

SKETCHING FOR PRINCIPAL COMPONENT REGRESSION*

LIRON MOR-YOSEF[†] AND HAIM AVRON[†]

Abstract. Principal component regression (PCR) is a useful method for regularizing least squares approximations. Although conceptually simple, straightforward implementations of PCR have high computational costs and so are inappropriate for large scale problems. In this paper, we propose efficient algorithms for computing approximate PCR solutions that, on one hand, are high quality approximations to the true PCR solutions (when viewed as minimizer of a constrained optimization problem) and, on the other hand, entertain rigorous risk bounds (when viewed as statistical estimators). In particular, we propose an input sparsity time algorithms for approximate PCR. We also consider computing an approximate PCR in the streaming model and kernel PCR. Empirical results demonstrate the excellent performance of our proposed methods.

Key words. sketching, randomized numerical linear algebra, linear regression, least squares, principal component regression, compressed least squares

AMS subject classifications. 65F20, 65F22, 65F05, 62J05, 68W20

DOI. 10.1137/18M1188860

1. Introduction. Least squares approximations of the form

$$\min_{\mathbf{x} \in \mathbb{R}^d} \|\mathbf{Ax} - \mathbf{b}\|_2$$

are fundamental building blocks in computational science, with applications ranging from statistical data analysis to inverse problems. However, it is well appreciated, especially in the aforementioned application areas, that regularization is often the key to achieving the best results.

One of the basic methods for regularizing least squares approximations is principal component regression (PCR) [23, 27, 2]. Given a data matrix \mathbf{A} , a right-hand side \mathbf{b} , and a target rank k , PCR is computed by first computing the coefficients $\mathbf{V}_{\mathbf{A},k}$ corresponding to the top k principal components of \mathbf{A} (i.e., to the dominant right invariant subspace of \mathbf{A}), then regressing on $\mathbf{AV}_{\mathbf{A},k}$ and \mathbf{b} , and finally projecting the solution back to the original space. In short, the PCR estimator is $\mathbf{x}_k = \mathbf{V}_{\mathbf{A},k}(\mathbf{AV}_{\mathbf{A},k})^+ \mathbf{b}$ and regularization is achieved via PCA-based dimensionality reduction. While there is some criticism of PCR in the statistical literature [2, 24], it is nevertheless a valuable tool in the toolbox of practitioners.

Up until recent breakthroughs on fast methods for least squares approximations, there was little penalty in terms of computational complexity when switching from ordinary least squares (OLS) to PCR. Indeed, the complexity of SVD-based computation of the dominant invariant subspace is $O(nd \min(n, d))$, and this matches the asymptotic complexity of straightforward computation of the OLS solution (i.e., via direct methods). However, recent progress on fast sketching-based algorithms for linear regression [17, 36, 31, 12, 44] has created a gap: exact computation of the principal components still requires SVD, so the overall complexity is still $O(nd \min(n, d))$, even

*Received by the editors May 22, 2018; accepted for publication (in revised form) March 1, 2019; published electronically April 30, 2019.

<http://www.siam.org/journals/simax/40-2/M118886.html>

Funding: This work was supported by Israel Science Foundation grant 1272/17 and by an IBM Faculty Award.

[†]School of Mathematical Sciences, Tel Aviv University, Tel Aviv, 6997801, Israel (lironmo2@mail.tau.ac.il, haimav@post.tau.ac.il).

though the OLS stage is faster. The gap is not insubstantial: when learning with large scale data (either large n or large d), $O(nd \min(n, d))$ is often not feasible, but modern sketching-based linear regression methods are.

1.1. Contributions. In this paper, we study the use of dimensionality reduction prior to computing PCR (so we can compute PCR on a smaller input matrix). In particular, for a data matrix \mathbf{A} , we relate the PCR solution of $\mathbf{A}\mathbf{R}$, where \mathbf{R} is any dimensionality reduction matrix, to the PCR solution of \mathbf{A} . To do so, we study the notion of approximate PCR both from an optimization perspective and from a statistical perspective, and we provide conditions on \mathbf{R} that guarantee that after projecting the solution back to the full space (by multiplying by \mathbf{R}^T) we have an approximate PCR solution with rigorous statistical risk bounds. These results are described in section 3.

We then leverage the aforementioned results to design fast, sketching-based algorithms for approximate PCR. We propose algorithms specialized for several cases (in the following, n is the number of data points, d is the dimension of the data): large n (using left sketching), large d (using right sketching), and both n and d large (using two-sided sketching). Furthermore, we propose an input-sparsity time algorithm for approximate PCR. These results are described in section 4.

We also consider computing approximate PCR in the streaming model, providing the first algorithm for computing approximate PCR in a stream. We also provide a fast algorithm for approximate kernel PCR (polynomial kernel only). These results are described in section 5.

Finally, empirical results (section 6) clearly demonstrate the ability of our proposed algorithms to compute approximate PCR solutions, the correctness of our theoretical analysis, and the advantages of using our techniques instead of simpler techniques like compressed least squares.

In general, unlike previous works on randomized methods for PCR (which we discuss in the next subsection), we analyze the use of sketching for PCR from a sketch-and-solve approach. We discuss the various advantages and disadvantages of the sketch-and-solve approach in comparison to iterative-based approaches in the next subsection.

1.2. Related work. Recently, matrix sketching, such as the use of random projections, has emerged as a powerful technique for accelerating and scaling many important statistical learning techniques. See recent surveys by Woodruff [44] and Yang, Meng, and Mahoney [46] for an extensive exposition on this subject. So far, there has been limited research on the use of matrix sketching in the context of principal component regression.

One natural strategy for leveraging sketching in the context of PCR is to use approximate principal components. Approximate principal components can be computed using a fast sketching-based algorithm for approximate PCA (also known as “randomized SVD”) [21, 44]. This was recently explored by Boutsidis and Magdon-Ismail [7]. The authors show that if the number of subspace iterations is sufficiently large, one can obtain a bound on the suboptimality of the approximate solution and on the error of the solution vector. We too bound the suboptimality of our solutions, but instead of bounding the error of the solution vector, we bound their distance to the right dominant subspace, or bound the distance of the projection to the left dominant subspace.

Frostig et al. leverage fast randomized algorithms for ridge regression to design iterative algorithms for PCR and principal component projection [18]. Frostig et al.’s

results were later improved upon by Allen-Zhu and Li [1]. Both of the aforementioned methods use iterations, while our work explores the use of a sketch-and-solve approach. While it is true that better accuracies can be achieved using iterative methods with sketching-based accelerators [36, 5, 31, 20, 3], there are some advantages to using a sketch-and-solve approach. In particular, sketch-and-solve algorithms are typically faster. However, this comes at a price: sketch-and-solve algorithms typically provide cruder approximations. Nevertheless, it is not uncommon for these cruder approximations to be sufficient in applications. Another advantage of the sketch-and-solve approach is that it is more amenable to streaming and kernelization; we consider both in this paper.

Closely related to our work is recent work on compressed least squares (CLS) [30, 25, 37, 38, 42]. In particular, our statistical analysis (section 3.2) is inspired by recent statistical analysis of CLS [37, 38, 25]. Additionally, CLS is sometimes considered as a computationally attractive alternative to PCR [38, 42]. While CLS certainly uses matrix sketching to compress the matrix, it also uses the compression to regularize the problem. The mix between compression for scalability and compression for regularization reduces the ability to fine-tune the method to the needs at hand, and thus obtain the best possible results. In contrast, our methods use sketching primarily to approximate the principal components and as such serve as a means for scalability only. We propose methods that are computationally as attractive as CLS, and are more faithful to the behavior of PCR (in fact, CLS is a special case of one of our proposed algorithms). These advantages over CLS are also evident in the experimental results reported in section 6.

PCR is a form of least squares regression with convex constraints (once the dominant subspace has been found). Pilanci and Wainwright recently explored the effect of regularization on the sketch size for least squares regression [34, 35]. In the aforementioned papers, sketching is applied only to the objective, while the constraint is enforced exactly. This is unsatisfactory in the context of PCR since for PCR the constraints are gleaned from the input, and enforcing them is as expensive as solving the problem exactly. In contrast, our methods use sketching not only to compress the objective function, but also to approximate the constraint set.

Ridge regression (also known as Tikhonov regularization) is another popular and well-studied method for regularizing least squares solutions. It is also closely related to PCR in the sense that the ridge term can be viewed as a soft damping of the singular values. Recently several sketching-based algorithms have been suggested to accelerate the solution of ridge regression [9, 4, 43, 10].

2. Preliminaries.

2.1. Notation and basic definitions. We denote scalars using Greek letters or using x, y, \dots . Vectors are denoted by $\mathbf{x}, \mathbf{y}, \dots$ and matrices by $\mathbf{A}, \mathbf{B}, \dots$. The $s \times s$ identity matrix is denoted by \mathbf{I}_s . We use the convention that vectors are column-vectors. $\text{nnz}(\mathbf{A})$ denotes the number of nonzeros in \mathbf{A} . The notation $\alpha = (1 \pm \gamma)\beta$ means that $(1 - \gamma)\beta \leq \alpha \leq (1 + \gamma)\beta$, and the notation $\alpha = \beta \pm \gamma$ means that $|\alpha - \beta| \leq \gamma$.

Given a matrix $\mathbf{X} \in \mathbb{R}^{m \times n}$, let $\mathbf{X} = \mathbf{U}_\mathbf{X} \Sigma_\mathbf{X} \mathbf{V}_\mathbf{X}^\top$ be a *thin SVD* of \mathbf{X} , i.e., $\mathbf{U}_\mathbf{X} \in \mathbb{R}^{m \times \min(m, n)}$ is a matrix with orthonormal columns, $\Sigma_\mathbf{X} \in \mathbb{R}^{\min(m, n) \times \min(m, n)}$ is a diagonal matrix with the nonnegative singular values on the diagonal, and $\mathbf{V}_\mathbf{X} \in \mathbb{R}^{n \times \min(m, n)}$ is a matrix with orthonormal columns. The thin SVD is not necessarily unique, so when we use this notation we mean that the statement is correct for any such decomposition. A thin SVD can be computed in $O(mn \min(m, n))$. We denote the singular values of \mathbf{X} by $\sigma_{\max}(\mathbf{X}) = \sigma_1(\mathbf{X}) \geq \dots \geq \sigma_{\min(m, n)}(\mathbf{X}) = \sigma_{\min}(\mathbf{X})$,

omitting the matrix from the notation if the relevant matrix is clear from the context. For $k \leq \min(m, n)$, we use $\mathbf{U}_{\mathbf{X},k}$ (respectively, $\mathbf{V}_{\mathbf{X},k}$) to denote the matrix consisting of the first k columns of $\mathbf{U}_{\mathbf{X}}$ (respectively, $\mathbf{V}_{\mathbf{X}}$), and we use $\Sigma_{\mathbf{X},k}$ to denote the leading $k \times k$ minor of $\Sigma_{\mathbf{X}}$. We use $\mathbf{U}_{\mathbf{X},k+}$ (respectively, $\mathbf{V}_{\mathbf{X},k+}$) to denote the matrix consisting of the last $\min(m, n) - k$ columns of $\mathbf{U}_{\mathbf{X}}$ (respectively, $\mathbf{V}_{\mathbf{X}}$), and we use $\Sigma_{\mathbf{X},k+}$ to denote the lower-right $(\min(m, n) - k) \times (\min(m, n) - k)$ block of $\Sigma_{\mathbf{X}}$. In other words,

$$\mathbf{U}_{\mathbf{X}} = \begin{bmatrix} \mathbf{U}_{\mathbf{X},k} & \mathbf{U}_{\mathbf{X},k+} \end{bmatrix}, \quad \Sigma_{\mathbf{X}} = \begin{bmatrix} \Sigma_{\mathbf{X},k} & 0 \\ 0 & \Sigma_{\mathbf{X},k+} \end{bmatrix}, \quad \mathbf{V}_{\mathbf{X}} = \begin{bmatrix} \mathbf{V}_{\mathbf{X},k} & \mathbf{V}_{\mathbf{X},k+} \end{bmatrix}.$$

The *Moore–Penrose pseudoinverse* of \mathbf{X} is $\mathbf{X}^+ := \mathbf{V}_{\mathbf{X}} \Sigma_{\mathbf{X}}^+ \mathbf{U}_{\mathbf{X}}^T$, where $\Sigma_{\mathbf{X}}^+ = \text{diag}(\sigma_1(\mathbf{X})^+, \dots, \sigma_{\min(m,n)}(\mathbf{X})^+)$ with $a^+ = a^{-1}$ when $a \neq 0$, and 0 otherwise.

The *stable rank* of a matrix \mathbf{X} is $\text{sr}(\mathbf{X}) := \|\mathbf{X}\|_F^2 / \|\mathbf{X}\|_2^2$. The k th *relative gap* of a matrix \mathbf{X} is

$$\text{gap}_k(\mathbf{X}) = \frac{\sigma_k^2 - \sigma_{k+1}^2}{\sigma_1^2}.$$

For a subspace \mathcal{U} , we use $\mathbf{P}_{\mathcal{U}}$ to denote the *orthogonal projection matrix* onto \mathcal{U} , and $\mathbf{P}_{\mathbf{X}}$ for the projection matrix on the column space of \mathbf{X} (i.e., $\mathbf{P}_{\mathbf{X}} = \mathbf{P}_{\text{range}(\mathbf{X})}$). We have $\mathbf{P}_{\mathbf{X}} = \mathbf{X}\mathbf{X}^+$. The *complementary projection matrix* is $\mathbf{P}_{\mathbf{X}}^\perp = \mathbf{I} - \mathbf{P}_{\mathbf{X}}$. A useful property of projection matrices is that if $\mathcal{S} \subseteq \mathcal{T}$, then $\mathbf{P}_{\mathcal{S}}\mathbf{P}_{\mathcal{T}} = \mathbf{P}_{\mathcal{T}}\mathbf{P}_{\mathcal{S}} = \mathbf{P}_{\mathcal{S}}$. Furthermore, we note the following result.

THEOREM 1 (Theorem 2.3 in [40]). *For any \mathbf{A} and \mathbf{B} with the same number of rows, the following statements hold:*

1. *If $\text{rank}(\mathbf{A}) = \text{rank}(\mathbf{B})$, then the singular values of $\mathbf{P}_{\mathbf{A}}\mathbf{P}_{\mathbf{B}}^\perp$ and $\mathbf{P}_{\mathbf{B}}\mathbf{P}_{\mathbf{A}}^\perp$ are the same, so*

$$\|\mathbf{P}_{\mathbf{A}}\mathbf{P}_{\mathbf{B}}^\perp\|_2 = \|\mathbf{P}_{\mathbf{B}}\mathbf{P}_{\mathbf{A}}^\perp\|_2.$$

2. *Moreover the nonzero singular values σ of $\mathbf{P}_{\mathbf{A}}\mathbf{P}_{\mathbf{B}}^\perp$ correspond to pairs $\pm\sigma$ of eigenvalues of $\mathbf{P}_{\mathbf{B}} - \mathbf{P}_{\mathbf{A}}$, so*

$$\|\mathbf{P}_{\mathbf{B}} - \mathbf{P}_{\mathbf{A}}\|_2 = \|\mathbf{P}_{\mathbf{A}}\mathbf{P}_{\mathbf{B}}^\perp\|_2.$$

3. *If $\|\mathbf{P}_{\mathbf{B}} - \mathbf{P}_{\mathbf{A}}\|_2 < 1$, then $\text{rank}(\mathbf{A}) = \text{rank}(\mathbf{B})$.*

2.2. Principal component regression and principal component projection. In the *principal component regression (PCR)* problem, we are given an input $n \times d$ data matrix \mathbf{A} , a right-hand side $\mathbf{b} \in \mathbb{R}^n$, and a rank parameter k which is smaller than or equal to the rank of \mathbf{A} . Furthermore, we assume that there is a nonzero eigengap at k : $\sigma_k > \sigma_{k+1}$. The goal is to find the PCR solution, \mathbf{x}_k , defined as

$$(1) \quad \mathbf{x}_k := \arg \min_{\mathbf{x} \in \text{range}(\mathbf{V}_{\mathbf{A},k})} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2.$$

It is easy to verify that $\mathbf{x}_k = \mathbf{V}_{\mathbf{A},k}(\mathbf{A}\mathbf{V}_{\mathbf{A},k})^+ \mathbf{b} = \mathbf{V}_{\mathbf{A},k} \Sigma_{\mathbf{A},k}^{-1} \mathbf{U}_{\mathbf{A},k}^T \mathbf{b}$. The *principal component projection (PCP)* of \mathbf{b} is $\mathbf{b}_k := \mathbf{A}\mathbf{x}_k = \mathbf{P}_{\mathbf{U}_{\mathbf{A},k}} \mathbf{b}$.

Straightforward computation of \mathbf{x}_k and \mathbf{b}_k via the SVD takes $O(nd \min(n, d))$ operations.¹ We are primarily interested in finding faster algorithms that compute an

¹The complexity when using iterative algorithms (e.g., Lanczos) to compute only the dominant invariant spaces depends on several additional facts and in particular on spectral properties of the matrix and sparsity level. Thus, to avoid overly complicating the discussion on computational complexity, we refrain from further discussion of iterative methods for computing dominant eigenspaces.

approximate PCR or PCP solution (we formalize the terms “approximate PCP/PCR” in section 3). Throughout the paper, we use \mathbf{A}, \mathbf{b} , and k as the arguments of the PCR/PCP problem to be solved.

2.3. Matrix perturbations and distance between subspaces. Our analysis uses matrix perturbation theory extensively. We now describe the basics of this theory and the results we use.

The *principal angles* $\theta_j \in [0, \pi/2]$ between two subspaces \mathcal{U} and \mathcal{W} are recursively defined by the identity

$$\cos(\theta_j) = \max_{\mathbf{u} \in \mathcal{U}} \max_{\mathbf{w} \in \mathcal{W}} \mathbf{u}^T \mathbf{w} \quad \text{s.t. } \|\mathbf{u}\|_2 = 1, \|\mathbf{w}\|_2 = 1 \quad \forall i < j. \mathbf{u}_i^T \mathbf{u} = 0, \mathbf{w}_i^T \mathbf{w} = 0.$$

We use \mathbf{u}_j and \mathbf{w}_j to denote the vectors for which $\cos(\theta_j) = \mathbf{u}_j^T \mathbf{w}_j$. Let $\Theta(\mathcal{U}, \mathcal{W})$ denote the $d \times d$ diagonal matrix whose j th diagonal entry is the j th principal angle, and as usual we allow writing matrices instead of subspaces as shorthand for the column space of the matrix. Henceforth, when we write a function on $\Theta(\cdot, \cdot)$, i.e., $\sin(\Theta(\mathbf{U}, \mathbf{W}))$, we mean evaluating the function entrywise on the diagonal only. It is well known [19, section 6.4.3] that if \mathbf{U} (respectively, \mathbf{W}) is a matrix with orthonormal columns whose column space is equal to \mathcal{U} (respectively, \mathcal{W}), then

$$\sigma_j(\mathbf{U}^T \mathbf{W}) = \cos(\theta_j).$$

The following lemma connects the tangent of the principal angles to the spectral norm of an appropriate matrix.

LEMMA 2 (Lemma 4.3 in [16]). *Let $\mathbf{Q} \in \mathbb{R}^{n \times s}$ have orthonormal columns, and let $\mathbf{W} = [\mathbf{W}_k \quad \mathbf{W}_{k+}] \in \mathbb{R}^{n \times n}$ be an orthogonal matrix where $\mathbf{W}_k \in \mathbb{R}^{n \times k}$ with $k \leq s$. If $\text{rank}(\mathbf{W}_k^T \mathbf{Q}) = k$, then*

$$\|\tan \Theta(\mathbf{Q}, \mathbf{W}_k)\|_2 = \|(\mathbf{W}_{k+}^T \mathbf{Q})(\mathbf{W}_k^T \mathbf{Q})\|_2.$$

Matrix perturbation theory studies how a perturbation of a matrix translates to perturbations of the matrix’s eigenvalues and eigenspaces. In order to bound the perturbation of an eigenspace, one needs some notion of distance between two subspaces. One common distance metric between two subspaces is

$$(2) \quad d_2(\mathcal{U}, \mathcal{W}) := \|\mathbf{P}_{\mathcal{U}} - \mathbf{P}_{\mathcal{W}}\|_2.$$

If \mathbf{U} and \mathbf{V} have the same number of columns, and both have orthonormal columns, then

$$d_2(\mathbf{U}, \mathbf{V}) = \sqrt{1 - \sigma_{\min}(\mathbf{U}^T \mathbf{V})^2} = \sin(\theta_{\max}) = \|\sin \Theta(\mathbf{U}, \mathbf{V})\|_2,$$

where θ_{\max} is the maximum principal angle between $\text{range}(\mathbf{U})$ and $\text{range}(\mathbf{V})$ [19, section 6.4.3].

A classical result that bounds the distance between the dominant subspaces of two symmetric matrices in terms of the spectral norm of difference between the two matrices is the Davis–Kahan $\sin(\Theta)$ theorem [15, section 2]. We need the following corollary of this theorem.

THEOREM 3 (corollary of Davis–Kahan $\sin \Theta$ theorem [15]). *Let $\mathbf{A}, \tilde{\mathbf{A}} \in \mathbb{R}^{n \times n}$ be two symmetric matrices, both of rank at least k . Suppose that $\lambda_k > \tilde{\lambda}_{k+1}$, where $\lambda_1 \geq \dots \geq \lambda_n$ and $\tilde{\lambda}_1 \geq \dots \geq \tilde{\lambda}_n$ are the eigenvalues of \mathbf{A} and $\tilde{\mathbf{A}}$. We have*

$$d_2(\mathbf{V}_{\mathbf{A},k}, \mathbf{V}_{\tilde{\mathbf{A}},k}) \leq \frac{\|\mathbf{A} - \tilde{\mathbf{A}}\|_2}{\lambda_k - \tilde{\lambda}_{k+1}}.$$

Proof. We use the following variant of the $\sin \Theta$ theorem (see [41, Theorem 2.16]): Suppose a symmetric matrix \mathbf{B} has a spectral representation

$$\mathbf{B} = \mathbf{X}\mathbf{L}\mathbf{X}^\top + \mathbf{Y}\mathbf{M}\mathbf{Y}^\top,$$

where $[\mathbf{X} \ \mathbf{Y}]$ is square orthonormal. Let the orthonormal matrix \mathbf{Z} be of the same dimensions as \mathbf{X} and suppose that

$$\mathbf{R} = \mathbf{B}\mathbf{Z} - \mathbf{Z}\mathbf{N},$$

where \mathbf{N} is symmetric. Furthermore, suppose that the spectrum of \mathbf{N} is contained in some interval $[\alpha, \beta]$ and that for some $\delta > 0$ the spectrum of \mathbf{M} lies outside of $[\alpha - \delta, \beta + \delta]$. Then

$$\|\sin \Theta(\mathbf{X}, \mathbf{Z})\|_2 \leq \frac{\|\mathbf{R}\|_2}{\delta}.$$

We prove Theorem 3 by applying the aforementioned variant of the $\sin \Theta$ theorem with $\mathbf{B} = \tilde{\mathbf{A}}$, $\mathbf{X} = \mathbf{V}_{\tilde{\mathbf{A}},k}$, $\mathbf{Y} = \mathbf{V}_{\tilde{\mathbf{A}},k+1}$, $\mathbf{L} = \mathbf{diag}(\tilde{\lambda}_1, \dots, \tilde{\lambda}_k)$, $\mathbf{M} = \mathbf{diag}(\tilde{\lambda}_{k+1}, \dots, \tilde{\lambda}_n)$, $\mathbf{Z} = \mathbf{V}_{\mathbf{A},k}$, $\mathbf{N} = \mathbf{diag}(\lambda_1, \dots, \lambda_k)$, and $\delta = \lambda_k - \tilde{\lambda}_{k+1}$. It is easy to verify that the conditions of the $\sin \Theta$ theorem hold, so

$$\|\sin \Theta(\mathbf{V}_{\mathbf{A},k}, \mathbf{V}_{\tilde{\mathbf{A}},k})\|_2 \leq \frac{\|\mathbf{R}\|_2}{\lambda_k - \tilde{\lambda}_{k+1}},$$

where $\mathbf{R} = \tilde{\mathbf{A}}\mathbf{V}_{\mathbf{A},k} - \mathbf{V}_{\mathbf{A},k}\mathbf{N}$. We have $\mathbf{A}\mathbf{V}_{\mathbf{A},k} = \mathbf{V}_{\mathbf{A},k}\mathbf{N}$, so $\|\mathbf{R}\|_2 = \|(\tilde{\mathbf{A}} - \mathbf{A})\mathbf{V}_{\mathbf{A},k}\|_2 \leq \|\tilde{\mathbf{A}} - \mathbf{A}\|_2$. Combining this inequality with the previous one and noting that $d_2(\mathbf{V}_{\mathbf{A},k}, \mathbf{V}_{\tilde{\mathbf{A}},k}) = \|\sin \Theta(\mathbf{V}_{\mathbf{A},k}, \mathbf{V}_{\tilde{\mathbf{A}},k})\|_2$ completes the proof. \square

Under the conditions of Theorem 3, since \mathbf{A} and $\tilde{\mathbf{A}}$ are symmetric matrices, Weyl's inequality implies that

$$d_2(\mathbf{V}_{\mathbf{A},k}, \mathbf{V}_{\tilde{\mathbf{A}},k}) \leq \frac{\|\mathbf{A} - \tilde{\mathbf{A}}\|_2}{\lambda_k - \lambda_{k+1} - \|\mathbf{A} - \tilde{\mathbf{A}}\|_2}$$

as long as $\|\mathbf{A} - \tilde{\mathbf{A}}\|_2 < \lambda_k - \lambda_{k+1}$. Thus, if $\|\mathbf{A} - \tilde{\mathbf{A}}\|_2 \ll \lambda_k - \lambda_{k+1}$, then we can compute an approximation to the k -dimensional dominant subspace of \mathbf{A} by computing the k -dimensional dominant subspace of $\tilde{\mathbf{A}}$.

3. PCR with dimensionality reduction. Our goal is to design algorithms which compute an approximate solution to the PCR or PCP problem. Our strategy for designing such algorithms is to reduce the dimensions of \mathbf{A} prior to computing the PCR/PCP solution. Specifically, let $\mathbf{R} \in \mathbb{R}^{d \times t}$ be some matrix where $t \leq d$, and define

$$(3) \quad \mathbf{x}_{\mathbf{R},k} := \mathbf{R}\mathbf{V}_{\mathbf{AR},k}(\mathbf{A}\mathbf{R}\mathbf{V}_{\mathbf{AR},k})^+ \mathbf{b}.$$

The rationale in (3) is as follows. First, \mathbf{A} is compressed by computing $\mathbf{A}\mathbf{R}$ (this is the dimensionality reduction step). Then we compute the rank k PCR solution of $\mathbf{A}\mathbf{R}$ and \mathbf{b} ; this is $(\mathbf{A}\mathbf{R}\mathbf{V}_{\mathbf{AR},k})^+ \mathbf{b}$. Finally, the solution is projected back to the original space by multiplying by $\mathbf{R}\mathbf{V}_{\mathbf{AR},k}$. Obviously, given \mathbf{R} we can compute $\mathbf{x}_{\mathbf{R},k}$ in $O(ndt)$ (and even faster if \mathbf{A} is sparse), so if $t \ll \min(n, d)$ there is a potential for significant gain in terms of computational complexity, provided it is possible to compute \mathbf{R} efficiently as well. Furthermore, if we design \mathbf{R} to have some special

structure that allows us to compute \mathbf{AR} in $O(nt^2)$ time, the overall complexity would reduce to $O(nt^2)$.

Of course, $\mathbf{x}_{\mathbf{R},k}$ is not the PCR solution \mathbf{x}_k (unless $\mathbf{R} = \mathbf{V}_{\mathbf{A},k}$). This suggests the following mathematical question (which, in turn, leads to an algorithmic question): under which conditions on \mathbf{R} is $\mathbf{x}_{\mathbf{R},k}$ a good approximation to the PCR solution \mathbf{x}_k ? In this section, we derive general conditions on \mathbf{R} that ensure deterministically that $\mathbf{x}_{\mathbf{R},k}$ is in some sense (which we formalize later in this section) a good approximation of \mathbf{x}_k . The results in this section are nonalgorithmic and independent of the method in which \mathbf{R} is computed. In the next section we address the algorithmic question: how can we compute such \mathbf{R} matrices efficiently?

We approach the mathematical question from two different perspectives: an optimization perspective and a statistical perspective. In the optimization perspective, we consider PCR/PCP as an optimization problem (equation (1)) and ask whether the value of the objective function of $\mathbf{x}_{\mathbf{R},k}$ is close to the optimal value of the objective function, while upholding the constraints approximately (see Definition 4). In the statistical perspective, we treat \mathbf{x}_k and $\mathbf{x}_{\mathbf{R},k}$ as statistical estimators and compare their excess risk under a fixed-design model. Interestingly, the conditions we derive for \mathbf{R} are the same for both perspectives.

Before proceeding, we remark that an important special case of (3) is when \mathbf{R} has exactly k columns. In that case, for brevity, we omit the subscript k from $\mathbf{x}_{\mathbf{R},k}$ and notice that

$$(4) \quad \mathbf{x}_{\mathbf{R}} = \mathbf{R}(\mathbf{AR})^+ \mathbf{b}.$$

Equation (4) is valid even if \mathbf{R} has more than k columns and/or the columns are not orthonormal. Thus, an established technique in the literature, frequently referred to as compressed least squares (CLS) [30, 25, 37, 38, 42], is to generate a random \mathbf{R} and compute $\mathbf{x}_{\mathbf{R}}$. To avoid confusion, we stress the difference between (3) and (4): in (3) we compute a PCR solution on the compressed matrix \mathbf{AR} , while in (4) ordinary least squares is used. These two strategies coincide when \mathbf{R} has k columns. In this paper, we focus on (3) and consider (4) only when it is a special case of (3) (when \mathbf{R} has exactly k columns). For an analysis of CLS from a statistical perspective, see recent work by Slawski [38].

3.1. Optimization perspective. The PCR solution can be written as the solution of a constrained least squares problem:

$$\mathbf{x}_k = \arg \min_{\substack{\|\mathbf{V}_{\mathbf{A},k+}^T \mathbf{x}\|_2 = 0 \\ \mathbf{x} \in \text{range}(\mathbf{A}^T)}} \|\mathbf{Ax} - \mathbf{b}\|_2.$$

In order to analyze a candidate solution $\tilde{\mathbf{x}}$ from an optimization perspective, we need to decide how to treat the constraints. One option is to require a candidate $\tilde{\mathbf{x}}$ to be inside the feasible set. Indeed, Pilanci and Wainwright recently considered sketching-based methods for constrained least squares regression [34]. However, there is no evident way to impose $\mathbf{V}_{\mathbf{A},k+}^T \mathbf{x} = 0$ without actually computing $\mathbf{V}_{\mathbf{A},k+}$, which is as expensive as computing $\mathbf{V}_{\mathbf{A},k}$. Thus, if we require an approximate solution to be inside the feasible set, we might as well compute the exact PCR solution. Thus, in our notion of approximate PCR, we relax the constraints and require only that the approximate solution is close to meeting the constraint; i.e., we seek a solution for which $\|\mathbf{A}\tilde{\mathbf{x}} - \mathbf{b}\|_2$ is close to $\|\mathbf{Ax}_k - \mathbf{b}\|_2$ and $\|\mathbf{V}_{\mathbf{A},k+}^T \tilde{\mathbf{x}}\|_2$ is small.

Similarly, if \mathbf{A} has full rank, the PCP solution can be written as the solution of a constrained least squares problem:

$$\mathbf{b}_k = \arg \min_{\substack{\|\mathbf{U}_{\mathbf{A},k}^T \tilde{\mathbf{b}}\|_2=0 \\ \tilde{\mathbf{b}} \in \text{range}(\mathbf{A})}} \|\tilde{\mathbf{b}} - \mathbf{b}\|_2.$$

Again, our notion of approximate PCP relaxes the constraint.

The discussion above motivates the following definition of approximate PCR/PCP.

DEFINITION 4 (approximate PCR and PCP). *An estimator $\tilde{\mathbf{x}}$ is an (ϵ, ν) -approximate PCR of rank k if*

$$\|\mathbf{A}\tilde{\mathbf{x}} - \mathbf{b}\|_2 = \|\mathbf{A}\mathbf{x}_k - \mathbf{b}\|_2 \pm \epsilon\|\mathbf{b}\|_2$$

and $\|\mathbf{V}_{\mathbf{A},k}^T \tilde{\mathbf{x}}\|_2 \leq \nu\|\mathbf{b}\|_2$. *An estimator $\tilde{\mathbf{b}}$ is an (ϵ, ν) -approximate PCP of rank k if*

$$\|\tilde{\mathbf{b}} - \mathbf{b}\|_2 = \|\mathbf{b}_k - \mathbf{b}\|_2 \pm \epsilon\|\mathbf{b}\|_2$$

and $\|\mathbf{U}_{\mathbf{A},k}^T \tilde{\mathbf{b}}\|_2 \leq \nu\|\mathbf{b}\|_2$.

Before proceeding, a few remarks are in order.

1. Imposing no constraints on $\tilde{\mathbf{x}}$ (or $\tilde{\mathbf{b}}$) does not make sense: we can always form or approximate the OLS solution, and it will demonstrate a smaller objective value. Indeed, the main motivation for using PCR is to impose some form of regularization, so it is crucial the definition of approximate PCR/PCP have some form of regularization built in.
2. We require only additive error on the objective function, while relative error bounds are usually viewed as more desirable. For approximate PCR, requiring relative error bounds is likely unrealistic: since it is possible that $\mathbf{b} = \mathbf{A}\mathbf{x}_k$, any algorithm that provides a relative error bound must search inside a space that contains $\text{range}(\mathbf{V}_{\mathbf{A},k})$. This is a strong restriction (and plausibly one that actually requires computing $\mathbf{V}_{\mathbf{A},k}$).
3. Approximate PCR implies approximate PCP: if $\tilde{\mathbf{x}}$ is an (ϵ, ν) -approximate PCR, then $\mathbf{A}\tilde{\mathbf{x}}$ is an $(\epsilon, \sigma_{k+1}\nu)$ -approximate PCP.
4. Our notion of approximate PCP is somewhat similar to the notion of approximate PCP proposed recently by Allen-Zhu and Li [1].
5. Yet another notion of approximate PCR appears in [7, Theorem 5]. They, too, consider an additive error on objective function, but instead of considering the distance to the dominant subspace they bound the distance of the approximate solution to the true solution. We remark that a bound on $\|\mathbf{x}_k - \tilde{\mathbf{x}}\|_2$ trivially implies a bound on $\|\mathbf{V}_{\mathbf{A},k}^T \tilde{\mathbf{x}}\|_2$.
6. Arguably, it would have been preferable to require the approximate PCR solution $\tilde{\mathbf{x}}$ to be such that $\|\mathbf{x}_k - \tilde{\mathbf{x}}\|_2$ is small (relative to $\|\mathbf{x}_k\|_2$). However, we believe that providing such guarantees with reasonable sketch sizes requires iterations. In this paper, we focus predominately on algorithms that do not require iterations (the only exception being the input sparsity algorithm in subsection 4.3).

We are now ready to state general conditions on \mathbf{R} that ensure deterministically that $\mathbf{x}_{\mathbf{R}}$ is an approximate PCR, and conditions on \mathbf{R} that ensure deterministically that $\mathbf{A}\mathbf{x}_{\mathbf{R},k}$ is an approximate PCP.

THEOREM 5. *Suppose that $\mathbf{R} \in \mathbb{R}^{d \times s}$, where $s \geq k$. Assume that $\nu \in (0, 1)$.*

1. If $d_2(\mathbf{U}_{\mathbf{A}\mathbf{R},k}, \mathbf{U}_{\mathbf{A},k}) \leq \nu$, then $\mathbf{A}\mathbf{x}_{\mathbf{R},k}$ is an (ν, ν) -approximate PCP.
2. If $s = k$, \mathbf{R} has orthonormal columns (i.e., $\mathbf{R}^T \mathbf{R} = \mathbf{I}_k$), and $d_2(\mathbf{R}, \mathbf{V}_{\mathbf{A},k}) \leq \nu(1 + \nu^2)^{-1/2}$, then $\mathbf{x}_{\mathbf{R}}$ is an $(\frac{\sigma_{k+1}}{\sigma_k} \nu, \frac{\nu}{(\sqrt{1-\nu^2}-\nu)\sigma_k})$ -approximate PCR.

Before proving this theorem, we state a theorem which is a corollary of a more general result proved recently by Drineas et al. [16], and then proceed to proving a couple of auxiliary lemmas.

THEOREM 6 (corollary of Theorem 2.1 in [16]). *Let \mathbf{A} be an $m \times n$ matrix with singular value decomposition $\mathbf{A} = \mathbf{U}_{\mathbf{A}} \Sigma_{\mathbf{A}} \mathbf{V}_{\mathbf{A}}^T$. Let $k \geq 0$, and let $\mathbf{R} \in \mathbb{R}^{d \times k}$ be any matrix such that $\mathbf{V}_{\mathbf{A},k}^T \mathbf{R}$ has full rank. Then*

$$\|\sin \Theta(\mathbf{A}\mathbf{R}, \mathbf{U}_{\mathbf{A},k})\|_2 \leq \|\Sigma_{\mathbf{A},k+}\|_2 \cdot \|\Sigma_{\mathbf{A},k}^{-1}\|_2 \cdot \|\tan \Theta(\mathbf{R}, \mathbf{V}_{\mathbf{A},k})\|_2.$$

LEMMA 7. *Assume $\text{rank}(\mathbf{A}) \geq k$. If $\mathbf{R} \in \mathbb{R}^{d \times k}$ has orthonormal columns and $d_2(\mathbf{R}, \mathbf{V}_{\mathbf{A},k}) \leq \nu$, then the following bounds hold:*

$$(5) \quad \|\mathbf{V}_{\mathbf{A},k+}^T \mathbf{R}\|_2 \leq \nu,$$

$$(6) \quad \sigma_{\min}(\mathbf{A}\mathbf{R}) \geq \sigma_k \left(\sqrt{1 - \nu^2} - \nu \right).$$

Furthermore, if $\nu < 1$, then $\text{rank}(\mathbf{A}\mathbf{R}) = k$.

Proof. Since both $\mathbf{V}_{\mathbf{A},k}$ and \mathbf{R} have orthonormal columns, $d_2(\mathbf{R}, \mathbf{V}_{\mathbf{A},k}) \leq \nu$ implies that the square of the singular values of $\mathbf{V}_{\mathbf{A},k}^T \mathbf{R}$ lies inside the interval $[1 - \nu^2, 1]$. The eigenvalues of $\mathbf{R}^T \mathbf{V}_{\mathbf{A},k} \mathbf{V}_{\mathbf{A},k}^T \mathbf{R}$ are exactly the square of the singular values of $\mathbf{V}_{\mathbf{A},k}^T \mathbf{R}$, so the eigenvalues of $\mathbf{I}_k - \mathbf{R}^T \mathbf{V}_{\mathbf{A},k} \mathbf{V}_{\mathbf{A},k}^T \mathbf{R}$ lie in $[0, \nu^2]$. Let \mathbf{Z} be any matrix with orthonormal columns that completes $\mathbf{V}_{\mathbf{A}}$ to a basis (i.e., $\mathbf{V}_{\mathbf{A}} \mathbf{V}_{\mathbf{A}}^T + \mathbf{Z} \mathbf{Z}^T = \mathbf{I}_d$) and is orthogonal to $\mathbf{V}_{\mathbf{A}}$ (i.e., $\mathbf{V}_{\mathbf{A}}^T \mathbf{Z} = 0$). Note that \mathbf{Z} can be an empty matrix if $d \leq n$. Denote $\mathbf{V}_{\mathbf{A},k\perp} = [\mathbf{V}_{\mathbf{A},k+} \quad \mathbf{Z}]$. We have

$$\begin{aligned} \|\mathbf{V}_{\mathbf{A},k\perp}^T \mathbf{R}\|_2 &= \sqrt{\|\mathbf{R}^T \mathbf{V}_{\mathbf{A},k\perp} \mathbf{V}_{\mathbf{A},k\perp}^T \mathbf{R}\|_2} \\ &= \sqrt{\|\mathbf{I}_k - \mathbf{R}^T \mathbf{V}_{\mathbf{A},k} \mathbf{V}_{\mathbf{A},k}^T \mathbf{R}\|_2} \\ &\leq \nu, \end{aligned}$$

where we used the fact that $\mathbf{V}_{\mathbf{A},k} \mathbf{V}_{\mathbf{A},k}^T + \mathbf{V}_{\mathbf{A},k\perp} \mathbf{V}_{\mathbf{A},k\perp}^T = \mathbf{I}_d$. We now note that $\mathbf{V}_{\mathbf{A},k+}^T \mathbf{R}$ is a submatrix of $\mathbf{V}_{\mathbf{A},k\perp}^T \mathbf{R}$ so $\|\mathbf{V}_{\mathbf{A},k+}^T \mathbf{R}\|_2 \leq \|\mathbf{V}_{\mathbf{A},k\perp}^T \mathbf{R}\|_2 \leq \nu$. This establishes the first part of the theorem.

As for the second part, recall the following identities: (1) for any matrix \mathbf{X} and \mathbf{Y} of the same size, $\sigma_{\min}(\mathbf{X} \pm \mathbf{Y}) \geq \sigma_{\min}(\mathbf{X}) - \sigma_{\max}(\mathbf{Y})$ [22, Theorem 3.3.19]; (2) if the number of rows in \mathbf{X} and \mathbf{Y} is at least as large as the number of columns, and $\mathbf{X}\mathbf{Y}$ is defined, then $\sigma_{\min}(\mathbf{X}\mathbf{Y}) \geq \sigma_{\min}(\mathbf{X}) \sigma_{\min}(\mathbf{Y})$. We have

$$\begin{aligned} \sigma_{\min}(\mathbf{A}\mathbf{R}) &= \sigma_{\min}(\mathbf{A} \mathbf{V}_{\mathbf{A},k} \mathbf{V}_{\mathbf{A},k}^T \mathbf{R} + \mathbf{A} \mathbf{V}_{\mathbf{A},k+} \mathbf{V}_{\mathbf{A},k+}^T \mathbf{R}) \\ &\geq \sigma_{\min}(\mathbf{A} \mathbf{V}_{\mathbf{A},k} \mathbf{V}_{\mathbf{A},k}^T \mathbf{R}) - \sigma_{\max}(\mathbf{A} \mathbf{V}_{\mathbf{A},k+} \mathbf{V}_{\mathbf{A},k+}^T \mathbf{R}) \\ &\geq \sigma_{\min}(\mathbf{A} \mathbf{V}_{\mathbf{A},k}) \sigma_{\min}(\mathbf{V}_{\mathbf{A},k}^T \mathbf{R}) - \sigma_{\max}(\mathbf{A} \mathbf{V}_{\mathbf{A},k+}) \sigma_{\max}(\mathbf{V}_{\mathbf{A},k+}^T \mathbf{R}) \\ &= \sigma_k \sigma_{\min}(\mathbf{V}_{\mathbf{A},k}^T \mathbf{R}) - \sigma_{k+1} \sigma_{\max}(\mathbf{V}_{\mathbf{A},k+}^T \mathbf{R}) \\ &\geq \sigma_k \sqrt{1 - \nu^2} - \sigma_{k+1} \nu \\ &\geq \sigma_k (\sqrt{1 - \nu^2} - \nu), \end{aligned}$$

where the first equality follows from the fact that $\mathbf{A}(\mathbf{V}_{\mathbf{A},k}\mathbf{V}_{\mathbf{A},k}^T + \mathbf{V}_{\mathbf{A},k+}\mathbf{V}_{\mathbf{A},k+}^T) = \mathbf{A}$. When $\nu < 1$ we have $\sigma_{\min}(\mathbf{A}\mathbf{R}) > 0$, so indeed the rank of $\mathbf{A}\mathbf{R}$ is k . \square

LEMMA 8. Assume $\text{rank}(\mathbf{A}) \geq k$. Suppose that $\mathbf{R} \in \mathbb{R}^{d \times k}$ has orthonormal columns, and that $d_2(\mathbf{R}, \mathbf{V}_{\mathbf{A},k}) \leq \nu(1 + \nu^2)^{-1/2} < 1$. We have

$$d_2(\mathbf{U}_{\mathbf{A}\mathbf{R}}, \mathbf{U}_{\mathbf{A},k}) \leq \frac{\sigma_{k+1}}{\sigma_k} \nu.$$

Proof. Since $\text{rank}(\mathbf{A}) \geq k$ and $\nu(1 + \nu^2)^{-1/2} < 1$, according to Lemma 7 the matrix $\mathbf{A}\mathbf{R}$ has full rank. According to Theorem 1 and the fact that $\mathbf{P}_{\mathbf{A}\mathbf{R}}$ and $\mathbf{P}_{\mathbf{U}_{\mathbf{A},k}}$ are orthogonal projections we have

$$(7) \quad d_2(\mathbf{U}_{\mathbf{A}\mathbf{R}}, \mathbf{U}_{\mathbf{A},k}) = d_2(\mathbf{A}\mathbf{R}, \mathbf{U}_{\mathbf{A},k}) = \|\mathbf{P}_{\mathbf{A}\mathbf{R}} - \mathbf{P}_{\mathbf{U}_{\mathbf{A},k}}\|_2 = \|\mathbf{P}_{\mathbf{A}\mathbf{R}}^\perp \mathbf{P}_{\mathbf{U}_{\mathbf{A},k}}\|_2.$$

Combining Theorem 6 and (7), we bound

$$\begin{aligned} d_2(\mathbf{A}\mathbf{R}, \mathbf{U}_{\mathbf{A},k}) &= \|\mathbf{P}_{\mathbf{A}\mathbf{R}}^\perp \mathbf{P}_{\mathbf{U}_{\mathbf{A},k}}\|_2 \\ &= \|(\mathbf{I} - \mathbf{P}_{\mathbf{A}\mathbf{R}})\mathbf{U}_{\mathbf{A},k}\|_2 \\ &= \|\sin \Theta(\mathbf{A}\mathbf{R}, \mathbf{U}_{\mathbf{A},k})\|_2 \\ &\leq \|\Sigma_{\mathbf{A},k+}\|_2 \cdot \|\Sigma_{\mathbf{A},k}^{-1}\|_2 \cdot \|\tan \Theta(\mathbf{R}, \mathbf{V}_{\mathbf{A},k})\|_2 \\ &= \frac{\sigma_{k+1}}{\sigma_k} \cdot \|\tan \Theta(\mathbf{R}, \mathbf{V}_{\mathbf{A},k})\|_2 \\ &\leq \frac{\sigma_{k+1}}{\sigma_k} \nu, \end{aligned}$$

where the last inequality follows from the fact that $\Theta(\mathbf{R}, \mathbf{V}_{\mathbf{A},k})$ is a diagonal matrix whose diagonal values are the inverse cosine of the singular values of $\mathbf{R}^T \mathbf{V}_{\mathbf{A},k}$, and these, in turn, are all larger than $\sqrt{1 - \nu^2(1 + \nu^2)^{-1}}$. \square

We are now ready to prove Theorem 5.

Proof of Theorem 5. We need to show both the additive error bounds on the objective function, and the error bound on the constraints. We start with the additive error bounds on the objective function, both for PCP (first part of the theorem) and for PCR (second part of the theorem). We have

$$\mathbf{A}\mathbf{x}_{\mathbf{R},k} = \mathbf{A}\mathbf{R}\mathbf{V}_{\mathbf{A},k}(\mathbf{A}\mathbf{R}\mathbf{V}_{\mathbf{A},k})^+ \mathbf{b} = \mathbf{P}_{\mathbf{U}_{\mathbf{A}\mathbf{R},k}} \mathbf{b}$$

and

$$\mathbf{A}\mathbf{x}_k = \mathbf{A}\mathbf{V}_{\mathbf{A},k}(\mathbf{A}\mathbf{V}_{\mathbf{A},k})^+ \mathbf{b} = \mathbf{P}_{\mathbf{U}_{\mathbf{A},k}} \mathbf{b}.$$

Thus,

$$\begin{aligned} \|\mathbf{A}\mathbf{x}_{\mathbf{R},k} - \mathbf{b}\|_2 &= \|\mathbf{A}\mathbf{x}_k - \mathbf{b} + \mathbf{A}\mathbf{x}_{\mathbf{R},k} - \mathbf{A}\mathbf{x}_k\|_2 \\ &= \|\mathbf{A}\mathbf{x}_k - \mathbf{b}\|_2 \pm \|\mathbf{A}\mathbf{x}_{\mathbf{R},k} - \mathbf{A}\mathbf{x}_k\|_2 \\ &= \|\mathbf{A}\mathbf{x}_k - \mathbf{b}\|_2 \pm \|(\mathbf{P}_{\mathbf{U}_{\mathbf{A}\mathbf{R},k}} - \mathbf{P}_{\mathbf{U}_{\mathbf{A},k}})\mathbf{b}\|_2 \\ &= \|\mathbf{A}\mathbf{x}_k - \mathbf{b}\|_2 \pm d_2(\mathbf{U}_{\mathbf{A}\mathbf{R},k}, \mathbf{U}_{\mathbf{A},k}) \cdot \|\mathbf{b}\|_2. \end{aligned}$$

In the first part of the theorem, we have $d_2(\mathbf{U}_{\mathbf{A}\mathbf{R},k}, \mathbf{U}_{\mathbf{A},k}) \leq \nu$, while in the second part of the theorem we have $\mathbf{U}_{\mathbf{A}\mathbf{R},k} = \mathbf{U}_{\mathbf{A}\mathbf{R}}$ (since $\mathbf{A}\mathbf{R}$ has k columns) and Lemma 8 ensures that $d_2(\mathbf{U}_{\mathbf{A}\mathbf{R},k}, \mathbf{U}_{\mathbf{A},k}) \leq \nu\sigma_{k+1}/\sigma_k$. Either way, the additive error bounds of the theorem are met.

We now bound the infeasibility of the approximate solution for the PCP guarantee (first part of the theorem):

$$\begin{aligned}
\|\mathbf{U}_{\mathbf{A},k+}^T \mathbf{A} \mathbf{x}_{\mathbf{R},k}\|_2 &= \|\mathbf{U}_{\mathbf{A},k+}^T \mathbf{A} \mathbf{x}_{\mathbf{R},k} - \mathbf{U}_{\mathbf{A},k+}^T \mathbf{A} \mathbf{x}_k + \mathbf{U}_{\mathbf{A},k+}^T \mathbf{A} \mathbf{x}_k\|_2 \\
&\leq \|\mathbf{U}_{\mathbf{A},k+}^T (\mathbf{A} \mathbf{x}_{\mathbf{R},k} - \mathbf{A} \mathbf{x}_k)\|_2 + \|\mathbf{U}_{\mathbf{A},k+}^T \mathbf{A} \mathbf{x}_k\|_2 \\
&\leq \|\mathbf{A} \mathbf{x}_{\mathbf{R},k} - \mathbf{A} \mathbf{x}_k\|_2 \\
&\leq d_2 (\mathbf{U}_{\mathbf{A},k}, \mathbf{U}_{\mathbf{A},k}) \|\mathbf{b}\|_2 \\
&\leq \nu \|\mathbf{b}\|_2,
\end{aligned}$$

where we used the fact that $\mathbf{A} \mathbf{x}_k \in \text{range}(\mathbf{U}_{\mathbf{A},k})$ so $\mathbf{U}_{\mathbf{A},k+}^T \mathbf{A} \mathbf{x}_k = 0$.

We now bound the infeasibility of the approximate solution for the PCR guarantee (second part of the theorem):

$$\begin{aligned}
\|\mathbf{V}_{\mathbf{A},k+}^T \mathbf{x}_{\mathbf{R}}\|_2 &= \|\mathbf{V}_{\mathbf{A},k+}^T \mathbf{R} (\mathbf{A} \mathbf{R})^+ \mathbf{b}\|_2 \\
&\leq \|\mathbf{V}_{\mathbf{A},k+}^T \mathbf{R}\|_2 \cdot \|(\mathbf{A} \mathbf{R})^+\|_2 \cdot \|\mathbf{b}\|_2 \\
&\leq \frac{\nu}{(\sqrt{1-\nu^2} - \nu) \sigma_k} \|\mathbf{b}\|_2,
\end{aligned}$$

where we used Lemma 7 to bound $\|\mathbf{V}_{\mathbf{A},k+}^T \mathbf{R}\|_2$ and $\|(\mathbf{A} \mathbf{R})^+\|_2$. \square

3.2. Statistical perspective. We now consider $\mathbf{x}_{\mathbf{R},k}$ from a statistical perspective. We use a similar framework to the one used in the literature to analyze CLS [37, 38, 42]. That is, we consider a fixed design setting in which the rows of \mathbf{A} , $\mathbf{a}_1, \dots, \mathbf{a}_n \in \mathbb{R}^d$, are considered as fixed, and \mathbf{b} 's entries, $b_1, \dots, b_n \in \mathbb{R}$, are

$$b_i = f_i + \xi_i,$$

where f_1, \dots, f_n are fixed values and the noise terms ξ_1, \dots, ξ_n are assumed to be independent random values with zero mean and σ^2 variance. We denote by $\mathbf{f} \in \mathbb{R}^n$ the vector whose i th entry is f_i . The goal is to recover \mathbf{f} from \mathbf{b} (i.e., denoise \mathbf{b}).

The optimal predictor $\mathbf{A} \mathbf{x}^*$ of \mathbf{f} given \mathbf{A} is a minimizer of

$$\min_{\mathbf{x} \in \mathbb{R}^d} \mathbb{E} [\|\mathbf{A} \mathbf{x} - \mathbf{b}\|_2^2 / n],$$

where here, and in subsequent expressions, the expectation is with respect to the noise ξ (if there are multiple minimizers, \mathbf{x}^* is the minimizer with minimum norm). It is easy to verify that $\mathbf{A} \mathbf{x}^* = \mathbf{P}_{\mathbf{A}} \mathbf{f}$.

Given an estimator $\theta = \theta(\mathbf{A}, \mathbf{b})$ of \mathbf{x}^* (which we assume is a random variable since \mathbf{b} is a random variable), its *excess risk* is defined as

$$\mathcal{E}(\theta) := \mathbb{E} [\|\mathbf{A} \theta - \mathbf{A} \mathbf{x}^*\|_2^2 / n].$$

The *ordinary least square estimator* (OLS) $\hat{\mathbf{x}}$ is simply a solution to $\min_{\mathbf{x} \in \mathbb{R}^d} \|\mathbf{A} \mathbf{x} - \mathbf{b}\|_2$: $\hat{\mathbf{x}} := \mathbf{A}^+ \mathbf{b}$. Simple calculations show that

$$\mathcal{E}(\hat{\mathbf{x}}) = \sigma^2 \text{rank}(\mathbf{A}) / n.$$

Thus, if the rank of \mathbf{A} is large, which is usually the case when $d \gg n$, then the excess risk might be large (and it does not asymptotically converge to 0 if $\text{rank}(\mathbf{A}) = \Omega(n)$).

This motivates the use of regularization (e.g., PCR). Indeed, the excess risk of the PCR estimator \mathbf{x}_k can be bounded [38]:

$$(8) \quad \mathcal{E}(\mathbf{x}_k) \leq \frac{\|\mathbf{V}_A^T \mathbf{x}^*\|_\infty^2 \cdot \sum_{i=k+1}^{\min(n,d)} \sigma_i^2}{n} + \frac{\sigma^2 k}{n}.$$

In many scenarios, \mathbf{x}_k has a significantly reduced excess risk in comparison to the excess risk of $\hat{\mathbf{x}}$ (see [38] for a discussion). This motivates the use of PCR when d is large.

In this section, we analyze the excess risk of $\mathbf{x}_{\mathbf{R},k}$ based on properties of \mathbf{R} . The bounds are based on the following identity [38]:² for any \mathbf{M} of appropriate size

$$(9) \quad \mathcal{E}(\mathbf{x}_M) = \mathcal{E}(\mathbf{M}(\mathbf{A}\mathbf{M})^+ \mathbf{b}) = \underbrace{\frac{1}{n} \|(\mathbf{I} - \mathbf{P}_{\mathbf{A}\mathbf{M}}) \mathbf{A} \mathbf{x}^*\|_2^2}_{\mathcal{B}(\mathbf{x}_M)} + \underbrace{\sigma^2 \frac{\text{rank}(\mathbf{A}\mathbf{M})}{n}}_{\mathcal{V}(\mathbf{x}_M)}.$$

In the above, $\mathcal{B}(\mathbf{x}_M)$ can be viewed as a bias term, and $\mathcal{V}(\mathbf{x}_M)$ can be viewed as a variance term. Equation (8) is obtained by bounding the bias term $\mathcal{B}(\mathbf{x}_k)$, although our results lead to a bound on $\mathcal{E}(\mathbf{x}_k)$ that is tighter in some cases (Corollary 11). An immediate corollary of (9) is the following bound for $\mathbf{x}_{\mathbf{R},k}$:

$$(10) \quad \mathcal{E}(\mathbf{x}_{\mathbf{R},k}) = \frac{1}{n} \|(\mathbf{I} - \mathbf{P}_{\mathbf{A}\mathbf{R}\mathbf{V}_{\mathbf{A},k}}) \mathbf{A} \mathbf{x}^*\|_2^2 + \frac{\sigma^2 k}{n}.$$

The following result addresses the case where \mathbf{R} has k orthonormal columns. The conditions are the same as the first part of Theorem 5 (optimization perspective analysis).

THEOREM 9. *Assume that $\text{rank}(\mathbf{A}) \geq k$. Suppose that $\mathbf{R} \in \mathbb{R}^{d \times k}$ has orthonormal columns, and that $d_2(\mathbf{R}, \mathbf{V}_{\mathbf{A},k}) \leq \nu(1 + \nu^2)^{-1/2} < 1$. Then*

$$\mathcal{E}(\mathbf{x}_{\mathbf{R}}) \leq \frac{(1 + \nu) \cdot \|\mathbf{x}^*\|_2^2 \cdot \sigma_{k+1}^2}{n} + \frac{\sigma^2 k}{n}.$$

For the proof, we need the following theorem due to Halko, Martinsson, and Tropp [21].

THEOREM 10 (Theorem 9.1 in [21]). *Let \mathbf{A} be an $m \times n$ matrix with singular value decomposition $\mathbf{A} = \mathbf{U}_A \Sigma_A \mathbf{V}_A^T$. Let $k \geq 0$, and let \mathbf{R} be any matrix such that $\mathbf{V}_{\mathbf{A},k}^T \mathbf{R}$ has full row rank. Then we have*

$$\|(\mathbf{I}_m - \mathbf{P}_{\mathbf{A}\mathbf{R}}) \mathbf{A}\|_2^2 \leq \|\Sigma_{\mathbf{A},k+}\|_2^2 + \|\Sigma_{\mathbf{A},k+} \mathbf{V}_{\mathbf{A},k+}^T \mathbf{R} (\mathbf{V}_{\mathbf{A},k}^T \mathbf{R})^+\|_2^2.$$

Proof of Theorem 9. The condition that $d_2(\mathbf{R}, \mathbf{V}_{\mathbf{A},k}) \leq \nu(1 + \nu^2)^{-1/2} < 1$ ensures that $\mathbf{V}_{\mathbf{A},k}^T \mathbf{R}$ has full rank, and that $\|\tan \Theta(\mathbf{R}, \mathbf{V}_{\mathbf{A},k})\|_2^2 \leq \nu$ (since $\Theta(\mathbf{R}, \mathbf{V}_{\mathbf{A},k})$ is a diagonal matrix whose diagonal values are the inverse cosine of the singular values of

²However, no proof of (9) appears in [38], so for completeness we include a proof in the appendix.

$\mathbf{R}^T \mathbf{V}$, and these, in turn, are all larger than $\sqrt{1 - \nu^2(1 + \nu^2)^{-1}}$. Thus we have

$$\begin{aligned}
\mathcal{B}(\mathbf{x}_R) &= \frac{1}{n} \|(\mathbf{I} - \mathbf{P}_{\mathbf{AR}}) \mathbf{A} \mathbf{x}^*\|_2^2 \\
&\leq \frac{1}{n} \|\mathbf{x}^*\|_2^2 \cdot \left(\|\Sigma_{\mathbf{A},k+}\|_2^2 + \|\Sigma_{\mathbf{A},k+} \mathbf{V}_{\mathbf{A},k+}^T \mathbf{R} (\mathbf{V}_{\mathbf{A},k}^T \mathbf{R})^+ \|_2^2 \right) \\
&\leq \frac{1}{n} \|\mathbf{x}^*\|_2^2 \left(\sigma_{k+1}^2 + \sigma_{k+1}^2 \|\mathbf{V}_{\mathbf{A},k+}^T \mathbf{R} (\mathbf{V}_{\mathbf{A},k}^T \mathbf{R})^+ \|_2^2 \right) \\
&= \frac{1}{n} \|\mathbf{x}^*\|_2^2 (\sigma_{k+1}^2 + \sigma_{k+1}^2 \|\tan \Theta(\mathbf{R}, \mathbf{V}_{\mathbf{A},k})\|_2^2) \\
&\leq \frac{(1 + \nu) \cdot \|\mathbf{x}^*\|_2^2 \cdot \sigma_{k+1}^2}{n},
\end{aligned}$$

where in the first inequality we used Theorem 10 and for the second equality we used Lemma 2. The result now follows from the fact that $\text{rank}(\mathbf{AR}) \leq k$. \square

COROLLARY 11. *For the PCR solution \mathbf{x}_k we have*

$$\mathcal{E}(\mathbf{x}_k) \leq \frac{\|\mathbf{x}^*\|_2^2 \cdot \sigma_{k+1}^2}{n} + \frac{\sigma^2 k}{n}.$$

Next, we consider the general case where \mathbf{R} does not necessarily have orthonormal columns, and potentially has more than k columns. The conditions are the same as the second part of Theorem 5 (optimization perspective).

THEOREM 12. *Suppose that $\mathbf{R} \in \mathbb{R}^{d \times s}$, where $s \geq k$. Assume that $\text{rank}(\mathbf{AR}) \geq k$. If $d_2(\mathbf{U}_{\mathbf{AR},k}, \mathbf{U}_{\mathbf{A},k}) \leq \nu < 1$, then*

$$\mathcal{E}(\mathbf{x}_{R,k}) \leq \mathcal{E}(\mathbf{x}_k) + \frac{(2\nu + \nu^2) \|\mathbf{f}\|_2^2}{n}.$$

Proof. Since \mathbf{AR} has rank at least k , we have $\mathbf{P}_{\mathbf{ARV}_{\mathbf{AR},k}} = \mathbf{P}_{\mathbf{U}_{\mathbf{AR},k}}$. From (10), the fact that $\mathbf{A} \mathbf{x}^* = \mathbf{P}_{\mathbf{A}} \mathbf{f}$, and $\mathbf{P}_{\mathbf{ARV}_{\mathbf{AR},k}} \mathbf{P}_{\mathbf{A}} = \mathbf{P}_{\mathbf{ARV}_{\mathbf{AR},k}}$ (since the range of $\mathbf{ARV}_{\mathbf{AR},k}$ is contained in the range of \mathbf{A}) we have

$$\begin{aligned}
\mathcal{B}(\mathbf{x}_{R,k}) &= \frac{1}{n} \|(\mathbf{I} - \mathbf{P}_{\mathbf{ARV}_{\mathbf{AR},k}}) \mathbf{A} \mathbf{x}^*\|_2^2 \\
&= \frac{1}{n} \|(\mathbf{P}_{\mathbf{A}} - \mathbf{P}_{\mathbf{ARV}_{\mathbf{AR},k}}) \mathbf{f}\|_2^2 \\
&= \frac{1}{n} \|(\mathbf{P}_{\mathbf{A}} - \mathbf{P}_{\mathbf{U}_{\mathbf{A},k}} + \mathbf{P}_{\mathbf{U}_{\mathbf{A},k}} - \mathbf{P}_{\mathbf{ARV}_{\mathbf{AR},k}}) \mathbf{f}\|_2^2 \\
&= \frac{1}{n} \left(\|(\mathbf{P}_{\mathbf{A}} - \mathbf{P}_{\mathbf{U}_{\mathbf{A},k}}) \mathbf{f}\|_2^2 + \|(\mathbf{P}_{\mathbf{U}_{\mathbf{A},k}} - \mathbf{P}_{\mathbf{U}_{\mathbf{AR},k}}) \mathbf{f}\|_2^2 \right. \\
&\quad \left. + 2(\mathbf{P}_{\mathbf{A}} \mathbf{f} - \mathbf{P}_{\mathbf{U}_{\mathbf{A},k}} \mathbf{f})^T (\mathbf{P}_{\mathbf{U}_{\mathbf{A},k}} \mathbf{f} - \mathbf{P}_{\mathbf{ARV}_{\mathbf{AR},k}} \mathbf{f}) \right) \\
&\leq \mathcal{B}(\mathbf{x}_k) + d_2(\mathbf{U}_{\mathbf{AR},k}, \mathbf{U}_{\mathbf{A},k})^2 \frac{\|\mathbf{f}\|_2^2}{n} + \frac{2}{n} |\mathbf{f}^T (\mathbf{P}_{\mathbf{A}} - \mathbf{P}_{\mathbf{U}_{\mathbf{A},k}})^T (\mathbf{P}_{\mathbf{U}_{\mathbf{A},k}} - \mathbf{P}_{\mathbf{ARV}_{\mathbf{AR},k}}) \mathbf{f}|.
\end{aligned}$$

For the cross-terms, we bound

$$\begin{aligned}
& |\mathbf{f}^\top (\mathbf{P}_\mathbf{A} - \mathbf{P}_{\mathbf{U}_{\mathbf{A},k}})^\top (\mathbf{P}_{\mathbf{U}_{\mathbf{A},k}} - \mathbf{P}_{\mathbf{A}\mathbf{R}\mathbf{V}_{\mathbf{A}\mathbf{R},k}}) \mathbf{f}| \\
&= |\mathbf{f}^\top (\mathbf{P}_\mathbf{A} \mathbf{P}_{\mathbf{U}_{\mathbf{A},k}} - \mathbf{P}_\mathbf{A} \mathbf{P}_{\mathbf{A}\mathbf{R}\mathbf{V}_{\mathbf{A}\mathbf{R},k}} - \mathbf{P}_{\mathbf{U}_{\mathbf{A},k}} \mathbf{P}_{\mathbf{U}_{\mathbf{A},k}} + \mathbf{P}_{\mathbf{U}_{\mathbf{A},k}} \mathbf{P}_{\mathbf{A}\mathbf{R}\mathbf{V}_{\mathbf{A}\mathbf{R},k}}) \mathbf{f}| \\
&= |\mathbf{f}^\top (\mathbf{P}_{\mathbf{U}_{\mathbf{A},k}} - \mathbf{P}_{\mathbf{U}_{\mathbf{A}\mathbf{R},k}} - \mathbf{P}_{\mathbf{U}_{\mathbf{A},k}} + \mathbf{P}_{\mathbf{U}_{\mathbf{A},k}} \mathbf{P}_{\mathbf{U}_{\mathbf{A}\mathbf{R},k}}) \mathbf{f}| \\
&= \mathbf{f}^\top (\mathbf{I} - \mathbf{P}_{\mathbf{U}_{\mathbf{A},k}}) \mathbf{P}_{\mathbf{U}_{\mathbf{A}\mathbf{R},k}} \mathbf{f} \\
&= \mathbf{f}^\top \mathbf{P}_{\mathbf{U}_{\mathbf{A},k}}^\perp \mathbf{P}_{\mathbf{U}_{\mathbf{A}\mathbf{R},k}} \mathbf{f} \\
&\leq \|\mathbf{P}_{\mathbf{U}_{\mathbf{A},k}}^\perp \mathbf{P}_{\mathbf{U}_{\mathbf{A}\mathbf{R},k}}\|_2 \cdot \|\mathbf{f}\|_2^2.
\end{aligned}$$

Since both $\mathbf{U}_{\mathbf{A}\mathbf{R},k}$ and $\mathbf{U}_{\mathbf{A},k}$ are full rank, we have (Theorem 1)

$$\|\mathbf{P}_{\mathbf{U}_{\mathbf{A},k}}^\perp \mathbf{P}_{\mathbf{U}_{\mathbf{A}\mathbf{R},k}}\|_2 = \|\mathbf{P}_{\mathbf{U}_{\mathbf{A},k}} - \mathbf{P}_{\mathbf{U}_{\mathbf{A}\mathbf{R},k}}\|_2 = d_2(\mathbf{U}_{\mathbf{A}\mathbf{R},k}, \mathbf{U}_{\mathbf{A},k}).$$

Thus, we find that

$$\mathcal{B}(\mathbf{x}_{\mathbf{R},k}) \leq \mathcal{B}(\mathbf{x}_k) + (2\nu + \nu^2) \frac{\|\mathbf{f}\|_2^2}{n}.$$

We reach the bound in the theorem statement by adding the variance $\mathcal{V}(\mathbf{x}_{\mathbf{R},k})$, which is equal to the variance of \mathbf{x}_k because the ranks are equal. \square

Discussion. Theorem 9 shows that if \mathbf{R} is a good approximation to $\mathbf{V}_{\mathbf{A},k}$, then there is a small relative increase to the bias term, while the variance term does not change. Since we are mainly interested in keeping the asymptotic behavior of the excess risk (as n goes to infinity), a fixed ν of modest value suffices. However, for this result to hold, \mathbf{R} has to have exactly k columns, and those columns should be orthonormal. Without these restrictions, we need to resort to Theorem 12. In that theorem, we get (if the conditions are met) only an additive increase in the bias term. Thus if, for example, $\|\mathbf{f}\|_2^2/n \rightarrow c$ as $n \rightarrow \infty$ for some constant c , then ν should tend to 0 as n goes to infinity, but a constant value should suffice if n is fixed.

4. Sketched PCR and PCP. In the previous section, we considered general conditions on \mathbf{R} which ensure that $\mathbf{x}_{\mathbf{R},k}$ is an approximate solution to the PCR/PCP problem. In this section, we propose algorithms to generate \mathbf{R} for which these conditions hold. The main technique we employ is *matrix sketching*. The idea is to first multiply the data matrix \mathbf{A} by some random transformation (e.g., a random projection) and extract an approximate subspace from the compressed matrix.

4.1. Dimensionality reduction using sketching. The compression (multiplication by a random matrix) alluded to in the previous paragraph can be applied from the left side, the right side, or both. In left sketching, which is more appropriate if the input matrix has many rows and a modest number of columns, we propose using $\mathbf{R} = \mathbf{V}_{\mathbf{S}\mathbf{A},k}$, where \mathbf{S} is some sketching matrix (we discuss a couple of options shortly). In right sketching, which is more appropriate if the input matrix has many columns and a modest number of rows, we propose using $\mathbf{R} = \mathbf{G}^\top$, where \mathbf{G} is some sketching matrix. Two-sided sketching, $\mathbf{R} = \mathbf{G}^\top \mathbf{V}_{\mathbf{S}\mathbf{A}\mathbf{G}^\top,k}$, is suited for the case in which the number of columns and the number of rows are large.

The sketching matrices, \mathbf{S} and \mathbf{G} , are randomized dimensionality reduction transformations. Quite a few sketching transforms have been proposed in the literature in recent years. For concreteness, we consider two specific cases, though our results hold for other sketching transformations as well (though some modifications in the bounds might be necessary). The first, which we refer to as “subgaussian map,” is a random

matrix in which every entry of the matrix is sampled i.i.d. from some subgaussian distribution (e.g., $N(0,1)$) and the matrix is appropriately scaled (however, scaling is not necessary in our case). The second transform is a sparse embedding matrix, in which each column is sampled uniformly and independently from the set of scaled identity vectors and multiplied by a random sign. We refer to such a matrix as a COUNTSKETCH matrix [8, 44].

Both transformations described above, and a few others, have, with high probability, provided enough rows are used, the following property, which we refer to as *approximate Gram property*.

DEFINITION 13. Let $\mathbf{X} \in \mathbb{R}^{m \times n}$ be a fixed matrix. For $\epsilon, \delta \in (0, 1/2)$, a distribution \mathcal{D} on matrices with m columns has the (ϵ, δ) -approximate Gram matrix property for \mathbf{X} if

$$\Pr_{\mathbf{S} \sim \mathcal{D}} \left(\|\mathbf{X}^T \mathbf{S}^T \mathbf{S} \mathbf{X} - \mathbf{X}^T \mathbf{X}\|_2 \geq \epsilon \|\mathbf{X}\|_2^2 \right) \leq \delta.$$

Recent results by Cohen, Nelson, and Woodruff [14]³ show that when \mathbf{S} has independent subgaussian entries, then as long as the number of rows in \mathbf{S} is $\Omega((\mathbf{sr}(\mathbf{X}) + \log(1/\delta))/\epsilon^2)$, we have the (ϵ, δ) -approximate Gram property for \mathbf{X} . If \mathbf{S} is a COUNTSKETCH matrix, then as long as the number of rows in \mathbf{S} is $\Omega(\mathbf{sr}(\mathbf{X})^2/(\epsilon^2\delta))$, we have the (ϵ, δ) -approximate Gram property for \mathbf{X} [14].

We first describe our results for the various modes of sketching, and then discuss algorithmic issues and computational complexity.

THEOREM 14 (left sketching). Let $\nu, \delta \in (0, 1/2)$ and denote

$$\epsilon = \frac{\nu(1 + \nu^2)^{-1/2}}{1 + \nu(1 + \nu^2)^{-1/2}} \cdot \mathbf{gap}_k(\mathbf{A}).$$

Suppose that \mathbf{S} is sampled from a distribution that provides a (ϵ, δ) -approximate Gram matrix for \mathbf{A} . Then for $\mathbf{R} = \mathbf{V}_{\mathbf{SA},k}$, with probability $1 - \delta$, the approximate solution \mathbf{x}_R is a $\left(\frac{\sigma_{k+1}}{\sigma_k} \nu, \frac{\nu}{(\sqrt{1-\nu^2}-\nu)\sigma_k}\right)$ -approximate PCR and

$$\mathcal{E}(\mathbf{x}_R) \leq \frac{(1 + \nu) \cdot \|\mathbf{x}^*\|_2^2 \cdot \sigma_{k+1}^2}{n} + \frac{\sigma^2 k}{n}.$$

Thus if, for example, \mathbf{S} is a COUNTSKETCH matrix, then the conditions are met when the number of rows in \mathbf{S} is

$$\Omega \left(\frac{\mathbf{sr}(\mathbf{A})^2}{\mathbf{gap}_k(\mathbf{A})^2 \nu^2 \delta} \right).$$

In another example, if \mathbf{S} is a subgaussian map, then the conditions are met when the number of rows in \mathbf{S} is

$$\Omega \left(\frac{\mathbf{sr}(\mathbf{A}) + \log(1/\delta)}{\mathbf{gap}_k(\mathbf{A})^2 \nu^2} \right).$$

Proof. Due to Theorems 5 and 9, it suffices to show that $d_2(\mathbf{R}, \mathbf{V}_{\mathbf{A},k}) \leq \nu(1 + \nu^2)^{-1/2}$. Under the conditions of the theorem, with probability of at least $1 - \delta$, we have $\|\mathbf{A}^T \mathbf{S}^T \mathbf{S} \mathbf{A} - \mathbf{A}^T \mathbf{A}\|_2 \leq \epsilon \|\mathbf{A}\|_2^2$. If that is indeed the case, $\mathbf{A}^T \mathbf{S}^T \mathbf{S} \mathbf{A}$ has rank at least k since $\mathbf{A}^T \mathbf{S}^T \mathbf{S} \mathbf{A}$ and $\mathbf{A}^T \mathbf{A}$ are symmetric matrices and we know that

³Theorem 1 in [14] with $k = \mathbf{sr}(\mathbf{X})$.

$\sigma_i^2(\mathbf{A}^\top \mathbf{S}^\top \mathbf{S} \mathbf{A}) = \sigma_i^2(\mathbf{A}^\top \mathbf{A}) \pm \|\mathbf{A}^\top \mathbf{S}^\top \mathbf{S} \mathbf{A} - \mathbf{A}^\top \mathbf{A}\|_2$ (Weyl's theorem and the fact that $\sigma_k^2 > \epsilon \sigma_1^2$). Furthermore, since $\nu > 0$ we have $\epsilon < \text{gap}_k(\mathbf{A})$, and Theorem 3 implies that

$$\begin{aligned} d_2(\mathbf{R}, \mathbf{V}_{\mathbf{A},k}) &\leq \frac{\|\mathbf{A}^\top \mathbf{A} - \mathbf{A}^\top \mathbf{S}^\top \mathbf{S} \mathbf{A}\|_2}{(\sigma_k^2 - \sigma_{k+1}^2) - \|\mathbf{A}^\top \mathbf{A} - \mathbf{A}^\top \mathbf{S}^\top \mathbf{S} \mathbf{A}\|_2} \\ &\leq \frac{\epsilon}{\text{gap}_k(\mathbf{A}) - \epsilon} \\ &\leq \nu(1 + \nu^2)^{-1/2}. \end{aligned}$$

Thus, we have shown that with probability $1 - \delta$ we have $d_2(\mathbf{R}, \mathbf{V}_{\mathbf{A},k}) \leq \nu(1 + \nu^2)^{-1/2}$, as required. \square

THEOREM 15 (right sketching). *Let $\nu, \delta \in (0, 1/2)$ and denote*

$$\epsilon = \frac{\nu}{1 + \nu} \cdot \text{gap}_k(\mathbf{A}).$$

Suppose that \mathbf{G} is sampled from a distribution that provides an (ϵ, δ) -approximate Gram matrix for \mathbf{A}^\top . Then for $\mathbf{R} = \mathbf{G}^\top$, with probability $1 - \delta$, the approximate solution $\mathbf{A}\mathbf{x}_{\mathbf{R},k}$ is a (ν, ν) -approximate PCP and

$$\mathcal{E}(\mathbf{x}_{\mathbf{R},k}) \leq \mathcal{E}(\mathbf{x}_k) + \frac{(2\nu + \nu^2)\|\mathbf{f}\|_2^2}{n}.$$

Thus if, for example, \mathbf{G} is a COUNTSKETCH matrix, then the conditions are met when the number of rows in \mathbf{G} is

$$\Omega\left(\frac{\text{sr}(\mathbf{A})^2}{\text{gap}_k(\mathbf{A})^2 \nu^2 \delta}\right).$$

In another example, if \mathbf{G} is a subgaussian map, then the conditions are met when the number of rows in \mathbf{G} is

$$\Omega\left(\frac{\text{sr}(\mathbf{A}) + \log(1/\delta)}{\text{gap}_k(\mathbf{A})^2 \nu^2}\right).$$

Proof. Due to Theorem 5 it suffices to show that $d_2(\mathbf{U}_{\mathbf{A}\mathbf{R},k}, \mathbf{U}_{\mathbf{A},k}) \leq \nu$. Under the conditions of the theorem, with probability at least $1 - \delta$, we have $\|\mathbf{A}\mathbf{G}^\top \mathbf{G} \mathbf{A}^\top - \mathbf{A}\mathbf{A}^\top\|_2 \leq \epsilon \|\mathbf{A}\|_2^2$. If that is indeed the case, $\mathbf{A}\mathbf{G}^\top \mathbf{G} \mathbf{A}^\top$ has rank at least k since $\mathbf{A}\mathbf{G}^\top \mathbf{G} \mathbf{A}^\top$ and $\mathbf{A}\mathbf{A}^\top$ are symmetric matrices and we know that $\sigma_i^2(\mathbf{A}\mathbf{G}^\top \mathbf{G} \mathbf{A}^\top) = \sigma_i^2(\mathbf{A}\mathbf{A}^\top) \pm \|\mathbf{A}\mathbf{G}^\top \mathbf{G} \mathbf{A}^\top - \mathbf{A}\mathbf{A}^\top\|_2$ (Weyl's theorem and the fact that $\sigma_k^2 > \epsilon \sigma_1^2$). Furthermore, since $\nu > 0$ we have $\epsilon < \text{gap}_k(\mathbf{A})$, and Theorem 3 implies

$$\begin{aligned} d_2(\mathbf{U}_{\mathbf{A}\mathbf{R},k}, \mathbf{U}_{\mathbf{A},k}) &\leq \frac{\|\mathbf{A}\mathbf{G}^\top \mathbf{G} \mathbf{A}^\top - \mathbf{A}\mathbf{A}^\top\|_2}{(\sigma_k^2 - \sigma_{k+1}^2) - \|\mathbf{A}\mathbf{G}^\top \mathbf{G} \mathbf{A}^\top - \mathbf{A}\mathbf{A}^\top\|_2} \\ &\leq \frac{\epsilon}{\text{gap}_k(\mathbf{A}) - \epsilon} \\ &\leq \nu. \end{aligned}$$

Thus, we have shown that with probability $1 - \delta$ we have $d_2(\mathbf{U}_{\mathbf{A}\mathbf{R},k}, \mathbf{U}_{\mathbf{A},k}) \leq \nu$, as required. \square

THEOREM 16 (two-sided sketching). *Let $\nu, \delta \in (0, 1/2)$ and denote*

$$\epsilon_2 = \frac{\nu}{2(1 + \nu/2)} \cdot \mathbf{gap}_k(\mathbf{A}).$$

Suppose that \mathbf{G} is sampled from a distribution that provides an $(\epsilon_2, \delta/2)$ -approximate Gram matrix for \mathbf{A}^\top . Denote

$$\epsilon_1 = \frac{\nu(1 + \nu^2/4)^{-1/2}/2}{1 + \nu(1 + \nu^2/4)^{-1/2}/2} \cdot \mathbf{gap}_k(\mathbf{A}\mathbf{G}^\top)$$

and suppose \mathbf{S} is sampled from a distribution that provides an $(\epsilon_1, \delta/2)$ -approximate Gram matrix for $\mathbf{A}\mathbf{G}^\top$. Then for $\mathbf{R} = \mathbf{G}^\top \mathbf{V}_{\mathbf{S}\mathbf{A}\mathbf{G}^\top, k}$ with probability $1 - \delta$ the approximate solution $\mathbf{A}\mathbf{x}_{\mathbf{R}, k}$ is a (ν, ν) -approximate PCP and

$$\mathcal{E}(\mathbf{x}_{\mathbf{R}, k}) \leq \mathcal{E}(\mathbf{x}_k) + \frac{(2\nu + \nu^2)\|\mathbf{f}\|_2^2}{n}.$$

Thus if, for example, \mathbf{S} is a COUNTSKETCH matrix and \mathbf{G} is a subgaussian map, then the conditions hold when the number of rows of \mathbf{S} is

$$\Omega\left(\frac{\mathbf{sr}(\mathbf{A}\mathbf{G}^\top)^2}{\mathbf{gap}_k(\mathbf{A}\mathbf{G}^\top)^2 \nu^2 \delta}\right)$$

and the number of rows in \mathbf{G} is

$$\Omega\left(\frac{\mathbf{sr}(\mathbf{A}) + \log(1/\delta)}{\mathbf{gap}_k(\mathbf{A})^2 \nu^2}\right).$$

Proof. Due to Theorem 5 it suffices to show that $d_2(\mathbf{U}_{\mathbf{A}\mathbf{R}, k}, \mathbf{U}_{\mathbf{A}, k}) \leq \nu$.

Under the conditions of the theorem, with probability at least $1 - \delta/2$, we have $\|(\mathbf{A}\mathbf{G}^\top)^\top \mathbf{S}^\top \mathbf{S} \mathbf{A}\mathbf{G}^\top - (\mathbf{A}\mathbf{G}^\top)^\top \mathbf{A}\mathbf{G}^\top\|_2 \leq \epsilon_1 \|\mathbf{A}\mathbf{G}^\top\|_2^2$, and with probability at least $1 - \delta/2$, we have $\|\mathbf{A}\mathbf{G}^\top \mathbf{G} \mathbf{A}^\top - \mathbf{A} \mathbf{A}^\top\|_2 \leq \epsilon_2 \|\mathbf{A}\|_2^2$. Thus, both inequalities hold with probability at least $1 - \delta$. If that is indeed the case, $\mathbf{A}\mathbf{G}^\top \mathbf{G} \mathbf{A}^\top$ has rank at least k since $\mathbf{A}\mathbf{G}^\top \mathbf{G} \mathbf{A}^\top$ and $\mathbf{A} \mathbf{A}^\top$ are symmetric matrices, and we know that $\sigma_i^2(\mathbf{A}\mathbf{G}^\top \mathbf{G} \mathbf{A}^\top) = \sigma_i^2(\mathbf{A} \mathbf{A}^\top) \pm \|\mathbf{A}\mathbf{G}^\top \mathbf{G} \mathbf{A}^\top - \mathbf{A} \mathbf{A}^\top\|_2$ (Weyl's theorem and the fact that $\sigma_k^2 > \epsilon_2 \sigma_1^2$). Moreover, $(\mathbf{A}\mathbf{G}^\top)^\top \mathbf{S}^\top \mathbf{S} \mathbf{A}\mathbf{G}^\top$ has rank at least k since $(\mathbf{A}\mathbf{G}^\top)^\top \mathbf{S}^\top \mathbf{S} \mathbf{A}\mathbf{G}^\top$ and $\mathbf{A}\mathbf{G}^\top \mathbf{G} \mathbf{A}^\top$ are symmetric matrices, and we know that $\sigma_i^2((\mathbf{A}\mathbf{G}^\top)^\top \mathbf{S}^\top \mathbf{S} \mathbf{A}\mathbf{G}^\top) = \sigma_i^2(\mathbf{A}\mathbf{G}^\top \mathbf{G} \mathbf{A}^\top) \pm \|(\mathbf{A}\mathbf{G}^\top)^\top \mathbf{S}^\top \mathbf{S} \mathbf{A}\mathbf{G}^\top - \mathbf{A}\mathbf{G}^\top \mathbf{G} \mathbf{A}^\top\|_2$ (Weyl's theorem and the fact that $\sigma_k^2(\mathbf{A}\mathbf{G}^\top) > \epsilon_1 \sigma_1^2(\mathbf{A}\mathbf{G}^\top)$). Since $\nu > 0$ we have $\epsilon_1 < \mathbf{gap}_k(\mathbf{A}\mathbf{G}^\top)$ and Theorem 3 implies

$$\begin{aligned} & d_2(\mathbf{V}_{\mathbf{A}\mathbf{G}^\top \mathbf{V}_{\mathbf{S}\mathbf{A}\mathbf{G}^\top, k}}, \mathbf{V}_{\mathbf{A}\mathbf{G}^\top, k}) \\ & \leq \frac{\|(\mathbf{A}\mathbf{G}^\top)^\top \mathbf{S}^\top \mathbf{S} \mathbf{A}\mathbf{G}^\top - (\mathbf{A}\mathbf{G}^\top)^\top \mathbf{A}\mathbf{G}^\top\|_2}{\sigma_k^2(\mathbf{A}\mathbf{G}^\top) - \sigma_{k+1}^2(\mathbf{A}\mathbf{G}^\top) - \|(\mathbf{A}\mathbf{G}^\top)^\top \mathbf{S}^\top \mathbf{S} \mathbf{A}\mathbf{G}^\top - (\mathbf{A}\mathbf{G}^\top)^\top \mathbf{A}\mathbf{G}^\top\|_2} \\ & \leq \frac{\epsilon_1 \|\mathbf{A}\mathbf{G}^\top\|_2^2}{\sigma_k^2(\mathbf{A}\mathbf{G}^\top) - \sigma_{k+1}^2(\mathbf{A}\mathbf{G}^\top) - \epsilon_1 \|\mathbf{A}\mathbf{G}^\top\|_2^2} \\ & = \frac{\epsilon_1}{\mathbf{gap}_k(\mathbf{A}\mathbf{G}^\top) - \epsilon_1} \\ & \leq \nu(1 + \nu^2/4)^{-1/2}/2. \end{aligned}$$

TABLE 1

Computational complexity of computing $\mathbf{x}_{\mathbf{R},k}$ and CLS for various options \mathbf{R} . For brevity, we omit the $O()$ from the notation.

		$\mathbf{x}_{\mathbf{R},k}$	CLS
Left sketching $\mathbf{R} = \mathbf{V}_{\mathbf{S}\mathbf{A},k}$	Subgaussian \mathbf{S}	$s_1 \cdot \text{nnz}(\mathbf{A}) + s_1 d \min(s_1, d) + nk^2$	N/A
	COUNTSKETCH \mathbf{S}	$k \cdot \text{nnz}(\mathbf{A}) + s_2 d \min(s_2, d) + nk^2$	N/A
Right sketching $\mathbf{R} = \mathbf{G}^T$	subgaussian \mathbf{G}	$t_1 \cdot \text{nnz}(\mathbf{A}) + nt_1 \min(n, t_1) + t_1 k \min(n, d)$	$t_1 \cdot \text{nnz}(\mathbf{A}) + nt_1 \min(n, t_1)$
	COUNTSKETCH \mathbf{G}	$\text{nnz}(\mathbf{A}) + nt_2 \min(n, t_2) + t_2 k \min(n, d)$	$\text{nnz}(\mathbf{A}) + nt_2 \min(n, t_2)$
Two sided $\mathbf{R} = \mathbf{G}^T \mathbf{V}_{\mathbf{S}\mathbf{A}\mathbf{G}^T,k}$	COUNTSKETCH \mathbf{G} and \mathbf{S}	$\text{nnz}(\mathbf{A}) + s_2 k^2 + k \min(nt_2, \text{nnz}(\mathbf{A})) + nk^2$	N/A

From Lemma 8 (with $\mathbf{A}\mathbf{G}^T$) we get that $d_2(\mathbf{U}_{\mathbf{A}\mathbf{G}^T \mathbf{V}_{\mathbf{S}\mathbf{A}\mathbf{G}^T,k}}, \mathbf{U}_{\mathbf{A}\mathbf{G}^T,k}) \leq \frac{\sigma_{k+1}(\mathbf{A}\mathbf{G}^T)}{\sigma_k(\mathbf{A}\mathbf{G}^T)}(\nu/2) \leq \nu/2$.

We now bound

$$\begin{aligned} d_2(\mathbf{U}_{\mathbf{A}\mathbf{R},k}, \mathbf{U}_{\mathbf{A},k}) &\leq d_2(\mathbf{U}_{\mathbf{A}\mathbf{R},k}, \mathbf{U}_{\mathbf{A}\mathbf{G}^T,k,k}) + d_2(\mathbf{U}_{\mathbf{A}\mathbf{G}^T,k,k}, \mathbf{U}_{\mathbf{A},k}) \\ &\leq \nu, \end{aligned}$$

where we similarly use Theorem 15 to bound $d_2(\mathbf{U}_{\mathbf{A}\mathbf{G}^T,k}, \mathbf{U}_{\mathbf{A},k}) \leq \nu/2$.

Thus, we have shown that with probability $1 - \delta$ we have $d_2(\mathbf{U}_{\mathbf{A}\mathbf{R},k}, \mathbf{U}_{\mathbf{A},k}) \leq \nu$, as required. \square

4.2. Fast approximate PCR/PCP. A prototypical algorithm for approximate PCR/PCP is to compute $\mathbf{x}_{\mathbf{R},k}$ with some choice of sketching-based \mathbf{R} . There are quite a few design choices that need to be made in order to turn this prototypical algorithm into a concrete algorithm, e.g., whether to use left, right, or two-sided sketching to form \mathbf{R} , and which sketch transform to use. There are various trade-offs; e.g., using COUNTSKETCH results in faster sketching, but usually requires larger sketch sizes. Furthermore, in computing $\mathbf{x}_{\mathbf{R},k}$ there are also algorithmic choices to be made with respect to choosing the order of matrix multiplications: in computing $(\mathbf{A}\mathbf{R}\mathbf{V}_{\mathbf{A}\mathbf{R},k})^+ \mathbf{b}$ should we first compute $\mathbf{A}\mathbf{R}$ and then multiply by $\mathbf{V}_{\mathbf{A}\mathbf{R},k}$, or vice versa? Likely, there is no one-size-fits-all algorithm, and different profiles of the input matrix (in particular, the size and sparsity level) call for a different variant of the prototypical algorithm.

Table 1 summarizes the running time complexity of several design options. In order to better make sense between these different choices, we first summarize the running time complexity of various design choices using the optimal implementation (from an asymptotic running time complexity perspective). To make the discussion manageable, we consider only subgaussian maps and COUNTSKETCH. Furthermore, for the sake of the analysis, we make some assumptions and adopt some notational conventions. First, we assume that computing $\mathbf{B}^+ \mathbf{c}$ for some $\mathbf{B} \in \mathbb{R}^{m \times n}$ and \mathbf{c} is done via straightforward methods based on QR or SVD factorizations, and as such takes $O(mn \min(m, n))$. We consider using fast sketch-based approximate least squares algorithms in the next subsection. Next, we let the sketch sizes be parameters in the complexity. In the discussion, we use our theoretical results to deduce reasonable assumptions on how these parameters are set, and thus to reason about the final complexity of sketched PCR/PCP. We denote the number of rows in the left sketch matrix \mathbf{S} by s_1 for a subgaussian map, and by s_2 for COUNTSKETCH. We denote the number of rows in the left sketch matrix \mathbf{G} by t_1 for a subgaussian map, and by t_2 for COUNTSKETCH. Finally, we assume $\text{nnz}(\mathbf{A}) \geq \max(n, d)$, and that all sketch sizes are greater than k .

Table 1 also lists, where relevant, the complexity of the CLS solution \mathbf{x}_R .

Discussion. We first compare the computational complexity of CLS to the computational complexity of our proposed right sketching algorithm. For both choices of \mathbf{G} we have for sketched PCP an additional term of $O(tk \min(n, d))$. However, close inspection reveals that this term is dominated by the term $O(nt \min(n, t))$. Thus our proposed algorithm has the same asymptotic complexity as CLS for the same sketch size. However, our algorithm does not mix regularization and compression and comes with stronger theoretical guarantees.

Next, in order to compare subgaussian maps to COUNTSKETCH, we first make some simplified assumptions on the required approximation quality ν , the relative eigengap $\mathbf{gap}_k(\mathbf{A})$, and the rank parameter k : ν is fixed, $\mathbf{gap}_k(\mathbf{A})$ is bounded from below by a constant, and we have $k = O(\mathbf{sr}(\mathbf{A}))$. The first assumption is justified if we are satisfied with fixed suboptimality in the objective (optimization perspective), or a small constant multiplicative increase in excess risk if left sketching is used, or n is fixed (statistical perspective). The first assumption is somewhat less justified from a statistical point of view when $n \rightarrow \infty$ and right sketching is used. The rationale behind the second assumption is that the PCR/PCP problem is in a sense ill-posed if there is a tiny eigengap. The third assumption is motivated by the fact that the stable rank is a measure of the number of large singular values, which are typically singular values that correspond to the signal rather than noise. With these assumptions, our theoretical results establish that $s_1, t_1 = O(k)$ and $s_2, t_2 = O(k^2)$ suffice. It is important to stress that we make these assumptions only for the sake of comparing the different sketching options, and we do not claim that these assumptions always hold, or that our proposed algorithms work only when these assumptions hold.

For left sketching, with these assumptions, we have a complexity of $O(k\mathbf{nnz}(\mathbf{A}) + k^2 \max(n, d))$ for subgaussian maps and $O(k\mathbf{nnz}(\mathbf{A}) + dk^2 \min(k^2, d) + nk^2)$ for COUNTSKETCH. Clearly, better asymptotic complexity is achieved with subgaussian maps. For right sketching, with these assumptions, we have a complexity of $O(k\mathbf{nnz}(\mathbf{A}) + nk^2)$ for subgaussian sketch and $O(\mathbf{nnz}(\mathbf{A}) + nk^2 \min(n, k^2))$ for COUNTSKETCH. The complexity in terms of the input sparsity $\mathbf{nnz}(\mathbf{A})$, which is arguably the dominant term, is better for COUNTSKETCH. For two-sided sketching, we have a complexity of $O(\mathbf{nnz}(\mathbf{A}) + k^4 + k \min(nk^2, \mathbf{nnz}(\mathbf{A})) + nk^2)$.

If $n \gg d$ and $\mathbf{nnz}(\mathbf{A}) = O(n)$ (sparse input matrix, and constant amount of nonzero features per data point), left sketching gives better asymptotic complexity. If $n \gg d$ and $\mathbf{nnz}(\mathbf{A}) = nd$ (full data matrix), left sketching has better complexity unless $d \gg k^3$. Furthermore, left sketching gives stronger theoretical guarantees. *Thus, for $n \gg d$ we advocate the use of left sketching.* If $d \gg n$ and $\mathbf{nnz}(\mathbf{A}) = O(d)$ (sparse input matrix), right sketching with a subgaussian maps always has better complexity than left sketching, and potentially (but not always) right sketching with COUNTSKETCH has even better complexity. If $d \gg n$ and $\mathbf{nnz}(\mathbf{A}) = nd$ (full data matrix), right sketching with subgaussian maps has the same complexity as left sketching, and potentially (but not always) right sketching with COUNTSKETCH has even better complexity (if d is sufficiently larger than n). *Thus, for $d \gg n$ we advocate the use of right sketching.* If $n \approx d$ (both very large) and $\mathbf{nnz}(\mathbf{A}) = n$, then it is possible to have $O(nk^2)$ with all three options (left, right, and two sided), as long as $k^2 \leq n$. A similar conclusion is achieved if $n \approx d$ and $\mathbf{nnz}(\mathbf{A}) = nd$, but if $k^2 \ll n$, then two-sided sketch is better.

4.3. Input sparsity approximate PCP. In this section, we propose an input sparsity algorithm for approximate PCP. By “input sparsity algorithm,” we mean an

Algorithm 1 Input Sparsity Approximate PCP.

-
- 1: **Input:** $\mathbf{A} \in \mathbb{R}^{n \times d}$, $\mathbf{b} \in \mathbb{R}^n$, $k \leq \min(n, d)$, $s, t \geq k$, $\epsilon \in (0, 1)$
 - 2: Generate two COUNTSKETCH matrices $\mathbf{S} \in \mathbb{R}^{s \times n}$ and $\mathbf{G} \in \mathbb{R}^{t \times n}$.
 - 3: Remove from \mathbf{G} any row that is zero.
 - 4: $\mathbf{C} \leftarrow \mathbf{A}\mathbf{G}^T$.
 - 5: $\mathbf{D} \leftarrow \mathbf{S}\mathbf{C}$.
 - 6: Compute $\mathbf{V}_{\mathbf{D},k}$, the k dominant right invariant space of \mathbf{G} (via SVD).
 - 7: For the analysis (no need to compute): $\mathbf{R} = \mathbf{G}^T \mathbf{V}_{\mathbf{D},k}$.
 - 8: Solve $\min_{\gamma} \|\mathbf{C}\mathbf{V}_{\mathbf{D},k}\gamma - \mathbf{b}\|_2$ to ϵ/d accuracy using input sparsity least squares regression (see [12, section 7.7]). (Do not compute $\mathbf{C}\mathbf{V}_{\mathbf{D},k}$. In each iteration, multiplying a vector by $\mathbf{C}\mathbf{V}_{\mathbf{D},k}$ is performed by first multiplying by $\mathbf{V}_{\mathbf{D},k}$ and then by \mathbf{C} .)
 - 9: Return $\mathbf{y} \leftarrow \mathbf{G}^T(\mathbf{V}_{\mathbf{D},k}\tilde{\gamma})$, where $\tilde{\gamma}$ is the output of the previous step.
-

algorithm whose running time is $O(\text{nnz}(\mathbf{A}) \log(d/\epsilon) + \text{poly}(k, s, t, \log(1/\epsilon)))$, where ϵ is some accuracy parameter (see formal theorem statement).

The basic idea is to use two-sided sketching, with an additional modification of using input sparsity algorithms to approximate $(\mathbf{A}\mathbf{R})^+ \mathbf{b} = \arg \min_{\gamma} \|\mathbf{A}\mathbf{R}\gamma - \mathbf{b}\|_2$. Specifically, we propose using the algorithm recently suggested by Clarkson and Woodruff [12]. A pseudocode description of our input sparsity approximate PCP algorithm is listed in Algorithm 1. We have the following statement about the algorithm.

THEOREM 17. *Run Algorithm 1 with ϵ, s, t, k as parameters. Under exact arithmetic,⁴ after*

$$O(\text{nnz}(\mathbf{A}) \log(d/\epsilon) + \log(d/\epsilon)tk + sk^2 + t^3k + k^3 \log^2 k)$$

operations, with probability $2/3$, the algorithm will return a \mathbf{y} such that

$$\|\mathbf{y} - \mathbf{x}_{\mathbf{R}}\|_2^2 \leq \epsilon \|\mathbf{x}_{\mathbf{R}}\|_2^2.$$

Proof. Denote $\mathbf{B} = \mathbf{A}\mathbf{R}$, and consider using the iterative method described in [12, section 7.7] to approximately solve $\min_{\gamma} \|\mathbf{B}\gamma - \mathbf{b}\|_2$. Denote the optimal solution by $\gamma_{\mathbf{R}}$, and the solution that our algorithm found by $\tilde{\gamma}$. Theorem 7.14 in [12] states that after the $O(\log(d/\epsilon))$ iterations the algorithm would have returned $\tilde{\gamma}$ such that

$$(11) \quad \|\mathbf{B}\mathbf{Z}(\tilde{\gamma} - \gamma_{\mathbf{R}})\|_2^2 \leq (\epsilon/d) \|\mathbf{B}\mathbf{Z}\gamma_{\mathbf{R}}\|_2^2$$

for some invertible \mathbf{Z} found by the algorithm. Furthermore, $\kappa(\mathbf{B}\mathbf{Z}) = O(1)$, where $\kappa(\cdot)$ is the condition number (ratio between the largest singular value and smallest). Equation (11) implies that $\|\tilde{\gamma} - \gamma_{\mathbf{R}}\|_2^2 \leq \kappa(\mathbf{B}\mathbf{Z})^2 (\epsilon/d) \|\gamma_{\mathbf{R}}\|_2^2 = O(\epsilon/d) \|\gamma_{\mathbf{R}}\|_2^2$. Now, noticing that $\mathbf{x}_{\mathbf{R}} = \mathbf{R}\gamma_{\mathbf{R}}$ and $\mathbf{y} = \mathbf{R}\tilde{\gamma}$, we find that

$$\|\mathbf{y} - \mathbf{x}_{\mathbf{R}}\|_2^2 \leq O(\epsilon/d) \kappa(\mathbf{R})^2 \|\mathbf{x}_{\mathbf{R}}\|_2^2.$$

We now need to bound $\kappa(\mathbf{R}) = \kappa(\mathbf{G}^T \mathbf{V}_{\mathbf{S}\mathbf{A}\mathbf{G}^T, k}) = \kappa(\mathbf{G}^T)$, where \mathbf{G} is a COUNTSKETCH matrix. Since \mathbf{G} has a single nonzero in each column, then $\|\mathbf{G}^T\|_2^2 \leq \|\mathbf{G}^T\|_F^2 \leq d$.

⁴The results are likely too optimistic for inexact arithmetic. We leave the numerical analysis to future work.

Furthermore, since we removed zero column from \mathbf{G}^T , for any \mathbf{x} the vector $\mathbf{G}^T \mathbf{x}$ has in one of its coordinates any coordinate of \mathbf{x} , so $\sigma_{\min}(\mathbf{G}^T) \geq 1$. So we found that $\kappa(\mathbf{G}^T)^2 \leq d$. We conclude that

$$\|\mathbf{y} - \mathbf{x}_R\|_2^2 \leq O(\epsilon) \|\mathbf{x}_R\|_2^2.$$

Adjusting ϵ to compensate for the constants completes the proof. \square

5. Extensions.

5.1. Streaming algorithm. We now consider computing an approximate PCR/PCP in the streaming model. We consider a one-pass row-insertion streaming model, in which the rows of \mathbf{A} , $\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_n$, and the corresponding entries in \mathbf{b} , b_1, b_2, \dots, b_n , are presented one by one and once only (i.e., in a stream). The goal is to use $o(n)$ memory (so \mathbf{A} cannot be stored in memory). The relevant resources to be bounded for numerical linear algebra in the streaming model are storage, update time (time spent per row), and final computation time (at the end of the stream) [11]. Our goal is to bound these by $O(\text{poly}(d))$.

Our proposed streaming algorithm for approximate PCP uses left sketching. It is easy to verify that if \mathbf{S} is a subgaussian map or COUNTSKETCH, then $\mathbf{R} = \mathbf{V}_{\mathbf{SA},k}$ can be computed in the streaming model: one has to update \mathbf{SA} as new rows are presented ($O(d)$ update for COUNTSKETCH, and $O(sd)$ for subgaussian map), and once the final row has been presented, factorizing \mathbf{SA} and extracting \mathbf{R} can be done in $O(sd \min(s, d))$, which is polynomial in d if s is polynomial in d . However, to compute \mathbf{x}_R one has to compute $(\mathbf{AR})^+ \mathbf{b}$, and storing \mathbf{AR} in memory requires $\Omega(n)$ memory. To circumvent this issue we propose introducing another sketching matrix \mathbf{T} , and approximating $(\mathbf{AR})^+ \mathbf{b}$ via $(\mathbf{TAR})^+ \mathbf{b}$. Thus, for $\mathbf{R} = \mathbf{V}_{\mathbf{SA},k}$ we approximate \mathbf{x}_R by $\tilde{\mathbf{x}}_R = \mathbf{R}(\mathbf{TAR})^+ \mathbf{b}$. It is easy to verify that $\tilde{\mathbf{x}}_R$ can be computed in the streaming model (by forming and updating \mathbf{TA} while computing \mathbf{R}).

More generally, for *any* \mathbf{R} which can be computed in the streaming model, we can also compute in the streaming model the following approximation of $\mathbf{x}_{R,k}$:

$$\tilde{\mathbf{x}}_{R,k} := \mathbf{R} \mathbf{V}_{\mathbf{AR},k} (\mathbf{TAR} \mathbf{V}_{\mathbf{AR},k})^+ \mathbf{T} \mathbf{b}.$$

The next theorem establishes conditions on \mathbf{T} that guarantee that $\tilde{\mathbf{x}}_{R,k}$ is an approximate PCR/PCP.

THEOREM 18. *Suppose that $\mathbf{R} \in \mathbb{R}^{d \times s}$ with $s \geq k$. Assume $\nu \in (0, 1)$. Suppose \mathbf{T} provides an $O(\nu)$ -distortion subspace embedding for $\text{range} \left(\begin{bmatrix} \mathbf{U}_{\mathbf{AR},k} & \mathbf{U}_{\mathbf{A},k} & \mathbf{b} \end{bmatrix} \right)$ that is*

$$\|\mathbf{U}_{\mathbf{AR},k} \mathbf{x}_1 + \mathbf{U}_{\mathbf{A},k} \mathbf{x}_2 + \mathbf{b} x_3\|_2^2 = (1 \pm O(\nu)) (\|\mathbf{x}_1\|_2^2 + \|\mathbf{x}_2\|_2^2 + x_3^2)$$

for every $\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{R}^k$ and $x_3 \in \mathbb{R}$. Then

1. if $d_2(\mathbf{U}_{\mathbf{AR},k}, \mathbf{U}_{\mathbf{A},k}) \leq \nu$, then $\mathbf{A} \tilde{\mathbf{x}}_{R,k}$ is an $(O(\nu), O(\nu))$ -approximate PCP;
2. if $s = k$ and \mathbf{R} has orthonormal columns (i.e., $\mathbf{R}^T \mathbf{R} = \mathbf{I}_k$) and $d_2(\mathbf{R}, \mathbf{V}_{\mathbf{A},k}) \leq \nu(1 + \nu^2)^{-1/2}$, then $\tilde{\mathbf{x}}_R$ is an $(O(\nu), O(\nu/\sigma_k))$ -approximate PCR.

The subspace embedding conditions on \mathbf{T} are met with probability at least $1 - \delta$ if, for example, \mathbf{T} is a COUNTSKETCH matrix with $O(k^2/\nu^2\delta)$ rows.

Proof. We need to show both the additive error bounds on the objective function and the error bound on the constraints. We start with the additive error bounds on the objective function for both for PCP (first part of the theorem) and for PCR

(second part of the theorem). The lower bound on $\|\mathbf{A}\tilde{\mathbf{x}}_{\mathbf{R},k} - \mathbf{b}\|_2$ follows immediately from the fact that $\tilde{\mathbf{x}}_{\mathbf{R},k} \in \text{range}(\mathbf{R}\mathbf{V}_{\mathbf{AR},k})$ and the fact that $\mathbf{x}_{\mathbf{R},k}$ is a minimizer of $\|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2$ subject to $\mathbf{x} \in \text{range}(\mathbf{R}\mathbf{V}_{\mathbf{AR},k})$. For the upper bound, we observe

$$\begin{aligned} \|\mathbf{A}\tilde{\mathbf{x}}_{\mathbf{R},k} - \mathbf{b}\|_2 &\leq (1 + O(\nu)) \|\mathbf{TARV}_{\mathbf{AR},k}(\mathbf{TARV}_{\mathbf{AR},k})^+ \mathbf{Tb} - \mathbf{Tb}\|_2 \\ &\leq (1 + O(\nu)) \|\mathbf{TARV}_{\mathbf{AR},k}(\mathbf{ARV}_{\mathbf{AR},k})^+ \mathbf{b} - \mathbf{Tb}\|_2 \\ &\leq (1 + O(\nu)) \|\mathbf{Ax}_{\mathbf{R},k} - \mathbf{b}\|_2 \\ &\leq \|\mathbf{Ax}_{\mathbf{R},k} - \mathbf{b}\|_2 + O(\nu)\|\mathbf{b}\|_2, \end{aligned}$$

where in the first and third inequalities we used the fact that \mathbf{T} provides a subspace embedding for $\text{range}([\mathbf{ARV}_{\mathbf{AR},k} \mathbf{b}])$, and in the second inequality we used the fact that $(\mathbf{TARV}_{\mathbf{AR},k})^+ \mathbf{Tb}$ is a minimizer of $\|\mathbf{TARV}_{\mathbf{AR},k} \mathbf{x} - \mathbf{Tb}\|_2$. Bounds on $\|\mathbf{Ax}_{\mathbf{R},k} - \mathbf{b}\|_2$ (Theorem 5) now imply the additive bound.

We now bound the constraint for the PCR guarantee (second part of the theorem). Let

$$\mathbf{C} = (\mathbf{TU}_{\mathbf{AR},k})^+ ((\mathbf{TU}_{\mathbf{AR},k})^T)^+.$$

Since $(\mathbf{TU}_{\mathbf{AR},k})^T$ and $(\mathbf{TU}_{\mathbf{AR},k})^+$ have the same row space, and $\mathbf{TU}_{\mathbf{AR},k}$ has more rows than columns, \mathbf{C} is nonsingular and we have $\mathbf{C}(\mathbf{TU}_{\mathbf{AR},k})^T = (\mathbf{TU}_{\mathbf{AR},k})^+$. Since \mathbf{T} provides a subspace embedding for $\mathbf{U}_{\mathbf{AR},k}$, all the singular values of $\mathbf{TU}_{\mathbf{AR},k}$ belong to the interval $[1 - O(\nu), 1 + O(\nu)]$. We conclude that $\|\mathbf{C} - \mathbf{I}_k\|_2 \leq O(\nu)$. We also have $(\mathbf{TU}_{\mathbf{AR},k} \Sigma_{\mathbf{AR},k})^+ = \Sigma_{\mathbf{AR},k}^{-1} (\mathbf{TU}_{\mathbf{AR},k})^+$ since $\mathbf{TU}_{\mathbf{AR},k}$ has linearly independent columns (since it provides a subspace embedding), and $\Sigma_{\mathbf{AR},k}$ has all linearly independent rows. Thus,

$$\begin{aligned} &\|\mathbf{U}_{\mathbf{A},k}^T \mathbf{A}\tilde{\mathbf{x}}_{\mathbf{R},k}\|_2 \\ &= \|\mathbf{U}_{\mathbf{A},k}^T \mathbf{A}\tilde{\mathbf{x}}_{\mathbf{R},k} - \mathbf{U}_{\mathbf{A},k}^T \mathbf{U}_{\mathbf{A},k} \mathbf{U}_{\mathbf{A},k}^T \mathbf{T}^T \mathbf{Tb}\|_2 \\ &\leq \|\mathbf{A}\tilde{\mathbf{x}}_{\mathbf{R},k} - \mathbf{U}_{\mathbf{A},k} \mathbf{U}_{\mathbf{A},k}^T \mathbf{T}^T \mathbf{Tb}\|_2 \\ &= \|\mathbf{U}_{\mathbf{AR},k} (\mathbf{TU}_{\mathbf{AR},k})^+ \mathbf{Tb} - \mathbf{U}_{\mathbf{A},k} \mathbf{U}_{\mathbf{A},k}^T \mathbf{T}^T \mathbf{Tb}\|_2 \\ &= \|\mathbf{U}_{\mathbf{AR},k} \mathbf{C} (\mathbf{TU}_{\mathbf{AR},k})^T \mathbf{Tb} - \mathbf{U}_{\mathbf{A},k} \mathbf{U}_{\mathbf{A},k}^T \mathbf{T}^T \mathbf{Tb}\|_2 \\ &\leq (1 + O(\nu)) \cdot \|\mathbf{U}_{\mathbf{AR},k} \mathbf{C} \mathbf{U}_{\mathbf{AR},k}^T \mathbf{T}^T - \mathbf{U}_{\mathbf{A},k} \mathbf{U}_{\mathbf{A},k}^T \mathbf{T}^T\|_2 \cdot \|\mathbf{b}\|_2 \\ &\leq (1 + O(\nu))^2 \cdot \|\mathbf{U}_{\mathbf{AR},k} \mathbf{C} \mathbf{U}_{\mathbf{AR},k}^T - \mathbf{U}_{\mathbf{A},k} \mathbf{U}_{\mathbf{A},k}^T\|_2 \|\mathbf{b}\|_2 \\ &\leq (1 + O(\nu)) \cdot \left(\|\mathbf{U}_{\mathbf{AR},k} (\mathbf{C} - \mathbf{I}_k) \mathbf{U}_{\mathbf{AR},k}^T\|_2 + \|\mathbf{U}_{\mathbf{AR},k} \mathbf{U}_{\mathbf{AR},k}^T - \mathbf{U}_{\mathbf{A},k} \mathbf{U}_{\mathbf{A},k}^T\|_2 \right) \cdot \|\mathbf{b}\|_2 \\ &\leq (1 + O(\nu)) \cdot \left(\|\mathbf{U}_{\mathbf{AR},k} (\mathbf{C} - \mathbf{I}_k) \mathbf{U}_{\mathbf{AR},k}^T\|_2 + \nu \right) \cdot \|\mathbf{b}\|_2 \\ &= (1 + O(\nu)) \cdot (\|\mathbf{C} - \mathbf{I}_k\|_2 + \nu) \cdot \|\mathbf{b}\|_2 \\ &\leq (1 + O(\nu)) \cdot (O(\nu) + \nu) \cdot \|\mathbf{b}\|_2 \\ &= O(\nu) \cdot \|\mathbf{b}\|_2. \end{aligned}$$

We now bound the constraint for the PCR guarantee (second part of the theorem).

To that end, we observe that

$$\begin{aligned}
\|\mathbf{V}_{\mathbf{A},k}^T \tilde{\mathbf{x}}_{\mathbf{R},k}\|_2 &\leq \|\mathbf{V}_{\mathbf{A},k}^T \mathbf{R} \mathbf{V}_{\mathbf{AR},k} (\mathbf{T} \mathbf{A} \mathbf{R} \mathbf{V}_{\mathbf{AR},k})^+ \mathbf{T} \mathbf{b}\|_2 \\
&\leq \|\mathbf{V}_{\mathbf{A},k}^T \mathbf{R}\|_2 \cdot \|(\mathbf{T} \mathbf{U}_{\mathbf{AR},k} \boldsymbol{\Sigma}_{\mathbf{AR},k})^+ \mathbf{T} \mathbf{b}\|_2 \\
&\leq \frac{\nu \cdot (1 + O(\nu)) \cdot \|\mathbf{b}\|_2}{\sigma_{\min}(\mathbf{T} \mathbf{U}_{\mathbf{AR},k} \boldsymbol{\Sigma}_{\mathbf{AR},k})} \\
&\leq \frac{\nu \cdot (1 + O(\nu)) \|\mathbf{b}\|_2}{(1 - O(\nu)) \sigma_{\min}(\mathbf{U}_{\mathbf{AR},k} \boldsymbol{\Sigma}_{\mathbf{AR},k})} \\
&\leq \frac{O(\nu) \cdot \|\mathbf{b}\|_2}{\sigma_{\min}(\mathbf{A} \mathbf{R})} \\
&\leq \frac{O(\nu)}{\sigma_k} \cdot \|\mathbf{b}\|_2,
\end{aligned}$$

where we used the fact that \mathbf{T} provides a subspace embedding for $\text{range}([\mathbf{U}_{\mathbf{AR},k} \ \mathbf{b}])$, and used Lemma 7 to bound $\|\mathbf{V}_{\mathbf{A},k}^T \mathbf{R}\|_2$ and $\|(\mathbf{A} \mathbf{R})^+\|_2$. \square

5.2. Approximate kernel PCR. For simplicity, we consider only the homogeneous polynomial kernel $\mathcal{K}(\mathbf{x}, \mathbf{z}) = (\mathbf{x}^T \mathbf{z})^q$. The results trivially extend to the nonhomogeneous polynomial kernel $\mathcal{K}_n(\mathbf{x}, \mathbf{z}) = (\mathbf{x}^T \mathbf{z} + c)^q$ by adding a single feature to each data point. We leave for future work the development of similar techniques for other kernels (e.g., Gaussian kernel).

Let $\phi: \mathbb{R}^d \rightarrow \mathbb{R}^{d^q}$ be the function that maps a vector $\mathbf{z} = (z_1, \dots, z_d)$ to the set of monomials formed by multiplying q entries of \mathbf{z} , i.e., $\phi(\mathbf{z}) = (z_{i_1} z_{i_2} \cdots z_{i_q})_{i_1, \dots, i_q \in \{1, \dots, d\}}$. For a data matrix $\mathbf{A} \in \mathbb{R}^d$ and a response vector $\mathbf{b} \in \mathbb{R}^n$, let $\Phi \in \mathbb{R}^{n \times d^q}$ be the matrix obtained by applying ϕ to the rows of \mathbf{A} , and consider computing the rank k PCR solution Φ and \mathbf{b} , which we denote by $\mathbf{x}_{\mathcal{K},k}$. The corresponding prediction function is $f_{\mathcal{K},k}(\mathbf{z}) := \phi(\mathbf{z})^T \mathbf{x}_{\mathcal{K},k}$. While $\mathbf{x}_{\mathcal{K},k}$ is likely a huge vector (since $\mathbf{x}_{\mathcal{K},k} \in \mathbb{R}^{d^q}$), and thus expensive to compute, in kernel PCR we are primarily interested in having an efficient method to compute $f_{\mathcal{K},k}(\mathbf{z})$ given a “new” \mathbf{z} . We can accomplish this via the kernel trick, as we now show.

We assume that Φ has full row rank (this holds if all data points are different). Let $\mathbf{a}_1, \dots, \mathbf{a}_n$ be the rows of \mathbf{A} . As usual with PCR, we have $\mathbf{x}_{\mathcal{K},k} = \mathbf{V}_{\Phi,k} \boldsymbol{\Sigma}_{\Phi,k}^{-1} \mathbf{U}_{\Phi,k}^T \mathbf{b}$. Since $\mathbf{V}_{\Phi,k} = \Phi^T \mathbf{U}_{\Phi,k} \boldsymbol{\Sigma}_{\Phi,k}^{-1}$ we have

$$(12) \quad f_{\mathcal{K},k}(\mathbf{z}) = \phi(\mathbf{z})^T \Phi^T \mathbf{U}_{\Phi,k} \boldsymbol{\Sigma}_{\Phi,k}^{-2} \mathbf{U}_{\Phi,k} \mathbf{b} = (\mathcal{K}(\mathbf{z}, \mathbf{a}_1) \cdots \mathcal{K}(\mathbf{z}, \mathbf{a}_n)) \alpha_{\mathcal{K},k},$$

where $\alpha_{\mathcal{K},k} := \mathbf{U}_{\Phi,k} \boldsymbol{\Sigma}_{\Phi,k}^{-2} \mathbf{U}_{\Phi,k}^T \mathbf{b}$. In the above, we used the fact that for any \mathbf{x} and \mathbf{z} we have $\phi(\mathbf{x})^T \phi(\mathbf{z}) = (\mathbf{x}^T \mathbf{z})^q = \mathcal{K}(\mathbf{x}, \mathbf{z})$. Let $\mathbf{K} \in \mathbb{R}^{n \times n}$ be the *kernel matrix* (also called *Gram matrix*) defined by $\mathbf{K}_{ij} = \mathcal{K}(\mathbf{a}_i, \mathbf{a}_j)$. It is easy to verify that $\mathbf{K} = \Phi \Phi^T$, so we can compute \mathbf{K} in $O(n^2(d+q))$ (and without forming Φ , which is a huge matrix). We also have $\mathbf{K} = \mathbf{U}_{\Phi} \boldsymbol{\Sigma}_{\Phi}^2 \mathbf{U}_{\Phi}^T$, so $\alpha_k = \mathbf{U}_{\mathbf{K},k} \boldsymbol{\Sigma}_{\mathbf{K},k}^{-1} \mathbf{U}_{\mathbf{K},k}^T \mathbf{b}$. Thus, we can compute $\alpha_{\mathcal{K},k}$ in $O(n^2(d+q+n))$ time. Once we have computed α_k , using (12) we can compute $f_{\mathcal{K},k}(\mathbf{z})$ for any \mathbf{z} in $O(ndq)$ time.

In order to compute an approximate kernel PCR, we introduce a right sketching matrix $\mathbf{R} \in \mathbb{R}^{d^q \times t}$. Such a matrix \mathbf{R} is frequently referred to, in the context of kernel learning, as a randomized feature map. We use the TENSORSKETCH feature map [32, 33]. The feature map is defined as follows. We first randomly generate q 3-wise independent hash functions $h_1, \dots, h_q \in \{1, \dots, d\} \rightarrow \{1, \dots, t\}$ and q 4-wise independent sign functions $g_1, \dots, g_q: \{1, \dots, d\} \rightarrow \{-1, +1\}$. Next, we define

$H : \{1, \dots, d\}^q \rightarrow \{1, \dots, t\}$ and $G : \{1, \dots, t\}^q \rightarrow \{-1, +1\}$:

$$H(i_1, \dots, i_q) := h_1(i_1) + \dots + h_q(i_q) \pmod{t},$$

$$G(i_1, \dots, i_q) = g_1(i_1) \cdot g_2(i_2) \cdot \dots \cdot g_q(i_q).$$

To define \mathbf{R} , we index the rows of \mathbf{R} by $\{1, \dots, d\}^q$ and set row (i_1, \dots, i_q) to be equal to $G(i_1, \dots, i_q) \cdot \mathbf{e}_{H(i_1, \dots, i_q)}$, where \mathbf{e}_j denotes the j th identity vector. A crucial observation that makes TENSORSKETCH useful is that via the representation using h_1, \dots, h_q and g_1, \dots, g_q we can compute $\mathbf{R}^T \phi(\mathbf{z})$ in time $O(q(\mathbf{nnz}(\mathbf{z}) + t \log t))$ (see Pagh [32] for details). Thus, we can compute $\Phi \mathbf{R}$ in time $O(q(\mathbf{nnz}(\mathbf{A}) + nt \log t))$.

Consider right sketching PCR on Φ and k with a TENSORSKETCH \mathbf{R} as the sketching matrix. The approximate solution is

$$\mathbf{x}_{\mathcal{K}, \mathbf{R}, k} := \mathbf{R} \mathbf{V}_{\Phi \mathbf{R}, k} (\Phi \mathbf{R} \mathbf{V}_{\Phi \mathbf{R}, k})^+ \mathbf{b} = \mathbf{R} \gamma_{\mathcal{K}, \mathbf{R}, k},$$

where $\gamma_{\mathcal{K}, \mathbf{R}, k} := \mathbf{V}_{\Phi \mathbf{R}, k} (\Phi \mathbf{R} \mathbf{V}_{\Phi \mathbf{R}, k})^+ \mathbf{b}$. We can compute $\gamma_{\mathbf{R}, k}$ in $O(q(\mathbf{nnz}(\mathbf{A}) + nt \log t) + nt^2)$ time. The prediction function is

$$f_{\mathcal{K}, \mathbf{R}, k}(\mathbf{z}) := \phi(\mathbf{z})^T \mathbf{x}_{\mathcal{K}, \mathbf{R}, k} = (\mathbf{R}^T \phi(\mathbf{z}))^T \gamma_{\mathcal{K}, \mathbf{R}, k},$$

so once we have $\gamma_{\mathcal{K}, \mathbf{R}, k}$ we can compute $f_{\mathcal{K}, \mathbf{R}, k}(\mathbf{z})$ in $O(q(\mathbf{nnz}(\mathbf{z}) + t \log t))$ time. Thus, the method is attractive from a computational complexity point of view if $t \ll n$ or $d \gg n$ and $d \gg t$. The following theorem bounds the excess risk of $\mathbf{x}_{\mathcal{K}, \mathbf{R}, k}$.

THEOREM 19. *Let $(\nu, \delta) \in (0, 1/2)$. Let $\lambda_1 \geq \dots \geq \lambda_n$ denote the eigenvalues of \mathbf{K} . If \mathbf{R} is a TENSORSKETCH matrix with*

$$t = \Omega \left(\frac{3^q \mathbf{Tr}(\mathbf{K})^2}{(\lambda_k - \lambda_{k+1})^2 \nu^2 \delta} \right)$$

columns, then with probability at least $1 - \delta$,

$$\mathcal{E}(\mathbf{x}_{\mathcal{K}, \mathbf{R}, k}) \leq \mathcal{E}(\mathbf{x}_{\mathcal{K}, k}) + \frac{(2\nu + \nu^2) \|\mathbf{f}\|_2^2}{n},$$

where \mathbf{f} is the expected value of \mathbf{b} (recall the statistical framework in section 3.2).

Before proving this theorem, we remark that the bound on the size of the sketch is somewhat disappointing in the sense that it is useful only if $d \gg n$ (since $\mathbf{Tr}(\mathbf{K})$ is likely to be large). However, this is only a bound, and possibly a pessimistic one. Furthermore, once the feature expanded data has been embedded in Euclidean space (via TENSORSKETCH), it can be further compressed using standard Euclidean space transforms like COUNTSKETCH and subgaussian maps (this is sometimes referred to as two-level sketching), or compression can be applied from the left. We leave the task of improving the bound and exploring additional compression techniques to future research.

Proof. The square singular values of Φ are exactly the eigenvalues of \mathbf{K} , so Theorem 15 asserts that the conclusions of the theorem hold if \mathbf{R}^T provides an (ϵ, δ) -approximate Gram matrix for Φ , where $\epsilon = O(\nu(\lambda_k - \lambda_{k+1})/\lambda_1)$. To that end, we combine the analysis of Avron, Nguyen, and Woodruff [6] of TENSORSKETCH with more recent results due to Cohen, Nelson, and Woodruff [14]. Although not stated as a formal theorem, as part of a larger proof, Avron, Nguyen, and Woodruff show that TENSORSKETCH has an OSE-moment property that together with the results of Cohen, Nelson, and Woodruff [14] implies that indeed the (ϵ, δ) -approximate Gram property holds for the specified numbers of columns in \mathbf{R} . \square

6. Experiments. In this section we report experimental results, on real data, that illustrate and support the main results of the paper, and demonstrate the ability of our algorithms to find appropriately regularized solutions.

Datasets. We experiment with three datasets: two regression datasets (*Twitter Buzz* and *E2006-tfidf*) and one classification dataset (*Gisette*).

Twitter Social Media Buzz [26] is a regression dataset in which the goal is to predict the popularity of topics as quantified by its mean number of active discussions given 77 predictor variables such as number of authors contributing to the topic over time, average discussion lengths, number of interactions between authors, etc. We preprocess the data in a manner similar to previous work [29, 37]. That is, several of the original predictor variables, as well as the response variable, are log-transformed prior to analysis. We then center and scale to unit norm. Finally, we add quadratic interactions, yielding a total of 3080 predictor variables (after preprocessing, the data matrix is 583250×3080). We used this dataset to explore only suboptimality of the objective and constraint satisfaction, as we have found that the generalization error is very sensitive to selection of the test set (when splitting a subset of the data into training and testing).

E2006-tfidf [28] is a regression dataset in which the features are extracted from SEC-mandated financial reports published annually by a publicly traded company, and the quantity to be predicted is volatility of stock returns, an empirical measure of financial risk. We use the standard training-test split available with the dataset.⁵ We use this dataset only for testing generalization. The only preprocessing we performed was subtracting the mean from the response variable, and reintroducing it when issuing predictions.

The *Gisette* dataset is a binary classification dataset that is constructed from the MNIST dataset. The goal is to separate the highly confusable digits “4” and “9”. The dataset has 6000 data points, each having 5000 features. We use the standard training-test split available with the dataset (this dataset was downloaded from the same website as the E2006-tfidf dataset). We convert the binary classification problem to a regression problem using standard techniques (Regularized Least Squares Classification). We use this dataset only for testing generalization.

Baselines. A first reference is the performance of plain PCR. For small problems, the dominant right subspace needed to compute the PCR solution can be computed via MATLAB’s dense SVD routine. For larger problems, we compute the dominant right subspace using PRIMME [39, 45], a state-of-the-art iterative algorithm for SVD. As additional reference, we also report results of two alternative algorithms: CLS and the iterative algorithm of Frostig et al. [18]. Both in the discussion and in the graphs, we refer to the Frostig et al. algorithm as “Iterative-PCR.” We use the implementation of Iterative-PCR supplied by the authors,⁶ for which we used the default parameters, except for the “tol” parameter, which we set to 10^{-6} instead of the default 10^{-3} . We found that the use of $\text{tol} = 10^{-3}$ produces results that generalize poorly, while the use of $\text{tol} = 10^{-6}$ produces much better results. However, the running time of Iterative-PCR with $\text{tol} = 10^{-6}$ is considerably higher than the running time for $\text{tol} = 10^{-3}$. Iterative-PCR controls singular vector truncation via a cut-off parameter λ , while in our experiments we set k (the number of principal components that are kept). We

⁵We downloaded the dataset from the LIBSVM website, <https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/>.

⁶<https://github.com/cpmusco/fast-pcr>

achieve this effect by setting $\lambda = (\sigma_k^2(\mathbf{A}) + \sigma_{k+1}^2(\mathbf{A}))/2$ (when we report running times, we do not include the time to compute the singular values). Finally, we remark that, based on the documentation, the algorithm analyzed by Frostig et al. [18] does not completely correspond to the default parameters of the implementation of Iterative-PCR supplied by the authors (e.g., the default parameter for the “method” parameter is “LANCZOS,” while “EXPLICIT” corresponds to the algorithm analyzed in [18]).

Suboptimality of objective and constraint satisfaction. We explore the Twitter Buzz dataset from the optimization perspective, namely, measure the suboptimality in the objective (vs. PCR) and constraint satisfaction. Since $n \gg d$, we use left sketching with subgaussian maps. We perform each experiment five times and report the median value. Error bars, when present, represent the minimum and maximum values of five runs. In the top panel of Figure 1, we use a fixed $k = 60$ and vary the sketch size (left sketching only), while in the bottom panel we vary k and set sketch size to be $s = 4k$. The left panel explores the value of the objective function, appropriately normalized (divided by $\|\mathbf{A}\mathbf{x}_k - \mathbf{b}\|_2$ for fixed k , and divided by $\|\mathbf{b}\|_2$ for varying k). The right panel explores the regularization effect by examining the value of the constraints $\|\mathbf{V}_{\mathbf{A},k}^T \mathbf{x}_k\|_2 / \|\mathbf{b}\|_2$.

In the left panel, we see that the value of the objective for the sketched PCR solution follows the value of the objective for the PCR solution. In general, as the sketch size increases, the variance in the objective value reduces (top left graph). The normalized value of the constraint for sketched PCR is rather small (for reference, we note that $\|\mathbf{V}_{\mathbf{A},k}^T \mathbf{x}_{OLS}\|_2 / \|\mathbf{b}\|_2 = 0.4165$) and generally decreases when the sketch size increases (top right graph), but increases with k for a fixed ratio between s and k (bottom right graph). Furthermore, the results of sketched PCR are very similar to the results of Iterative-PCR (bottom panel), while running time is considerably shorter (see Table 2).

The role of the constraints as a regularizer is illustrated by the results for CLS (for fixed k we use $t = 4k$). As expected, CLS achieves lower objective value at the price of larger constraint infeasibility. The values of $\|\mathbf{V}_{\mathbf{A},k}^T \mathbf{x}_{CLS}\|_2 / \|\mathbf{b}\|_2$ are much smaller than the OLS value, but much larger than the values for sketched PCR. Furthermore, it is hard to control the regularization effect for CLS: when sketch size increases, the objective decreases and the constraint increases (compare to PCR and sketched PCR, top panel).

Generalization results. We also explored the prediction error and the trade-offs between compression and regularization. We perform each experiment five times and report the median value. Error bars, when present, represent the minimum and maximum value of those five runs.

We report the mean squared error (MSE) of predictions for the E2006.tfidf dataset in Figure 2. We compare CLS, Iterative-PCR, right sketching, and two-sided sketching (the matrix is too large for exact PCR, and $d \gg n$, so right sketching is more appropriate). In the left panel we fix $k = 600$ and vary the sketch size. The MSE decreases as the sketch size increases for both sketching methods. For CLS, initially the MSE decreases and is close to the MSE of the two sketching methods, but for large sketch sizes the MSE starts to go up, likely due to decreased level of regularization. We note that the minimum MSE achieved by CLS is larger than that achieved by both sketching methods. A similar phenomenon is observed when we vary the value of k in the right panel.

We report the classification error for the Gisette dataset in Figure 3. In the left panel we fix $k = 400$ and vary the sketch size. For reference, the error rate of exact

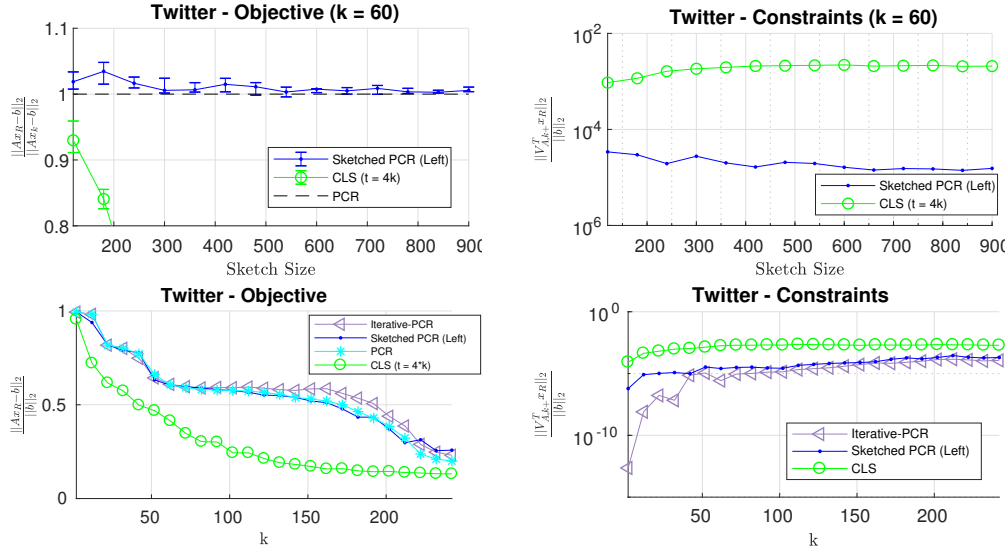


FIG. 1. Suboptimality of objective and constraint satisfaction for the Twitter Buzz dataset.




FIG. 2. Mean squared error of predictions for the E2006.tfidf dataset.

PCR (with $k = 400$) is 2.8% and the error rate for OLS is 9.3%. Left sketching has an error rate very close to the error rate of exact PCR, especially when s is large enough. Right sketching does not perform as well as left sketching, but it too achieves a low error rate for large s . For both methods, the error rate drops as the sketch sizes increase, and the variance reduces. For CLS the error rate and variance initially drop as the sketch size increases, but eventually when the sketch size is large enough, the error rate and the variance increase. This is hardly surprising: as the sketch size increases, CLS approaches OLS. This is due to the fact that CLS uses the compression to regularize, and when the sketch size is large, there is little regularization. In the right panel, we vary the value of k and set $s = 4k$ (left sketching) and $t = 4k$ (right sketching and CLS). Left sketch and PCR consistently achieve about the same error rate. For small values of t , CLS performs well, but when t is too large, the error starts to increase. In contrast, right sketching continues to perform well with large values of k . Again, we see that CLS mixes compression and regularization, and one cannot use a large sketch size and modest amount of regularization with CLS.

Running time. In Table 2 we report a sample of the various running times of the different algorithms. All experiments were conducted using MATLAB, although

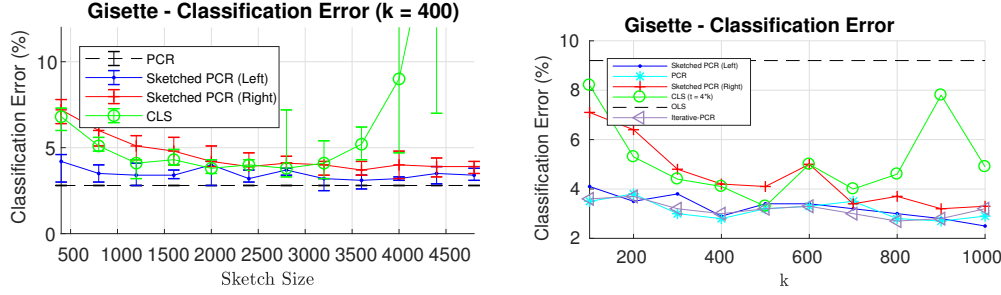


FIG. 3. Classification error for the Gisette dataset.

TABLE 2

Running times (in seconds). For sketched PCR, we report in parentheses the type of sketching used (left, right, or two sided).

	CLS ($t = 400$)	PRIMME-PCR	Iter-PCR (tol= 10^{-3} , iter=10)	Iter-PCR (tol= 10^{-6})	Sketched PCR
Twitter, $k = 82$	21.2	1730	742	4067	4.7 (left)
Twitter, $k = 152$	48.9	5907	1278	7759	8.4 (left)
E2006, $k = 1000$	100	14694	140	601	150 (two sided)
E2006, $k = 2000$	270	FAIL	320	791	815 (two sided)
Gisette, $k = 400$	2.0	497	7.1	30.3	0.2 (left)
Gisette, $k = 1000$	9.3	FAIL	9.3	39.4	0.9 (left)

the sketching routines were written in C. Running times were measured on a machine with a 6-core Intel Xeon Processor E5-1650 v4 CPU and 128 GB of main memory, running Ubuntu 16.04. For plain PCR, we report running time using PRIMME, which we ran with default parameters and no preconditioner. For PRIMME, we cap the number of iterations at 100,000, and write “FAIL” in the table if the PRIMME failed to convergence within that cap. For Iterative-PCR we also report running times when we set tol to the default value, and reduce the max number of iteration from 40 to 10. This results in much faster running time, but much degraded generalization (not reported); for example, for E2006 the test MSE for Iterative-PCR (tol = 10^{-3} , iter = 10) is 0.32. With respect to running time, Iterative-PCR is competitive with sketched PCR only for E2006, but with worse classification error. Using PRIMME for PCR is not competitive with sketched PCR. However, we stress that we experimented with only three datasets, so the comparison is not comprehensive.

7. Conclusions and future work. In this paper, we studied the use of sketching to accelerate the solution of PCR and PCP. In particular, for a data matrix \mathbf{A} , we relate the PCR/PCP solution of \mathbf{AR} , where \mathbf{R} is any dimensionality reduction matrix, to the PCR/PCP solution of \mathbf{A} . We presented a notion of approximate PCR/PCP, motivated both from an optimization perspective and from a statistical perspective, and provide conditions on \mathbf{R} that guarantee rigorous theoretical bounds. We then leverage the aforementioned results to design fast, sketching-based algorithms for approximate PCR/PCP and demonstrate empirically the utility of our proposed algorithms. Throughout, our focus in this paper has been on algorithms that use the “sketch-and-solve” approach.

There are multiple ways in which the current work can be extended and the theoretical results improved. We have presented two notions of approximation: ap-

proximate PCR and approximate PCP. Our results for approximate PCR use only dimensionality reduction matrices \mathbf{R} whose number of columns is equal to the target rank. It is natural to conjecture that the use of dimensionality reduction matrices with a higher number of columns will lead to stronger PCR bounds, but we prove only PCP bounds. The underlying reason is that our bounds for PCR are based on analyzing the distance between the column space of \mathbf{R} and the column space of $\mathbf{V}_{\mathbf{A},k}$. However, once the number of columns in \mathbf{R} is different from the number of columns in $\mathbf{V}_{\mathbf{A},k}$, the definition of $d_2(\mathbf{R}, \mathbf{V}_{\mathbf{A},k})$ is no longer applicable. One possible strategy for analyzing PCR when \mathbf{R} has more than k columns might be to use a generalization of the distance between two subspaces that allows subspaces of different size; see [47] for such generalizations. Another crucial component will then be to generalize the Davis–Kahan theorem to bound such distances. We conjecture it is possible to derive algorithms that depend on gaps between σ_k and σ_{k+l} , where l is some oversampling parameter, as opposed to the smaller gap between σ_k and σ_{k+1} . We leave this for future work.

Another interesting direction is in finding other ways to identify a valid approximate dominant subspace. If we consider the statistical perspective and inspect (9), we see that all we need is to find a subspace $\mathcal{S} \subseteq \text{range}(\mathbf{A})$ of rank k such that $\|(\mathbf{I} - \mathbf{P}_{\mathcal{S}})\mathbf{A}\|_F$ is small, while our theoretical results try to achieve a stronger bound: having the dominant subspaces align. One possible way for finding such an \mathcal{S} is using so-called projection-cost preserving sketches [13]. We leave this for future work.

Appendix A. Bias-variance decomposition for $\mathcal{E}(\mathbf{x}_{\mathbf{R}})$. The following appears, without proof, in [38]. For completeness, we include a proof.

CLAIM 20. *The excess risk of $\mathbf{x}_{\mathbf{R}}$ can be bounded as follows:*

$$\mathcal{E}(\mathbf{x}_{\mathbf{R}}) = \underbrace{\frac{1}{n} \|(\mathbf{I} - \mathbf{P}_{\mathbf{AR}}) \mathbf{A} \mathbf{x}^*\|_2^2}_{\mathcal{B}(\mathbf{x}_{\mathbf{R}})} + \underbrace{\sigma^2 \frac{\text{rank}(\mathbf{AR})}{n}}_{\mathcal{V}(\mathbf{x}_{\mathbf{R}})}.$$

Proof. The column space of \mathbf{AR} is contained in the column space of \mathbf{A} , so we have $\mathbf{P}_{\mathbf{AR}} = \mathbf{P}_{\mathbf{AR}} \mathbf{P}_{\mathbf{A}}$. We now observe

$$\begin{aligned} \mathcal{E}(\mathbf{x}_{\mathbf{R}}) &= \frac{1}{n} \mathbb{E} [\|\mathbf{A} \mathbf{x}_{\mathbf{R}} - \mathbf{A} \mathbf{x}^*\|_2^2] \\ &= \frac{1}{n} \mathbb{E} [\|\mathbf{P}_{\mathbf{AR}} \mathbf{b} - \mathbf{P}_{\mathbf{A}} \mathbf{f}\|_2^2] \\ &= \frac{1}{n} \mathbb{E} [\|\mathbf{P}_{\mathbf{AR}} \mathbf{f} - \mathbf{P}_{\mathbf{A}} \mathbf{f}\|_2^2] + \frac{1}{n} \mathbb{E} [\|\mathbf{P}_{\mathbf{AR}} \xi\|_2^2] \\ &= \frac{1}{n} \mathbb{E} [\|\mathbf{P}_{\mathbf{AR}} \mathbf{P}_{\mathbf{A}} \mathbf{f} - \mathbf{P}_{\mathbf{A}} \mathbf{f}\|_2^2] + \sigma^2 \frac{\text{rank}(\mathbf{AR})}{n} \\ &= \frac{1}{n} \|(\mathbf{I} - \mathbf{P}_{\mathbf{AR}}) \mathbf{A} \mathbf{x}^*\|_2^2 + \sigma^2 \frac{\text{rank}(\mathbf{AR})}{n}, \end{aligned}$$

where in the third line we used the fact that the expected value of ξ is 0, and in the fourth line we used that fact that for any matrix \mathbf{M} and random vector \mathbf{y} with independent entries with 0 mean and σ^2 variance we have $\mathbb{E} [\mathbf{y}^T \mathbf{M} \mathbf{y}] = \text{Tr}(\mathbf{M})$. \square

REFERENCES

- [1] Z. ALLEN-ZHU AND Y. LI, *Faster principal component regression and stable matrix Chebyshev approximation*, in Proceedings of the 34th International Conference on Machine Learning (ICML), 2017, pp. 107–115, <http://proceedings.mlr.press/v70/allen-zhu17c.html>.
- [2] A. ARTEMIOU AND B. LI, *On principal components and regression: A statistical explanation of a natural phenomenon*, Statist. Sinica, 19 (2009), pp. 1557–1565.
- [3] H. AVRON, K. L. CLARKSON, AND D. P. WOODRUFF, *Faster kernel ridge regression using sketching and preconditioning*, SIAM J. Matrix Anal. Appl., 38 (2017), pp. 1116–1138, <https://doi.org/10.1137/16M1105396>.
- [4] H. AVRON, K. L. CLARKSON, AND D. P. WOODRUFF, *Sharper bounds for regularized data fitting*, in Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2017), K. Jansen, J. D. P. Rolim, D. Williamson, and S. S. Vempala, eds., LIPIcs. Leibniz Internat. Proc. Inform. 81, Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2017, pp. 27:1–27:22, <https://doi.org/10.4230/LIPIcs.APPROX-RANDOM.2017.27>.
- [5] H. AVRON, P. MAYMOUNKOV, AND S. TOLEDO, *Blendenpik: Supercharging LAPACK’s least-squares solver*, SIAM J. Sci. Comput., 32 (2010), pp. 1217–1236, <https://doi.org/10.1137/090767911>.
- [6] H. AVRON, H. NGUYEN, AND D. WOODRUFF, *Subspace embeddings for the polynomial kernel*, in Neural Information Processing Systems (NIPS), 2014, pp. 2258–2266.
- [7] C. BOUTSIDIS AND M. MAGDON-ISMAIL, *Faster SVD-truncated regularized least-squares*, in 2014 IEEE International Symposium on Information Theory, 2014, pp. 1321–1325, <https://doi.org/10.1109/ISIT.2014.6875047>.
- [8] M. CHARIKAR, K. CHEN, AND M. FARACH-COLTON, *Finding frequent items in data streams. Automata, languages and programming*, Theoret. Comput. Sci., 312 (2004), pp. 3–15.
- [9] S. CHEN, Y. LIU, M. R. LYU, I. KING, AND S. ZHANG, *Fast relative-error approximation algorithm for ridge regression*, in Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence, UAI’15, AUAI Press, Arlington, VA, 2015, pp. 201–210, <http://dl.acm.org/citation.cfm?id=3020847.3020869>.
- [10] A. CHOWDHURY, J. YANG, AND P. DRINEAS, *An iterative, sketching-based framework for ridge regression*, in Proceedings of the 35th International Conference on Machine Learning, J. Dy and A. Krause, eds., Proc. Mach. Learn. Res. 80, Stockholm, Sweden, PMLR, 2018, pp. 989–998, <http://proceedings.mlr.press/v80/chowdhury18a.html>.
- [11] K. L. CLARKSON AND D. P. WOODRUFF, *Numerical linear algebra in the streaming model*, in Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing, STOC ’09, ACM, 2009, pp. 205–214, <https://doi.org/10.1145/1536414.1536445>.
- [12] K. L. CLARKSON AND D. P. WOODRUFF, *Low-rank approximation and regression in input sparsity time*, J. ACM, 63 (2017), pp. 54:1–54:45, <https://doi.org/10.1145/3019134>.
- [13] M. B. COHEN, S. ELDER, C. MUSCO, C. MUSCO, AND M. PERSU, *Dimensionality reduction for k-means clustering and low rank approximation*, in Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing, STOC 2015, Portland, OR, 2015, R. A. Servedio and R. Rubinfeld, eds., ACM, 2015, pp. 163–172, <https://doi.org/10.1145/2746539.2746569>.
- [14] M. B. COHEN, J. NELSON, AND D. P. WOODRUFF, *Optimal approximate matrix product in terms of stable rank*, in Proceedings of the 43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, Rome, Italy, 2016, pp. 11:1–11:14, <https://doi.org/10.4230/LIPIcs.ICALP.2016.11>.
- [15] C. DAVIS AND W. M. KAHAN, *The rotation of eigenvectors by a perturbation. III*, SIAM J. Numer. Anal., 7 (1970), pp. 1–46, <https://doi.org/10.1137/0707001>.
- [16] P. DRINEAS, I. C. F. IPSEN, E.-M. KONTOPOULOU, AND M. MAGDON-ISMAIL, *Structural convergence results for approximation of dominant subspaces from block Krylov spaces*, SIAM J. Matrix Anal. Appl., 39 (2018), pp. 567–586, <https://doi.org/10.1137/16M1091745>.
- [17] P. DRINEAS, M. W. MAHONEY, S. MUTHUKRISHNAN, AND T. SARLÓS, *Faster least squares approximation*, Numer. Math., 117 (2011), pp. 219–249, <https://doi.org/10.1007/s00211-010-0331-6>.
- [18] R. FROSTIG, C. MUSCO, C. MUSCO, AND A. SIDFORD, *Principal component projection without principal component analysis*, in International Conference on Machine Learning (ICML), 2016, pp. 2349–2357.
- [19] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, JHU Press, 2012.
- [20] A. GONEN, F. ORABONA, AND S. SHALEV-SHWARTZ, *Solving ridge regression using sketched preconditioned SVRG*, in Proceedings of the 33rd International Conference on International

- Conference on Machine Learning, ICML'16, JMLR.org, 2016, pp. 1397–1405, <http://dl.acm.org/citation.cfm?id=3045390.3045538>.
- [21] N. HALKO, P.-G. MARTINSSON, AND J. A. TROPP, *Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions*, SIAM Rev., 53 (2011), pp. 217–288, <https://doi.org/10.1137/090771806>.
 - [22] R. A. HORN AND C. R. JOHNSON, *Matrix Analysis*, Cambridge University Press, 1990.
 - [23] H. HOTELLING, *Analysis of a complex of statistical variables into principal components*, J. Educ. Psych., 24 (1933), pp. 417–441.
 - [24] I. T. JOLLIFFE, *A note on the use of principal components in regression*, Appl. Statist., 31 (1982), pp. 300–303.
 - [25] A. KABAN, *New bounds on compressive linear least squares regression*, in Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics, S. Kaski and J. Corander, eds., Proc. Mach. Learn. Res. 33, Reykjavik, Iceland, 2014, PMLR, 2014, pp. 448–456, <http://proceedings.mlr.press/v33/kaban14.html>.
 - [26] F. KAWALA, A. DOUZAL-CHOUAKRIA, E. GAUSSIER, AND E. DIMERT, *Prédictions d'activité dans les réseaux sociaux en ligne*, in 4ième conférence sur les modèles et l'analyse des réseaux: Approches mathématiques et informatiques, 2013.
 - [27] M. G. KENDALL, *A Course in Multivariate Analysis*, C. Griffin, 1957.
 - [28] S. KOGAN, D. LEVIN, B. R. ROUTLEDGE, J. S. SAGI, AND N. A. SMITH, *Predicting risk from financial reports with regression*, in Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics, Association for Computational Linguistics, 2009, pp. 272–280.
 - [29] Y. LU AND D. P. FOSTER, *Fast ridge regression with randomized principal component analysis and gradient descent*, in Proceedings of the Thirtieth Conference on Uncertainty in Artificial Intelligence, UAI'14, AUAI Press, 2014, pp. 525–532.
 - [30] O. MAILLARD AND R. MUNOS, *Compressed least-squares regression*, in Neural Information Processing Systems (NIPS), 2009, pp. 1213–1221.
 - [31] X. MENG, M. A. SAUNDERS, AND M. W. MAHONEY, *LSRN: A parallel iterative solver for strongly over- or underdetermined systems*, SIAM J. Sci. Comput., 36 (2014), pp. C95–C118, <https://doi.org/10.1137/120866580>.
 - [32] R. PAGH, *Compressed matrix multiplication*, ACM Trans. Comput. Theory, 5 (2013), pp. 9:1–9:17, <https://doi.org/10.1145/2493252.2493254>.
 - [33] N. PHAM AND R. PAGH, *Fast and scalable polynomial kernels via explicit feature maps*, in Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '13, ACM, 2013, pp. 239–247, <https://doi.org/10.1145/2487575.2487591>.
 - [34] M. PILANCI AND M. J. WAINWRIGHT, *Randomized sketches of convex programs with sharp guarantees*, IEEE Trans. Inform. Theory, 61 (2015), pp. 5096–5115.
 - [35] M. PILANCI AND M. J. WAINWRIGHT, *Iterative Hessian sketch: Fast and accurate solution approximation for constrained least-squares*, J. Mach. Learn. Res., 17 (2016), pp. 1842–1879.
 - [36] V. ROKHLIN AND M. TYGERT, *A fast randomized algorithm for overdetermined linear least-squares regression*, Proc. Natl. Acad. Sci. USA, 105 (2008), pp. 13212–13217, <https://doi.org/10.1073/pnas.0804869105>.
 - [37] M. SLAWSKI, *Compressed least squares regression revisited*, in Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS), 2017, pp. 1207–1215.
 - [38] M. SLAWSKI, *On Principal Components Regression, Random Projections, and Column Subsampling*, preprint, <https://arxiv.org/abs/1709.08104>, 2017.
 - [39] A. STATHOPOULOS AND J. R. MCCOMBS, *PRIMME: PREconditioned Iterative MultiMethod Eigensolver: Methods and software description*, ACM Trans. Math. Softw., 37 (2010), pp. 21:1–21:30.
 - [40] G. W. STEWART, *On the perturbation of pseudo-inverses, projections and linear least squares problems*, SIAM Rev., 19 (1977), pp. 634–662, <https://doi.org/10.1137/1019104>.
 - [41] G. W. STEWART, *Matrix Algorithms: Volume II: Eigensystems*, SIAM, 2001, <https://doi.org/10.1137/1.9780898718058>.
 - [42] G.-A. THANEI, C. HEINZE, AND N. MEINSHAUSEN, *Random projections for large-scale regression*, in Big and Complex Data Analysis, Springer, 2017, pp. 51–68, <https://doi.org/10.1007/978-3-319-41573-4.3>.
 - [43] S. WANG, A. GITTENS, AND M. W. MAHONEY, *Sketched ridge regression: Optimization perspective, statistical perspective, and model averaging*, J. Mach. Learn. Res., 18 (2017), pp. 8039–8088, <http://dl.acm.org/citation.cfm?id=3122009.3242075>.
 - [44] D. P. WOODRUFF, *Sketching as a tool for numerical linear algebra*, Found. Trends Theor.

- Comput. Sci., 10 (2014), pp. 1–157.
- [45] L. WU, E. ROMERO, AND A. STATHOPOULOS, *Primme_SVDs: A high-performance preconditioned SVD solver for accurate large-scale computations*, SIAM J. Sci. Comput., 39 (2017), pp. S248–S271, <https://doi.org/10.1137/16M1082214>.
 - [46] J. YANG, X. MENG, AND M. W. MAHONEY, *Implementing randomized matrix algorithms in parallel and distributed environments*, Proc. IEEE, 104 (2016), pp. 58–92, <https://doi.org/10.1109/JPROC.2015.2494219>.
 - [47] K. YE AND L.-H. LIM, *Schubert varieties and distances between subspaces of different dimensions*, SIAM J. Matrix Anal. Appl., 37 (2016), pp. 1176–1197, <https://doi.org/10.1137/15M1054201>.