

On principal components regression, random projections, and column subsampling

Martin Slawski

*Department of Statistics
George Mason University
Fairfax, VA 22030, USA
e-mail: mslawsk3@gmu.edu*

Abstract: Principal Components Regression (PCR) is a traditional tool for dimension reduction in linear regression that has been both criticized and defended. One concern about PCR is that obtaining the leading principal components tends to be computationally demanding for large data sets. While random projections do not possess the optimality properties of the leading principal subspace, they are computationally appealing and hence have become increasingly popular in recent years. In this paper, we present an analysis showing that for random projections satisfying a Johnson-Lindenstrauss embedding property, the prediction error in subsequent regression is close to that of PCR, at the expense of requiring a slightly large number of random projections than principal components. Column sub-sampling constitutes an even cheaper way of randomized dimension reduction outside the class of Johnson-Lindenstrauss transforms. We provide numerical results based on synthetic and real data as well as basic theory revealing differences and commonalities in terms of statistical performance.

Received October 2017.

1. Introduction

Principal Components Regression (PCR), first introduced in [24, 32], is perhaps the most basic approach to dimension reduction in linear regression. In PCR, the design matrix $X \in \mathbb{R}^{n \times d}$ containing the original predictor variables is replaced by $XR \in \mathbb{R}^{n \times r}$, $r < d \wedge n$, where $R \in \mathbb{R}^{r \times d}$ reduces X to its top r principal components. From a statistical point of view, PCR can be motivated as a way of dealing with multi-collinearity and reducing estimation variance at the expense of additional bias. From the point of view of computation, PCR potentially achieves a reduction from a large number of variables to a parsimonious model, which can be beneficial for both model fitting and prediction of future observations. The use of PCR is debated in the literature as it does not need to be case that principal components corresponding to small singular values do not significantly contribute in predicting the response variable [6, 29]. Herein, we mostly avoid touching upon this issue. Instead, the purpose of this paper is to establish a connection between PCR and the use of randomized methods of dimension reduction in linear regression in which the matrix R above is sampled from a

suitable distribution. The latter approach, typically referred to as the method of random projections, is motivated from large-scale datasets in which both the number of samples n and the dimension d are large; in this case, computation of principal components via the SVD can be demanding. Random projections only require a matrix multiplication which can be easily parallelized. Having its roots in the celebrated Johnson-Lindenstrauss-Lemma [28], the idea has meanwhile a long history in computer science [46], and has recently attracted considerable interest in statistics (see [14] and the references therein).

Contributions and related work It is critical to understand the statistical properties of the dimension reduction provided by random projections. Regarding linear regression, there are many more papers (e.g., [2, 23, 39, 40, 47, 49]) on the scenario in which X is reduced to RX , i.e., R is multiplied from left instead of from the right with X being reduced to XR . The latter scenario was first analyzed in [36] under the term “compressed least squares” (CLS) which will also be employed herein. For a fixed design setting, refinements appear in [30], and very recently in [44]. Together with a preliminary version [43] of the present paper, the paper [44] is the first to make a connection between PCR and CLS. However, as we show below, while improving over the main result in [30], the upper bounds on the prediction error of CLS in [44] still leave a considerable gap to PCR. In the present paper, we try to close this gap. In brief, our main result states that CLS can roughly match the performance of PCR with respect to prediction at the expense of a moderate increase (at most by a logarithmic factor) in the reduced dimension. This property is shown to hold for a certain class of matrices comprising those that are typically considered in the literature on randomized dimensionality reduction and compressed sensing. We leave it as an open problem whether a similar result can be established for column subsampling, i.e., the columns of R are chosen uniformly at random from the canonical basis vectors. Results are provided indicating that more stringent conditions are required in that case. Finally, we note that the very recent work [31] presents an analysis of CLS when the goal is to recover the vector of regression coefficients under sparsity. Such an assumption is not made herein, and accordingly settings, goals and results are not comparable.

Outline In Section §2, we provide background and review the results of prior work in more detail. Our main result is contained in §3, while §4 discusses extensions and open questions. We conclude with a brief summary in §5. The appendix contains all proofs.

Notation For a positive integer m , we write $[m] := \{1, \dots, m\}$. For a matrix M , we write P_M for the projection operator onto the column space of M . Its Frobenius norm is denoted by $\|M\|_F = \sqrt{\text{tr}(M^\top M)}$, where “tr” is the trace of a diagonal matrix. The Gaussian distribution with zero mean and variance s^2 is denoted by $N(0, s^2)$, and $\text{Unif}(S)$ denotes the uniform probability distribution on a set S . For $a, b \in \mathbb{R}$, we write $a \wedge b = \min\{a, b\}$ and $a \vee b = \max\{a, b\}$.

Positive constants are denoted by C, C_1, c, c_1 etc. We make use of the usual Big-O notation in terms of O, o, Ω and Θ .

2. Background

We start by providing some context for our main result. After fixing the setup, we derive bounds on the prediction error of PCR and put them into relation to existing results on CLS. This will point to a significant gap that motivates our analysis of CLS in §3.

2.1. General setting

We consider fixed design linear regression for data (y_i, x_i) , with y_i taking values in \mathbb{R} and x_i taking values in \mathbb{R}^d , $i \in [n]$. The predictors x_i are considered as fixed, and

$$y_i = f_i + \xi_i,$$

with $f_i = \mathbf{E}[y_i]$ and ξ_i following a distribution with mean zero and variance σ^2 , $i \in [n]$. Moreover, the $\{\xi_i\}_{i=1}^n$ are assumed to be uncorrelated. More concisely, we write $y = f + \xi$, where $y = (y_i)_{i=1}^n$, etc. We denote by $X \in \mathbb{R}^{n \times d}$ the design matrix whose rows are given by x_i^\top , $i \in [n]$. The optimal linear predictor Xw^* of y given X with respect to squared loss is defined by the optimization problem

$$\min_{w \in \mathbb{R}^d} \mathbf{E}[\|y - Xw\|_2^2/n],$$

where the expectation is with respect to the noise ξ . Any minimizer w^* of the above problem satisfies $Xw^* = P_X f$ with P_X defined as in the paragraph on notation above; if there are multiple such w^* we choose the one with minimum ℓ_2 -norm. Accordingly, we define the excess risk of an estimator $\hat{\theta} = \hat{\theta}(X, y)$ of w^* by

$$\mathcal{E}(\hat{\theta}) = \mathbf{E}[\|Xw^* - X\hat{\theta}\|_2^2/n],$$

If the linear model holds exactly (i.e., $P_X f = f$), $\mathcal{E}(\hat{\theta})$ equals the in-sample mean squared prediction error that measures how well the $\{x_i^\top \hat{\theta}\}_{i=1}^n$ predict the “denoised” observations $\{x_i^\top w^*\}_{i=1}^n$ on average. An ordinary least squares (OLS) estimator \hat{w} satisfies $X\hat{w} = P_X y$. Its excess risk is given by

$$\mathcal{E}(\hat{w}) = \sigma^2 \text{rank}(X)/n. \tag{1}$$

To keep matters simple, we assume that X has full rank $d \wedge n$ unless otherwise stated. In this paper, we are interested in a high-dimensional setup in which $\text{rank}(X)$ is of the same order of magnitude as n . In this situation, OLS does not yield satisfactory statistical performance. Moreover, if both n and d are large, obtaining \hat{w} or making predictions based on \hat{w} becomes computationally costly.

In light of these issues, it makes sense to consider alternatives that aim at leveraging some sort of low-dimensional structure. Scenarios in which w^* exhibits one of various forms of sparsity are dominating in the literature, see the

monographs of [10, 22] for an overview. In the present paper, we follow another direction in which the predictors $\{x_i\}_{i=1}^n$ are linearly mapped into a lower-dimensional space, and linear least squares regression is then performed based on the subspace obtained in this way. Put differently, one considers a new design matrix $X_R = XR$ with R being a d -by- k matrix, $k \ll d$. On the statistical side, one potentially achieves a substantial reduction of variance at the expense of an increased bias as made precise below. The excess risk of the approach is given by

$$\mathcal{E}(R) = \mathbf{E} [\|Xw^* - X_R\hat{w}_R\|_2^2/n], \quad (2)$$

where \hat{w}_R is a least squares solution based on the reduced design matrix X_R , i.e., \hat{w}_R satisfies $X_R\hat{w}_R = P_{X_R}y$, and the expectation is with respect to ξ (in later sections, R will be random, and we will then also take the expectation with respect to R). Straightforward calculations show that $\mathcal{E}(R)$ can be decomposed into a bias and a variance term:

$$\mathcal{E}(R) = \underbrace{\|(I - P_{X_R})Xw^*\|_2^2/n}_{\text{Bias}} + \underbrace{\sigma^2 \text{rank}(X_R)/n}_{\text{Variance}}. \quad (3)$$

We commonly have $\text{rank}(X_R) = k$. The choice of k determines the bias-variance trade-off. If k can be chosen much smaller than $d \wedge n$ while at the same time the magnitude of the bias can be controlled, an improvement over the excess risk of OLS in (1) is obtained.

The approach can as well be motivated from the computational side given that we only need to solve a least squares problem of dimension k instead of $d \wedge n$. In addition, having a smaller number of predictors yields savings in storage and when making predictions.

2.2. Excess risk of PCR

The traditional choice of constructing X_R is in terms of the leading principal components of X . Subsequent use of this new set of predictors in regression is known as principal components regression (PCR). Let $X = U\Sigma V^\top$ be the singular value decomposition (SVD) of X , where $U \in \mathbb{R}^{n \times d \wedge n}$, $U^\top U = I$, is the matrix of left singular vectors, $\Sigma \in \mathbb{R}^{d \wedge n \times d \wedge n}$ is the diagonal matrix whose diagonal contains the decreasingly ordered sequence of singular values $\sigma_1 \geq \dots \geq \sigma_{d \wedge n}$, and $V \in \mathbb{R}^{d \times d \wedge n}$, $V^\top V = I$, is the matrix of right singular vectors. For $r \in \{1, \dots, d \wedge n\}$, consider

$$U = [U_r \ U_{r+}], \quad \Sigma = \begin{bmatrix} \Sigma_r & 0 \\ 0 & \Sigma_{r+} \end{bmatrix}, \quad V = [V_r \ V_{r+}], \quad (4)$$

where U_r and $V_r \in \mathbb{R}^{d \times r}$ contain the top r left respectively right singular vectors, and Σ_r contains the corresponding singular values. The remaining singular vectors respectively singular values are contained in U_{r+} , V_{r+} and Σ_{r+} . The top r principal components are extracted from X by setting $R = V_r$:

$$X_R = XV_r = (U_r \Sigma_r V_r^\top + U_{r+} \Sigma_{r+} V_{r+}^\top) V_r = U_r \Sigma_r.$$

The corresponding projection is given by $P_{X_R} = U_r U_r^\top$ and the bias term in (3) results as

$$(I - P_{X_R})Xw^* = (I - U_r U_r^\top)Xw^* = U_{r+} \Sigma_{r+} V_{r+}^\top w^*. \quad (5)$$

Let us define $\alpha^* \in \mathbb{R}^{d \wedge n}$ by

$$\alpha^* = V^\top w^* = \begin{bmatrix} V_r^\top w^* \\ V_{r+}^\top w^* \end{bmatrix} = \begin{bmatrix} \alpha_r^* \\ \alpha_{r+}^* \end{bmatrix}. \quad (6)$$

Combining (2), (5) and (6), the excess risk of PCR can then be expressed as follows.

$$\begin{aligned} \mathcal{E}(V_r) &= \|U_{r+} \Sigma_{r+} V_{r+}^\top w^*\|_2^2 / n + \sigma^2 r / n \\ &= \|\Sigma_{r+} \alpha_{r+}^*\|_2^2 / n + \sigma^2 r / n \\ &= \sum_{j=r+1}^{d \wedge n} \sigma_j^2 (\alpha_j^*)^2 / n + \sigma^2 r / n \end{aligned} \quad (7)$$

We see from (7) that the excess risk of PCR behaves favorably if (i) the tail of the squared singular values at truncation level r is small (i.e., X can be well approximated by a matrix of rank r) and (ii) if there are no large coefficients in α^* outside its top r entries corresponding to the leading singular vectors. Condition (ii) constitutes the main source of criticism of PCR: if nature is malicious, then α^* has most of its mass in α_{r+}^* . Under generic random sampling, however, this is not a concern: if V is sampled uniformly at random from its respective Stiefel manifold, then $\alpha^* / \|\alpha^*\|_2$ is uniformly distributed on the unit sphere in $\mathbb{R}^{d \wedge n}$ so that the entries of α^* are roughly homogeneous in magnitude.

In the sequel, we derive a series of bounds on the excess risk of PCR depending on the decay of the squared singular values $\{\sigma_j^2\}_{j=1}^{d \wedge n}$. For this purpose, we use the following simple upper bound on $\mathcal{E}(V_r)$:

$$\mathcal{E}(V_r) \leq \|\alpha^*\|_\infty^2 \frac{\|\Delta_r\|_F^2}{n} + \sigma^2 \frac{r}{n}, \quad \Delta_r := X - \mathcal{T}_r(X), \quad (8)$$

where $\mathcal{T}_r(X) = U_r \Sigma_r V_r^\top$ equals the best rank r -approximation to X with respect to the Frobenius norm. For what follows, we assume that X is scaled such that $\|X\|_F^2 = \sum_{j=1}^{d \wedge n} \sigma_j^2 = n \cdot d$. For $1 \leq s \leq d \wedge n$, we define

$$\gamma(s) = \sum_{j=1}^s \sigma_j^2, \quad \tau(s) = \frac{\gamma(d \wedge n) - \gamma(s)}{\gamma(d \wedge n)}. \quad (9)$$

The quantity $\|\Delta_r\|_F^2 / n$ in (8) can then be expressed as

$$\|\Delta_r\|_F^2 / n = \{\gamma(d \wedge n) - \gamma(r)\} / n = \tau(r) \cdot \gamma(d \wedge n) / n = \tau(r) \cdot d. \quad (10)$$

After these preparations, we study the excess risk of PCR in three basic scenarios.

Scenario (F): perfectly flat spectrum

A flat spectrum means that $\sigma_j^2 = n \vee d$, $j \in [d \wedge n]$. We obtain that $\|\Delta_r\|_F^2/n = d \cdot \tau(r) = \frac{d}{n} \vee 1 (d \wedge n - r)$. The choice of r minimizing the bound (8) is given by $r^* = d \wedge n$ if $\|\alpha^*\|_\infty > \sigma/\sqrt{d \vee n}$, and $r^* = 0$ otherwise. We note that when using the exact expression (7) for the excess risk, the optimal r^* would result as the largest value of r such that $\alpha_r > \sigma/\sqrt{d \vee n}$. Eventually, this does not make much of a difference if the entries of α^* are of a comparable magnitude, and does not affect the conclusion that in general, we cannot hope for improvements over OLS when the spectrum is constant.

Scenario (P): polynomial decay

Suppose that $\sigma_j^2 = C \cdot j^{-q}$, $j \in [d]$, for $q \geq 2$ and a constant C determined by the relation $\sum_{j=1}^{d \wedge n} \sigma_j^2 = n \cdot d$. Comparing series and integrals, we obtain $\gamma(d \wedge n) - \gamma(r) \leq C(q-1)^{-1} r^{-(q-1)}$. Moreover, $\gamma(d \wedge n) \geq C$, so that $\tau(r) \leq (q-1)^{-1} r^{-(q-1)}$. By (8)

$$\mathcal{E}(V_r) \leq (q-1)^{-1} r^{-(q-1)} \cdot d \cdot \|\alpha^*\|_\infty^2 + \sigma^2 r/n.$$

Minimizing the right hand side w.r.t. r , we obtain

$$r^* = \{\|\alpha^*\|_\infty^2 (n \cdot d) / \sigma^2\}^{1/q}, \quad \mathcal{E}(V_{r^*}) \leq 2 (d \|\alpha^*\|_\infty^2)^{1/q} (\sigma^2/n)^{(q-1)/q}. \quad (11)$$

To get some insight into (11), fix $q = 2$ and consider the case of generic random sampling of V as discussed above so that $\alpha^*/\|\alpha^*\|_2$ follows a uniform distribution on the unit sphere in $\mathbb{R}^{d \wedge n}$. In this situation, $\|\alpha^*\|_\infty/\|\alpha^*\|_2$ scales as $O(\sqrt{\log(d)/d})$ as d gets large. Assuming further that $\|\alpha^*\|_2 = O(1)$ yields $r^* = O(\{\log(d)n\}^{1/2})$ and $\mathcal{E}(V_{r^*}) = O(\sqrt{\log(d)/n})$. We have hence identified a regime in which PCR achieves better statistical and computational performance than OLS if $d = \Omega(n)$. Clearly, the improvements get amplified as q increases.

Scenario (E): exponential decay

Suppose that $\sigma_j^2 = C_0 \theta^j$ for $\theta \in (0, 1)$. Then, $\tau(r) \leq \frac{\theta^r}{1-\theta} = C_1 \exp(-cr)$, say. The optimal choice of r^* and the corresponding bound on $\mathcal{E}(V_{r^*})$ result as

$$r^* = \frac{1}{c} \log(C_2 \|\alpha^*\|_\infty^2 n d / \sigma^2), \quad \mathcal{E}(V_{r^*}) \leq \frac{2}{c} \{\log(C_2 \|\alpha^*\|_\infty^2 n d / \sigma^2) \vee 1\} \sigma^2/n. \quad (12)$$

Cases (P) and (E) show that PCR may improve significantly over OLS in terms of achievable dimension reduction and excess risk depending on the decay of the spectrum of X .

2.3. Existing bounds for CLS

We now consider the case of dimension reduction via a *random* matrix R . We refer to the columns of R as “random projections” as R maps the predictors $\{x_i\}_{i=1}^n$ to a random linear subspace, typically of dimension k . Regarding the distribution of R , sampling its entries i.i.d. from a Gaussian distribution with expectation zero and variance $1/k$ constitutes the basic case in the literature on randomized dimensionality reduction [46]. The column space of R then follows the uniform distribution on the Grassmannian $\mathbb{G}(d, k)$. Random Gaussian matrices of this form are the canonical example of Johnson-Lindenstrauss transforms [28] (henceforth JLTs for short), cf. Definition 1 below. This class of matrices extends to i.i.d. sub-Gaussian matrices [1, 38], the fast JLT of [3], and certain row-subsampled orthonormal matrices [4, 33, 45] as they are also used in compressed sensing [13].

Maillard & Munos [36] were the first to study the use of JLTs for randomized dimension reduction in least squares regression with random design under the name “compressed least squares” (CLS). They show a bound on a corresponding notion of excess risk of the order

$$O\left(\sigma\|w^*\|_2\left(\mathbf{E}[\|x_1\|_2^2]\right)^{1/2}\sqrt{\log(n)/n}\right) \quad (13)$$

for $k = \Theta(\sqrt{n \log n} \|w^*\|_2 \mathbf{E}[\|x_1\|_2^2]^{1/2} / \sigma)$ random projections. For fixed design, Kaban [30] (specializing Theorem 1 therein to the case where R has i.i.d. $N(0, 1/k)$ entries) shows that

$$\begin{aligned} \mathbf{E}[\mathcal{E}(R)] &\leq \frac{(w^*)^\top (\text{tr}(\Gamma)I + \Gamma)w^*}{k} + \sigma^2 \frac{k}{n}, & \Gamma &:= X^\top X/n, \\ &\leq c \frac{\text{tr}(\Gamma)\|w^*\|_2^2}{k} + \sigma^2 \frac{k}{n} \\ &= c \cdot \|w^*\|_2^2 \left(\frac{1}{n} \sum_{j=1}^{d \wedge n} \sigma_j^2 \right) + \sigma^2 \frac{k}{n} \text{ for some } c \in [1, 2], \end{aligned} \quad (14)$$

where $\mathcal{E}(R)$ is defined in (2) and the expectation in (14) is with respect to R .

Optimizing the bound (14) with respect to k , we obtain

$$k^* = c' \sqrt{n \text{tr}(\Gamma)} / \sigma, \quad \mathbf{E}[\mathcal{E}(R)] \leq 2\sigma c' \|w^*\|_2 \sqrt{\text{tr}(\Gamma)/n}. \quad (15)$$

with $c' = \sqrt{c}$. Comparing the bound of Maillard and Munos (13) with (15), we essentially observe an agreement apart from a $\sqrt{\log n}$ factor, noting that for random design with $\{x_i\}_{i=1}^n \sim x$ such that $\mathbf{E}[xx^\top] = \Gamma_0$, we have $\sqrt{\mathbf{E}[\|x_1\|_2^2]} = \sqrt{\text{tr}(\Gamma_0)}$.

In [36, 42], the excess risk bound in (15) is interpreted as being of the order $O(1/\sqrt{n})$. This would mean that the performance of CLS is comparable to that of PCR in scenario (P) with exponent $q = 2$ above, independent of the $\{\sigma_j^2\}_{j=1}^{d \wedge n}$. However, such interpretation is not valid in general. In a fixed design setting,

it is common to assume that the columns $\{X_j\}_{j=1}^d$ of X are scaled such that $\|X_j\|_2^2 = n$, $j \in [d]$, whereas for standard random designs, e.g. X with i.i.d. rows from a zero-mean Gaussian distribution with unit variances, this scaling holds in expectation. In this case, $\text{tr}(\Gamma)$ respectively $\mathbf{E}[\|x_1\|_2^2]$ evaluate as d which makes the bounds (13) and (15) of rather limited use. For (15), we obtain

$$k^* = c' \sqrt{nd}/\sigma, \quad \mathbf{E}[\mathcal{E}(R)] \leq 2\sigma c' \|w^*\|_2 \sqrt{d/n}.$$

This means that k^* is of the same order (or may even exceed) $d \wedge n$, while the bound on the excess risk is generally inferior to that of OLS (1).

It turns out that the outcome (15) is the consequence of a crude bound on the bias of CLS. From (14), we find that the correct variance term $\sigma^2 k/n$ is present. On the other hand, a simple argument shows that the bias term $\text{tr}(\Gamma) \|w^*\|_2^2/k$ is improvable: if X has rank r , $1 \leq r \leq d \wedge n$, CLS with R as a matrix with $N(0, 1)$ -entries yields $P_X = P_{X_R}$ and in turn zero bias in (2) with probability one as long as $k \geq r$. By contrast, according to (14) basically $k/d \rightarrow \infty$ is required for the bias to vanish.

In [30], the expected bias (where the expectation is w.r.t. R) is bounded as

$$\begin{aligned} \mathbf{E}[\|(I - P_{X_R})Xw^*\|_2^2/n] &= \mathbf{E} \left[\min_{v \in \mathbb{R}^k} \|Xw^* - XRv\|_2^2/n \right] \\ &\leq \mathbf{E}[\|Xw^* - XRR^\top w^*\|_2^2/n] \end{aligned} \quad (16)$$

Evaluation of (16) then inevitably leads to the term $\text{tr}(\Gamma) \|w^*\|_2^2/k$ as it appears in (14).

In a recent paper by Thanei et al. [44], the following bound is used instead of (16):

$$\mathbf{E} \left[\min_{v \in \mathbb{R}^k} \|Xw^* - XRv\|_2^2/n \right] \leq \min_{v \in \mathbb{R}^k} \mathbf{E}[\|Xw^* - XRv\|_2^2/n] \quad (17)$$

The authors evaluate the above expectation for R with i.i.d. $N(0, 1/k)$ entries, and then minimize with respect to v . This yields the bound (cf. Theorem 2 in [44])

$$\mathbf{E}[\mathcal{E}(R)] \leq \frac{1}{n} \sum_{j=1}^{d \wedge n} (\alpha_j^*)^2 \sigma_j^2 \omega_j + \sigma^2 \frac{k}{n}, \quad (18)$$

$$\omega_j = \frac{(1 + 1/k)\sigma_j^4 + (1 + 2/k)\sigma_j^2 \text{tr}(\Gamma) + \text{tr}(\Gamma)^2/k}{(k + 2 + 1/k)\sigma_j^4 + 2(1 + 1/k)\sigma_j^2 \text{tr}(\Gamma) + \text{tr}(\Gamma)^2/k}, \quad j \in [d \wedge n]. \quad (19)$$

In order to assess (18) for improvements over earlier bounds, one needs to gain more insights into the weights $\{\omega_j\}$. In Appendix A, we show that $\min_{1 \leq j \leq d \wedge n} \omega_j \geq 2/(2 + k)$, which implies that

$$\sum_{j=1}^{d \wedge n} (\alpha_j^*)^2 \sigma_j^2 \omega_j + \sigma^2 \frac{k}{n} \geq \frac{2}{2 + k} \frac{1}{n} \sum_{j=1}^{d \wedge n} (\alpha_j^*)^2 \sigma_j^2 + \sigma^2 \frac{k}{n} = \frac{2}{2 + k} (w^*)^\top \Gamma w^* + \sigma^2 \frac{k}{n}. \quad (20)$$

Compared to (14), this means that at best, the term $\text{tr}(\Gamma)\|w^*\|_2^2$ gets replaced by $(w^*)^\top \Gamma w^*$. However, even the lower bound in (20) does not yield satisfactory results if X is (approximately) of low rank as the bias term again scales as $O(1/k)$ independent of the spectrum. As a consequence, no matter how small the rank of X is, according to (20) we still obtain an upper bound on the bias of $O(1/\sqrt{n})$ and $k^* = \Omega(\sqrt{n})$ for the optimal number of random projections.

3. Improved Analysis (of CLS)

In this section we present and discuss the main result of the paper, a bound on the excess risk of CLS that is of a similar flavor of that of PCR in §2.2. In this manner, we establish a substantially stronger connection between CLS and PCR as could be made based on other analyses reviewed in §2.3.

3.1. Assumptions on the random projections

Our analysis basically requires R to be a Johnson-Lindenstrauss transform (JLT). The precise conditions are given as follows. The first one or slight variations thereof are conventionally used in the literature on random projections and randomized numerical linear algebra (see, e.g., Definition 1 in [41] or Definition 2.3 in [48]).

Definition 1. Fix an arbitrary set of points of m points $\{v_1, \dots, v_m\} \subset \mathbb{R}^d$. A random d -by- k matrix R is said to be an (m, ε, δ) -JLT for $\varepsilon, \delta \in (0, 1)$ if

$$(1 - \varepsilon)\|v_i\|_2^2 \leq \|R^\top v_i\|_2^2 \leq (1 + \varepsilon)\|v_i\|_2^2, \quad i \in [m],$$

holds with probability at least $1 - \delta$.

The next definition is akin to the restricted isometry property in the theory of sparse estimation [7] with the difference that approximate norm preservation is required only for a single subspace (as opposed to the union of subspaces of sparse vectors).

Definition 2. Let $\mathcal{V} \subset \mathbb{R}^d$ be an arbitrary subspace of dimension s . A random d -by- k matrix R is said to be an (s, ε, δ) -restricted isometry for $\varepsilon, \delta \in (0, 1)$ if

$$(1 - \varepsilon)\|v\|_2 \leq \|R^\top v\|_2 \leq (1 + \varepsilon)\|v\|_2 \quad \forall v \in \mathcal{V}$$

holds with probability at least $1 - \delta$.

Remark. The above two definitions are related in the following way: it is shown in [7] that if R is a $(\{\frac{12}{\varepsilon}\}^k, \varepsilon/2, \delta)$ -JLT, then R is also an (s, ε, δ) -restricted isometry; cf. also Theorem 2.3 in [48]. We here state two separate definitions for ease of reference.

3.2. Main result

We are now in position to state our main result.

Theorem 1. For $r \in \{1, \dots, d \wedge n\}$, let the following conditions be satisfied:

- (C1) R is an $(2nr, \varepsilon_1/\sqrt{r}, \delta_1)$ -JLT.
 (C2) R is an $(r, \varepsilon_2, \delta_2)$ -restricted isometry.

Then with probability at least $1 - \delta_1 - \delta_2$,

$$\|(I - P_{X_R})X\|_F^2 \leq \left(1 + \frac{\varepsilon_1^2}{(1 - \varepsilon_2)^4}\right) \|\Delta_r\|_F^2. \quad (21)$$

Conditional on the event (21), the excess risk of CLS (3) can be bounded as

$$\mathcal{E}(R) \leq \left(1 + \frac{\varepsilon_1^2}{(1 - \varepsilon_2)^4}\right) \|w^*\|_2^2 \frac{\|\Delta_r\|_F^2}{n} + \sigma^2 \frac{k}{n}. \quad (22)$$

A meaningful interpretation of Theorem 1 and the bound (22) requires an understanding of how large the number of random projections need to be so that the conditions (C1) and (C2) are satisfied. The next statement addresses this key point.

Proposition 1. Let R have entries drawn i.i.d. from a zero-mean sub-Gaussian distribution and variance $1/k$. If $k = \Omega(\varepsilon_1^{-2} r \{\log(r) + \log(n)\} \vee \varepsilon_2^{-2} \log(\varepsilon_2^{-1}) r^1)$, then R^\top satisfies conditions (C1), (C2) with $\delta_1 = \exp(-c' \log(nr))$ and $\delta_2 = \exp(-c \log(\varepsilon_2^{-1}) r^1)$ for absolute constants $c, c' > 0$.

According to Proposition 1, the bound on the excess risk (22) holds with high probability for $k = \Omega(r \log n)$ sub-Gaussian random projections. The logarithmic factor can potentially be removed, in view of results in Halko et al. [21] for Gaussian random projections. Specifically, in this case Halko et al. show that the expectation of the left hand side of (21) with respect to R can be bounded as

$$\mathbf{E}[\|(I - P_{X_R})X\|_F^2] \leq \left(1 + \frac{r}{k - r - 1}\right) \|\Delta_r\|_F^2, \quad (23)$$

for $k \geq r + 2$. In particular, for $k = 2r + 1$, the error in Frobenius norm for approximating the matrix X by $P_{X_R}X$ is within a factor two of the r -truncated SVD (4).

Bounds similar to (21) and (23) appear in various places in the literature on randomized linear algebra. For example, the proof of Lemma 4.2 in [48] yields the bound

$$\|(I - P_{X_R})X\|_F^2 \leq (1 + \epsilon) \|\Delta_r\|_F^2 \quad (24)$$

with probability at least 0.9 under two conditions that can be related to (C1) and (C2), respectively; the conditions in [48] are shown to hold for Gaussian R as long as $k = \Omega(r/\epsilon)$. The paper [41] establishes (24) for matrices

¹ If $r = O(1)$, r can be replaced by $r \vee \log n$ to make the failure probability small as n grows.

R having i.i.d. ± 1 zero-mean entries with probability at least $1/2$ as long as $k = \Omega(r/\epsilon + r \log r)$. An attractive feature of the specific approach taken herein is that verification of the underlying conditions is fairly manageable at a technical level for a wide array of random matrix ensembles.

3.3. Comparison of PCR and CLS

Statistical performance Comparing the bounds on the excess risk (8) and (22) for PCR and CLS, respectively, we see an agreement in their structure up to constant factors (ignoring the potentially spurious log-factor in the required number of projections k) and the change of $\|\alpha^*\|_\infty^2$ to $\|w^*\|_2^2$. As a result, CLS profits from a rapid decay in the sequence of singular values of X as does PCR, enjoying similarly favorable excess risk bounds under scenarios **(P)** and **(E)** given at the end of §2.2; changes in those bounds are only in terms of constants and the replacement of $\|\alpha^*\|_\infty^2$ by $\|w^*\|_2^2$. The latter may amount to a factor of $d \wedge n$ in the worst case². However, in light of (11) and (12) this difference does not have much of an effect as long as the spectrum of X exhibits strong decay. To be fair, it is worth pointing out that the bound (22) need not always be an improvement over those reviewed in §2.3, but it yields qualitatively a much better fit if the singular values decay rapidly.

In spite of this, there are cases in which the ratio of the excess risk of CLS and PCR can be arbitrarily large as we elaborate on in §4.3. In fact, it can be shown that CLS can be arbitrarily worse than ridge regression. This sharply contrasts with a result in [15] which states that the PCR risk is always within a factor of 4 of the corresponding ridge regression risk, while the converse is not true. Related observations have been made in the literature on non-parametric regression using reproducing kernel Hilbert space methods, e.g., [16, 17]. In that sense, PCR is superior to ridge regression, but this superiority does not hold for CLS.

Computational cost PCR requires access to the top r left singular vectors of X which is typically done via Krylov subspace methods like Lanczos' algorithm [19] in $O(ndr)$ flops on average. Assuming that $k = O(r)$, this is comparable to CLS whose dominating operation is given by the matrix-matrix multiplication XR . However, as discussed in [21], reducing PCR to an $O(ndr)$ operation is problematic as the actual computational complexity can vary significantly depending on subtle spectral properties of X : while the nominal complexities of the two approaches are comparable, CLS often requires less runtime. In addition, CLS provides several other advantages. The matrix-matrix multiplication is trivially parallelizable, and can be computed in a single pass over the data. The latter property becomes beneficial once X is too large to fit into the main memory since accessing external memory is slow. Moreover, Theorem 1 and

²if $n < d$, we may assume without loss of generality that w^* is contained in the orthogonal complement of the null space of X so that $\|\alpha^*\|_2^2 = \|w^*\|_2^2$

Proposition 1 are likely extendable to structured JLTs [3, 4, 33, 45] for which the cost of forming XR is considerably reduced.

The use of both Krylov subspace methods and CLS has recently been challenged by techniques developed in [5, 18] that approximate PCR by iterated ridge regression, thereby enabling the use of highly efficient large-scale optimization methods [8] for the latter and avoiding explicit dependence of the computational cost on the number of principal components. Despite promising empirical results in [5], the performance relative to CLS does not seem to have been extensively studied yet, and hence deserves further investigation.

4. Further topics

We here discuss several miscellaneous topics that naturally arise from the analysis of the preceding section.

4.1. Assessing the applicability of CLS when being on a computational budget

The analysis above indicates that CLS can achieve reasonable statistical performance while using a substantially reduced number of predictors provided that the singular values of X decay at a fast rate. Verifying such decay seems to require the SVD though, which would eliminate potential computational advantages of CLS over PCR. Direct evaluation of the quantity $\delta_R^2 = \|(I - P_{X_R})X\|_F^2$ in (21) is computationally demanding as well: finding the left singular vectors $U \in \mathbb{R}^{n \times k}$ of X_R can be done in $O(nk^2)$ flops; however, forming $P_{X_R}X = UU^\top X$ requires $O(ndk)$ flops which is of the same order of magnitude as computing X_R , the dominating operation in CLS. Hence, we would like to circumvent this operation. We here suggest to recycle the method of random projections in order to get an accurate estimate of δ_R^2 while achieving a reduction to $O(nd)$ flops. The basic idea is to apply P_{X_R} to a small number L of random elements from the range of X rather than to all its columns.

Proposition 2. *Consider a collection of L i.i.d. d -dimensional standard Gaussian random vectors $\{\omega_l\}_{l=1}^L$ independent of R . Conditional on R , consider the following estimator of $\delta_R^2 = \|(I - P_{X_R})X\|_F^2$:*

$$\widehat{\delta}_R^2 = \frac{1}{L} \sum_{l=1}^L \|X\omega_l - P_{X_R}X\omega_l\|_2^2.$$

Then, for any $c \in (0, 1)$ and any $C > 1$, as long as

$$L \geq \max \left\{ \frac{16}{(1-c)^2}, \frac{144}{(C-1)^2} \right\}$$

it holds that $\mathbf{P} \left(c\delta_R^2 \leq \widehat{\delta}_R^2 \leq C\delta_R^2 \right) \geq 0.96$, where the probability is w.r.t. $\{\omega_l\}_{l=1}^L$ and conditional on R .




FIG 1. Estimation of δ_R^2 for 100 different realizations of R with *i.i.d.* Gaussian entries and a fixed matrix X whose singular values decay polynomially (setting (\mathbf{P}) in §2.2, $q = 2$). Here, $\delta_R^2(k)$ is estimated for multiple k simultaneously using the same $\{\omega_l\}_{l=1}^L$ (here $L = 10$). Left: trajectories of $\{\delta_R^2(k)\}_{k=1}^{300}$ and $\{\hat{\delta}_R^2(k)\}_{k=1}^{300}$ for all 100 realizations of R . Right: Trajectories of $\{\hat{\delta}_R^2(k)\}_{k=1}^{300}$ for the first three realizations of R .

For example, setting $C = 3$, $c = 1/3$, we would need $L = 36$ to estimate δ_R^2 within a multiplicative factor of 3 with probability near 1. Note that computing $P_{X_R} X \omega_l = \mathcal{U}(\mathcal{U}^\top(X \omega_l))$ for a single l only amounts to $O(nd)$ flops. The constants in Proposition 2 may not necessarily be optimal. In the example of Figure 1, we use $L = 10$ random vectors to estimate $\delta_R^2(k)$ simultaneously for $1 \leq k \leq 300$.

4.2. Column Subsampling vs. Dense Random Projections

We can think of CLS as a scheme that picks k random elements from the subspace spanned by the columns of X , and subsequently uses these random elements to predict the response y . The analysis of the previous section asserts that if k is in proper relation to r , then we will do roughly as good as when using the top r principal components as predictors. For this result to hold, Theorem 1 imposes certain restrictions on the matrix R , or equivalently, the way the random elements are generated. It is natural to ask whether one can be more flexible in this regard. The simplest approach one can possibly think of would be to select columns from X uniformly at random without replacement, i.e.,

$$R = [e_{i_1} \dots e_{i_k}], \quad i_1 \sim \text{Unif}([d]), \quad i_\ell \sim \text{Unif}([d] \setminus \{i_1, \dots, i_{\ell-1}\}), \quad \ell = 2, \dots, k, \tag{25}$$

where $\{e_1, \dots, e_d\}$ are the canonical basis vectors of \mathbb{R}^d , and $\text{Unif}(\dots)$ denotes the uniform probability distribution, so that $XR = [X_{i_1} \dots X_{i_k}]$ is a random column submatrix of X . Note that a random matrix R generated according to (25) fails to be a $(1, \varepsilon, \delta)$ -JLT for any $\varepsilon \in (0, 1)$, any $\delta \in (0, 1/2)$ and any $k \leq \frac{d^3}{2}$, hence the framework used for Theorem 1 does not yield useful results.

³To see this, consider the first canonical basis vector e_1 . Then $\mathbf{P}(R^\top e_1 = 0) \geq 1/2$ for all $k \leq d/2$.

In general, the excess risk of column subsampling can be considerably worse than when using a suitable JLT. Both approaches are equivalent in terms of excess risk if i) $n \geq d$ and the spectrum is flat (scenario **(F)** in §2.2), or ii) X is a random Gaussian matrix with i.i.d. $N(0, 1)$ entries.

Proposition 3. *Let S be a random $d \times k$ matrix generated according to (25) and let R be an $d \times k$ random matrix with i.i.d. $N(0, 1/k)$ entries.*

i) *If $n \geq d$ and X has flat spectrum, i.e., $\sigma_1 = \dots = \sigma_d = \sqrt{n}$, then*

$$\mathbf{E}[\mathcal{E}(R)] = \mathbf{E}[\mathcal{E}(S)] = \left(1 - \frac{k}{d}\right) \|\alpha^*\|_2^2 + \frac{k}{n} \sigma^2.$$

ii) *If $n < d$, there exists an X with flat spectrum, i.e., $\sigma_1 = \dots = \sigma_n = \sqrt{d}$ such that*

$$\begin{aligned} \mathbf{E}[\mathcal{E}(R)] &= \|\alpha^*\|_2^2 \frac{d}{n} \left(1 - \frac{k}{n}\right) \|\alpha^*\|_2^2 + \frac{k}{n} \sigma^2 \\ &< \mathbf{E}[\mathcal{E}(S)] = \|\alpha^*\|_2^2 \frac{d}{n} \left(1 - \frac{k}{d}\right) + \frac{k}{n} \sigma^2. \end{aligned}$$

iii) *If X is a random $n \times d$ matrix with $N(0, 1)$ entries, then $\mathbf{E}[\mathcal{E}(R)] = \mathbf{E}[\mathcal{E}(S)]$, where the expectations are w.r.t. the randomness of both X and R respectively X and S .*

We note that property iii) crucially relies on Gaussianity of the entries of X . Proposition 3 will be complemented with numerical results presented in §5 below.

4.3. Averaging and Comparison to Ridge Regression

An idea in the spirit of bagging [9] considered in [44] is to generate an ensemble of i.i.d. random projections $\{R_b\}_{b=1}^B$ and then average the resulting predictions

$$\frac{1}{B} \sum_{b=1}^B X R_b \widehat{w}_{R_b} = \frac{1}{B} \sum_{b=1}^B P_{X R_b},$$

where \widehat{w}_{R_b} is the least squares estimator given response y and design matrix X_{R_b} , $b \in [B]$, and $\{P_{X R_b}\}_{b=1}^B$ are the projections on the respective column spaces of $\{X R_b\}_{b=1}^B$. As $B \rightarrow \infty$, the average of projectors can be expected to behave similarly to $\mathcal{P}_k = \mathbf{E}[P_{X R}]$. The next proposition summarizes basic properties of averaging and the operator \mathcal{P}_k .

Proposition 4. *Let R and $\{R_b\}_{b=1}^B$ be i.i.d. random projections.*

i) *Reduction in bias: $\forall B \geq 1$, $\mathbf{E}[\|(I - \frac{1}{B} \sum_{b=1}^B P_{X R_b}) X w^*\|_2^2 / n] \leq \mathbf{E}[\|(I - P_{X R}) X w^*\|_2^2 / n]$.*

ii) Let R have i.i.d. $N(0, 1)$ entries and let $\Delta(k) := \{z \in \mathbb{R}^{d \wedge n} : \sum_{j=1}^{d \wedge n} z_j = k, z_j \geq 0, j \in [d \wedge n]\}$. Then there exists $\{\eta_j\}_{j=1}^{d \wedge n} \in \Delta(k)$ depending only on $\{\sigma_j\}_{j=1}^{d \wedge n}$ s.t.

$$\begin{aligned} \frac{1}{n} \|(I - \mathcal{P}_k)Xw^*\|_2^2 &= \frac{1}{n} \sum_{j=1}^{d \wedge n} \{\alpha_j^*\}^2 \sigma_j^2 (1 - \eta_j)^2 \\ &\leq \frac{1}{n} \mathbf{E}[\|(I - P_{XR})Xw^*\|_2^2] = \frac{1}{n} \sum_{j=1}^{d \wedge n} \{\alpha_j^*\}^2 \sigma_j^2 (1 - \eta_j). \end{aligned}$$

iii) Reduction in variance: $\forall B \geq 1, \mathbf{E} \left[\left\| \frac{1}{B} \sum_{b=1}^B P_{XR_b} \xi \right\|_2^2 / n \right] \leq \sigma^2 k / n$.
 iv) In the setting of ii), we additionally have

$$\begin{aligned} \frac{1}{n} \mathbf{E}[\|\mathcal{P}_k \xi\|_2^2] &= \frac{\sigma^2}{n} \text{tr}(\mathcal{P}_k^2) = \frac{\sigma^2}{n} \sum_{j=1}^{d \wedge n} \eta_j^2 \\ &= \mathbf{E} \left[\sum_{\ell=1}^k \cos^2 \theta_\ell(\text{range}(XR), \text{range}(XR')) \right], \end{aligned}$$

where the last expectation is w.r.t. R and R' drawn independently, and $\{\theta_\ell(\mathcal{L}, \mathcal{L}')\}_{\ell=1}^k$ denote the canonical angles between two k -dimensional subspaces \mathcal{L} and \mathcal{L}' [19].

Parts ii) and iv) already appear in [44] for the $n \geq d$ case with slight differences in presentation. Parts ii) and iv) provide a rough quantification of the reduction in bias and variance by averaging in the limit $B \rightarrow \infty$ and Gaussian random projections. Since the relationship between the singular values $\{\sigma_j\}_{j=1}^{d \wedge n}$ and the coefficients $\{\eta_j\}_{j=1}^{d \wedge n}$ is not well understood, a more precise connection is yet to be made. In light of Theorem 1, we know that when choosing $k = \Omega(r \log n)$ large enough, the bias essentially depends only on the tail $\{\sigma_j^2\}_{j>r}$. Accordingly, the $\{\eta_j\}_{j=1}^r$ must be close to one, and accordingly the variance term in iv) is at least about $\sigma^2 r / n$. For a choice of k that makes the tail $\{\sigma_j^2\}_{j>r}$ (and hence the bias) negligible, averaging thus does not yield significant benefits. However, when using averaging, the optimal choice of k is guaranteed to be lower than when using a single random projection.

For fixed k , the maximum possible reduction in variance is seen to be a factor $k / (d \wedge n)$, from $\sigma^2 k / n$ to $\sigma^2 k^2 / \{(d \wedge n) \cdot n\}$, which follows immediately from the representation of the variance in terms of the $\{\eta_j\}_{j=1}^{d \wedge n}$ in iv) and the fact that the minimum ℓ_2 -norm over the simplex $\Delta(k)$ is attained at its barycenter. It is not hard to see that this maximum reduction is attained precisely when the spectrum of X is flat. This case remains of limited interest though as the bias is of the same order as without averaging.

The last identity in iv) provides a geometric interpretation of the variance after averaging in terms of the expected sum of squared cosines of the principal angles between to independently generated subspaces $\text{range}(XR)$ and $\text{range}(XR')$. Again, we see that the bias is low, i.e., if $\text{range}(XR)$ well approximates $\text{range}(X)$,

two randomly chosen subspaces will be essentially aligned so that the squared cosines evaluate all about one, which implies that averaging will not achieve any substantial reduction in variance.

Part ii) also yields a link between CLS and regularization methods such as ridge regression that limit the influence of small values among $\{\sigma_j^2\}$: given a regularization parameter $\lambda \geq 0$, ridge regression and PCR correspond to the maps $x \mapsto x/(x + \lambda)$ and $x \mapsto I(x \geq \lambda)x$, respectively, where $I(\cdot)$ denotes the indicator function. Accordingly, consider

$$\begin{aligned}\hat{w}_\lambda^{\text{ridge}} &= \operatorname{argmin}_{w \in \mathbb{R}^d} \frac{1}{n} \|y - Xw\|_2^2 + \lambda \|w\|_2^2, \\ \hat{w}_\lambda^{\text{PCR}} &= V_{r(\lambda)} \Sigma_{r(\lambda)}^{-1} U_{r(\lambda)}^\top y, \quad r(\lambda) = \max\{1 \leq j \leq d \wedge n : \sigma_j^2 \geq \lambda\}\end{aligned}$$

It has been shown in [15] that for all $\lambda \geq 0$

$$0 \leq \mathcal{E}(\hat{w}_\lambda^{\text{PCR}}) / \mathcal{E}(\hat{w}_\lambda^{\text{ridge}}) \leq 4,$$

with the left hand inequality being sharp. In the statement below, we provide an example showing that while PCR can be arbitrarily better than ridge regression in view of [15], CLS can be arbitrarily worse than the latter. Moreover, in that example, *averaged CLS* performs just on par with ridge regression, but can still be arbitrarily worse than PCR.

Proposition 5. *Suppose that $d \geq 6$ and that $X^\top X/n = \Delta I_d + (1 - \Delta) \mathbf{1}_d \mathbf{1}_d^\top$ for $0 < \Delta < 1$, so that $\sigma_1^2/n = d(1 - \Delta) + \Delta$, $\sigma_2^2/n = \dots = \sigma_d^2/n = \Delta$, and suppose further that $1 = (\alpha_1^*)^2 > \tau_n^2 := \sigma^2/n$, and $\alpha_2^* = \dots = \alpha_d^* = 0$. In the sequel, let $R = R_{d \times k}$ be a random d -by- k matrix with i.i.d. $N(0, 1)$ entries. Then*

$$\begin{aligned}(\text{PCR}) \quad & \min_{0 \leq r \leq d} \mathcal{E}(V_r) = \tau_n^2 \\ (\text{ridge}) \quad & \frac{1}{4} \left(\tau_n^2 + 2\tau_n \Delta (1 - \Delta)^{-1/2} \wedge \tau_n^2 d \wedge 1 \right) \\ & \leq \min_{\lambda \geq 0} \mathcal{E}(\hat{w}_\lambda) \leq \tau_n^2 + 2\tau_n \Delta (1 - \Delta)^{-1/2} \\ (\text{CLS}) \quad & \sqrt{d - 2} \Delta^{1/2} \tau_n \wedge \frac{1}{2} (d - 2) \tau_n^2 \wedge \frac{1}{2} \leq \min_{0 \leq k \leq d} \mathbf{E}[\mathcal{E}(R_{d \times k} \hat{w}_R)] \\ & \leq 2\sqrt{2} \tau_n \sqrt{d} \Delta^{1/2} \\ (\emptyset\text{-CLS}) \quad & 2^{-3/2} \tau_n \Delta (1 - \Delta)^{-1/2} \wedge \ell(\tau_n, d) \leq \min_{0 \leq k \leq d} \mathcal{E}(\mathbf{E}[R_{d \times k} \hat{w}_R]) \\ & \leq \tau_n^2 + 2^{5/2} \tau_n \Delta (1 - \Delta)^{-1/2}, \\ & \text{where } \ell(\tau, d) := \frac{1}{4} \left(\tau_n \sqrt{d} \wedge \frac{d - 4}{d} \tau_n^2 d \wedge 1 \right).\end{aligned}$$

The statement provides lower and upper bounds on the (expected) excess risk of PCR, Ridge Regression, CLS and averaged CLS (\emptyset -CLS) in the sense of Proposition 4, iv) in the situation that the Gram matrix $X^\top X/n$ has equi-correlation (or compound symmetry) structure, a specific spiked covariance

model with a single spike parameterized by $\Delta \in (0, 1)$; as $\Delta \rightarrow 0$, the spike dominates. Under the condition that $1 = (\alpha_1^*)^2 > \tau_n^2 = \sigma^2/n$, i.e., that the signal dominates the noise, and that there is no signal outside the spike, PCR with $r = 1$ achieves the optimum risk of $\tau_n^2 = O(n^{-1})$. In a high-dimensional regime with $d/n = \Omega(1)$, the excess risks of all other approaches exhibit slower rates than PCR. Specifically, in that regime the risk of both Ridge Regression and averaged CLS is seen to be proportional to $\tau_n \Delta (1 - \Delta)^{-1/2} = \Theta(n^{-1/2})$ as long as Δ is bounded away from zero and one. CLS behaves even worse with an expected risk of the order $\Omega(1)$. It is also noteworthy that the risk of CLS decays at the rate $\Theta(\Delta^{1/2})$ as $\Delta \rightarrow 0$, while the rate of decay is $\Theta(\Delta)$ for Ridge Regression and averaged CLS.

5. Experiments

We present the results of experiments with synthetic and real data in order to illustrate and support the main results of the paper, as well as pointing to some open questions. Specifically, the goal is to study the following two questions: (i) How does the CLS bias behave relative to that of PCR? (ii) How does CLS perform relative to column subsampling? Regarding (i), we observe overall good agreement with the theory, while for (ii) it is confirmed that there can indeed be substantial differences among the two approaches.

5.1. Synthetic data

We start by generating a random n -by- d matrix X_0 with $n = 1000$, $d = 500$, where the entries of X_0 are drawn i.i.d. from the $N(0, 1)$ distribution. The SVD of X_0 is given by

$$X_0 = U_0 \Sigma_0 V_0^\top, \quad U_0 \in \mathbb{R}^{n \times d}, \quad \Sigma_0 \in \mathbb{R}^{d \times d}, \quad V_0 \in \mathbb{R}^{d \times d}. \quad (26)$$

We then replace Σ_0 by a diagonal matrix Σ whose diagonal elements $\{\sigma_j\}_{j=1}^d$ are chosen in a deterministic fashion according to one of the following regimes:

$$\begin{aligned} \text{polynomial} &: \sigma_j \propto j^{-q}, \quad q \in \{.5, .75, 1, 1.5, 2, 4\}, \quad j \in [d], \\ \text{exponential} &: \sigma_j \propto 0.9^j, \quad j \in [d], \end{aligned}$$

where the constant of proportionality is determined by the scaling $\sum_{j=1}^d \sigma_j^2 = n \cdot d$. We subsequently work with $X = U_0 \Sigma V_0^\top$, generating data from the model

$$y = X w^* + \sigma \xi, \quad (27)$$

where w^* is drawn uniformly from the unit sphere in \mathbb{R}^d , $\sigma \in 2^p$, $p \in \{-1, -0.5, \dots, 1\}$, and ξ has i.i.d. standard Gaussian entries.

Given data (X, y) , we then perform PCR with ten different choices of r , using an equi-spaced grid of values depending on the regime according to which X has been generated. For CLS, R is chosen as a standard d -by- k Gaussian matrix with

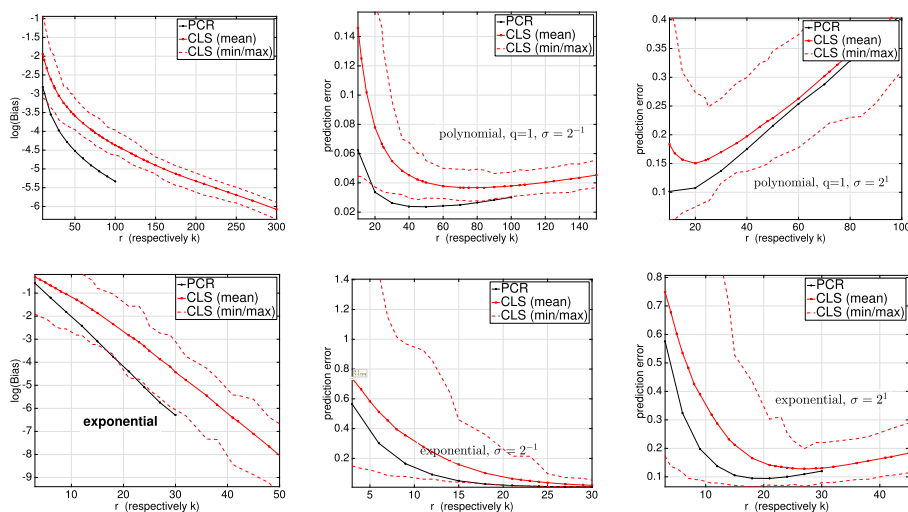


FIG 2. Results of the synthetic data experiment. From left to right: the bias $\|(I - P_{X_R})Xw^*\|_2^2/n$ (log scale), and the mean squared prediction error $\|Xw^* - X_R\hat{w}_R\|_2^2/n$ for $\sigma = 1/2$ (middle) and $\sigma = 2$ (right) in dependence of k (horizontal axis) for CLS in relation to PCR. Solid curves are averages, dashed curves minima and maxima (only CLS) over 100 replications. Top: spectrum with polynomial decay ($q = 1$), Bottom: exponential decay. For comparison, the mean squared prediction error of OLS $\sigma^2 d/n$ equals $1/8 = .125$ for the middle column and 2 for the right column.

$k = \alpha r$, where the over-sampling factor $\alpha \in \{1, 1.2, 1.5, 2, 2.5, 3\}$. We conduct 100 independent replications for each regime. Our main interest is in the bias and the prediction error of PCR and CLS:

$$\begin{aligned} \|(I - P_{U_r})Xw^*\|_2^2/n & \quad \text{vs.} \quad \|(I - P_{X_R})Xw^*\|_2^2/n, \\ \|Xw^* - XV_r\hat{w}_{V_r}\|_2^2/n & \quad \text{vs.} \quad \|Xw^* - X_R\hat{w}_R\|_2^2/n, \end{aligned}$$

where \hat{w}_{V_r} and \hat{w}_R denote the least squares estimator for data (XV_r, y) and (X_R, y) , respectively. In order to compare CLS and column subsampling, we also generate R as a $d \times k$ column submatrix of the identity chosen uniformly at random, cf. (25).

A subset of the results involving two different regimes of decay is shown in Figure 2. In the regime of polynomial decay ($q = 1$), we observe that the bias of CLS is roughly proportional to that of PCR (or alternatively, we need to choose k as a suitable multiple of r to achieve the same bias). Accordingly, the dip in the prediction error curve occurs for $k = 2r^*$ with $r^* = 40$ yielding the smallest prediction error for PCR. In both low and high noise settings, PCR and CLS improve significantly over OLS in terms of prediction error (≈ 0.02 and ≈ 0.04 vs. 0.125 and ≈ 0.1 and ≈ 0.15 vs. 2). In the regime of exponential decay, the bias of CLS is not quite proportional to that of PCR for small r , but this improves once r reaches 20. Overall, the results agree well with what is suggested by Theorem 1.




FIG 3. Bias $\|(I - P_{X_R})Xw^*\|_2^2/n$ (log scale) when using Gaussian random projections respectively column subsampling. Left: the entries of X_0 (26) are Gaussian. Right: the entries of X_0 (26) are Cauchy. Solid curves are averages, dashed curves minima and maxima over 100 replications.

Figure 3 shows that if X_0 in (26) is generated from a Gaussian distribution, Gaussian random projections and column subsampling perform the same on average as asserted by Proposition 3. Interestingly, the performance of column sampling degrades considerably when the entries of X_0 are drawn i.i.d. from the standard Cauchy distribution, whereas Gaussian random projections are not affected by this change.

5.2. Real data

We here consider two data sets from the UCI machine learning repository for illustration purposes.

Twitter social media buzz Our data analysis is inspired by that in [35]. This is a regression problem 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 here only work with the first 8000 observations. Several of the original predictor variables as well as the response variable are log-transformed prior to analysis. Following [35], we add quadratic interactions which yields $d = 3080$ predictors in total. We consider 50 random partitions into a training set of size 6000 and a test set of size 2000 which is used to evaluate the prediction error. The design matrices of the training sets are centered and subsequently scaled to unit norm before performing least squares regression with a centered response and a reduced design matrix obtained from i) a truncated SVD with $r \in \{5, 10, \dots, 50, 60, \dots, 100, 120, \dots, 200\}$, ii) random Gaussian projections with $k = r\alpha$, where the grid for the factor α is as for the synthetic data, iii) subsets of k columns sampled uniformly at random without replacement. The thus obtained regression coefficients are back-transformed to account for the preliminary centering/scaling step before using them to make predictions on the test set. Approaches i) to iii) are compared in terms of the mean squared

prediction error on the test set. For each of the 50 partitions into training and test set, we obtain ten i.i.d. sets of random projections respectively subsampled columns for ii) respectively iii) and perform regression with each of the resulting reduced matrices, and compute the average as well as the maximum error over each of those ten runs.

Blog feedback The task associated with the data set is the prediction of the number comments on blog posts [12] within a 24-hour time window after a certain base time. The training set consists of $n = 52,397$ observations and originally 280 predictors including meta data about the blogs in which the posts were made, the number of comments received within specific time windows before the base time, number of comments on related posts, bag-of-words data, and the weekday of the post. After eliminating redundant and non-informative predictors, we end up with 114 predictors, several of which are log-transformed. A subset of those are expanded in terms of quadratic and interaction terms which eventually yields $d = 2,589$. The target variable is log-transformed as well. As distinguished from the first case study, this data set comes with a fixed test set of size 7,624. Evaluation then proceeds as described above, with the only difference that $r \in \{10, 20, 50, 100, \dots, 500\}$ in order to adjust for a slower rate of decay of the singular values.

The main results are summarized in Figure 4. The left panels show that the singular values exhibit different rates of decay for the first respectively second data set. For the Twitter data, the decay is not far from linear on a log-scale, whereas the decay is noticeably slower for the Blog Feedback data. As already seen for the synthetic data experiments, CLS requires only a moderate amount of oversampling to achieve the approximation error of PCR $\|X - P_{U_r} X\|_F^2$; we here use this quantity as a surrogate for the bias since w^* is unknown, and in order to establish a connection to result (21) in Theorem 1. The corresponding quantity of column subsampling is essentially to that of CLS for the Blog Feedback data, but is markedly worse for the Twitter data set. Turning to the right panels, we observe that the test error of PCR dips for $r = 70$ (Twitter) respectively $r = 150$ (Blog Feedback). The errors of CLS and column subsampling are not far off, but again an increase of k relative to r is indicated to achieve optimal results. Comparing CLS and column sampling, we see that the latter overall performs better particularly for small k , but is also less stable in the sense that the gap between the average error and the maximum error (taken over ten different realizations of the random matrix R) is substantially larger than for CLS. In particular, for the Twitter data, the maximum error of CLS is uniformly smaller than that of column subsampling. Eventually, the relative performance will depend on properties of X (as illustrated by Figure 3), and likely on possible sparsity of w^* and its interaction with properties of X .

6. Conclusion

Regarding the use of random projections in linear regression, the literature has mainly focused on the setting in which the random matrix R is applied from

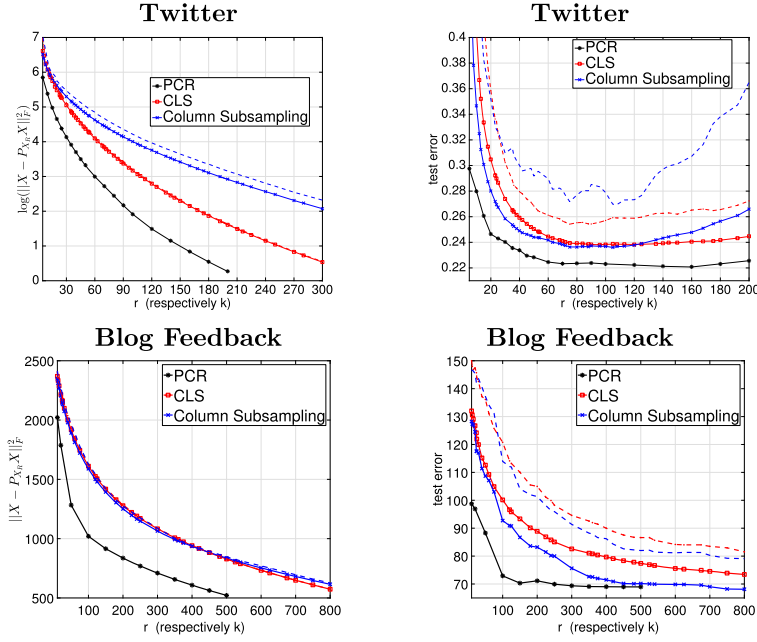


FIG 4. Left: Approximation errors $\|X - P_{U_r} X\|_F^2$ (PCR) respectively $\|X - P_{X_R} X\|_F^2$ (CLS and column subsampling) in dependence of r (or k); for “Twitter”, we use a log-scale and averages over the 50 training/test set partitions. Right: Mean squared prediction errors on the test data vs. r or k . For CLS and column subsampling, we plot both the mean and the maximum over the 10 realizations of R .

the left, i.e., X is reduced to RX . The converse setting with X being reduced to XR has been studied in several earlier papers as well, however, without establishing a tight link to PCR at the level of achievable excess risk as made herein. Towards the end of our paper, we raise the question how much randomness is needed in generating a random subspace so that such connection holds true. Gaussian random projections induce “maximum randomness”, whereas in column sub-sampling we only consider random subspaces spanned by subsets of the canonical basis vectors. Given the dramatic computational advantages of the latter over “dense” random projections it is worth elaborating general conditions under which column sub-sampling can be shown to exhibit similar statistical performance as classical Johnson-Lindenstrauss transforms.

Appendix A: Proof of the lower bound (20)

Introducing a new set of weights $\{a_j\}_{j=1}^{d \wedge n}$ by the relation $\sigma_j^2 = a_j \text{tr}(\Sigma)$, $\sum_{j=1}^{d \wedge n} a_j = 1$, the $\{\omega_j\}_{j=1}^{d \wedge n}$ (19) can be re-expressed as

$$\omega_j = \frac{(1 + 1/k)a_j^2 \text{tr}(\Sigma)^2 + (1 + 2/k)a_j \text{tr}(\Sigma)^2 + \text{tr}(\Sigma)^2/k}{(k + 2 + 1/k)a_j^2 \text{tr}(\Sigma)^2 + 2(1 + 1/k)a_j \text{tr}(\Sigma)^2 + \text{tr}(\Sigma)^2/k}, \quad j \in [d \wedge n].$$

Factoring out the term $\text{tr}(\Sigma)^2$, we arrive at (cf. Equation (12) in [44])

$$\omega_j = \frac{(1 + 1/k)a_j^2 + (1 + 2/k)a_j + 1/k}{(k + 2 + 1/k)a_j^2 + 2(1 + 1/k)a_j + 1/k}, \quad j \in [d \wedge n]$$

After some algebra, one obtains the following more compact representation:

$$\omega_j = \frac{1 + a_j}{1 + a_j + a_j k} = \frac{1/a_j + 1}{1/a_j + 1 + k}, \quad j \in [d \wedge n].$$

which is a monotonically decreasing function in $0 \leq a_j \leq 1$, whose minimum is achieved for $a_j = 1$. Consequently, $\omega_j \geq \frac{2}{2+k}$, $j \in [d \wedge n]$, and we conclude (20).

Appendix B: Proof of Proposition 1

Let R be a random matrix of i.i.d. sub-Gaussian random variables with zero mean and variance $1/k$.

Regarding **(C1)**, it is shown in [38] (see Theorem 3.1 and the proof therein; cf. also [26]) that for any fixed $v \in \mathbb{R}^d$ and $\varepsilon' \in (0, 1)$

$$\mathbf{P} \left((1 - \varepsilon') \|v\|_2^2 \leq \|R^\top v\|_2^2 \leq (1 + \varepsilon') \|v\|_2^2 \right) \leq 2 \exp(-c_0(\varepsilon')^2 k).$$

for some absolute constant $c_0 > 0$. It hence follows from the union bound that for any set \mathcal{S} of vectors in \mathbb{R}^d , $|\mathcal{S}| = 2n \cdot r$,

$$\mathbf{P}(\forall v \in \mathcal{S} : (1 - \varepsilon') \|v\|_2^2 \leq \|R^\top v\|_2^2 \leq (1 + \varepsilon') \|v\|_2^2) \leq \exp(-c_0(\varepsilon')^2 k + \log(4nr)).$$

Setting $\varepsilon' = \varepsilon_1/\sqrt{r}$ for $\varepsilon_1 \in (0, 1)$, it follows that for $k = \Omega(\varepsilon_1^{-2} r \log(nr))$, condition **(C1)** holds with $\delta_1 = \exp(-c' \log(nr))$.

Turning to **(C2)**, it follows from arguments in [7] (cf. Lemma 5.1 therein) that for any fixed subspace \mathcal{V} of dimension r in \mathbb{R}^d , $r < k$,

$$(1 - \varepsilon_2) \|v\|_2 \leq \|R^\top v\|_2 \leq (1 + \varepsilon_2) \|v\|_2 \quad \text{for all } v \in \mathcal{V},$$

with probability at least

$$1 - 2(12/\varepsilon_2)^r \exp(-c_0 \varepsilon_2^2 k) = 1 - \exp(-c_0 \varepsilon_2^2 k + r \log(12/\varepsilon_2) + \log(2)),$$

Hence, for $k = \Omega(\varepsilon_2^{-2} \log(\varepsilon_2^{-1}) r)$, **(C2)** holds with $\delta_2 = \exp(-c \log(\varepsilon_2^{-1}) r)$. This concludes the proof of the proposition.

Appendix C: Proof of Theorem 1

Before going into the proof, let us introduce a bit more of notation. Below, $\mathcal{T}_s(M)$ denotes the best rank- s approximation of a matrix M with respect to Frobenius norm which can be obtained from a truncated SVD, cf. (8). Moreover, we write M^- for the Moore-Penrose pseudoinverse of a matrix M . The j -th column of M is denoted by $M_{:,j}$. $\|M\|_2$ denotes the spectral norm.

Note that in view of (3)

$$\mathcal{E}(R) = \|(I - P_{X_R})Xw^*\|_2^2/n + \sigma^2 \text{rank}(X_R)/n,$$

we have

$$\begin{aligned} \mathcal{E}(R) &\leq (\|(I - P_{X_R})\|_2^2/n) \|w^*\|_2^2 + \sigma^2 k/n \\ &\leq (\|(I - P_{X_R})\|_F^2/n) \|w^*\|_2^2 + \sigma^2 k/n, \end{aligned}$$

so that (22) immediately follows from (21). In the sequel, we hence prove (21), following the strategy of the proof of Theorem 14 in [41]. The key idea is a reduction to a series of *sketched* least squares regression problems of the form

$$\min_{\lambda_i} \|R^\top (X^\top)_{:,i} - R^\top \mathcal{T}_r(X)^\top \lambda_i\|_2^2, \quad i \in [n],$$

where $(X^\top)_{:,i}$ denotes the i -th column of X^\top , $i \in [n]$. Problems of this form are a cornerstone in the randomized linear algebra literature, where they have been studied extensively. The importance of the above regression problems can be seen from the intuition that

$$\begin{aligned} \sum_{i=1}^n \min_{\lambda_i} \|R^\top (X^\top)_{:,i} - R^\top \mathcal{T}_r(X)^\top \lambda_i\|_2^2 &\approx \sum_{i=1}^n \min_{\lambda_i} \|(X^\top)_{:,i} - \mathcal{T}_r(X)^\top \lambda_i\|_2^2 \\ &= \|X - \mathcal{T}_r(X)\|_F^2. \end{aligned}$$

With that in mind, we turn to the proof which is partitioned into three basic steps.

Step 1

Lemma C.1. *We have*

$$\|X - P_{X_R}X\|_F^2 \leq \|X - \mathcal{T}_r(P_{X_R}X)\|_F^2. \quad (28)$$

Proof. Observe that according to the definition of P_{X_R} , we have

$$\min_{B \in \mathbb{R}^{k \times d}} \|X - X_R B\|_F^2 = \|X - P_{X_R}X\|_F^2. \quad (29)$$

Let $B^* \in \mathbb{R}^{k \times d}$ denote a minimizer of the optimization problem on the l.h.s. of (29) such that $P_{X_R}X = X_R B^*$. Let the SVD of that matrix be given by

$$X_R B^* = \begin{matrix} \Upsilon & \Xi & \Psi^\top \\ n \times d & d \times d & d \times d \end{matrix}.$$

Denote by $M_r \in \mathbb{R}^{d \times d}$ the diagonal matrix whose first r diagonal entries are equal to one and zero else. Then $\mathcal{T}_r(P_{X_R}X) = X_R B^* M_r = X_R \tilde{B}$. Since $\tilde{B} = B^* M_r$ is a feasible solution for the minimization problem (29), we conclude (28). \square

Lemma C.2. *We have*

$$\|X - \mathcal{T}_r(P_{X_R}X)\|_F^2 \leq \|X - \Pi X\|_F^2, \tag{30}$$

where Π is the orthogonal projection on the subspace spanned by the columns of $\Phi_r = P_{X_R}\mathcal{T}_r(X)$, i.e.

$$\Pi = P_{\Phi_r} = P_{P_{X_R}\mathcal{T}_r(X)}. \tag{31}$$

Proof. Consider the following optimization problem:

$$\min_{\text{rank}(B) \leq r} \|X - X_R B\|_F^2.$$

Then any minimizer B^* of the above problem satisfies $X_R B^* = \mathcal{T}_r(P_{X_R}X)$ (see Proposition 1 and Lemma 14 in [11]). Noting that $\Pi = X_R M$ for some matrix $M \in \mathbb{R}^{k \times d}$ with $\text{rank}(M) \leq r$ (as $\mathcal{T}_r(X)$ has rank no more than r), M is feasible for the above optimization problem, and we conclude (30). \square

We conclude **Step 1**. by combining Lemmas C.1 and C.2:

$$\|X - P_{X_R}X\|_F^2 \leq \|X - \Pi X\|_F^2 \tag{32}$$

with Π defined in (31).

Step 2

In the second step, we decompose $\|X - \Pi X\|_F^2$ into two parts: an “easy” part and one more delicate part that requires sophisticated analysis. Recalling (4), we have

$$\begin{aligned} \|X - \Pi X\|_F^2 &= \|U\Sigma V^\top - \Pi U\Sigma V^\top\|_F^2 \\ &= \|U\Sigma - \Pi U\Sigma\|_F^2 \\ &= \|U_r \Sigma_r - \Pi U_r \Sigma_r\|_F^2 + \|U_{r+} \Sigma_{r+} - \Pi U_{r+} \Sigma_{r+}\|_F^2 \\ &= \|U_r \Sigma_r - \Pi U_r \Sigma_r\|_F^2 + \|(I - \Pi)U_{r+} \Sigma_{r+}\|_F^2 \\ &\leq \|U_r \Sigma_r - \Pi U_r \Sigma_r\|_F^2 + \|U_{r+} \Sigma_{r+}\|_F^2 \\ &= \underbrace{\|U_r \Sigma_r - \Pi U_r \Sigma_r\|_F^2}_{\text{part requiring special treatment}} + \underbrace{\|U_{r+} \Sigma_{r+}\|_F^2}_{\text{part that we need (up to constant)}} \end{aligned} \tag{33}$$

where the inequality follows from the fact that $I - \Pi$ is an orthogonal projection.

Step 3

It remains to bound

$$\|U_r \Sigma_r - \Pi U_r \Sigma_r\|_F^2 = \|\mathcal{T}_r(X) - \Pi \mathcal{T}_r(X)\|_F^2.$$

Let us write $C^* = X_R^-$ and $\tilde{C} = (\mathcal{T}_r(X)R)^-$. Note that for any matrix M of appropriate dimension, we have

$$\|M - P_{X_R} M\|_F^2 = \min_{C \in \mathbb{R}^{k \times n}} \|M - X_R C\|_F^2 = \|M - X_R C^*\|_F^2 \leq \|M - X_R \tilde{C}\|_F^2. \quad (34)$$

Moreover, observe that according to the definition of Π in (31)

$$\Pi \mathcal{T}_r(X) = P_{P_{X_R} \mathcal{T}_r(X)} \mathcal{T}_r(X) = P_{X_R} \mathcal{T}_r(X). \quad (35)$$

Using (34) and (35), we obtain that

$$\begin{aligned} \|\mathcal{T}_r(X) - \Pi \mathcal{T}_r(X)\|_F^2 &= \|\mathcal{T}_r(X) - X_R (X_R)^- \mathcal{T}_r(X)\|_F^2 \\ &\leq \|\mathcal{T}_r(X) - X_R \{\mathcal{T}_r(X)R\}^- \mathcal{T}_r(X)\|_F^2 \\ &= \|\mathcal{T}_r(X)^\top - \mathcal{T}_r(X)^\top \{R^\top \mathcal{T}_r(X)^\top\}^- R^\top X^\top\|_F^2 \end{aligned} \quad (36)$$

Define

$$b_i = (X^\top)_{:,i} \in \mathbb{R}^d, \quad i \in [n], \quad \text{and} \quad A = \mathcal{T}_r(X)^\top \in \mathbb{R}^{d \times n}, \quad (37)$$

and consider the least squares problems

$$\min_{\lambda_i} \|b_i - A\lambda_i\|_2^2$$

with minimizer λ_i^* , $i = 1, \dots, n$, and the corresponding *sketched regression problems* with sketching matrix R^\top :

$$\min_{\lambda_i} \|R^\top b_i - R^\top A\lambda_i\|_2^2,$$

with minimizer $\tilde{\lambda}_i$, $i = 1, \dots, n$. It is straightforward to show that

$$A\lambda_i^* = (\mathcal{T}_r(X)^\top)_{:,i}, \quad i \in [n].$$

For the sketched regression problems, an optimal set of coefficients is given by

$$\tilde{\lambda}_i = \{R^\top \mathcal{T}_r(X)^\top\}^- R^\top (X^\top)_{:,i}, \quad i \in [n],$$

so that

$$A\tilde{\lambda}_i = \mathcal{T}_r(X)^\top \{R^\top \mathcal{T}_r(X)^\top\}^- R^\top (X^\top)_{:,i}, \quad i \in [n].$$

Identifying terms, we see that the right hand side in (36) can be written as

$$\begin{aligned} &\|\mathcal{T}_r(X)^\top - \mathcal{T}_r(X)^\top \{R^\top \mathcal{T}_r(X)^\top\}^- R^\top X^\top\|_F^2 \\ &= \sum_{i=1}^n \|(\mathcal{T}_r(X)^\top)_{:,i} - \mathcal{T}_r(X)^\top \{R^\top \mathcal{T}_r(X)^\top\}^- R^\top (X^\top)_{:,i}\|_2^2 \\ &= \sum_{i=1}^n \|A(\lambda_i^* - \tilde{\lambda}_i)\|_2^2 \\ &= \sum_{i=1}^n \|\beta_i\|_2^2, \quad \beta_i = A(\lambda_i^* - \tilde{\lambda}_i), \quad i \in [n]. \end{aligned} \quad (38)$$

Consider the residuals

$$w_i = b_i - A\lambda_i^* = (X^\top)_{:,i} - (\mathcal{T}_r(X)^\top)_{:,i}. \quad (39)$$

By analyzing the structure of (general) sketched regression problems, it can be shown that

$$V_r^\top R R^\top V_r \beta_i = V_r^\top R R^\top w_i, \quad (40)$$

where V_r is the same matrix as in (4). The analysis leading to property (40) will be given at the end of this proof. In the sequel, we use this property in combination with conditions **(C1)** and **(C2)** to deduce the final result. We will first derive a lower bound on the l.h.s. of (40) with the help of **(C2)**, and then we derive an upper bound on the r.h.s. by means of **(C1)**. Combining both, we obtain an upper bound on $\sum_{i=1}^n \|\beta_i\|_2^2$ and in turn on the quantity $\|\mathcal{T}_r(X) - \Pi \mathcal{T}_r(X)\|_F^2$ that we eventually need to bound.

Let $\mathcal{V}_r \subset \mathbb{R}^d$ denote the column space of V_r . Invoking **(C2)** with $\mathcal{V} = \mathcal{V}_r$, the following event holds with probability at least $1 - \delta_2$:

$$\|R^\top V_r v\|_2^2 \geq (1 - \varepsilon_2)^2 \|v\|_2^2 \quad \forall v \in \mathbb{R}^d,$$

or equivalently,

$$\lambda_{\min}(\Omega) \geq (1 - \varepsilon_2)^2,$$

where $\Omega = V_r^\top R R^\top V_r$ and $\lambda_{\min}(\cdot)$ denotes the smallest eigenvalue. Conditional on that event, we have that

$$\begin{aligned} \|V_r^\top R R^\top V_r \beta_i\|_2^2 &= \beta_i^\top \Omega^2 \beta_i \\ &\geq \lambda_{\min}(\Omega^2) \|\beta_i\|_2^2 \\ &\geq (1 - \varepsilon_2)^4 \|\beta_i\|_2^2. \end{aligned} \quad (41)$$

Next, observe that $V_{:,j}^\top w_i = 0$, $j = 1, \dots, r$, $i = 1, \dots, n$, as follows immediately from the definition of the $\{w_i\}_{i=1}^n$ in (39). We now apply **(C1)** with the following set of vectors:

$$\mathcal{S} = \{V_{:,j} + \tilde{w}_i, V_{:,j} - \tilde{w}_i, i \in [n], j \in [r]\},$$

where $\tilde{w}_i = w_i / \|w_i\|_2$, $i = 1, \dots, n$. Note that $|\mathcal{S}| = 2rn$. In the next step, we will establish that with the specified probability, the inner products between $V_{:,j}^\top w_i$, are preserved up to an additive term of $\varepsilon'_1 \|w_i\|_2$, $i \in [n]$, $j \in [r]$, where $\varepsilon'_1 = \varepsilon_1 / \sqrt{r}$ according to **(C1)**.

Recall that for arbitrary x, y , it holds that $\langle x, y \rangle = \frac{1}{4} (\|x + y\|_2^2 - \|x - y\|_2^2)$. With R^\top being a $(2nr, \varepsilon'_1, \delta_1)$ JLT, we therefore have with probability at least $1 - \delta_1$

$$\begin{aligned} 4 \langle R^\top V_{:,j}, R^\top \tilde{w}_i \rangle &= \|R^\top V_{:,j} + R^\top \tilde{w}_i\|_2^2 - \|R^\top V_{:,j} - R^\top \tilde{w}_i\|_2^2 \\ &\geq (1 - \varepsilon'_1) \|V_{:,j} + \tilde{w}_i\|_2^2 - (1 + \varepsilon'_1) \|V_{:,j} - \tilde{w}_i\|_2^2 \\ &= 4 \langle V_{:,j}, \tilde{w}_i \rangle - 2\varepsilon'_1 (\|V_{:,j}\|_2^2 + \|\tilde{w}_i\|_2^2) \\ &= 4 \langle V_{:,j}, \tilde{w}_i \rangle - 4\varepsilon'_1. \end{aligned}$$

It follows that $\langle R^\top V_{:,j}, R^\top \tilde{w}_i \rangle \geq \langle V_{:,j}, \tilde{w}_i \rangle - \varepsilon'_1$ and in turn also $\langle R^\top V_{:,j}, R^\top w_i \rangle \geq \langle V_{:,j}, w_i \rangle - \varepsilon'_1 \|w_i\|_2$ by homogeneity.

Regarding the upper bound, we argue analogously:

$$\begin{aligned} 4 \langle R^\top V_{:,j}, R^\top \tilde{w}_i \rangle &= \|R^\top V_{:,j} + R^\top \tilde{w}_i\|_2^2 - \|R^\top V_{:,j} - R^\top \tilde{w}_i\|_2^2 \\ &\leq (1 + \varepsilon'_1) \|V_{:,j} + \tilde{w}_i\|_2^2 - (1 - \varepsilon'_1) \|V_{:,j} - \tilde{w}_i\|_2^2 \\ &= 4 \langle V_{:,j}, \tilde{w}_i \rangle + 2\varepsilon'_1 (\|V_{:,j}\|_2^2 + \|\tilde{w}_i\|_2^2) \\ &= 4 \langle V_{:,j}, \tilde{w}_i \rangle + 4\varepsilon'_1. \end{aligned}$$

and thus $\langle R^\top V_{:,j}, R^\top \tilde{w}_i \rangle \leq \langle V_{:,j}, \tilde{w}_i \rangle + \varepsilon'_1$ and in turn $\langle R^\top V_{:,j}, R^\top w_i \rangle \leq \langle V_{:,j}, w_i \rangle + \varepsilon'_1 \|w_i\|_2$. We now use these bounds as follows (recall that $\langle V_{:,j}, w_i \rangle = 0$, $j \in [r]$, $i \in [n]$):

$$\begin{aligned} \sum_{i=1}^n \|V_r^\top R R^\top w_i\|_2^2 &= \sum_{i=1}^n \sum_{j=1}^r \langle R^\top V_{:,j}, R^\top w_i \rangle^2 \\ &\leq \sum_{i=1}^n \sum_{j=1}^r (\varepsilon'_1)^2 \|w_i\|_2^2 \\ &= r(\varepsilon'_1)^2 \sum_{i=1}^n \|w_i\|_2^2 \\ &= \varepsilon_1^2 \|X - \mathcal{T}_r(X)\|_F^2 \end{aligned} \tag{42}$$

where the last line is immediate from the definition of the $\{w_i\}_{i=1}^n$ in (39). Combining (33), (36), (38), (40), (41), (42), we obtain (21) and the assertion of the theorem follows.

In order to finish the proof, it remains to establish (40) as is done below.

For $A \in \mathbb{R}^{d \times n}$, $b \in \mathbb{R}^d$, consider the least squares problem of the form

$$\min_{\lambda \in \mathbb{R}^n} \|A\lambda - b\|_2^2$$

and the corresponding sketched regression problem with sketching matrix R^\top

$$\min_{\lambda} \|R^\top A\lambda - R^\top b\|_2^2,$$

Let λ^* denote a minimizer of the original least squares problem and let $\tilde{\lambda}$ denote the minimizer of the sketched least squares problem. Furthermore, we write \mathcal{U} for the matrix of left singular vectors of A .

We then have the following properties:

- (P1) $A\lambda^* = \mathcal{U}\alpha$,
- (P2) $b = A\lambda^* + w$, with w orthogonal to the columns of \mathcal{U} .
- (P3) $A\tilde{\lambda} - A\lambda^* = \mathcal{U}\beta$,

for certain vectors α and β . We now decompose the least squares error when using $\tilde{\lambda}$:

$$\|b - A\tilde{\lambda}\|_2^2 = \|b - A\lambda^* + A(\lambda^* - \tilde{\lambda})\|_2^2$$

$$\begin{aligned}
&= \|b - A\lambda^*\|_2^2 + \|A(\lambda^* - \tilde{\lambda})\|_2^2 \\
&= \|w\|_2^2 + \|\mathcal{U}\beta\|_2^2 \\
&= \|w\|_2^2 + \|\beta\|_2^2
\end{aligned}$$

Bringing the sketching matrix R^\top into play, we have

$$\begin{aligned}
R^\top \mathcal{U}(\alpha + \beta) &= R^\top A\lambda^* + R^\top (A\tilde{\lambda} - A\lambda^*) \\
&= R^\top A\tilde{\lambda} \\
&= P_{R^\top A} R^\top b \\
&= P_{R^\top \mathcal{U}} R^\top b.
\end{aligned}$$

Furthermore, we have

$$\begin{aligned}
P_{R^\top \mathcal{U}} R^\top b &= P_{R^\top \mathcal{U}} R^\top (\mathcal{U}\alpha + w) \\
&= R^\top \mathcal{U}\alpha + P_{R^\top \mathcal{U}} R^\top w.
\end{aligned}$$

Combining the previous displays, we obtain that

$$R^\top \mathcal{U}(\alpha + \beta) = R^\top \mathcal{U}\alpha + P_{R^\top \mathcal{U}} R^\top w$$

and thus

$$R^\top \mathcal{U}\beta = P_{R^\top \mathcal{U}} R^\top w.$$

Multiplying both sides with $\mathcal{U}^\top R$, this implies

$$\begin{aligned}
\mathcal{U}^\top R R^\top \mathcal{U}\beta &= \mathcal{U}^\top R P_{R^\top \mathcal{U}} R^\top w \\
&= \mathcal{U}^\top R R^\top w.
\end{aligned} \tag{43}$$

Note that (43) has the form as claimed in (40) with V_r playing the role of \mathcal{U} : according to (37), this is as it should be since V_r contains the left singular vectors of $\mathcal{T}_r(X)^\top$. The proof is thus complete.

Appendix D: Proof of Proposition 2

Let us recall that the statement is conditional on R , and for what follows only $\{\omega_l\}_{l=1}^L$ is considered as random. We first verify that $\|X\omega_l - P_{X_R} X\omega_l\|_2^2$ is an unbiased estimator of δ_R^2 , $l \in [L]$. We have

$$\begin{aligned}
\mathbf{E}[\|X\omega_l - P_{X_R} X\omega_l\|_2^2] &= \mathbf{E}[\|(I - P_{X_R})X\omega_l\|_2^2] \\
&= \mathbf{E}[\text{tr}(\omega_l^\top X^\top (I - P_{X_R})X\omega_l)] \\
&= \text{tr}(X^\top (I - P_{X_R})X \mathbf{E}[\omega_l \omega_l^\top]) \\
&= \text{tr}(X^\top (I - P_{X_R})X) \\
&= \|X - P_{X_R} X\|_F^2.
\end{aligned}$$

Concentration We now establish concentration for the estimator $\widehat{\delta}_R^2$ by invoking results in [25, 34]. Let $\omega \in \mathbb{R}^{dL}$ be the vector one obtains when stacking $\omega_1, \dots, \omega_L$ vertically. Let us also introduce $\Psi = X^\top(I - P_{X_R})X$ and let $\mathbf{\Psi} = \frac{1}{L}I_L \otimes \Psi$, where \otimes denotes the Kronecker product. Then $\widehat{\delta}_R^2$ can be re-written in the following way:

$$\begin{aligned} \omega^\top \mathbf{\Psi} \omega &= \omega^\top \frac{1}{L} \begin{bmatrix} \Psi & 0 & \dots & \dots & 0 \\ 0 & \Psi & \dots & \dots & 0 \\ \vdots & 0 & \ddots & \ddots & \vdots \\ \vdots & \vdots & \ddots & \ddots & 0 \\ 0 & 0 & \dots & 0 & \Psi \end{bmatrix} \omega \\ &= \frac{1}{L} \sum_{l=1}^L \omega_l^\top \Psi \omega_l \\ &= \frac{1}{L} \sum_{l=1}^L \omega_l^\top X^\top (I - P_{X_R}) X \omega_l \\ &= \frac{1}{L} \sum_{l=1}^L \|(I - P_{X_R}) X \omega_l\|_2^2 = \widehat{\delta}_R^2. \end{aligned}$$

In other words, $\widehat{\delta}_R^2$ can be expressed as a quadratic form in a Gaussian random vector of dimension dL and a positive definite matrix. We can thus use the following tail inequalities [25, 34]

$$\begin{aligned} \mathbf{P}(\omega^\top \mathbf{\Psi} \omega > \text{tr}(\mathbf{\Psi}) + 2\sqrt{t \text{tr}(\mathbf{\Psi}^2)} + 2\|\mathbf{\Psi}\|_2 t) &\leq \exp(-t), \quad t > 0. \\ \mathbf{P}(\omega^\top \mathbf{\Psi} \omega < \text{tr}(\mathbf{\Psi}) - 2\sqrt{t \text{tr}(\mathbf{\Psi}^2)}) &\leq \exp(-t), \quad t > 0. \end{aligned}$$

This can be re-written using the following relations:

$$\begin{aligned} \text{tr}(\mathbf{\Psi}) &= \text{tr}(\Psi) = \mathbf{E}[\widehat{\delta}_R^2] = \delta_R^2, \quad \sqrt{\text{tr}(\mathbf{\Psi}^2)} = \|\mathbf{\Psi}\|_F = \frac{\|\Psi\|_F}{\sqrt{L}} \leq \frac{\text{tr}(\Psi)}{\sqrt{L}}, \\ \|\mathbf{\Psi}\|_2 &\leq \|\mathbf{\Psi}\|_F \leq \text{tr}(\mathbf{\Psi}), \end{aligned}$$

$$\begin{aligned} \mathbf{P}\left(\widehat{\delta}_R^2 > \delta_R^2 \left(1 + \frac{2(t + \sqrt{t})}{\sqrt{L}}\right)\right) &\leq \exp(-t), \\ \mathbf{P}\left(\widehat{\delta}_R^2 < \delta_R^2 \left(1 - \frac{2\sqrt{t}}{\sqrt{L}}\right)\right) &\leq \exp(-t), \end{aligned}$$

Setting $t = 4$

$$\mathbf{P}\left(\left(1 - \frac{4}{\sqrt{L}}\right) \delta_R^2 \leq \widehat{\delta}_R^2 \leq \delta_R^2 \left(1 + \frac{12}{\sqrt{L}}\right)\right) \geq 1 - 2\exp(-4) \geq 0.96.$$

As a result, for any $0 < c < 1$ and any $C > 1$, as long as

$$L \geq \max \left\{ \frac{16}{(1-c)^2}, \frac{144}{(C-1)^2} \right\}$$

it holds that

$$\mathbf{P} \left(c\delta_R^2 \leq \widehat{\delta}_R^2 \leq C\delta_R^2 \right) \geq 1 - 2 \exp(-4) \geq 0.96.$$

Appendix E: Proof of Proposition 3

We start with a basic observation to be used several times. Let \widetilde{R} be an $(d \wedge n) \times k$ random matrix with $N(0, 1)$ entries. From the rotational invariance of the Gaussian distribution, we have that

$$V^\top R \stackrel{\mathcal{D}}{=} \widetilde{R} \quad (44)$$

where $\stackrel{\mathcal{D}}{=}$ denotes equality in distribution. Turning to property i), using that $\Sigma = \sqrt{n}I$ we have $X_R = \sqrt{n}UV^\top R \stackrel{\mathcal{D}}{=} \sqrt{n}U\widetilde{R}$ according to (44). We then compute

$$P_{X_R} = X_R(X_R^\top X_R)^{-1} X_R^\top \stackrel{\mathcal{D}}{=} U\widetilde{R}(\widetilde{R}^\top \widetilde{R})^{-1} \widetilde{R}^\top U = UP_{\widetilde{R}}U^\top,$$

where the inverse exists with probability one. Accordingly,

$$\begin{aligned} \mathbf{E} [\|Xw^* - P_{X_R}Xw^*\|_2^2/n] &= \mathbf{E}[\|UV^\top - UP_{\widetilde{R}}U^\top X\|_2^2] \\ &= \mathbf{E}[\|U(I - P_{\widetilde{R}})V^\top w^*\|_2^2] \\ &= \mathbf{E}[(\alpha^*)^\top (I - P_{\widetilde{R}})\alpha^*], \end{aligned} \quad (45)$$

With $\alpha_u^* = \alpha^*/\|\alpha^*\|_2$ and $U_{\widetilde{R}^\perp}$ as a matrix containing a set of orthonormal basis vectors of $\text{range}(\widetilde{R})^\perp$ as its columns, we have

$$\mathbf{E}[(\alpha^*)^\top (I - P_{\widetilde{R}})\alpha^*] = \|\alpha^*\|_2^2 \mathbf{E}[(\alpha_u^*)^\top U_{\widetilde{R}^\perp}U_{\widetilde{R}^\perp}^\top (\alpha_u^*)], \quad (46)$$

Since \widetilde{R} is Gaussian, $\text{range}(U_{\widetilde{R}^\perp}) \sim \text{Unif}(\mathbf{G}(d, d-k))^4$ [27]. By rotational invariance,

$$(\alpha_u^*)^\top U_{\widetilde{R}^\perp}U_{\widetilde{R}^\perp}^\top (\alpha_u^*) = \|U_{\widetilde{R}^\perp}^\top (\alpha_u^*)\|_2^2 \stackrel{\mathcal{D}}{=} \|E_{d-k}u\|_2^2 = 1 - k/d, \quad (47)$$

where $E_{d-k} \in \mathbb{R}^{(d-k) \times d}$ contains the first $d-k$ canonical basis vectors as its rows and $u \sim \text{Unif}(\mathbb{S}^{d-1})$. Combining (45), (46) and (47) concludes the derivation of the bias. The expression for $\mathbf{E}[\mathcal{E}(R)]$ given in the proposition is obtained by adding the variance term $\sigma^2 k/n$. Similarly, we evaluate $\mathbf{E}[\mathcal{E}(S)]$ by computing its bias. Expanding P_{X_S} , we get that

$$P_{X_S} = nUV^\top S(S^\top S)^{-1} S^\top VU^\top = nUV^\top S S^\top VU^\top, \quad (48)$$

⁴We recall that $\mathbf{G}(m, l)$ denotes the set of l -dimensional subspaces of \mathbb{R}^m .

where we have used that $VV^\top = I$ and $S^\top S = I$, where the latter property results from the fact that column sampling is done without replacement. It remains to evaluate $\mathbf{E}[SS^\top]$. The entries of SS^\top are given $(\langle S_{i,:}, S_{j,:} \rangle)_{1 \leq i, j \leq d}$, where $S_{l,:}$ denotes the l -th row of S , $l \in [d]$. We have

$$\mathbf{E}[\langle S_{i,:}, S_{j,:} \rangle] = \mathbf{E} \left[\sum_{l=1}^k S_{il} S_{jl} \right] = \sum_{l=1}^k \mathbf{P}(S_{il} = 1, S_{jl} = 1) = \begin{cases} \frac{k}{d} & \text{if } i = j, \\ 0 & \text{if } i \neq j. \end{cases}$$

Putting together the pieces, we obtain that

$$\mathbf{E}[\|(I - P_{X_S})Xw^*\|_2^2/n] = (1 - k/d)\|\alpha^*\|_2^2.$$

Turning to property ii), the arguments for $\mathcal{E}(R)$ parallel those used for i), with the difference that $\Sigma = \sqrt{d}I_n$. The subsequent steps are as in i) and are thus omitted. The situation is different for $\mathcal{E}(S)$ because the expansion (48) is no longer valid since $VV^\top \neq I$ as $n < d$. Consider a matrix X with dimensions $n < d$ whose matrix of right singular vectors $V \in \mathbb{R}^{d \times n}$ takes the form

$$\begin{pmatrix} I_n \\ \mathbf{O}_{n-d,d} \end{pmatrix},$$

where $\mathbf{O}_{n-d,d}$ denotes an $(n-d) \times d$ matrix of zeroes. Note that we still have $V^\top V = I_n$, hence this is a valid choice. With this specific form for V , we obtain that

$$\begin{aligned} P_{X_S} &= UV^\top S(S^\top VV^\top S)^{-1} S^\top VU^\top \\ &= US_{1:n,:}([S_{1:n,:}]^\top S_{1:n,:})^{-1} (S_{1:n,:})^\top U^\top = US_{1:n,:}(S_{1:n,:})^\top U^\top, \end{aligned}$$

where $S_{1:n,:}$ denotes the submatrix of S consisting of its first n rows. It follows that

$$\begin{aligned} \mathbf{E} \left[\frac{1}{n} \|(I - P_{X_S})Xw^*\|_2^2 \right] &= \frac{d}{n} (\alpha^*)^\top (I - \mathbf{E}[S_{1:n,:}(S_{1:n,:})^\top]) \alpha^* \\ &= \frac{d}{n} \|\alpha^*\|_2^2 \left(1 - \frac{k}{d} \right). \end{aligned}$$

Regarding property iii), observe that by the rotational invariance according to (44)

$$XR = U\Sigma V^\top R \stackrel{\mathcal{D}}{=} U\Sigma \tilde{R}.$$

Moreover, $XS = U\Sigma(V^\top S)$ and $\text{range}(\tilde{R}) \stackrel{\mathcal{D}}{=} \text{range}(V^\top S) \sim \text{Unif}(\mathbf{G}(d \wedge n, k))$ according to [27] since the entries of both R and X are i.i.d. zero-mean Gaussian, thus $\text{range}(XS) \stackrel{\mathcal{D}}{=} \text{range}(XR)$.

Appendix F: Proof of Proposition 4

For property i), observe that the map $A \mapsto \phi(A) := \|(I - A)Xw^*\|_2^2/n$ from $\mathbb{R}^{n \times n}$ to \mathbb{R}_+ is convex, hence $\phi\left(\frac{1}{B} \sum_{b=1}^B P_{X_{R_b}}\right) \leq \frac{1}{B} \sum_{b=1}^B \phi(P_{X_{R_b}})$. Taking

expectations then yields the assertion. Likewise, regarding property iii), we have

$$\begin{aligned} & \mathbf{E} \left[\left\| \frac{1}{B} \sum_{b=1}^B P_{XR_b} \xi \right\|_2^2 \middle| \{P_{XR_b}\}_{b=1}^B \right] \\ &= \mathbf{E} \left[\frac{1}{B^2} \sum_{b=1}^B \sum_{b'=1}^B \xi^\top P_{XR_b} P_{XR_{b'}} \xi \middle| \{P_{XR_b}\}_{b=1}^B \right] \end{aligned}$$

The claim then follows by noting that for any pair (b, b') , we have $\text{tr}(P_{XR_b} P_{XR_{b'}}) \leq \|P_{XR_b}\|_F \|P_{XR_{b'}}\|_F = k$.

We finally turn to properties ii) and iv). Consider the operator $\mathcal{P}_k = \mathbf{E}[P_{XR}]$. We first show that $\text{range}(\mathcal{P}_k) = \text{range}(X)$. The inclusion $\text{range}(\mathcal{P}_k) \subseteq \text{range}(X)$ holds trivially. For the other direction, since \mathcal{P}_k is symmetric positive definite, it suffices to show that $v^\top \mathcal{P}_k v > 0 \forall v \in \text{range}(X)$. Suppose by contradiction that there exists $v \in \text{range}(X)$ s.t.

$$v^\top \mathcal{P}_k v = v^\top \mathbf{E}[P_{XR}]v = \mathbf{E}_R[\|P_{XR}v\|_2^2] = 0,$$

which would imply that v is contained in the orthogonal complement of $\text{range}(XR)$ with probability one, i.e., $v \in \text{null}((XR)^\top) \Leftrightarrow R^\top X^\top v = 0$ with probability one. This contradicts the fact that the entries of R are from a distribution that is absolutely continuous with respect to the Lebesgue measure. In particular, the fact that $\text{range}(\mathcal{P}_k) = \text{range}(X)$ implies that \mathcal{P}_k has exactly $d \wedge n$ positive eigenvalues $\{\eta_j\}_{j=1}^{d \wedge n}$ contained in the simplex $\Delta(k) = \{z : \sum_{j=1}^{d \wedge n} z_j = k, 0 \leq z_j \leq 1\}$, noting that \mathcal{P}_k is an expectation over orthogonal projections onto k -dimensional subspaces. The last property of \mathcal{P}_k to be established in order to arrive at ii) and iv) is the fact that $U^\top \mathcal{P}_k U = \text{diag}(\eta_1, \dots, \eta_{d \wedge n})$, where U is the matrix of left singular vectors of X as its columns. We have

$$\begin{aligned} \mathcal{P}_k &= \mathbf{E}[P_{XR}] = \mathbf{E}[XR(R^\top X^\top XR)^{-1}R^\top X^\top] \\ &= \mathbf{E}[U\Sigma V^\top R(R^\top V\Sigma^2 V^\top R)^{-1}R^\top V\Sigma U^\top] \\ \Rightarrow U^\top \underbrace{\mathbf{E}[P_{XR}]}_{\mathcal{P}_k} U &= \mathbf{E}[\Sigma V^\top \tilde{R}(\tilde{R}^\top V\Sigma^2 V^\top \tilde{R})^{-1} \tilde{R}^\top V\Sigma] \\ &= \mathbf{E}[\Sigma \tilde{R}(\tilde{R}^\top \Sigma^2 \tilde{R})^{-1} \tilde{R}^\top \Sigma], \end{aligned} \tag{49}$$

where the last identity uses that $V^\top R \stackrel{\mathcal{D}}{=} \tilde{R}$ by rotational invariance (44). It remains to show that the matrix

$$\mathbf{E}[\Sigma \tilde{R}(\tilde{R}^\top \Sigma^2 \tilde{R})^{-1} \tilde{R}^\top \Sigma] \tag{50}$$

is diagonal. This has been shown in [37], noting that a matrix A is diagonal if and only if $DAD = A$ for all diagonal matrices D with diagonal elements ± 1 ; the claim then follows from the fact that $D\Sigma \tilde{R} \stackrel{\mathcal{D}}{=} \Sigma \tilde{R}$ and that $\tilde{R}^\top D\Sigma^2 D\tilde{R} = \tilde{R}^\top \Sigma^2 \tilde{R}$ for all such D . We note that the diagonal elements $\{\eta_j\}_{j=1}^{d \wedge n}$ of the diagonal matrix (50) depend only on the singular values $\{\sigma_j\}_{j=1}^{d \wedge n}$ but not on U or V as follows again from rotational invariance. Equipped with the property

$U^\top \mathcal{P}_k U = \text{diag}(\eta_1, \dots, \eta_{d \wedge n})$, we compute

$$\begin{aligned} \mathbf{E} [\|Xw^* - P_{X_R} Xw^*\|_2^2/n] &= \frac{1}{n} (w^*)^\top X^\top \mathbf{E}[I - P_{X_R}] Xw^* \\ &= \frac{1}{n} (w^*)^\top X^\top (I - \mathcal{P}_k) Xw^* \\ &= \frac{1}{n} \sum_{j=1}^{d \wedge n} \sigma_j^2 \{\alpha_j^*\}^2 (1 - \eta_j), \end{aligned}$$

after expanding X in its singular value composition and recalling that $\alpha^* = V^\top w^*$. In the same vein, we obtain that

$$\begin{aligned} &\|Xw^* - \mathcal{P}_k Xw^*\|_2^2/n \\ &= \frac{1}{n} (\|Xw^*\|_2^2 - 2(w^*)^\top X^\top \mathcal{P}_k Xw^* + (w^*)^\top X^\top \mathcal{P}_k^2 Xw^*) \\ &= \frac{1}{n} \left(\sum_{j=1}^{d \wedge n} \sigma_j^2 \{\alpha_j^*\}^2 - 2 \sum_{j=1}^{d \wedge n} \sigma_j^2 \{\alpha_j^*\}^2 \eta_j + \sum_{j=1}^{d \wedge n} \sigma_j^2 \{\alpha_j^*\}^2 \eta_j^2 \right) \\ &= \frac{1}{n} \sum_{j=1}^{d \wedge n} \sigma_j^2 \{\alpha_j^*\}^2 (1 - \eta_j)^2. \end{aligned}$$

From $U^\top \mathcal{P}_k^2 U = \text{diag}(\eta_1^2, \dots, \eta_{d \wedge n}^2)$, we immediately obtain the first identity in property iv). For the second identity, we let R' be an i.i.d. copy of R and note that

$$\begin{aligned} \text{tr}(\mathcal{P}_k^2) &= \text{tr}(\mathbf{E}[P_{X_R}] \mathbf{E}[P_{X_{R'}}]) = \text{tr}(\mathbf{E}[P_{X_R} P_{X_{R'}}]) \\ &= \mathbf{E}[\text{tr}(P_{X_R} P_{X_{R'}})] \\ &= \mathbf{E} \left[\sum_{\ell=1}^k \cos^2 \theta_\ell(\text{range}(XR), \text{range}(XR')) \right], \end{aligned}$$

where the last identity is obtained directly from the definition of canonical angles between subspaces [19].

Appendix G: Proof of Proposition 5

In the sequel, we use the shortcuts

$$L := \Delta + (1 - \Delta)d, \quad \kappa := L/\Delta = \frac{\Delta + (1 - \Delta)d}{\Delta}$$

for the maximum eigenvalue and condition number of $X^\top X/n$, respectively.

Lemma G.1. *Consider the coefficients $\{\eta_j\}_{j=1}^{d \wedge n}$ of Proposition 4 ii), and let $\nu_j = 1 - \eta_j$, $j \in [d]$. Then in the situation of Proposition 5*

$$\nu_1 = \mathbf{E} \left[\frac{1}{1 + \kappa Z G^{-1}} \right], \quad \eta_j = 1 - \nu_j = \frac{k - \eta_1}{d - 1}, \quad 2 \leq j \leq d,$$

for $1 \leq k < d$, where $Z \sim \chi^2(k)$ and $G \sim \chi^2(d - k)$ are Chi-square random variables with k and $d - k$ degrees of freedom, respectively.

Proof. From the proof of Proposition 4, Eq. (49), we know that η_1 can be computed as

$$\begin{aligned}\eta_1 &= \mathbf{E}[e_1^\top \Sigma R (R^\top \Sigma^2 R)^{-1} R^\top \Sigma e_1] \\ &= \mathbf{E} \left[L g_1^\top \left(L g_1 g_1^\top + \Delta \sum_{j=2}^d g_j g_j^\top \right)^{-1} g_1 \right],\end{aligned}$$

where $\{g_j\}_{j=1}^d$ are i.i.d. standard Gaussian random vectors of dimension k corresponding to the rows of R . Note that $W := \sum_{j=2}^d g_j g_j^\top$ follows a Wishart distribution of dimension $k \times k$ with $d-1$ degrees of freedom. With the help of the Sherman-Woodbury-Morrison formula, one computes

$$\begin{aligned}\eta_1 &= \mathbf{E} \left[\frac{L}{\Delta} g_1^\top W^{-1} g_1 - \frac{(L/\Delta)^2 g_1^\top W^{-1} g_1 g_1^\top W^{-1} g_1}{1 + (L/\Delta) g_1^\top W^{-1} g_1} \right] \\ &= \mathbf{E} \left[\kappa Q - \frac{\kappa^2 Q^2}{1 + \kappa Q} \right] = \mathbf{E} \left[\frac{\kappa Q}{1 + \kappa Q} \right], \quad Q := g_1^\top W^{-1} g_1 \\ \Rightarrow \nu_1 = 1 - \eta_1 &= \mathbf{E} \left[\frac{1}{1 + \kappa Q} \right] = \mathbf{E} \left[\frac{1}{1 + \kappa Z u^\top W^{-1} u} \right] = \mathbf{E} \left[\frac{1}{1 + \kappa Z e_1^\top W^{-1} e_1} \right] \\ &= \mathbf{E} \left[\frac{1}{1 + \kappa Z G^{-1}} \right],\end{aligned}$$

with $Z \sim \chi^2(k)$ and $G \sim \chi^2(d-k)$ as claimed. In the previous chain of identities, $u \sim \text{Unif}(\mathbb{S}^{k-1})$; the replacement of u by the first canonical basis vector e_1 can be done because of rotational invariance. Finally, in order to conclude the result, we have used that the diagonal entries of W^{-1} follow an inverse Chi-squared distribution with $d-k$ degrees of freedom ([20], Corollary 3.4.2.1). The result for $\{\nu_j\}_{j>1}$ is clear since $\sum_{j=1}^d \eta_j = k$ and $\sigma_2^2 = \dots = \sigma_d^2$ which implies $\eta_2 = \dots = \eta_d$.

CLS, upper bound:

Equipped with Lemma G.1, we have that for any $k \geq 3$, $\frac{G}{Z} \frac{k}{d-k}$ follows an F-distribution with $(d-k, k)$ degrees of freedom. Consequently,

$$\begin{aligned}\nu_1 &= \mathbf{E} \left[\frac{1}{1 + \kappa Z G^{-1}} \right] = \mathbf{E} \left[\frac{G}{G + \kappa Z} \right] \\ &\leq \frac{1}{\kappa} \mathbf{E} \left[\frac{G}{Z} \right] = \frac{1}{\kappa} \frac{d-k}{k} \frac{k}{k-2} = \frac{1}{\kappa} \frac{d-k}{k-2} \quad (k \geq 3).\end{aligned}$$

In the situation of Proposition 5, the bias term is hence upper bounded by

$$\frac{1}{n} \sum_{j=1}^d (\alpha_j^*)^2 \sigma_j^2 \underbrace{(1 - \eta_j)}_{=\nu_j} \leq \frac{\sigma_1^2}{n} \frac{1}{\kappa} \frac{d-k}{k-2} = \underbrace{\frac{\kappa \Delta}{\sigma_1^2/n}}_{=\sigma_1^2/n} \frac{1}{\kappa} \frac{d-k}{k-2} \leq \Delta \frac{2 \cdot d}{k} \quad (k \geq 4). \quad (51)$$

This yields $\mathbf{E}[\mathcal{E}(R)] \leq k\tau_n^2 + \Delta \frac{2d}{k}$. Minimizing this bound with respect to k , we obtain $\mathbf{E}[\mathcal{E}(R)] \leq 2\sqrt{2}\sqrt{d}\Delta^{1/2}$ as claimed.

CLS, lower bound:

We continue with the lower bound for CLS. Jensen's inequality combined with convexity of the map $x \mapsto \frac{1}{1+\kappa x}$ on $[0, \infty)$ yields

$$\begin{aligned} \nu_1 &= \mathbf{E} \left[\frac{1}{1 + \kappa G^{-1}Z} \right] \geq \frac{1}{1 + \kappa \mathbf{E}[G^{-1}Z]} = \frac{1}{1 + \kappa \cdot \frac{d-k}{d-k-2} \cdot \frac{k}{d-k}} \\ &= \frac{(d-k-2)}{(d-k-2) + \kappa k} \end{aligned} \quad (52)$$

for $k < d-2$ (note that if that is not the case, the claimed lower bound holds). Analogously to (51), the CLS bias is lower bounded by

$$\underbrace{\frac{(d-k-2)}{(d-k-2) + \kappa k}}_{\text{lower bound on } \nu_1} \cdot \underbrace{\kappa \Delta}_{=\sigma_1^2/n} \geq \begin{cases} \frac{1}{2}\kappa\Delta \geq \frac{1}{2} & \text{if } \kappa k < d-k-2, \\ \frac{d-k-2}{2k}\Delta & \text{otherwise.} \end{cases} \quad (53)$$

In the first case, the claimed lower bound holds. We hence turn to the second case, in which case $\mathbf{E}[\mathcal{E}(R)]$ is lower bounded by

$$\frac{d-k-2}{2k}\Delta + \tau_n^2 k.$$

If $k > (d-2)/2$, then we know that the risk is lower bounded by $\frac{d-2}{2}\tau_n^2$. Supposing conversely that $k < (d-2)/2$, we obtain the lower bound

$$\frac{d-2}{4k}\Delta + \tau_n^2 k.$$

Minimizing this expression with respect to k , we obtain the lower bound $\sqrt{d-2}\sqrt{\Delta}\tau_n$. This concludes the derivation of the lower bound for CLS.

\emptyset -CLS, upper bound:

According to Proposition 4 ii), we have

$$\mathcal{E}(\mathbf{E}[R_{d \times k} \widehat{w}_R]) = \frac{1}{n} \sum_{j=1}^d (\alpha_j^*)^2 \sigma_j^2 \underbrace{(1 - \eta_j)^2}_{=\nu_j^2} + \tau_n^2 \sum_{j=1}^d \eta_j^2 \quad (54)$$

With arguments similar as for CLS above, one shows that the bias term in (54) is upper bounded by ($k \geq 4$)

$$\frac{\Delta^2}{(1-\Delta)d} \left(\frac{d-k}{k-2} \right)^2 \leq \frac{\Delta^2}{(1-\Delta)} \frac{4d}{k^2} \quad (55)$$

In virtue of Lemma G.1, the variance term in (54) can be upper bounded as

$$\tau_n^2(\eta_1^2 + (d-1)\eta_2^2) \leq \tau_n^2 \left(1 + \frac{k^2}{(d-1)^2} \cdot (d-1) \right) \leq \tau_n^2 + \frac{2k^2}{d} \tau_n^2. \tag{56}$$

Combining (55) and (56), and minimizing the result w.r.t. k , the upper bound follows.

0-CLS, lower bound:

We start by lower bounding the variance term from (54). Similar to (56), we have

$$\tau_n^2 \left\{ \frac{(k - \eta_1)^2}{(d-1)^2} \cdot (d-1) + \eta_1^2 \right\} = \tau_n^2 \left\{ \frac{k^2 - 2\eta_1 k + d\eta_1^2}{d-1} \right\}.$$

The quadratic term in η_1 the numerator is minimized by $\eta_1 = k/d$. This yields the lower bound

$$\tau_n^2 \left\{ \frac{k^2 - 2\eta_1 k + d\eta_1^2}{d-1} \right\} \geq \tau_n^2 \left\{ \frac{k^2(1 - 1/d)}{d-1} \right\} = \tau_n^2 \frac{k^2}{d}. \tag{57}$$

For the bias term, we adopt the lower bound for ν_1 in (52). Accordingly,

$$\nu_1^2 \geq \left(\frac{d - k - 2}{(d - k - 2) + \kappa \cdot k} \right)^2$$

for $k < d - 2$, otherwise the claimed lower bound follows from the lower bound on the variance term (57). As a result, analogously to (53), the bias term is lower bounded by

$$\left(\frac{(d - k - 2)}{(d - k - 2) + \kappa k} \right)^2 \cdot \kappa \Delta \geq \begin{cases} \frac{1}{4} \kappa \Delta \geq \frac{1}{4} & \text{if } \kappa k < d - k - 2, \\ \frac{(d - k - 2)^2}{4\kappa k^2} \Delta & \text{otherwise.} \end{cases}$$

In the first case, the claimed lower bound follows. In the second case, we may assume that $k < d/2 - 2$, otherwise the lower bound on the variance (57) yields

$$\mathcal{E}(\mathbf{E}[R_{d \times k} \widehat{w}_R]) \geq \tau_n^2 \frac{(d/2 - 2)^2}{d} = \frac{1}{4} \frac{(d - 4)}{d} \tau_n^2 (d - 4).$$

On the other hand, if $k < d/2 - 2$,

$$\begin{aligned} \frac{(d - k - 2)^2}{4\kappa k^2} \Delta &\geq \frac{d^2}{16\{(1 - \Delta) \cdot d + \Delta\}k^2} \Delta^2 \\ &\geq \begin{cases} \frac{d^2}{64k^2} & \text{if } \Delta \geq d(1 - \Delta) \Leftrightarrow \Delta \geq \frac{d}{d+1}, \\ \frac{d}{32k^2} \frac{\Delta^2}{1 - \Delta} & \text{otherwise.} \end{cases} \end{aligned}$$

Combined with (57), this yields

$$\mathcal{E}(\mathbf{E}[R_{d \times k} \widehat{w}_R]) \geq \frac{1}{64} \frac{d^2}{k^2} + \tau_n^2 \frac{k^2}{d}$$

in the first case. Minimizing this lower bound with respect to k , we obtain

$$\mathcal{E}(\mathbf{E}[R_{d \times k} \hat{w}_R]) \geq \frac{1}{4} \tau_n \sqrt{d}.$$

On the other hand, in the second case

$$\mathcal{E}(\mathbf{E}[R_{d \times k} \hat{w}_R]) \geq \frac{d}{32k^2} \frac{\Delta^2}{1 - \Delta} + \tau_n^2 \frac{k^2}{d}.$$

Minimization with respect to k yields

$$\mathcal{E}(\mathbf{E}[R_{d \times k} \hat{w}_R]) \geq 2^{-3/2} \tau_n \sqrt{\Delta^2 / (1 - \Delta)},$$

which completes the demonstration of the lower bound for \emptyset -CLS.

Ridge regression, upper bound:

The excess risk of Ridge regression can be shown to have the form

$$\mathcal{E}(\hat{w}_\lambda) = \tau_n^2 \sum_{j=1}^d \frac{(\sigma_j^2/n)^2}{(\sigma_j^2/n + \lambda)^2} + \lambda^2 \sum_{j=1}^d (\alpha_j^*)^2 \frac{\sigma_j^2/n}{(\sigma_j^2/n + \lambda)^2}.$$

Specializing to the setting of Proposition 5, we obtain

$$\begin{aligned} \mathcal{E}(\hat{w}_\lambda) &= \tau_n^2 \frac{L^2}{(L + \lambda)^2} + (d - 1) \tau_n^2 \frac{\Delta^2}{(\Delta + \lambda)^2} + \frac{L}{(L + \lambda)^2} \lambda^2 \\ &\leq \tau_n^2 + (d - 1) \tau_n^2 \frac{\Delta^2}{\lambda^2} + \frac{1}{L} \lambda^2. \end{aligned}$$

Minimizing the resulting expression with respect to λ , we obtain the following optimal choice of λ :

$$\lambda^* = (\tau_n^2 \Delta^2 (d - 1) L)^{1/4}.$$

Substituting this choice back into the upper bound for the risk, we obtain

$$\mathcal{E}(\hat{w}_{\lambda^*}) \leq \tau_n^2 + 2\tau_n \sqrt{\frac{d-1}{L}} \Delta \leq \tau_n^2 + 2\tau_n \sqrt{\frac{\Delta^2}{1-\Delta}}.$$

Ridge regression, lower bound:

The lower bounded can be shown with arguments very similar to the upper bound; the derivation is hence omitted. \square

Acknowledgments

The author acknowledges extremely value input of two reviewers and an associate editor, which among numerous improvements and clarifications, has led to the addition of Proposition 5.

References

- [1] D. ACHLIOPTAS, *Database-friendly random projections: Johnson-Lindenstrauss with binary coins*, Journal of Computer and System Sciences, 66 (2003), pp. 671–687. [MR2005771](#)
- [2] D. AHFOCK, W. ASTLE, AND S. RICHARDSON, *Statistical Properties of Sketching Algorithms*. arXiv:[1706.03665](#), 2017.
- [3] N. AILON AND B. CHAZELLE, *Approximate nearest neighbors and the fast Johnson-Lindenstrauss transform*, in Proceedings of the Symposium on Theory of Computing (STOC), 2006, pp. 557–563. [MR2277181](#)
- [4] N. AILON AND E. LIBERTY, *Almost optimal unrestricted fast Johnson-Lindenstrauss transform*, in Symposium on Discrete Algorithms (SODA), 2011. [MR2857120](#)
- [5] Z. ALLEN-ZHU AND Y. LI, *Faster Principal Component Regression and Stable Matrix Chebyshev Approximation*, in International Conference on Machine Learning, 2017.
- [6] A. ARTEMIU AND B. LI, *On principal components and regression: a statistical explanation of a natural phenomenon*, Statistica Sinica, 19 (2009), pp. 1557–1565. [MR2589197](#)
- [7] R. BARANIUK, M. DAVENPORT, R. DEVORE, AND M. WAKIN, *A Simple Proof of the Restricted Isometry Property for Random Matrices*, Constructive Approximation, 28 (2006), pp. 253–263. [MR2453366](#)
- [8] L. BOTTOU, F. CURTIS, AND J. NOCEDAL, *Optimization Methods for Large-Scale Machine Learning*. arXiv:[1606.04838](#), 2018. [MR3797719](#)
- [9] L. BREIMAN, *Bagging predictors*, Machine Learning, 24 (1996), pp. 123–140.
- [10] P. BÜHLMANN AND S. VAN DE GEER, *Statistics for High-Dimensional Data*, Springer, 2011. [MR2807761](#)
- [11] F. BUNEA, Y. SHE, AND M. WEGKMAP, *Optimal selection of reduced rank estimators of high-dimensional matrices*, The Annals of Statistics, 39 (2011), pp. 1282–1309. [MR2816355](#)
- [12] K. BUZA, *Data Analysis, Machine Learning and Knowledge Discovery*, Springer, 2014, ch. Feedback Prediction for Blogs, pp. 145–152.
- [13] E. CANDÈS AND M. WAKIN, *An introduction to compressive sampling.*, IEEE Signal Processing Magazine, 25 (2008), pp. 21–30. [MR2275736](#)
- [14] T. CANNINGS AND R. SAMWORTH, *Random-projection ensemble classification*, Journal of the Royal Statistical Society Series B, 79 (2017), pp. 959–1035. [MR3689307](#)
- [15] P. DHILLON, D. FOSTER, S. KAKADE, AND L. UNGAR, *A risk comparison of ordinary least squares vs. ridge regression*, Journal of Machine Learning Research, 14 (2013), pp. 1505–1511. [MR3092989](#)
- [16] L. DICKER, D. FOSTER, AND D. HSU, *Kernel ridge vs. principal component regression: Minimax bounds and the qualification of regularization operators*, Electronic Journal of Statistics, 11 (2017), pp. 1022–1047. [MR3629418](#)

- [17] B. F. S. PEREVERZEV, AND L. ROSASCO, *On regularization algorithms in learning theory*, Journal of Complexity, 23, pp. 52–72. [MR2297015](#)
- [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.
- [19] G. GOLUB AND C. V. LOAN, *Matrix Computations*, Johns Hopkins University Press, 3rd ed., 1996. [MR1417720](#)
- [20] A. GUPTA AND D. NAGAR, *Matrix Variate Distributions*, CRC Press, 1999. [MR1738933](#)
- [21] N. HALKO, P. MARTINSSON, AND J. TROPP, *Finding Structure with Randomness: Probabilistic Algorithms for Constructing Approximate Matrix Decompositions*, SIAM Review, (2011), pp. 217–288. [MR2806637](#)
- [22] T. HASTIE, R. TISBHIRANI, AND M. WAINWRIGHT, *Statistical Learning with Sparsity*, CRC Press, 2015. [MR3616141](#)
- [23] D. HOMRIGHAUSEN AND D. MCDONALD, *Compressed and Penalized Linear Regression*. arXiv:[1705.08036](#), 2017.
- [24] H. HOTELLING, *Analysis of a complex of statistical variables into principal components*, Journal of Educational Psychology, 24 (1933), pp. 417–441.
- [25] D. HSU, S. KAKADE, AND T. ZHANG, *A tail inequality for quadratic forms of sub-Gaussian random vectors*, Electronic Communications in Probability, 52 (2012), pp. 1–6. [MR2994877](#)
- [26] P. INDYK AND A. NAOR, *Nearest-neighbor-preserving embeddings*, ACM Transactions on Algorithms, 3 (2007), p. 31. [MR2344022](#)
- [27] A. JAMES, *Normal Multivariate Analysis and the Orthogonal Group*, The Annals of Mathematical Statistics, 25 (1954), pp. 40–75. [MR0060779](#)
- [28] W. JOHNSON AND J. LINDENSTRAUSS, *Extensions of Lipschitz mappings into a Hilbert space*, Contemporary Mathematics, (1984), pp. 189–206. [MR0737400](#)
- [29] I. JOLIFFE, *A note on the use of principal components in regression*, Journal of the Royal Statistical Society Series C, 31 (1982), pp. 300–303.
- [30] A. KABAN, *New Bounds on Compressive Linear Least Squares Regression*, in Artificial Intelligence and Statistics (AISTATS), 2014, pp. 448–456.
- [31] S. KASIVISWANATHAN AND M. RUDELSON, *Compressed sparse linear regression*. arXiv:[1707.08902](#).
- [32] M. KENDALL, *A course in Multivariate Analysis*, Griffith, London, 1957. [MR0092297](#)
- [33] F. KRAHMER AND R. WARD, *New and improved Johnson-Lindenstrauss embeddings via the Restricted Isometry Property*, SIAM Journal on Mathematical Analysis, 43 (2011), pp. 1269–1281. [MR2821584](#)
- [34] B. LAURENT AND P. MASSART, *Adaptive estimation of a quadratic functional by model selection*, The Annals of Statistics, 28 (2000), pp. 1302–1338. [MR1805785](#)
- [35] Y. LU AND D. FOSTER, *Fast Ridge Regression with Randomized Principal Component Analysis and Gradient Descent*, in Uncertainty in Artificial Intelligence (UAI), 2014.

- [36] O. MAILLARD AND R. MUNOS, *Compressed least-squares regression*, in Advances in Neural Information Processing Systems (NIPS), 2009, pp. 1213–1221.
- [37] T. MARZETTA, G. TUCCI, AND S. SIMON, *A Random Matrix-Theoretic Approach to Handling Singular Covariance Estimates*, IEEE Transactions on Information Theory, 57 (2011), pp. 6256–6271. [MR2857971](#)
- [38] J. MATOUSEK, *On variants of the Johnson-Lindenstrauss lemma*, Random Structures and Algorithms, 33 (2008), pp. 142–156. [MR2436844](#)
- [39] M. PILANCI AND M. WAINWRIGHT, *Randomized Sketches of Convex Programs With Sharp Guarantees*, IEEE Transactions on Information Theory, 61 (2015), pp. 5096–5115. [MR3386504](#)
- [40] G. RASKUTTI AND M. MAHONEY, *A Statistical Perspective on Randomized Sketching for Ordinary Least-Squares*, Journal of Machine Learning Research, 17 (2016), pp. 1–31. [MR3595148](#)
- [41] T. SARLOS, *Improved approximation algorithms for large matrices via random projections*, in Foundations of Computer Science (FOCS), 2006, pp. 143–152.
- [42] R. SHAH AND N. MEINSHAUSEN, *On b-bit min-wise hashing for large-scale regression and classification with sparse data*, Journal of Machine Learning Research, 18 (2018), pp. 1–42. [MR3827066](#)
- [43] M. SLAWSKI, *Compressed least squares regression revisited*, in Artificial Intelligence and Statistics (AISTATS), 2017, pp. 1207–1215.
- [44] G.-A. THANEI, C. HEINZE, AND N. MEINSHAUSEN, *Big and Complex Data Analysis*, Springer, 2017, ch. Random Projections for Large-Scale Regression, pp. 51–68. [MR3644120](#)
- [45] J. TROPP, *Improved analysis of the subsampled randomized Hadamard transform*, Adaptive Data Analysis, (2011), pp. 115–126. [MR2835584](#)
- [46] S. VEMPALA, *The Random Projection Method*, American Mathematical Society, 2005. [MR2073630](#)
- [47] S. WANG, A. GITTENS, AND M. MAHONEY, *Sketched Ridge Regression: Optimization Perspective, Statistical Perspective, and Model Averaging*, Journal of Machine Learning Research, 18 (2017), pp. 1–50. [MR3827106](#)
- [48] M. WOODRUF, *Sketching as a Tool for Numerical Linear Algebra*, Foundations and Trends in Theoretical Computer Science, 10 (2014), pp. 1–157. [MR3285427](#)
- [49] S. ZHOU, J. LAFFERTY, AND L. WASSERMAN, *Compressed and privacy-sensitive regression*, IEEE Transactions on Information Theory, 55 (2009), pp. 846–866. [MR2597272](#)