0 such that $\forall n > n_0$, $a(n) \geq kb(n)$. For this case, SA is computable in $O(\mathbf{nnz}(A))$ operations.

- ii. S is a sparse subspace embedding with

$$s = \Omega(\log_\alpha(\mathbf{sd}_\lambda(A)/\delta)/\epsilon)$$

non-zero elements in each column where $\alpha > 2$, $\delta < 1/2$, $\epsilon < 1/2$, [38, 39], with a sketch size

$$m = \Omega(\alpha \mathbf{sd}_\lambda(A) \log(\mathbf{sd}_\lambda(A)/\delta) / \epsilon^2).$$

For this case, SA is computable in $O(s \mathbf{nnz}(A))$ operations.

- iii. S is a SRHT sketching matrix [34, 24] with a sketch size

$$m = \Omega((\mathbf{sd}_\lambda(A) + \log(1/\epsilon\delta)) \log(\mathbf{sd}_\lambda(A)/\delta) / \epsilon^2).$$

For this case, SA is computable in $O(nd \log(m))$ operations.

- iv. S is a Sub-Gaussian sketching matrix [22, 34] with a sketch size

$$m = \Omega(\mathbf{sd}_\lambda(A) / \epsilon^2).$$

For this case, SA is computable in $O(ndm)$ operations.

Proof. The following identities will be used on U_1 where U, Σ , and D are defined in the proof

of Theorem 1:

$$\|U_1\|_F^2 = \|U\Sigma D\|_F^2 = \|\Sigma(\Sigma^2 + \lambda I_r)^{-1/2}\|_F^2 = \sum_{i=1}^r \frac{\sigma_i^2}{\sigma_i^2 + \lambda} = \mathbf{sd}_\lambda(A),$$

and $\|U_1\|_2^2 = \frac{\sigma_1^2}{\sigma_1^2 + \lambda} \leq 1$ for any λ . If the sketch matrix S is drawn from a randomized distribution \mathcal{D} over matrices $\mathbb{R}^{m \times n}$, then by using the Approximate Matrix Property (AMM) which is given below, it will be proven that the condition in (10) can be met with a desired level of probability.

As proven in [34], if a distribution \mathcal{D} over $S \in \mathbb{R}^{m \times n}$ has the $(\epsilon, \delta, 2k, \ell)$ -OSE moment property for some $\delta < 1/2$ and $\ell \geq 2$, then it has (ϵ, δ, k) -AMM Property for any A, B , i.e.,

$$\mathbb{P}_{S \sim \mathcal{D}} \left(\|A^T S^T S B - A^T B\|_2 > \epsilon \sqrt{\left(\|A\|_2^2 + \frac{\|A\|_F^2}{k} \right) \left(\|B\|_2^2 + \frac{\|B\|_F^2}{k} \right)} \right) < \delta. \quad (15)$$

The definition of the OSE-moment property can be found in [34]. As it will be detailed next, using the AMM property in (15), the sketch sizes in the statement of Lemma 3 can be found relative to the embedding size k to satisfy the condition in (10) of Theorem 1.

For case (i) of Lemma 3, Count Sketch with a single nonzero element in each column and size $m \geq 2/(\epsilon'^2 \delta)$ has $(\epsilon', \delta, 2)$ -JL moment property [40]. JL-Moment Property can be found in Definition 6.1 of [38]. By Theorem 6.2 in [38]:

$$\|U_1 S^T S U_1 - U_1^T U_1\|_F < 3\epsilon' \|U_1\|_F^2 = 3\epsilon' \mathbf{sd}_\lambda(A) \leq \epsilon$$

for $\epsilon' = \epsilon/(3\mathbf{sd}_\lambda(A))$. So, condition in (10) holds with probability at least $1 - \delta$, if $m = O(\mathbf{sd}_\lambda(A)^2/(\epsilon^2 \delta))$.

For case (ii) of Lemma 3, combining Theorem 4.2 of [41] and Remark 2 of [34] implies that any sketch matrix drawn from an OSNAP [39] with the conditions given in case (ii) of Lemma 3 satisfies the $(\epsilon', \delta, k, \log(k/\delta))$ -OSE moment property thus the $(\epsilon', \delta, k/2)$ -AMM Property. Setting $A = B = U_1$ and $k = \mathbf{sd}_\lambda(A)/2$ in (15) gives:

$$\|U_1^T S^T S U_1 - U_1^T U_1\|_2 \leq \epsilon' (\|U_1\|_2^2 + 2) \leq 3\epsilon' \leq \epsilon$$

with probability of at least $(1 - \delta)$.

Remark 1. Based on the lower bounds established for any OSE in [42], the Conjecture 14 in [39] states that any OSNAP with $m = \Omega((k + \log(1/\delta))/\epsilon^2)$ and $s = \Omega(\log(k/\delta)/\epsilon)$ have the $(\epsilon, \delta, k, \ell)$ -OSE moment property for $\ell = \Theta(\log(k/\delta))$, an even integer. If this conjecture is proved, then by the AMM property in (15), the condition in (10) can be satisfied with probability at least $(1 - \delta)$ by using an OSNAP matrix with size $m = \Omega((\mathbf{sd}_\lambda(A) + \log(1/\delta))/\epsilon^2)$ and sparsity $s = \Omega(\log(\mathbf{sd}_\lambda(A)/\delta)/\epsilon)$.

For case (iii) of Lemma 3, by Theorem 9 of [34], SRHT with the sketch size given in case (iii) has the $(\epsilon', \delta, 2\mathbf{sd}_\lambda(A), \log(\mathbf{sd}_\lambda(A)/\delta))$ -OSE moment property and thus it provides

$(\epsilon', \delta, \mathbf{sd}_\lambda(A))$ -AMM property. Again, setting $A = B = U_1$ and $k = \mathbf{sd}_\lambda(A)$ in (15) produces the desired result.

For case (iv) of Lemma 3, the Subgaussian matrices having entries with mean zero and variance $1/m$ satisfy the JL Lemma [17] with optimal sketch size [38]. Also, they have the $(\epsilon/2, \delta, \Theta(\log(1/\delta)))$ -JL moment property [43]. Thus by Lemma 4 of [34] such matrices have $(\epsilon, \delta, k, \Theta(k + \log(1/\delta)))$ -OSE moment property for $\delta < 9^{-k}$, which means $m = \Omega(k/\epsilon^2)$. Again, by setting $A = B = U_1$ and $k = \mathbf{sd}_\lambda(A)$ in (15) produces the desired result. \square

Lemma 3 suggests that in order to satisfy the condition in Theorem 1, the sketch size can be chosen proportional to the statistical dimension of the coefficient matrix which can be considerably smaller than its rank. Moreover, to obtain a solution approximation, the second condition in Lemma 11 of [23] is not a requirement, hence we obtained better results for the sparse subspace embeddings in the cases of (i) and (ii) of Lemma 3. In the following corollary, we obtained substantially simplified empirical versions of the convergence rate, momentum parameters and required sketch size by using the MPL and approximating the filtering coefficients of Tikhonov regularization with binary coefficients. Corollary 4 suggests that the ratio between the statistical dimension and the sketch size determines the convergence rate of the proposed algorithms, which interestingly seems valid even for the sketch matrices with a single non-zero element in each column as we demonstrate in Figure 1(b). Moreover, an empirical algorithm to estimate statistical dimension $\mathbf{sd}_\lambda(A)$ using a Hutchinson-like estimator is proposed in Section 3.3. As demonstrated through numerical experiments in Section 4 (in particular Figure 7), the complexity of estimating $\mathbf{sd}_\lambda(A)$ using the proposed approach is insignificant compared to the iteration complexity of the M-IHS algorithm.

Corollary 4. *If the entries of the sketching matrix are independent, zero mean, unit variance with bounded higher order moments, and the Truncated SVD regularization with truncation parameter $\lceil \mathbf{sd}_\lambda(A) \rceil$ is used, then the M-IHS and the Dual M-IHS with the following momentum parameters*

$$\beta = \frac{\mathbf{sd}_\lambda(A)}{m}, \quad \alpha = (1 - \beta)^2$$

converge to the optimal solutions x^ and v^* , respectively, with a convergence rate of $\sqrt{\beta}$ as $m \rightarrow \infty$ while the ratio $\mathbf{sd}_\lambda(A)/m$ remains constant. Any sketch size $m > \mathbf{sd}_\lambda(A)$ can be chosen to obtain an η -optimal solution approximation in most $\log(\eta)/\log(\sqrt{\beta})$ iterations.*

Proof. Consider the regularized LS solution with parameter λ and the Truncated SVD solution with parameter $\lceil \mathbf{sd}_\lambda(A) \rceil$:

$$x^* = \sum_{i=1}^r \frac{\sigma_i^2}{\sigma_i^2 + \lambda} \frac{u_i^T b}{\sigma_i} v_i \quad \text{and} \quad x^\dagger = \sum_{i=1}^{\lceil \mathbf{sd}_\lambda(A) \rceil} \frac{u_i^T b}{\sigma_i} v_i \quad (16)$$

where u_i 's and v_i 's are columns of U and V matrices in the SVD. The Tikhonov regularization with the closed form solution is preferred in practice to avoid the high computational cost of the SVD. The filtering coefficients of the Tikhonov regularization, $\frac{\sigma_i^2}{\sigma_i^2 + \lambda}$, become very close

to the binary filtering coefficients of the TSVD, as the decay rate of the singular values of A increases. In these cases, x^* and x^\dagger in (16) are very close to each other (Section 4 and 5 of [44]). Thereby, the diagonal matrix ΣD which is used in the proof of Lemma 3 can be approximated by the diagonal matrix Π where

$$\Pi_{ii} = \begin{cases} 1 & \text{if } i \leq \mathbf{sd}_\lambda(A) \leq r \\ 0 & \text{otherwise} \end{cases},$$

which is equivalent to replacing the Tikhonov coefficients by the binary coefficients. Then, we have the following close approximation:

$$\begin{aligned} (\widehat{A}^T \widehat{S}^T \widehat{S} \widehat{A})^{-1} &= (D \Sigma U^T S^T S U \Sigma D + \lambda D^2)^{-1} \\ &\approx (\Pi (S U)^T (S U) \Pi + I_r - \Pi)^{-1} = \begin{bmatrix} \overline{S}^T \overline{S} & \mathbf{0} \\ \mathbf{0} & I_{(r - \mathbf{sd}_\lambda(A))} \end{bmatrix}^{-1}, \end{aligned}$$

where $\overline{S} = S U \Pi \in \mathbb{R}^{m \times \mathbf{sd}_\lambda(A)}$ has the same distribution as S , since $U \Pi$ is an orthonormal transformation. By the MPL, the minimum and the maximum eigenvalues of this approximation converge to $(1 \pm \sqrt{\mathbf{sd}_\lambda(A)/m})^{-2}$ as $m \rightarrow \infty$ and while the ratio $\mathbf{sd}_\lambda(A)/m$ converges to a constant [31]. The rest of the proof follows from the analysis given in the proof of Theorem 1. \square

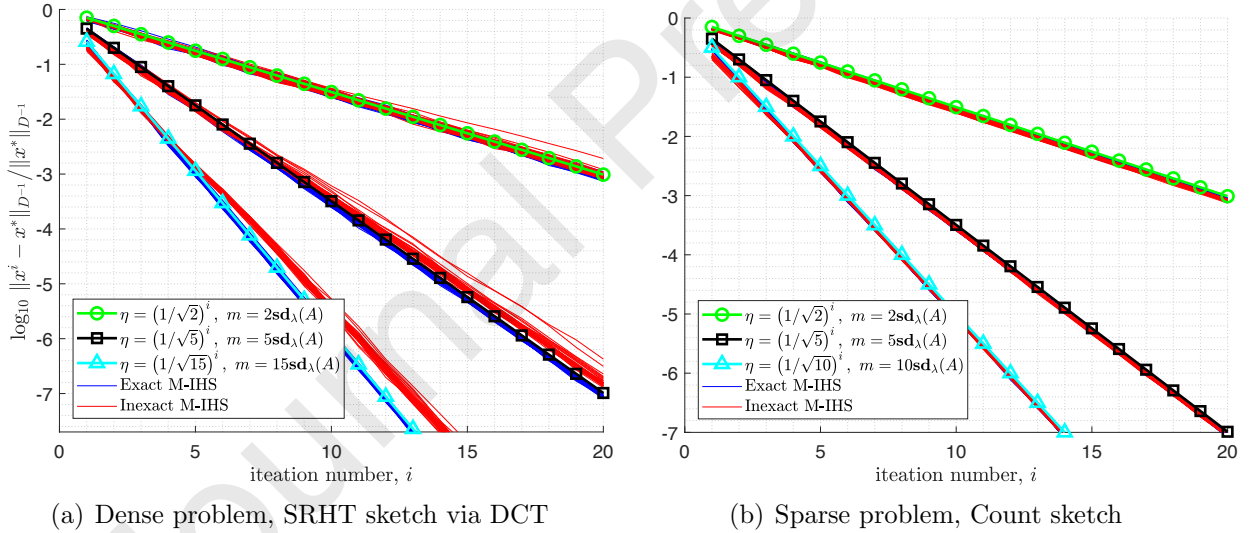


Figure 1: Comparison of the theoretical rate given in Corollary 4 and the empirical convergence rate. The lines with different markers show the theoretical convergence rate for different sketch sizes. Both the *exact* and the *inexact* (given in Algorithm 1) versions of the M-IHS were run 32 times and the result of each run is plotted as a separate line. Except for a small degradation in the dense case, setting the forcing term to a small constant such as $\epsilon_{sub} = 0.1$ is sufficient for the inexact scheme to achieve the same rate as the exact version in these experiments.

Although the MPL provides bounds for the singular values of the sketching matrix S

in the asymptotic regime, i.e., as $m \rightarrow \infty$; these bounds become very good estimators of the actual bounds when m takes finite values as demonstrated in Figure 1. In Figure 1(a), $A \in \mathbb{R}^{32768 \times 1000}$ with $\kappa(A) = 10^8$ was generated as described in Section 4.1. In Figure 1(b), $A \in \mathbb{R}^{24336 \times 1296}$ was generated by using *sprand* command of MATLAB. We first created a sparse matrix with size $\tilde{A} \in \mathbb{R}^{20 \times 6}$ and sparsity of 15%, then the final form was obtained by taking $A = \tilde{A}^{\otimes 4}$ and deleting the all-zero rows. The final form of A has a sparsity ratio of 0.1% and the condition number of $\kappa(A) = 10^7$. The noise level was set to 1% and the regularization parameter λ that minimizes the error $\|x_0 - x(\lambda)\|_2$ was used in both experiments. The resulting statistical dimensions were 119 and 410, respectively. The rate of $\sqrt{\beta}$ in Corollary 4 creates a remarkable fit to the numerical convergence rate of the M-IHS variants when the momentum parameters given in Corollary 4 are used even for the Tikhonov regularization. This is because the sigmoid-like filtering coefficients in the Tikhonov regularization can be thought of as the smoothed version of the binary coefficients in the TSVD solution and therefore the binary coefficients constitute a good approximation for the filtering coefficients of the Tikhonov regularization.

Remark 2. The momentum parameters given in Corollary 4 maximizes the convergence rate when the statistical dimension is known. If $\mathbf{sd}_\lambda(A)$ is overestimated and thus β is chosen larger than the ratio $\mathbf{sd}_\lambda(A)/m$ and $\alpha = (1 - \beta)^2$, then the convergence rate is still $\sqrt{\beta}$ since all modes of the dynamical system driven by the characteristic equation given in (12) will be still in the under-damped regime. In Section 3.3, we propose an empirical algorithm to estimate $\mathbf{sd}_\lambda(A)$.

3.1. Efficient M-IHS sub-solvers

In practice, the M-IHS and the Dual M-IHS eliminate the dominant term $O(md^2)$ in the complexity expression of well known solvers such as the Blendenpik and the LSRN by approximately solving the lower dimensional linear systems in (6) and (9) avoiding matrix decompositions or inversions. This *inexact sub-solver* approach provides a trade-off opportunity between the computational complexity and the convergence rate, that is highly desirable in very large dimensional problems. Unfortunately, such a trade-off is not possible for the Blendenpik and the LSRN techniques which require a full matrix decomposition of the sketched matrix. Inexact sub-solvers have been known to be a good heuristic way to create this trade-off and they are widely used in the algorithms that are based on the Newton method to solve the large scale normal equations [45]. In these inexact (or truncated) Newton methods, inner iterations are terminated at the moment that the relative residual error is lower than an iteration-dependent threshold, named as the *forcing terms* [46]. In the literature, there are various techniques to choose these forcing terms that guarantee a global convergence [47], but the number of iterations suggested by these techniques are significantly higher than the total number of iterations used in practice. Therefore, in this work the heuristic constant threshold ϵ_{sub} , that checks the relative residual error of the linear system, is used [48].

Efficient but approximate solutions to the sub-problems in (6) and (9) can be obtained by Krylov subspace based first-order methods. However, LSQR-like solvers that are adapted

for the normal equations would require computations of 4 matrix-vector multiplications per iteration. On the other hand, due to the explicit calculation of $(SA)^T(SA)z$, the symmetric CG, that would require only 2 matrix-vector multiplications, might be unstable for the ill-conditioned problems [35]. Therefore, in Section 3.5, we propose a stable sub-solver, referred to as **AAb_Solver**, which is particularly designed for the problems in the form of $A^T Ax = b$. The **AAb_Solver** is based on the Golub Kahan Bidiagonalization and it uses a similar approach that the LSQR uses on the LS problem. The inexact versions of the **M-IHS** and **Dual M-IHS** that use **AAb_Solver** are given in Algorithm 1 and Algorithm 2, where **RP_fun** represents the function that generates the desired sketched matrix such that $\mathbb{E}[S^T S] = I_m$ whose implementation details can be found in the relevant references in Lemma 3. Number of operations required at each step is stated at the right most column of the algorithms, where $C(\cdot)$ represents the complexity of constructing the sketching matrix as given in Lemma 3. Setting the forcing term ϵ_{sub} , for instance, to 0.1 for all iterations is enough for the inexact **M-IHS** variants to converge at the same rate $\sqrt{\beta}$ as the exact versions as demonstrated in Figure 1.

Algorithm 1 **M-IHS** (for $n \geq d$)

1: <i>Input</i> : $A, b, m, \lambda, N, x^1, \mathbf{sd}_\lambda(A), \epsilon_{sub}$	<i>complexity</i>
2: $SA = \mathbf{RP_fun}(A, m)$	$C(m, n, d)$
3: $\beta = \mathbf{sd}_\lambda(A)/m$	$O(1)$
4: $\alpha = (1 - \beta)^2$	$O(1)$
5: for $i = 1 : N$ do	
6: $g^i = A^T(b - Ax^i) - \lambda x^i$	$4nd + 3d$
7: $\Delta x^i = \mathbf{AAb_Solver}(SA, g^i, \lambda, \epsilon_{sub})$	$O(md)$
8: $x^{i+1} = x^i + \alpha \Delta x^i + \beta(x^i - x^{i-1})$	$5d$
9: end for	

3.2. Two-stage sketching for the M-IHS variants

Lemma 3 suggests that if the statistical dimension is several times smaller than the dimensions of A , then for the **M-IHS** techniques, it is possible to choose a substantially smaller sketch size than $\min(n, d)$. If this is the case, then the quadratic objective functions in (6) and (9) become strongly under-determined problems, which makes it possible to approximate the Hessian of the objective functions one more time by taking their convex dual as it has been done in the **Dual M-IHS**. This approach is similar to the approach where the problems in (6) and (9) are approximately solved by using the **AAb_Solver**, with an additional dimension reduction. As a result of two Hessian sketching, the linear sub-problem whose dimensions are reduced from both sides can be efficiently solved by the **AAb_Solver** for a pre-determined tolerance as before. For the details of this two-step approach, consider

Algorithm 2 Dual M-IHS (for $n \leq d$)

1: <i>Input:</i> $A, b, m, \lambda, N, \mathbf{sd}_\lambda(A), \epsilon_{sub}$	<i>complexity</i>
2: $SA^T = \text{RP_fun}(A^T, m)$	$C(m, n, d)$
3: $\beta = \mathbf{sd}_\lambda(A)/m$	$O(1)$
4: $\alpha = (1 - \beta)^2$	$O(1)$
5: $\nu^1 = 0$	$O(1)$
6: for $i = 1 : N$ do	
7: $g^i = b - AA^T \nu^i - \lambda \nu^i$	$4nd + 3n$
8: $\Delta \nu^i = \text{AAb_Solver}(SA^T, g^i, \lambda, \epsilon_{sub})$	$O(mn)$
9: $\nu^{i+1} = \nu^i + \alpha \Delta \nu^i + \beta(\nu^i - \nu^{i-1})$	$5n$
10: end for	
11: $x^{N+1} = A^T \nu^{N+1}$	$2nd$

the following dual of the sub-problem in (6)

$$z^{i,*} = \underset{z \in \mathbb{R}^m}{\operatorname{argmin}} \underbrace{\frac{1}{2} \|A^T S^T z + \nabla f(x^i)\|_2^2 + \frac{\lambda}{2} \|z\|_2^2}_{h(z, x^i)}, \quad (17)$$

where $\nabla_z h(z, x) = SA(A^T S^T z + \nabla f(x)) + \lambda z$. Problem (17) is a strongly over-determined if $m \ll \min(n, d)$. Hence, it can be approximately solved by the M-IHS updates as

$$\begin{aligned} \Delta z^{i,j} &= \underset{z \in \mathbb{R}^m}{\operatorname{argmin}} \|W A^T S^T z\|_2^2 + \lambda \|z\|_2^2 + 2 \langle \nabla_z h(z^{i,j}, x^i), z \rangle, \\ z^{i,j+1} &= z^{i,j} + \alpha_2 \Delta z^{i,j} + \beta_2 (z^{i,j} - z^{i,j-1}). \end{aligned} \quad (18)$$

After M iterations, the solution of (6) can be recovered by using the relation in (8) as $\Delta x^i = (\nabla f(x^i) - A^T S^T z^{i,M})/\lambda$. The same strategy can be applied on the sub-problem in (9) by replacing SA with SA^T and $\nabla f(x^i)$ with $\nabla g(\nu^i)$. The resulting algorithms, referred to as **Primal Dual M-IHS**, are given in Algorithm 3 and Algorithm 4, respectively.

The primal-dual idea presented here is first suggested by Zhang et al. in [28]. They used the A-IHS technique to solve the sub-problems that arise during the iterations of the Accelerated Iterative Dual Random Projection (A-IDRP) which is a dual version of the A-IHS. However, since both of the A-IHS and the A-IDRP are based on the CG technique, the convergence rate of the proposed A-IHS, A-IDRP and the primal dual algorithm called as Accelerated Iterative Primal Dual Sketch (A-IPDS) are all degraded in the LS problems with high condition numbers due to the instability issue of the symmetric CG technique [35]. Even if the regularization is used, still the performance of the solvers proposed in [28] are considerably deteriorated compared to the other randomized preconditioning techniques as shown in Section 4. Further, applying the preconditioning idea of IHS to the stable techniques such as the LSQR that are adapted for the LS problem is not so efficient as the

Algorithm 3 Primal Dual M-IHS (for $n \leq d$)

1: <i>Input:</i> $A, b, m_1, m_2, \lambda, N, M, \mathbf{sd}_\lambda(A), \epsilon_{sub}$	<i>complexity</i>
2: $SA^T = \text{RP_fun}(A^T, m_1)$	$C(m_1, n, d)$
3: $WAS^T = \text{RP_fun}(SA^T, m_2)$	$C(m_1, m_2, n)$
4: $\beta_\ell = \mathbf{sd}_\lambda(A)/m_\ell, \quad \ell = 1, 2$	$O(1)$
5: $\alpha_\ell = (1 - \beta_\ell)^2, \quad \ell = 1, 2$	$O(1)$
6: $\nu^1 = 0, z^{1,1} = 0$	$O(1)$
7: for $i = 1 : N$ do	
8: $b^i = b - AA^T \nu^i - \lambda \nu^i$	$4nd + 3n$
9: for $j = 1 : M$ do	
10: $g^{i,j} = SA^T(b^i - AS^T z^{i,j}) - \lambda z^{i,j}$	$4nm_1 + 3m_1$
11: $\Delta z^{i,j} = \text{AAb_Solver}(WAS^T, g^{i,j}, \lambda, \epsilon_{sub})$	$O(m_1 m_2)$
12: $z^{i,j+1} = z^{i,j} + \alpha_2 \Delta z^{i,j} + \beta_2 (z^{i,j} - z^{i,j-1})$	$5m_1$
13: end for	
14: $\Delta \nu^i = (b^i - AS^T z^{i,M+1})/\lambda, \quad z^{i+1,1} = z^{i,M+1}$	$2nm_1 + 2n$
15: $\nu^{i+1} = \nu^i + \alpha_1 \Delta \nu^i + \beta_1 (\nu^i - \nu^{i-1})$	$5n$
16: end for	
17: $x^{N+1} = A^T \nu^{N+1}$	$2nd$

Algorithm 4 Primal Dual M-IHS (for $n \geq d$)

1: Input: $A, b, m_1, m_2, N, M, \lambda, x^1, \mathbf{sd}_\lambda(A), \epsilon_{sub}$	<i>complexity</i>
2: $SA = \text{RP_fun}(A, m_1)$	$C(m_1, n, d)$
3: $WAS^T = \text{RP_fun}((SA)^T, m_2)$	$C(m_1, m_2, d)$
4: $\beta_\ell = \mathbf{sd}_\lambda(A)/m_\ell, \quad \ell = 1, 2$	$O(1)$
5: $\alpha_\ell = (1 - \beta_\ell)^2, \quad \ell = 1, 2$	$O(1)$
6: $z^{1,1} = z^{1,0} = 0$	$O(1)$
7: for $i = 1 : N$ do	
8: $b^i = A^T(b - Ax^i) - \lambda x^i$	$4nd + 3d$
9: for $j = 1 : M$ do	
10: $g^{i,j} = SA(b^i - A^T S^T z^{i,j}) - \lambda z^{i,j}$	$4dm_1 + 3m_1$
11: $\Delta z^{i,j} = \text{AAb_Solver}(WAS^T, g^{i,j}, \lambda, \epsilon_{sub})$	$O(m_1 m_2)$
12: $z^{i,j+1} = z^{i,j} + \alpha_2 \Delta z^{i,j} + \beta_2 (z^{i,j} - z^{i,j-1})$	$5m_1$
13: end for	
14: $\Delta x^i = (b^i - A^T S^T z^{i,M+1})/\lambda, \quad z^{i+1,1} = z^{i,M+1}$	$2dm_1 + 2d$
15: $x^{i+1} = x^i + \alpha_1 \Delta x^i + \beta_1 (x^i - x^{i-1})$	$5d$
16: end for	

M-IHS variants, because they require two preconditioning systems to be solved per iteration.

The computational saving when we apply a second dimension reduction as in the **Primal Dual M-IHS** may not be significant due to the second gradient computations in *Line 10* of the given algorithms, but the lower dimensional sub-problems that we obtain at the end of the second sketching can be used to estimate several parameters including the regularization parameter itself, if it is unknown. As we have shown in [10], such two-stage sketching approach is particularly effective to estimate the unknown regularization parameter when the coefficient matrix is square or close to be square.

The **Primal Dual M-IHS** techniques are extensions of the inexact schemes. Therefore, their convergence rates depend on their forcing terms that are used to stop the inner iterations [47]. In [28], an upper bound for the error of the primal dual updates is proposed. However, as detailed in Section Appendix A, there are several inaccuracies in the development of the bound. Therefore, finding a provably valid lower bound on the number of inner loop iterations, that guarantee a certain rate of convergence at the main loop, is still an open problem for the primal dual algorithms.

3.3. Estimation of the statistical dimension

The statistical dimension $\mathbf{sd}_\lambda(A)$ in Algorithms 1, 2, 3 and 4 can be estimated by using a Hutchinson-like randomized trace estimator [49]. Alternatively, $\mathbf{sd}_\lambda(A)$ can be estimated by using the algorithm proposed in [23] within a constant factor in $\mathbf{nnz}(A)$ time with a constant probability, if $\mathbf{sd}_\lambda(A) \leq \xi$ where:

$$\xi = \min\{n, d, \lfloor (n + d)^{1/3} / \text{poly}(\log(n + d)) \rfloor\}.$$

However, due to the third order root and the division by typically higher than a sixth order polynomial, ξ becomes very small and the proposed algorithm in [23] can only be used when the singular values of A decay severely/exponentially. Therefore, we preferred to use the heuristic trace estimator in Algorithm 5, where the input matrix SA can be replaced with SA^T or even with $WA^T S^T$ and WAS^T according to the requirements of the algorithm used. Any estimator in [49] can be substituted for the Hutchinson estimator and the number of samples T can be chosen accordingly. In the conducted experiments with various singular value profiles, a small sample size such as 2 and $\epsilon_{tr} = 0.5$ was sufficient to obtain satisfactory estimates for $\mathbf{sd}_\lambda(A)$ used in Corollary 4. Note that, as long as $\mathbf{sd}_\lambda(A)$ is overestimated, the convergence rates of the proposed algorithms will be strictly controlled by β as in Corollary 4.

3.4. Complexity analyses of the proposed algorithms

The iterations of both the exact and inexact **M-IHS** and **Dual M-IHS** consist of 4 stages with the computational complexities given in Table 1 where T is the number of samples used in Hutchinson-like estimators, $\kappa(\lambda)$ is the condition number of $A^T A + \lambda I$, $r = \min(n, d)$, and $C(\cdot)$ is the complexity of generating the sketched matrix which is noted in Lemma 3. For the proposed techniques, sketch size m can be chosen proportional to the statistical dimension $\mathbf{sd}_\lambda(A)$ which is always smaller than r . We assumed that the sub-problems in (6) and (9) are solved by using the QR decomposition for the exact schemes. The tolerance

Algorithm 5 Inexact Hutchinson Trace Estimator

1: Input: $SA \in \mathbb{R}^{m \times d}$, λ , T , ϵ_{tr}	<i>complexity</i>
2: $v^\ell = \{-1, +1\}^d$, $\ell = 1, \dots, T$	$O(Td)$
3: $\tau = 0$	$O(1)$
4: for $i = 1:T$ do	
5: $z^i = \text{AAb_Solver}(SA, v^i, \lambda, \epsilon_{tr})$	$O(md)$
6: $\tau = \tau + \lambda \langle v^i, z^i \rangle$	$2d$
7: end for	
8: Output: $\widehat{\text{sd}}_\lambda = d - \tau/T$	$O(1)$

parameters ϵ_{tr} and ϵ_{sub} , that are used to terminate the sub-solvers, are typically chosen around 0.5 or 0.1 as noted in Sections 3.1 and 3.3. The proposed techniques provide, in

Table 1: Computational complexity of each stage in the M-IHS techniques

Stage	Exact schemes	Inexact schemes
generation of SA or SA^T	$C(n, d, \text{sd}_\lambda(A))$	$C(n, d, \text{sd}_\lambda(A))$
QR (R – factor only)	$O(r \text{sd}_\lambda(A)^2)$	<i>N.A.</i>
$\text{sd}_\lambda(A)$ estimation	$O(T \text{sd}_\lambda(A)^2)$	$O(T \sqrt{\kappa(\lambda)} \log(\epsilon_{tr}^{-1}) \text{sd}_\lambda(A)r)$
1 iteration	$O(nd + \text{sd}_\lambda(A)^2)$	$O(nd + \sqrt{\kappa(\lambda)} \log(\epsilon_{sub}^{-1}) \text{sd}_\lambda(A)r)$

addition to several others as described earlier, two major computational advantages over the current randomized preconditioning solvers: the first is the capability to use the sketch sizes that are smaller than r , and the second is the ability of avoiding the complexity of the QR step. In the inexact M-IHS variants, the cubic complexity $O(\min(mr^2, rm^2))$ of matrix decomposition or inversion are avoided by solving sub-problems in each iteration via Krylov subspace-based solvers. Although solving these sub-problems iteratively incur a complexity of $O(\sqrt{\kappa(\lambda)}mr)$, the overall complexity can be reduced significantly in large scale problems where the dimensions n, d and hence r, m can be considerably higher than the condition number $\kappa(\lambda)$. For the applications where m grows larger, this saving becomes even more critical as shown in Section 4.5. The memory space required by all the above techniques is $O(\text{sd}_\lambda(A)r)$. In Table 1, it is assumed that A is a dense matrix. For those cases where A is sparse, further computational savings can be achieved for the proposed techniques. In the following Table 2, the overall computational complexities of the algorithms to obtain an η -optimal solution approximation are given for both dense and sparse A 's. In Table 2³, we assumed that the SRHT sketch matrices are used for a dense A while the sparse subspace embeddings with s nonzero entry in each column are used for a sparse A with $\text{nnz}(A)$ number of nonzero entries. Depending on the type of choice of the sketching used, the complexity of the proposed techniques vary significantly. For dense coefficient matrices while the SRHT

³Note that for sparse A , the QR algorithm that is based on Givens rotation does not provide significant complexity reduction in theory since sequentially applied Givens rotation causes fill-in in the data matrix.

Table 2: Overall computational complexity against baselines

	Dense A	Sparse A
Inex. M-IHS	$O\left(nd \log(\mathbf{sd}_\lambda(A)) + \log(\eta^{-1}) \left(nd + \sqrt{\kappa(\lambda)} \mathbf{sd}_\lambda(A)r\right)\right)$	$O\left(\text{nnz}(A) + \log(\eta^{-1}) \left(\text{nnz}(A) + \sqrt{\kappa(\lambda)} \mathbf{sd}_\lambda(A)r\right)\right)$
Exact M-IHS	$O\left(nd \log(\mathbf{sd}_\lambda(A)) + r \mathbf{sd}_\lambda(A)^2 + \log(\eta^{-1}) nd\right)$	$O\left(\text{nnz}(A) + r \mathbf{sd}_\lambda(A)^2 + \log(\eta^{-1}) \text{nnz}(A)\right)$
Blendenpik	$O\left(nd \log(r) + r^3 + \log(\eta^{-1}) nd\right)$	$O\left(nd \log(r) + r^3 + \log(\eta^{-1}) \text{nnz}(A)\right)$
LSQR/CG	$O\left(\sqrt{\kappa(\lambda)} \log(\eta^{-1}) nd\right)$	$O\left(\sqrt{\kappa(\lambda)} \log(\eta^{-1}) \text{nnz}(A)\right)$
QR	$O(ndr)$	$O(ndr)$

matrices has lower run time in sequential environments, Gaussian matrices would be more efficient in parallel computing. If the coefficient matrix is sparse, then the data oblivious sketching types such as OSNAP or CountSketch matrices would be effective choices with run time of $O(s \text{nnz}(A))$ where s is proportional to $\log(\mathbf{sd}_\lambda(A))$ as noted in Lemma 3. The proposed techniques can be still used even if the coefficient matrix is an operator, in this case Gaussian or sparse embeddings can be utilized. If the coefficient matrix is sparse or an operator that allows fast matrix-vector computations, then both the exact and inexact schemes are automatically sped up due to the saving in the gradient computation. For instance, in the sparse case, complexity of the gradient computation is reduced from $O(nd)$ to $O(\text{nnz}(A))$ as shown in Table 2.

In a similar manner, the complexity of each stage in the **Primal Dual M-IHS** variants is given in Table 3. Here, M denote the number of inner iterations, and both sketch size m_1 and

Table 3: Computational complexity of each stage in the Primal Dual M-IHS techniques

Stage	Exact schemes	Inexact schemes
generation of WAS^T	$C(n, d, \mathbf{sd}_\lambda(A)) + C(r, \mathbf{sd}_\lambda(A), \mathbf{sd}_\lambda(A))$	$C(n, d, \mathbf{sd}_\lambda(A)) + C(r, \mathbf{sd}_\lambda(A), \mathbf{sd}_\lambda(A))$
QR or SVD	$O(\mathbf{sd}_\lambda^3(A))$	<i>N.A.</i>
$\mathbf{sd}_\lambda(A)$ estimation	$O(T \mathbf{sd}_\lambda^2(A))$	$O\left(T \sqrt{\kappa(\lambda)} \log(\epsilon_{tr}^{-1}) \mathbf{sd}_\lambda^2(A)\right)$
1 outer iteration	$O\left(nd + M(r \mathbf{sd}_\lambda(A) + \mathbf{sd}_\lambda^2(A))\right)$	$O\left(nd + M \mathbf{sd}_\lambda(A) \left(r + \sqrt{\kappa(\lambda)} \log(\epsilon_{sub}^{-1}) \mathbf{sd}_\lambda(A)\right)\right)$

m_2 can be chosen proportional to the statistical dimension $\mathbf{sd}_\lambda(A)$, e.g., $m_1 = m_2 = 2\mathbf{sd}_\lambda(A)$ as demonstrated in Figure 5. Unless $M \mathbf{sd}_\lambda(A) \ll r$, the **Primal Dual M-IHS** does not provide significant saving over the **M-IHS** or the **Dual M-IHS**. However, when n and d scale similar and the ratio $\mathbf{sd}_\lambda(A)/r$ is very small, if the decomposition of the sketched matrix is required for parameter estimation purpose as discussed earlier (see Chapter 4 of [10]), then due to the decomposition of $m_2 \times m_1$ -dimensional doubly sketched matrix, the **Primal Dual M-IHS** variants require far fewer operations than any exact schemes which need to compute the decomposition of $r \times m_1$ dimensional sketched matrix. Such conditions are prevalent, for example, in image de-blurring or seismic travel-time tomography problems [50]. The memory space required by **Primal Dual M-IHS** techniques is $O(\mathbf{sd}_\lambda(A)r + \mathbf{sd}_\lambda^2(A))$.

3.5. A solver for linear systems in the form of $A^T A x = b$

The linear sub-problems in the form of $(A^T A + \lambda I)x = b$, whose solutions are required by all four of the proposed **M-IHS** variants, can be approximately solved by using the *bidiag2* procedure described in [35], which produces an upper bidiagonal matrix as:

$$P_k^T A V_k = R_k = \begin{bmatrix} \rho_1 & \theta_2 & & \\ & \ddots & \ddots & \\ & & \rho_{k-1} & \theta_k \\ & & & \rho_k \end{bmatrix} \in \mathbb{R}^{k \times k},$$

where $P_k \in \mathbb{R}^{n \times k}$, $V_k \in \mathbb{R}^{d \times k}$ and $P_k^T P_k = V_k^T V_k = I_k$. The upper bidiagonal decomposition R_k is computed by using the Lanczos-like three term recurrence:

$$\begin{aligned} AV_k = P_k R_k & \implies Av^1 = \rho_1 p^1, \\ A^T P_k = V_k R_k^T + \theta_{k+1} v^{k+1} e_k^T & \implies A^T p^j = \rho_j v^j + \theta_{j+1} v^{j+1} \quad j \leq k, \\ & \implies Av^j = \theta_j p^{j-1} + \rho_j p^j, \quad j \leq k, \end{aligned}$$

where θ_j 's and ρ_j 's are chosen so that $\|v^j\|_2 = \|p^j\|_2 = 1$, respectively. Note that P_k and V_k are not needed to be orthogonal in **AAb_Solver**, therefore we do not need any reorthogonalization steps. Unlike the LSQR, we choose $\theta_1 v^1 = b$ with $\theta_1 = \|b\|_2$ so that the columns of the matrix V_k constitute an orthonormal basis for the k -th order Krylov Subspace:

$$\text{span}\{v^1, \dots, v^k\} = \mathcal{K}_k(A^T A, b) = \mathcal{K}_k(A^T A + \mu I_d, b), \quad \forall \mu \in \mathbb{R}_+.$$

Since the Krylov Subspace is invariant under a constant shift, regularization does not affect this property. In the k -th iteration of the proposed **AAb_Solver**, let the solution estimate of the linear system be $x^k = V_k y^k$ for some vector $y^k \in \mathbb{R}^k$, i.e., $x^k \in \mathcal{K}_k(A^T A, b)$, then we have $(A^T A + \lambda I_d)V_k y^k = b$ which implies

$$\bar{R}_k y^k = \bar{R}_k^{-T} V_k^T b \stackrel{(a)}{=} \theta_1 \bar{R}_k^{-T} e_1,$$

where (a) is due to the choice of v^1 and \bar{R}_k is obtained by applying a sequence of Givens rotation on $[R_k^T \sqrt{\lambda} I_k]^T$ in order to eliminate the sub-diagonal elements due to the regularization [51]. One instance of this elimination procedure is

$$\begin{bmatrix} \bar{\rho}_k & \theta_{k+1} \\ 0 & \rho_{k+1} \\ 0 & 0 \\ 0 & \sqrt{\lambda} \end{bmatrix} \rightarrow \begin{bmatrix} \bar{\rho}_k & c_k \theta_{k+1} \\ 0 & \rho_{k+1} \\ 0 & 0 \\ 0 & \bar{\lambda}_{k+1} \end{bmatrix} \rightarrow \begin{bmatrix} \bar{\rho}_k & \bar{\theta}_{k+1} \\ 0 & \bar{\rho}_{k+1} \\ 0 & 0 \\ 0 & 0 \end{bmatrix} \xrightarrow[\text{iteration}]{\text{next}} \begin{bmatrix} \bar{\rho}_{k+1} & \theta_{k+2} \\ 0 & \rho_{k+2} \\ 0 & 0 \\ 0 & \sqrt{\lambda} \end{bmatrix},$$

where $c_k = \rho_k / \bar{\rho}_k$, $s_k = \bar{\lambda}_k / \bar{\rho}_k$, $\bar{\theta}_{k+1} = c_k \theta_{k+1}$, $\bar{\lambda}_{k+1}^2 = \lambda + (s_k \theta_{k+1})^2$ and $\bar{\rho}_{k+1} = \sqrt{\rho_{k+1}^2 + \bar{\lambda}_{k+1}^2}$. Since \bar{R}_k is an upper bidiagonal matrix, the inverse always exists and $f^k := \bar{R}_k^{-T} e_1$ can be

computed analytically as:

$$\phi_1 = \frac{\theta_1}{\bar{\rho}_1} \quad \text{and} \quad \phi_k = -\phi_{k-1} \frac{\bar{\theta}_k}{\bar{\rho}_k} \quad \text{where} \quad f^k = [\phi_1, \dots, \phi_k]^T. \quad (19)$$

Furthermore, the solution at the k -th iteration, $x^k = V_k \bar{R}_k^{-1} f^k$, can be obtain without computing any inversions by using the forward substitution. Define $D_k = V_k \bar{R}_k^{-1}$:

$$[D_{k-1}, d^k] \left\{ \begin{array}{l} \left[\begin{array}{cc} \bar{R}_{k-1} & e_{k-1} \bar{\theta}_k \\ 0 & \bar{\rho}_k \end{array} \right] = [V_{k-1}, v^k] \\ D_{k-1} \bar{R}_{k-1} = V_{k-1} \\ \bar{\theta}_k d^{k-1} + \bar{\rho}_k d^k = v^k \end{array} \right. \left. \begin{array}{l} d^k = (v^k - \bar{\theta}_k d^{k-1}) / \bar{\rho}_k \\ x^k = x^{k-1} + \phi_k d^k, \end{array} \right.$$

and the relative residual error that will be used as a stopping criterion can be found as:

$$\begin{aligned} \|A^T A x^k + \lambda x^k - b\|_2^2 &= \|A^T A V_k y^k + \lambda V_k y^k - b\|_2^2 = \|A^T P_k \bar{R}_k y^k - b\|_2^2 \\ &= \|(V_k \bar{R}_k^T + \theta_{k+1} v^{k+1} e_k^T) \bar{R}_k y^k - b\|_2^2 \\ &\stackrel{(i)}{=} \|\bar{R}_k^T \bar{R}_k y^k - V_k^T b\|_2^2 + \|\theta_{k+1} v^{k+1} e_k^T \bar{R}_k y^k - (I - V_k V_k^T) b\|_2^2 \\ &= |\phi_k \bar{\theta}_{k+1}| = |\phi_{k+1} \bar{\rho}_{k+1}|. \end{aligned}$$

The first norm in (i) is zero since the linear system is always consistent. The second term in the second norm is also zero, since $b \in \text{span}(V_k)$ by the initial choice of $\theta_1 v^1 = b$. By definition, $f^k = \bar{R}_k y^k$ gives the final results. The overall algorithm is given in Algorithm 6. The **AAb_Solver** is also a Krylov subspace method, therefore, it finds the solution in at most $\min(n, d, m)$ iterations in the exact arithmetic, but far fewer number of iterations is sufficient for our purpose.

Efficient solutions for linear systems in the form of $(A + \lambda I)x = b$ for a symmetric matrix or $(A^T A + \lambda I)x = b$ for a rectangular matrix have been well studied subject. In the first case, Lanczos tridiagonalization algorithm can be used for deriving a stable solver [52]. In the second case, which is our main concern, if the lower bidiagonalization processes (*bidiag1* in [35]) is used such as in [53], then a tridiagonal system in the form of $B_k^T B_k y^k = \theta_1 e_1$ must be solved where $B_k \in \mathbb{R}^{(k+1) \times k}$ is a lower bidiagonal matrix. This system can be solved by first eliminating the lower diagonal elements in the tridiagonal matrix $B_k^T B_k$ and then by using forward substitution. However, the condition number of $B_k^T B_k$ is the square of the condition number of B_k , and thus, increases the instability of the operations in the inexact arithmetic. Therefore, in the proposed **AAb_Solver**, we use upper bidiagonalization process to solve a tridiagonal system in the form of $R_k^T R_k y^k = \theta_1 e_1$. The major advantage of this form over the one obtained by lower bidiagonal matrix B_k is that $R_k^{-T} e_1$ can be calculated analytically as in (19). Then the solution y^k can be obtained via forward substitution. In this way, we avoid both squaring the condition number and the elimination process of the lower diagonal entries. As a result, we obtain a solver with better stability properties and with slightly lower computational requirements.

Algorithm 6 AAb.Solver (for problems in the form of $(A^T A + \lambda I)x = b$)

1: Input: $A \in \mathbb{R}^{m \times n}, b, \lambda, \epsilon$	▷	choose ρ and θ to make $\ p\ _2 = \ v\ _2 = 1$	<i>complexity</i>
2: $\theta_1 v = b$			$3n$
3: $\rho p = Av$			$2mn + 3m$
4:			
5: $\bar{\rho} = \sqrt{\rho^2 + \lambda}, \quad c = \rho/\bar{\rho}, \quad s = \sqrt{\lambda/\bar{\rho}}, \quad \phi = \theta_1/\bar{\rho}, \quad t = \infty$			$O(1)$
6: $d = v/\bar{\rho}$			n
7: $x = \phi d$			n
8: while $t \geq \epsilon$ do			
9: $\theta v := A^T p - \rho v$			$2mn + 5n$
10: $\rho p := Av - \theta p$			$2mn + 5m$
11:			
12: $\bar{\lambda}^2 := \lambda + (s\theta)^2, \quad \bar{\theta} = c\theta$			$O(1)$
13: $\bar{\rho} := \sqrt{\rho^2 + \bar{\lambda}^2}, \quad c = \rho/\bar{\rho}, \quad s = \bar{\lambda}/\bar{\rho}$			$O(1)$
14:			
15: $d := (v - \bar{\theta}d)/\bar{\rho}$			$3m$
16: $\phi := -\phi\bar{\theta}/\bar{\rho}$			$O(1)$
17: $x := x + \phi d$			$2n$
18: $t = \phi\bar{\rho} /\theta_1$			$O(1)$
19: end while			

4. Numerical Experiments and Comparisons

We compare the operation counts required by the algorithms to obtain a certain level of accuracy in the solution approximation metric. For a fair comparison, we have implemented all the proposed algorithms in this manuscript as well as those that are used for the comparisons in MATLAB which can be found in the provided in <https://github.com/ibrahimkurban/M-IHS>.

4.1. Experimental setting

The coefficient matrix $A \in \mathbb{R}^{n \times d}$ was generated for various sizes as follows: we first sampled the entries of A from the distribution $\mathcal{N}(1_d, \Gamma)$ where $\Gamma_{ij} = 5 \cdot 0.9^{|i-j|}$ so that the columns are highly correlated with each other. Then by using the SVD, we replaced the singular values with *philips* profile provided in RegTool [54]. We scaled the singular values to set the condition number $\kappa(A)$ to 10^8 and we used the same input signal provided by RegTool. In this way, we have obtained a challenging setup for any first order iterative solvers to compare their performances. In all the experiments, the same setup has been used unless indicated. We counted the number of operations according to Hunger's report [55]. All the reported results have been obtained by averaging over 32 MC simulations.

4.2. Compared methods and their implementation details

We compared the proposed algorithms with the state of the art randomized preconditioning techniques which can reach any level of desired accuracy within a bounded number

of iterations. In the conducted comparison study we used a total of 5 previously proposed techniques that can be briefly described as follows. The Blendenpik uses the R matrix in the QR decomposition of the sketched matrix SA as the preconditioning matrix for the LSQR algorithm just like the method proposed by Rokhlin et al. [26, 25] and it uses Randomized Orthonormal System (ROS) to generate the sketched matrix [22]. The LSRN uses the V matrix in the SVD similar to the Blendenpik. In spite of its high running complexity, for parallelization purposes, the Gaussian sketch matrices are preferred in the LSRN. In addition to the LSQR, also the CS can be preferred in the LSRN as the core solver in distributed computational environments [27]. The IHS uses the sketched Hessian as the preconditioning matrix for the gradient descent. The Accelerated IHS (A-IHS) uses this idea for the CG algorithm in over-determined problems. The dual counter-part of the A-IHS algorithm, A-IDRP, is shown to be faster than the Dual Random Projection algorithm proposed in [4], so we did not include the DRP in the simulations. Additionally, we include a CS variant of the IHS (IHS-CS) to the comparisons: we combined the randomized preconditioning idea of the IHS with the preconditioned CS method [14]. We found the bounds for the eigenvalues in the same way as in the LSRN. We have solved the low dimensional sub-problems required by all the IHS variants by taking the QR decomposition, but for *inexact* schemes, we have used the proposed `AAb_Solver` with a constant forcing term. Although the inexact approach is also applicable for the accelerated algorithms proposed in [28], we did not include them in the simulations since their exact versions are outperformed by the *Exact M-IHS* variants in all settings. Except for the LSRN variants which use Gaussian sketch matrices, we used Discrete Cosine Transform in the ROS for all the compared techniques.

4.3. Linear systems with noiseless measurements

In the first experiment, we did not include noise in the linear system to emphasize the convergence rate that the algorithms can provide in such severely ill posed problems. To make the problem more challenging, for this experiment only we sampled the input vector x_0 from uniform distribution $\text{Uni}(-1, 1)$. In such scenarios, convergence rates of Krylov subspace-based iterative solvers without preconditioning fall to its minimum value since the energy of the input is distributed equally over the range space of A . The obtained results are shown in Figure 2.

Due to high running time of the Gaussian sketches, $O(mnd)$, the LSRN variants require more operations (for the size of the problems considered here approximately 10 times larger) than the others. Due to the lack of inner product calculations, the *M-IHS* requires slightly fewer operations than the Blendenpik, nonetheless, it reaches to the same accuracy with the LSRN-LSQR. The A-IHS algorithm has the worst performance which is expected in the un-regularized problems, since it is adapted on the CG technique that can be unstable for the un-regularized LS problems due to the high condition number [35]. The convergence of the CS-based techniques, both of the IHS and the LSRN variants, are substantially slower than the *M-IHS*, which suggests that the *M-IHS* algorithm can take the CS's place in those applications where parallel computation is an option. A similar comparison of the *M-IHS* with the Accelerated Randomized Kaczmarz (ARK) and the CGLS without preconditioning has been shown in Figure 2 of [30].

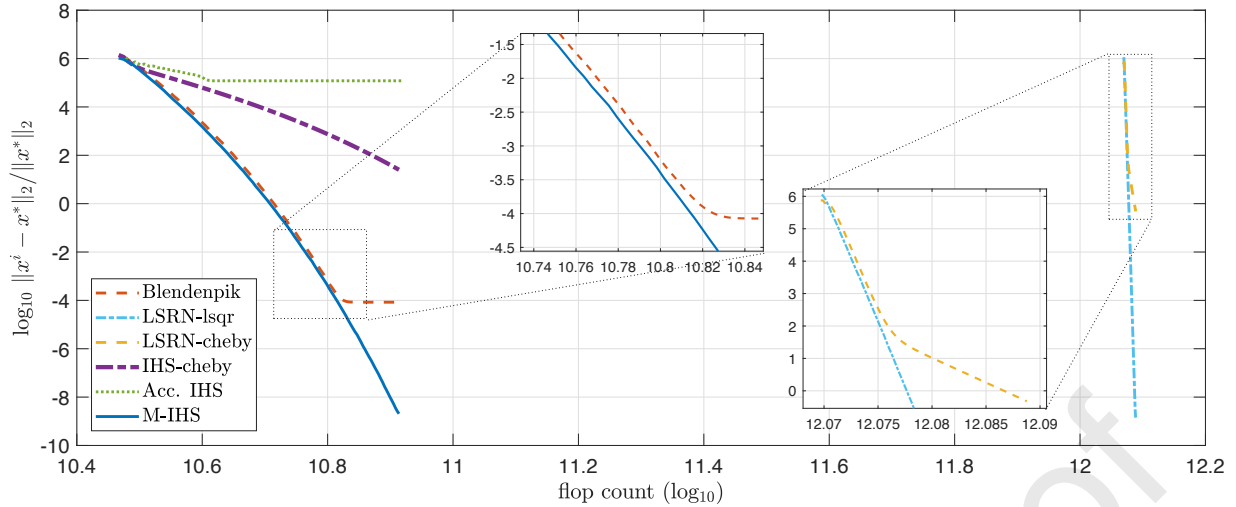


Figure 2: Performance comparison on an un-regularized LS problem with size $2^{16} \times 2000$. In order to compare the convergence rates, number of iterations for all solvers are set to $N = 100$ with the same sketch size: $m = 4000$. According to Corollary 4, we expect the M-IHS to reach an accuracy: $\|x^N - x_0\|_2 \leq \kappa(A)\|x_0\|_2 (1/\sqrt{2})^N = 9 \cdot 10^{-8}$, which closely fits to the observed case.

4.4. Linear systems with noisy measurements

We tested robustness of the methods against noise on regularized LS problems by using an additive i.i.d. Gaussian noise at level $\|\omega\|_2/\|Ax_0\|_2 = 1\%$. For this purpose, the optimal regularization parameter that minimizes the error $\|x^*(\lambda) - x_0\|_2$ is provided to all techniques. Each technique is allowed to conduct a total of 20 iterations. Results for strongly over-determined and strongly under-determined cases can be seen in Figures 3 and 4, respectively.

We used a sketch size of $m = \min(n, d)$ to emphasize the promise of the RP techniques although such sizes are not applicable for the LSRN variants. Even if the sketch size has been increased further, the convergence of the LSRN variants were considerably slower than the others; so we leave out the LSRN variants from the comparison set in the regularized settings. Also, in the regularized setup, the A-IHS and A-IDRP methods are slower than the Blendenpik, IHS-CS and M-IHS variants. Besides, the inexact schemes proposed for the M-IHS and Dual M-IHS require significantly less operations to reach to the same level of accuracy as their exact versions. Although the inexact schemes require approximately 10 times less operations than their exact versions in these setups; the saving gets larger as the sketch size increases as examined in Section 4.5, because while any full decomposition requires $O(mr^2)$ operations, approximately solving the sub-problem requires only $O(mr)$ operations.

As long as the statistical dimension of the problem is small with respect to the dimensions of coefficient matrix A , Lemma 3 implies eligibility of sketch sizes that are smaller than the rank, $m \leq \min(n, d)$. This implication can be verified in Figure 5 on which we showed the performance of the Primal Dual M-IHS techniques.

Here, the *inexact* schemes of the M-IHS and Dual M-IHS use a sketch size $m = 2 \cdot \text{sd}_\lambda(A)$. The primal dual schemes use $m_1 = m_2 = 2 \cdot \text{sd}_\lambda(A)$ except for the Primal Dual M-IHS




Figure 3: Performance comparison on a regularized LS problem ($n \gg d$) with dimensions $(n, d, m, \mathbf{sd}_\lambda(A)) = (2^{16}, 4000, 4000, 443)$. According to Corollary 4, M-IHS is expected to satisfy: $\|x^N - x^*\|_2 \leq \|x^*\|_2 \sqrt{\kappa(A^T A + \lambda I_d)} (\sqrt{443/4000})^N = 6 \cdot 10^{-9}$ which is almost exactly the case. The *Inexact* M-IHS requires significantly fewer operations to reach the same accuracy as others. For example to obtain an ($\eta = 10^{-4}$)-optimal solution approximation, the *Inexact* M-IHS requires approximately 10 times less operations than any techniques that need factorization or inversion of the sketched matrix.

shown as a green curve which uses $m_1 = m_2 = 8 \cdot \mathbf{sd}_\lambda(A)$. All the methods are allowed to conduct $N = 60$ iterations except the **Primal Dual** M-IHS with larger sketch size is allowed to conduct only 20 iterations. The number of inner iterations are restricted by $M = 25$ for all the primal dual schemes. Lastly, a fixed forcing term $\epsilon_{sub} = 0.1$ is used in the **AAb_Solver** for all the inexact schemes. Applying a second dimension reduction may not seem to create significant computational saving, but this approach produces smaller sub-problems than the M-IHS and the **Dual** M-IHS techniques therefore enables estimation of parameters such as λ with far fewer number of operations. Lastly, the **Primal Dual** M-IHS variants have a noticeably higher rate of convergence than the A-IPDS algorithm which is based on the CG technique.

4.5. Scalability to larger size problems

In this section, as the size of the coefficient matrix and the sketch size increase we show that the saving gained by the inexact schemes become critically more important. For this purpose, the algorithms were run on the over-determined problems with size $5 \cdot 10^4 \times \gamma \cdot 500$ where $\gamma \in \{1, 2, 4, 8, 16\}$. The sketch size was chosen as $m = d = \gamma \cdot 500$ and the regularization parameter was set to 0.1453 for all the experiments so that $\mathbf{sd}_\lambda(A) = d/10$ remains the same for all the experiments. The data was generated by using the setup described in Section 4.1. Note that the convergence properties of the proposed techniques depend only on the statistical dimension but not directly to the decay rate of the singular values. To show this, for these experiments, we used *heat* singular value profile that has significantly lower decay rate than the *philips* profile used earlier. The experiments were realized on a desktop with 4Ghz i7-4790K CPU processor and 32Gb RAM. The flop




Figure 4: Performance comparison on a regularized LS problem ($n \ll d$) with dimensions $(n, d, m, \text{sd}_\lambda(A)) = (4000, 2^{16}, 4000, 462)$. The comments in Figure 3 are also valid for this case. The *Inexact* scheme for **Dual M-IHS** is capable of significantly reducing the complexity.

count and wall clock time of the algorithms to reach to an ($\eta = 10^{-4}$)-optimal solution approximation are shown in Figure 6.

As d and m reach thousands, the number of operations required by the exact schemes (Blendepik and M-IHS) becomes larger than 100 times of the operation count required by the inexact scheme. Moreover, the exact schemes need 25 time longer time than the inexact scheme to reach the desired accuracy. Additionally, the operation counts and elapsed time in each stage of the algorithms can be seen in Figure 7 which shows that even the cost of the decomposition applied on the sketched matrices reaches to prohibitive levels for large scale problems. Hence the use of solvers such as M-IHS variants that allow inexact schemes is the only practical choice in these regimes. In these experiments, for the estimation of the statistical dimension, we set $T = 2$ and $\epsilon_{tr} = 0.5$. The additional cost of the $\text{sd}_\lambda(A)$ estimation for the proposed M-IHS variants becomes negligibly small when R -factor is utilized; for the inexact schemes, still it has a low cost, around the cost of one M-IHS-inexact iteration, that does not cause an issue unlike a matrix decomposition.

Remark 3. Benchmarking of the exact and inexact schemes by using wall clock time in MATLAB is not a fair comparison because for-loops in the interpreted languages such as MATLAB is well known to be much slower than the loops in compiled languages such as C. Most of the decompositions in MATLAB have C-based implementation with professional use of BLAS operations, while the inexact schemes are based on a for-loop. Therefore, we prefer to rely on the operation counts. However, to give an opinion, in spite of the disadvantages we demonstrate timing as well.

4.6. Effect of the statistical dimension on the performance of the inexact schemes

The inexact schemes become more efficient as the statistical dimension decreases since the sub-problems are solved in less iterations. To show the effect of varying statistical

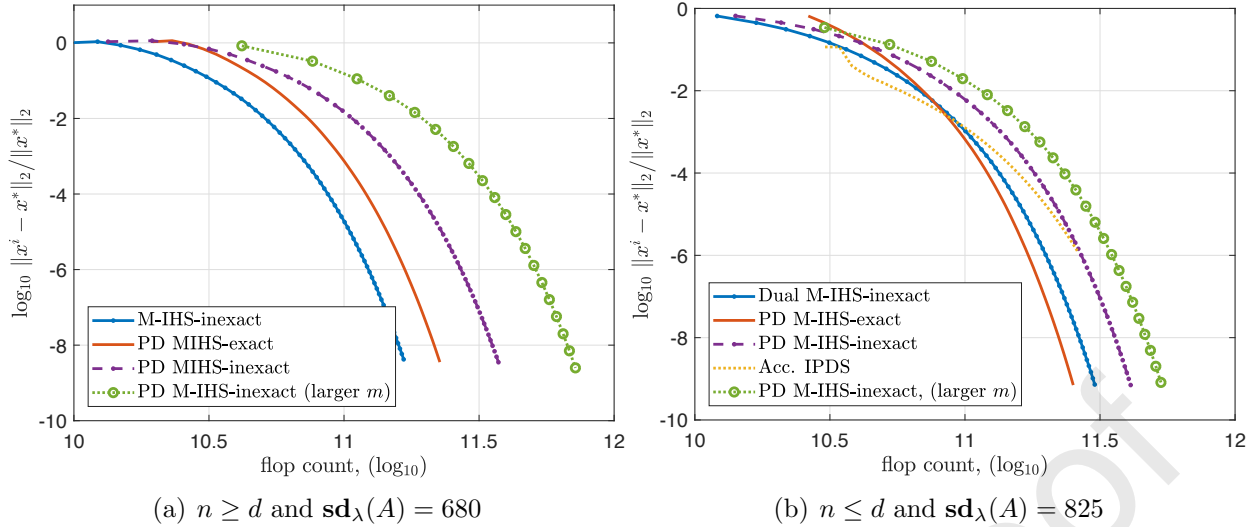


Figure 5: Performance comparison on regularized LS problems with square-like dimensions. The problem dimensions are set to $\max(n, d) = 5 \cdot 10^4$ and $\min(n, d) = 10^4$ with a noise level of 10%. The results verifies two hypothesis: first the sketch size for the M-IHS variants can be chosen proportional to the statistical dimension even if it becomes smaller than the size of the coefficient matrix. Second, the coefficient matrix can be sketched from both sides to reduce computational complexity.

dimension on the complexity of the algorithms, we used over-determined problems with size $5 \cdot 10^4 \times 16 \cdot 10^3$ and varied the regularization parameter to obtain different $\rho = \mathbf{sd}_\lambda(A)/d$ ratios where $\rho \in \{0.5\%, 1\%, 2\%, 5\%, 10\%, 20\%, 50\%\}$. The sketch size was chosen as $m = d$ and *heat* profile was used. As in Section 4.5, the flop count and wall clock time of the algorithms to reach to an ($\eta = 10^{-4}$)-optimal solution approximation for the problems with different statistical dimensions are shown in Figure 8.

Complexity of the exact schemes increases by the increasing statistical dimension since convergence rate $\mathbf{sd}_\lambda(A)/m$ decreases as m remains constant. Complexity of the inexact scheme increase faster since sub-problems require more iterations as the effective range space gets larger. The effect of the increasing statistical dimension over the different stages of the algorithms are shown in Figure 9. For the estimation of $\mathbf{sd}_\lambda(A)$, same parameters $T = 2$ and $\epsilon_{tr} = 0.5$ were used as early. Even for the large ρ ratios, utilizing a sub-solver is still more effective than computing a matrix decomposition.

5. Conclusions

Based on the IHS and the Heavy Ball Acceleration we proposed and analyzed the M-IHS solvers for large scale LS problems. We examined the effect of ℓ_2 norm regularization on the optimal momentum parameters. We obtained lower bounds on the sketch size for several randomized distributions in order to get a pre-determined convergence rate with a constant probability. The bounds suggest that the sketch size can be chosen proportional to the statistical dimension of the problem. Hence, the M-IHS variants can be used for any dimension regimes if the statistical dimension is sufficiently smaller than the dimensions of

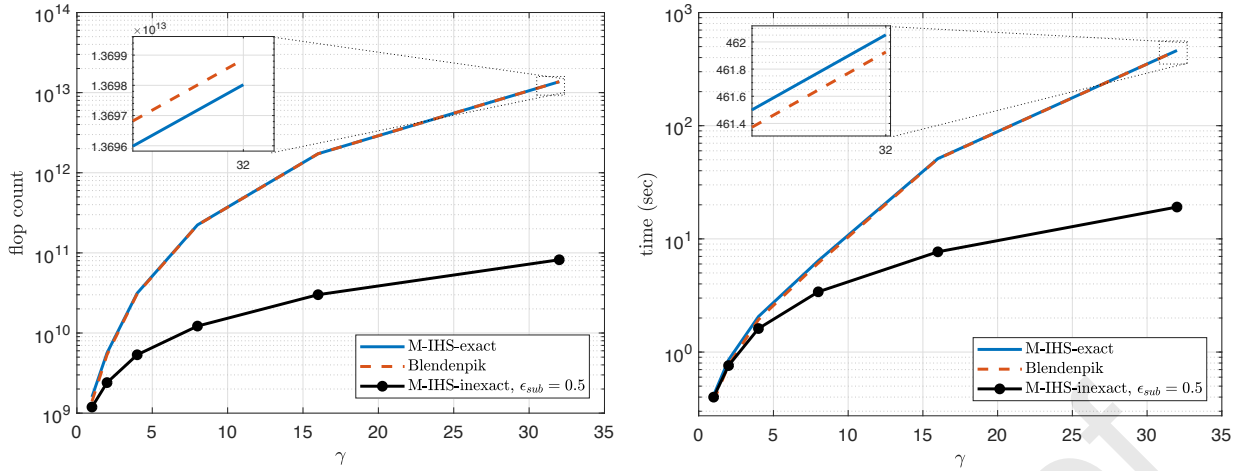


Figure 6: Complexity of the algorithms in terms of operation count and computation time on a set of $5 \cdot 10^4 \times 500 \cdot \gamma$ dimensional over-determined problems with $m = d$ and $\text{sd}_\lambda(A) = d/10$.

the coefficient matrix. We empirically showed the ratio between statistical dimension and the sketch size determines the convergence rate of the proposed M-IHS variants. The main advantage of the proposed M-IHS variants over the state of the art randomized preconditioning techniques such as the Blendenpik, A-IHS and LSRN is their ability to use inexact schemes that avoids matrix decompositions or inversions. As demonstrated in a wide range of numerical experiments, computational saving provided by the proposed solver becomes decidedly significant in large scale problems. Lastly, the proposed M-IHS variants avoid using any inner products in their iterations and they are shown to be faster than CS-based randomized preconditioning algorithms. Therefore the proposed M-IHS variants are strong candidates to be the techniques of choice in parallel or distributed architectures.

Appendix A. Discussion on the Proposed Error Upper Bound for Iterations of Primal Dual Algorithms in [28]

In this appendix, we provide details of a critical discussion on the steps of the derivation that leads to an error upper bound for the iterations of the primal dual algorithms given in [28]. First, we provide a short list of minor issues that can easily be corrected.

1. During the initialization stage in *Line 2* of both *Algorithm 4* in page 4097 and *Algorithm 5* in page 4098, the residual error vector $\mathbf{r}^{(0)}$ must be set to $-\lambda\mathbf{y}$ instead of $-\mathbf{y}$, otherwise iterates of the both of the algorithms diverge from the optimal solution.
2. During the initialization stage in *Line 2* of *Algorithm 7* in 4912, the dual residual error vector $\mathbf{r}_{\text{Dual}}^{(0)}$ must be set to $-\lambda\mathbf{y}$ instead of $-\mathbf{y}$ and during the initialization stage of the inner loop iterations in *Line 15*, the primal residual error vector $\mathbf{r}_p^{(0)}$ must be set to $-\mathbf{R}^T \mathbf{X}^T \mathbf{r}_D^{(t+1)}$; otherwise iterates of the algorithm diverges from the optimal solution. The MATLAB codes provided in the <https://github.com/ibrahimkurban/M-IHSlink> includes these corrections.

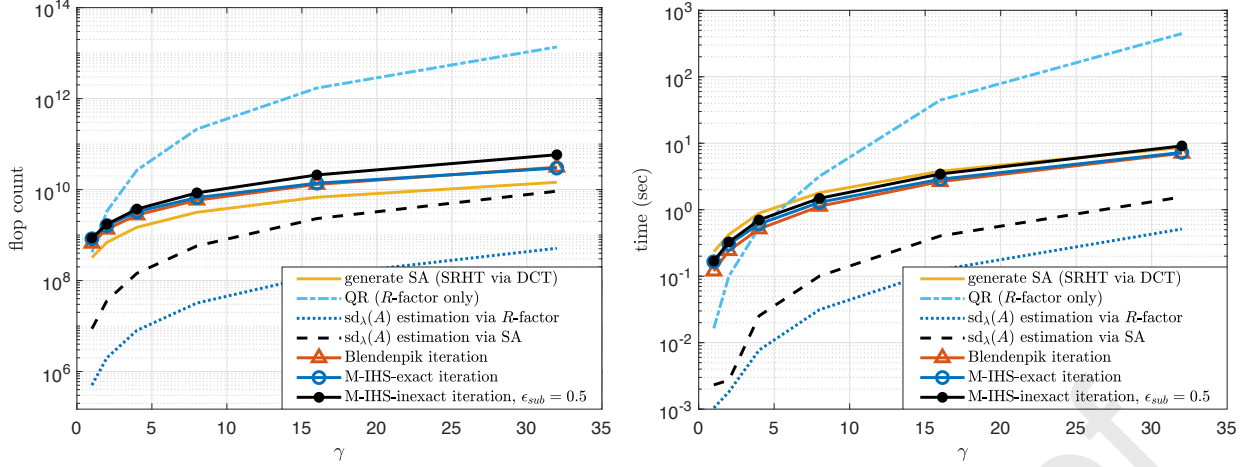


Figure 7: Complexity of the each stage in terms of operation count and computation time on a set of $5 \cdot 10^4 \times 500 \cdot \gamma$ dimensional over-determined problems with $m = d$ and $\text{sd}_\lambda(A) = d/10$. All methods contain SA generation stage. The Blendenpik and M-IHS-exact contain also QR decomposition stage but M-IHS-inexact does not. The M-IHS-exact estimates $\text{sd}_\lambda(A)$ by using the R -factor while the M-IHS-inexact uses directly SA matrix and the `AAb_Solver` as proposed in Algorithm 5. The results show that the matrix decompositions are the main computational bottleneck for the exact schemes in large scale problems where the advantage of the inexact schemes becomes more significant.

In addition to the above mentioned minor issues, there are some major issues as well. Unfortunately, we could not obtain corrective actions on these major issues as we could have done on the minor issues mentioned above. Therefore, a lower bound on the number of inner loop iterations, that guarantee a certain rate of convergence at the main loop, is still an open question for the primal dual algorithms. In the remaining of this appendix, we will provide steps of the derivation presented in [28], along with our critical remarks on their validity.

Consider the A-IHS updates $\widehat{\mathbf{w}}^{t+1} = \widehat{\mathbf{w}}^t + \widehat{\mathbf{u}}^t$. We are going to use exactly the same notation as [28] except for that HS subscript for the A-IHS iterates are omitted. In the primal dual algorithms, instead of exact sequence $\{\widehat{\mathbf{w}}^t\}$, a sequence $\{\widetilde{\mathbf{w}}^t\}$ is obtained due to the approximate minimizers that are used in place of $\widehat{\mathbf{u}}^t$. Consequently while the sequence $\{\widehat{\mathbf{w}}^t\}$ is obtained after t exact iterations of the A-IHS algorithm, sequence $\{\widetilde{\mathbf{w}}^t\}$ is obtained after t primal dual iterations in each of which k inner loop updates are used to approximate $\widehat{\mathbf{u}}^t$'s. The details of the inner and outer loops can be found in Algorithm 7 of [28]. The aim of the *Theorem 9* is to establish an upper bound for $\|\widetilde{\mathbf{w}}^{t+1} - \mathbf{w}^*\|_{\mathbf{X}}$ where \mathbf{w}^* is the true minimizer of the primal objective function. The triangle inequality and the convergence rate of the A-IHS that is established in *Theorem 2* of [28] is used to find an upper bound for this error norm:

$$\|\widetilde{\mathbf{w}}^{t+1} - \mathbf{w}^*\|_{\mathbf{X}} \leq \|\widehat{\mathbf{w}}^{t+1} - \mathbf{w}^*\|_{\mathbf{X}} + \|\widetilde{\mathbf{w}}^{t+1} - \widehat{\mathbf{w}}^{t+1}\|_{\mathbf{X}} \leq \alpha^t \|\mathbf{w}\|_{\mathbf{X}} + \|\widetilde{\mathbf{w}}^{t+1} - \widehat{\mathbf{w}}^{t+1}\|_{\mathbf{X}},$$

where $\alpha = \frac{C_0 \sqrt{\overline{\mathbf{W}^2(\mathbf{X}\mathbb{R}^p \cap \mathcal{S}^{n-1})} \log(1/\delta)}}{1 - C_0 \sqrt{\overline{\mathbf{W}^2(\mathbf{X}\mathbb{R}^p \cap \mathcal{S}^{n-1})} \log(1/\delta)}}$. At this point a new iterate, $\overline{\mathbf{w}}^{t+1}$, is introduced, which is the result of one exact step of the IHS initialized at $\widetilde{\mathbf{w}}^t$. The inner loop iterations at the




Figure 8: Complexity of the algorithms in terms of operation count and computation time on a $5 \cdot 10^4 \times 4 \cdot 10^3$ dimensional problem for different $\rho = \text{sd}_\lambda(A)/d$ ratios.

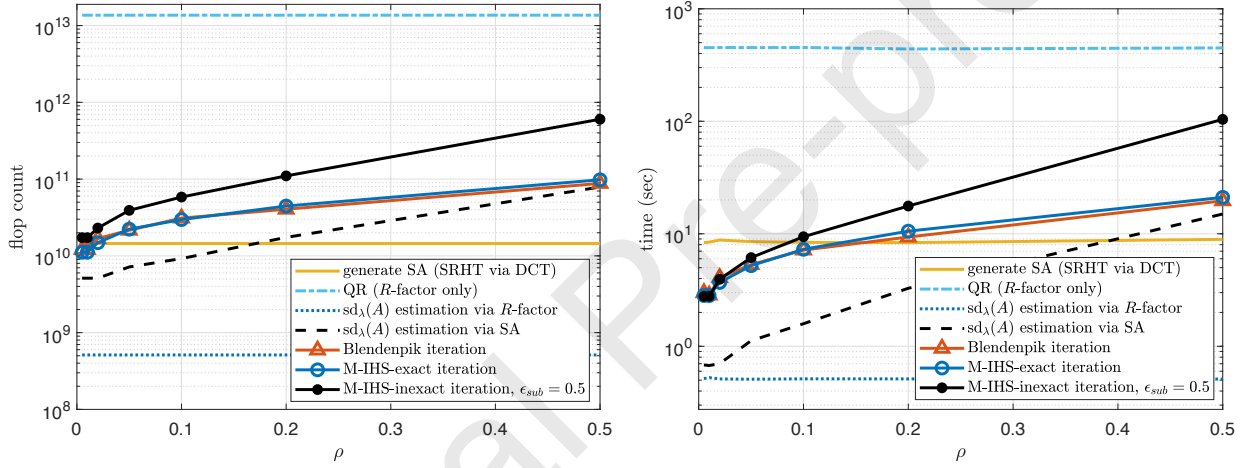


Figure 9: Complexity of each stage in terms of operation count and computation time on a $5 \cdot 10^4 \times 4 \cdot 10^3$ dimensional problem for different $\rho = \text{sd}_\lambda(A)/d$ ratios. Stages of each algorithm is given in Figure 7.

t -th outer (main) loop iteration of the primal dual iterations are expected to converge $\bar{\mathbf{w}}^{t+1}$. Therefore,

$$\begin{aligned}
 \|\tilde{\mathbf{w}}^{t+1} - \hat{\mathbf{w}}^{t+1}\|_{\mathbf{X}} &\leq \|\tilde{\mathbf{w}}^{t+1} - \bar{\mathbf{w}}^{t+1}\|_{\mathbf{X}} + \|\bar{\mathbf{w}}^{t+1} - \hat{\mathbf{w}}^{t+1}\|_{\mathbf{X}}, \\
 \|\tilde{\mathbf{w}}^{t+1} - \bar{\mathbf{w}}^{t+1}\|_{\mathbf{X}} &\leq \lambda_{\max} \left(\frac{\mathbf{X}^T \mathbf{X}}{n} \right) \beta^k \|\bar{\mathbf{w}}^{t+1}\|_2 \leq \lambda_{\max} \left(\frac{\mathbf{X}^T \mathbf{X}}{n} \right) \beta^k (\|\bar{\mathbf{w}}^{t+1} - \mathbf{w}^*\|_2 + \|\mathbf{w}^*\|_2) \\
 &\leq 2\lambda_{\max} \left(\frac{\mathbf{X}^T \mathbf{X}}{n} \right) \beta^k \|\mathbf{w}^*\|_2, \tag{A.1}
 \end{aligned}$$

where $\beta = \frac{C_0 \sqrt{\mathbb{W}^2(\mathbf{X}^T \mathbb{R}^p \cap \mathcal{S}^{p-1}) \log(1/\delta)}}{1 - C_0 \sqrt{\mathbb{W}^2(\mathbf{X} \mathbb{R}^p \cap \mathcal{S}^{p-1}) \log(1/\delta)}}$. The last inequality is not valid unless $\|\bar{\mathbf{w}}^{t+1} - \mathbf{w}^*\|_2 \leq \|\mathbf{w}^*\|_2$. However, particularly during the initial phases of the main iterations this condition can be violated. Therefore, this step of the proof requires a major revision. Assuming that such revision is possible, up to this point, the following is obtained:

$$\|\tilde{\mathbf{w}}^{t+1} - \mathbf{w}^*\|_{\mathbf{X}} \leq \alpha^t \|\mathbf{w}\|_{\mathbf{X}} + 2\lambda_{\max} \left(\frac{\mathbf{X}^T \mathbf{X}}{n} \right) \beta^k \|\mathbf{w}^*\|_2 + \|\bar{\mathbf{w}}^{t+1} - \hat{\mathbf{w}}^{t+1}\|_{\mathbf{X}}. \quad (\text{A.2})$$

To proceed for the final form of the upper bound, the following upper bound on the last term of (A.2) is given in [28]:

$$\|\bar{\mathbf{w}}^{t+1} - \hat{\mathbf{w}}^{t+1}\|_{\mathbf{X}} \leq \|\tilde{\mathbf{H}}^{-1}\|_2 \|\tilde{\mathbf{H}} - \mathbf{H}\|_2 \|\tilde{\mathbf{w}}^t - \hat{\mathbf{w}}^t\|_{\mathbf{X}} \leq \frac{4\lambda_{\max} \left(\frac{\mathbf{X}^T \mathbf{X}}{n} \right)}{\lambda} \|\tilde{\mathbf{w}}^t - \hat{\mathbf{w}}^t\|_{\mathbf{X}},$$

which is a valid bound. Then in [28] the following upper bound is given without necessary justification:

$$\|\tilde{\mathbf{w}}^t - \hat{\mathbf{w}}^t\|_{\mathbf{X}} \leq 2\lambda_{\max} \left(\frac{\mathbf{X}^T \mathbf{X}}{n} \right) \beta^k \|\mathbf{w}^*\|_2. \quad (\text{A.3})$$

to reach the final form of the error upper bound:

$$\|\tilde{\mathbf{w}}^{t+1} - \mathbf{w}^*\|_{\mathbf{X}} \leq \alpha^t \|\mathbf{w}^*\|_{\mathbf{X}} + \frac{10\lambda_{\max}^2 \left(\frac{\mathbf{X}^T \mathbf{X}}{n} \right)}{\lambda} \beta^k \|\mathbf{w}^*\|_2.$$

However, this final form of the upper bound is not supported in detail as part of the presented proof. Because of the following major issue, we conclude that the proposed bound remains an unproven conjecture. The bound established for $\|\tilde{\mathbf{w}}^{t+1} - \bar{\mathbf{w}}^{t+1}\|_{\mathbf{X}}$ in (A.1) is used to upper bound $\|\tilde{\mathbf{w}}^t - \hat{\mathbf{w}}^t\|_{\mathbf{X}}$ in (A.3). This is not justified as part of the proof in [28].

References

- [1] Å. Björck, Numerical methods for least squares problems, SIAM, 1996.
- [2] P. C. Hansen, Discrete inverse problems: insight and algorithms, Vol. 7, SIAM, 2010.
- [3] D. L. Donoho, et al., Compressed sensing, IEEE Trans. Inform. Theory 52 (4) (2006) 1289–1306.
- [4] L. Zhang, M. Mahdavi, R. Jin, T. Yang, S. Zhu, Recovering the optimal solution by dual random projection, in: Conf. Learn. Theory, 2013, pp. 135–157.
- [5] S. Boyd, N. Parikh, E. Chu, B. Peleato, J. Eckstein, Distributed optimization and statistical learning via the alternating direction method of multipliers, Found. Trends Mach. Learn. 3 (1) (2011) 1–122.
- [6] S. Bubeck, et al., Convex optimization: Algorithms and complexity, Found. Trends Mach. Learn. 8 (3-4) (2015) 231–357.
- [7] M. Pilanci, M. J. Wainwright, Newton sketch: A near linear-time optimization algorithm with linear-quadratic convergence, SIAM J. Optim. 27 (1) (2017) 205–245.
- [8] C. R. Vogel, Computational methods for inverse problems, Vol. 23, SIAM, 2002.
- [9] M. E. Kilmer, D. P. O’Leary, Choosing regularization parameters in iterative methods for ill-posed problems, SIAM J. Matrix Anal. Appl. 22 (4) (2001) 1204–1221.

- [10] I. K. Ozaslan, Fast and robust solution techniques for large scale linear least squares problems, Master's thesis, Bilkent University (2020).
- [11] D. G. Luenberger, Introduction to linear and nonlinear programming, Vol. 28, Addison-Wesley Reading, MA, 1973.
- [12] B. Bartan, M. Pilanci, Polar coded distributed matrix multiplication, arXiv preprint arXiv:1901.06811 (2019).
- [13] E. Jonas, Q. Pu, S. Venkataraman, I. Stoica, B. Recht, Occupy the cloud: Distributed computing for the 99%, in: Proc. Symp. Cloud Comput., ACM, 2017, pp. 445–451.
- [14] R. Barrett, M. W. Berry, T. F. Chan, J. Demmel, J. Donato, J. Dongarra, V. Eijkhout, R. Pozo, C. Romine, H. Van der Vorst, Templates for the solution of linear systems: Building blocks for iterative methods, Vol. 43, SIAM, 1994.
- [15] M. Benzi, Preconditioning techniques for large linear systems: a survey, *J. Comput. Phys.* 182 (2) (2002) 418–477.
- [16] M. H. Gutknecht, S. Röllin, The chebyshev iteration revisited, *Parallel Comput.* 28 (2) (2002) 263–283.
- [17] W. B. Johnson, J. Lindenstrauss, Extensions of lipschitz mappings into a hilbert space, *Contemp. Math.* 26 (189-206) (1984) 1.
- [18] P. Drineas, M. W. Mahoney, RandNLA: Randomized numerical linear algebra, *Comm. ACM* 59 (6) (2016) 80–90.
- [19] M. W. Mahoney, et al., Randomized algorithms for matrices and data, *Found. Trends Theor. Mach. Learn* 3 (2) (2011) 123–224.
- [20] D. P. Woodruff, et al., Sketching as a tool for numerical linear algebra, *Found. Trends Theor. Comput. Sci.* 10 (1–2) (2014) 1–157.
- [21] P. Drineas, M. W. Mahoney, S. Muthukrishnan, T. Sarlós, Faster least squares approximation, *Numer. Math.* 117 (2) (2011) 219–249.
- [22] M. Pilanci, M. J. Wainwright, Randomized sketches of convex programs with sharp guarantees, *IEEE Trans. Inform. Theory* 61 (9) (2015) 5096–5115.
- [23] H. Avron, K. L. Clarkson, D. P. Woodruff, Sharper bounds for regularized data fitting, in: Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, 2017, pp. 27:1–27:22.
- [24] M. Pilanci, M. J. Wainwright, Iterative hessian sketch: Fast and accurate solution approximation for constrained least-squares, *J. Mach. Learn. Res.* 17 (1) (2016) 1842–1879.
- [25] V. Rokhlin, M. Tygert, A fast randomized algorithm for overdetermined linear least-squares regression, *Proc. Nat. Acad. Sci. India Sect. A* 105 (36) (2008) 13212–13217.
- [26] H. Avron, P. Maymounkov, S. Toledo, Blendenpik: Supercharging lapack's least-squares solver, *SIAM J. Sci. Comput.* 32 (3) (2010) 1217–1236.
- [27] X. Meng, M. A. Saunders, M. W. Mahoney, Lsrn: A parallel iterative solver for strongly over-or underdetermined systems, *SIAM J. Sci. Comput.* 36 (2) (2014) C95–C118.
- [28] J. Wang, J. D. Lee, M. Mahdavi, M. Kolar, N. Srebro, et al., Sketching meets random projection in the dual: A provable recovery algorithm for big and high-dimensional data, *Electron. J. Stat.* 11 (2) (2017) 4896–4944.
- [29] M. Bertero, M. Piana, Inverse problems in biomedical imaging: modeling and methods of solution, Springer Milan, Milano, 2006, pp. 1–33.
- [30] I. K. Ozaslan, M. Pilanci, O. Arikan, Iterative hessian sketch with momentum, in: Int. Conf. Acoust., Speech, Signal Process., IEEE, 2019, pp. 7470–7474.
- [31] A. Edelman, Y. Wang, Random matrix theory and its innovative applications, in: Advances in Applied Mathematics, Modeling, and Computational Science, Springer, 2013, pp. 91–116.
- [32] H. Avron, K. L. Clarkson, D. P. Woodruff, Faster kernel ridge regression using sketching and preconditioning, *SIAM J. Matrix Anal. Appl.* 38 (4) (2017) 1116–1138.
- [33] A. Chowdhury, J. Yang, P. Drineas, An iterative, sketching-based framework for ridge regression, in: Int. Conf. Mach. Learn., 2018, pp. 988–997.
- [34] M. B. Cohen, J. Nelson, D. P. Woodruff, Optimal approximate matrix product in terms of stable rank,

- in: *Int. Colloquium on Automata, Languages, and Programming*, Vol. 55, 2016, pp. 11:1–11:14.
- [35] C. C. Paige, M. A. Saunders, LSQR: An algorithm for sparse linear equations and sparse least squares, *ACM Trans. Math. Software* 8 (1) (1982) 43–71.
 - [36] B. T. Polyak, Some methods of speeding up the convergence of iteration methods, *Comput. Math. Math. Phys* 4 (5) (1964) 1–17.
 - [37] B. O’donoghue, E. Candes, Adaptive restart for accelerated gradient schemes, *Found. Comput. Math.* 15 (3) (2015) 715–732.
 - [38] D. M. Kane, J. Nelson, Sparser Johnson-Lindenstrauss transforms, *J. ACM.* 61 (1) (2014) 4.
 - [39] J. Nelson, H. L. Nguyen, OSNAP: Faster numerical linear algebra algorithms via sparser subspace embeddings, in: *Annu. Symp. Found. Comput. Sci.*, IEEE, 2013, pp. 117–126.
 - [40] M. Thorup, Y. Zhang, Tabulation-based 5-independent hashing with applications to linear probing and second moment estimation, *SIAM J. Comput.* 41 (2) (2012) 293–331.
 - [41] M. B. Cohen, Nearly tight oblivious subspace embeddings by trace inequalities, in: *ACM/SIAM Symp. Discr. Algorithms*, SIAM, 2016, pp. 278–287.
 - [42] J. Nelson, H. L. Nguyen, Lower bounds for oblivious subspace embeddings, in: *Int. Colloquium Automata, Languages, and Program.*, Springer, 2014, pp. 883–894.
 - [43] D. Kane, R. Meka, J. Nelson, Almost optimal explicit Johnson-Lindenstrauss families, in: *Int. Workshop on Approximation Algorithms for Combinatorial Optimization*, Springer, 2011, pp. 628–639.
 - [44] P. C. Hansen, Rank-deficient and discrete ill-posed problems: numerical aspects of linear inversion, Vol. 4, SIAM, 2005.
 - [45] J. Nocedal, S. Wright, *Numerical optimization*, Springer Science & Business Media, 2006.
 - [46] S. C. Eisenstat, H. F. Walker, Choosing the forcing terms in an inexact newton method, *SIAM J. Sci. Comput.* 17 (1) (1996) 16–32.
 - [47] S. G. Nash, A survey of truncated-newton methods, *J. Comput. Appl. Math.* 124 (1-2) (2000) 45–59.
 - [48] A. S. Berahas, R. Bollapragada, J. Nocedal, An investigation of newton-sketch and subsampled newton methods, *Optimization Methods and Software* 35 (4) (2020) 661–680.
 - [49] H. Avron, S. Toledo, Randomized algorithms for estimating the trace of an implicit symmetric positive semi-definite matrix, *J. ACM* 58 (2) (2011) 8.
 - [50] S. Gazzola, P. C. Hansen, J. G. Nagy, Ir tools: a matlab package of iterative regularization methods and large-scale test problems, *Numer. Algorithms* 81 (3) (2019) 773–811.
 - [51] L. Elden, Algorithms for the regularization of ill-conditioned least squares problems, *BIT* 17 (2) (1977) 134–145.
 - [52] H. D. Simon, Analysis of the symmetric lanczos algorithm with reorthogonalization methods, *Linear Algebra Appl.* 61 (1984) 101–131.
 - [53] S. Gazzola, P. Novati, M. R. Russo, On krylov projection methods and tikhonov regularization, *Electron. Trans. Numer. Anal* 44 (1) (2015) 83–123.
 - [54] P. C. Hansen, Regularization tools: a matlab package for analysis and solution of discrete ill-posed problems, *Numer. Algorithms* 6 (1) (1994) 1–35.
 - [55] R. Hunger, *Floating point operations in matrix-vector calculus*, Technical University of Munich, 2007.