

Faster least squares approximation

Petros Drineas · Michael W. Mahoney ·
S. Muthukrishnan · Tamás Sarlós

Received: 26 May 2009 / Revised: 6 August 2010 / Published online: 2 October 2010
© Springer-Verlag 2010

Abstract Least squares approximation is a technique to find an approximate solution to a system of linear equations that has no exact solution. In a typical setting, one lets n be the number of constraints and d be the number of variables, with $n \gg d$. Then, existing exact methods find a solution vector in $O(nd^2)$ time. We present two randomized algorithms that provide accurate relative-error approximations to the optimal value and the solution vector of a least squares approximation problem more rapidly than existing exact algorithms. Both of our algorithms preprocess the data with the Randomized Hadamard transform. One then uniformly randomly samples constraints and solves the smaller problem on those constraints, and the other performs a sparse random projection and solves the smaller problem on those projected coordinates. In both cases, solving the smaller problem provides relative-error approximations, and, if n is sufficiently larger than d , the approximate solution can be computed in $O(nd \ln d)$ time.

Mathematics Subject Classification (2000) 65F99

P. Drineas (✉)

Department of Computer Science, Rensselaer Polytechnic Institute, Troy, NY, USA
e-mail: drinep@cs.rpi.edu

M. W. Mahoney

Department of Mathematics, Stanford University, Stanford, CA, USA
e-mail: mmahoney@cs.stanford.edu

S. Muthukrishnan

Google, Inc., New York, NY, USA
e-mail: muthu@google.com

T. Sarlós

Yahoo! Research, Sunnyvale, CA, USA
e-mail: stamas@yahoo-inc.com

1 Introduction

In many applications in mathematics and statistical data analysis, it is of interest to find an approximate solution to a system of linear equations that has no exact solution. For example, let a matrix $A \in \mathbb{R}^{n \times d}$ and a vector $b \in \mathbb{R}^n$ be given. If $n \gg d$, there will not in general exist a vector $x \in \mathbb{R}^d$ such that $Ax = b$, and yet it is often of interest to find a vector x such that $Ax \approx b$ in some precise sense. The method of least squares, whose original formulation is often credited to Gauss and Legendre [26], accomplishes this by minimizing the sum of squares of the elements of the residual vector, i.e., by solving the optimization problem

$$\mathcal{Z} = \min_{x \in \mathbb{R}^d} \|Ax - b\|_2. \quad (1)$$

It is well-known that the minimum ℓ_2 -norm vector among those satisfying Eq. (1) is

$$x_{\text{opt}} = A^\dagger b, \quad (2)$$

where A^\dagger denotes the Moore-Penrose generalized inverse of the matrix A [6, 16]. This solution vector has a very natural statistical interpretation as providing an optimal estimator among all linear unbiased estimators, and it has a very natural geometric interpretation as providing an orthogonal projection of the vector b onto the span of the columns of the matrix A .

Recall that to minimize the quantity in Eq. (1), we can set the derivative of $\|Ax - b\|_2^2 = (Ax - b)^T(Ax - b)$ with respect to x equal to zero, from which it follows that the minimizing vector x_{opt} is a solution of the so-called normal equations

$$A^T Ax_{\text{opt}} = A^T b. \quad (3)$$

Geometrically, this means that the residual vector $b^\perp = b - Ax_{\text{opt}}$ is required to be orthogonal to the column space of A , i.e., $b^{\perp T} A = 0$. While solving the normal equations squares the condition number of the input matrix (and thus is not recommended in practice), direct methods (such as the QR decomposition [16]) solve the problem of Eq. (1) in $O(nd^2)$ time assuming that $n \geq d$. Finally, an alternative expression for the vector x_{opt} of Eq. (2) emerges by leveraging the singular value decomposition (SVD) of A . If $A = U_A \Sigma_A V_A^T$ denotes the SVD of A , then

$$x_{\text{opt}} = V_A \Sigma_A^{-1} U_A^T b.$$

1.1 Our results

In this paper, we describe two randomized algorithms that will provide accurate relative-error approximations to the minimal ℓ_2 -norm solution vector x_{opt} of Eq. (2) faster than existing exact algorithms for a large class of overconstrained least-squares problems. In particular, we will prove the following theorem.

Theorem 1 Suppose $A \in \mathbb{R}^{n \times d}$, $b \in \mathbb{R}^n$, and let $\epsilon \in (0, 1)$. Then, there exists a randomized algorithm that returns a vector $\tilde{x}_{\text{opt}} \in \mathbb{R}^d$ such that, with probability at least 0.8, the following two claims hold: first, \tilde{x}_{opt} satisfies

$$\|A\tilde{x}_{\text{opt}} - b\|_2 \leq (1 + \epsilon)\mathcal{Z}; \tag{4}$$

and, second, if $\kappa(A)$ is the condition number of A and if we assume that $\gamma \in [0, 1]$ is the fraction of the norm of b that lies in the column space of A (i.e., $\gamma = \|U_A U_A^T b\|_2 / \|b\|_2$, where U_A is an orthogonal basis for the column space of A), then \tilde{x}_{opt} satisfies

$$\|x_{\text{opt}} - \tilde{x}_{\text{opt}}\|_2 \leq \sqrt{\epsilon} \left(\kappa(A) \sqrt{\gamma^{-2} - 1} \right) \|x_{\text{opt}}\|_2. \tag{5}$$

Finally, the solution \tilde{x}_{opt} can be computed in $O(nd \ln d)$ time if n is sufficiently larger than d and less than e^d .

We will provide a precise statement of the running time for our two algorithms (including the ϵ -dependence) in Theorems 2 (Sect. 4) and 3 (Sect. 5), respectively. It is worth noting that the claims of Theorem 1 can be made to hold with probability $1 - \delta$, for any $\delta > 0$, by repeating the algorithm $\lceil \ln(1/\delta) / \ln(5) \rceil$ times. For example, one could run ten independent copies of the algorithm and keep the vector \tilde{x}_{opt} that minimizes the residual. This clearly does not increase the running time of the algorithm by more than a constant factor, while driving the failure probability down to (approximately) 10^{-7} . Also, we will assume that n is a power of two and that the rank of the $n \times d$ matrix A equals d . (We note that padding A and b with all-zero rows suffices to remove the first assumption.)

We now provide a brief overview of our main algorithms. Let the matrix product HD denote the $n \times n$ Randomized Hadamard transform (see also Sect. 2.4). Here the $n \times n$ matrix H denotes the (normalized) matrix of the Hadamard transform and the $n \times n$ diagonal matrix D is formed by setting its diagonal entries to $+1$ or -1 with equal probability in n independent trials. This transform has been used as one step in the development of a “fast” version of the Johnson-Lindenstrauss lemma [1, 18]. Our first algorithm is a random sampling algorithm. After premultiplying A and b by HD , this algorithm samples uniformly at random r constraints from the preprocessed problem. (See Eq. (22), as well as the remarks after Theorem 2 for the precise value of r .) Then, this algorithm solves the least squares problem on just those sampled constraints to obtain a vector $\tilde{x}_{\text{opt}} \in \mathbb{R}^d$ such that Theorem 1 is satisfied. Note that applying the randomized Hadamard transform to the matrix A and vector b only takes $O(nd \ln r)$ time. This follows since we will actually sample only r of the constraints from the Hadamard-preprocessed problem [2]. Then, exactly solving the $r \times d$ sampled least-squares problem will require only $O(rd^2)$ time. Assuming that ϵ is a constant and $n \leq e^d$, it follows that the running time of this algorithm is $O(nd \ln d)$ when $\frac{n}{\ln n} = \Omega(d^2)$.

In a similar manner, our second algorithm also initially premultiplies A and b by HD . This algorithm then multiplies the result by a $k \times n$ sparse projection matrix T ,

where $k = O(d/\epsilon)$. This matrix T is described in detail in Sect. 5.2. Its construction depends on a sparsity parameter, and it is identical to the “sparse projection” matrix in Matoušek’s version of the Ailon–Chazelle result [1, 18]. Finally, our second algorithm solves the least squares problem on just those k coordinates to obtain $\tilde{x}_{\text{opt}} \in \mathbb{R}^d$ such that the three claims of Theorem 1 are satisfied. Assuming that ϵ is a constant and $n \leq e^d$, it follows that the running time of this algorithm is $O(nd \ln d)$ when $n = \Omega(d^2)$.

It is worth noting that our second algorithm has a (marginally) less restrictive assumption on the connection between n and d . However, the first algorithm is simpler to implement and easier to describe. Clearly, an interesting open problem is to relax the above constraints on n for either of the proposed algorithms.

1.2 Related work

We should note several lines of related work.

- First, techniques such as the “method of averages” [10] preprocess the input into the form of Eq. (6) of Sect. 3 and can be used to obtain exact or approximate solutions to the least squares problem of Eq. (1) in $o(nd^2)$ time under strong statistical assumptions on A and b . To the best of our knowledge, however, the two algorithms we present and analyze are the first algorithms to provide nontrivial approximation guarantees for overconstrained least squares approximation problems in $o(nd^2)$ time, while making no assumptions at all on the input data.
- Second, Ibarra et al. [17] provide a reduction of the least squares approximation problem to the matrix multiplication problem. In particular, they show that $MM(d)O(n/d)$ time, where $MM(d)$ is the time needed to multiply two $d \times d$ matrices, is sufficient to solve this problem. All of the running times we report in this paper assume the use of standard matrix multiplication algorithms, since $o(d^3)$ matrix multiplication algorithms are almost never used in practice. Moreover, even with the current best value for the matrix multiplication exponent, $\omega \approx 2.376$ [9], our algorithms are still faster.
- Third, motivated by our preliminary results as reported in [12] and [24], both Rokhlin and Tygert [22] as well as Avron et al. [4, 5] have empirically evaluated numerical implementations of variants of one of the algorithms we introduce. We describe this in more detail below in Sect. 1.3.
- Fourth, very recently, Clarkson and Woodruff proved space lower bounds on related problems [8]; and Nguyen et al. [20] achieved a small improvement in the sampling complexity for related problems.

1.3 Empirical performance of our randomized algorithms

In prior work we have empirically evaluated randomized algorithms that rely on the ideas that we introduce in this paper in several large-scale data analysis tasks. Nevertheless, it is a fair question to ask whether our “random perspective” on linear algebra will work well in numerical implementations of interest in scientific computation.

We address this question here. Although we do *not* provide an empirical evaluation in this paper, in the wake of the original Technical Report version of this paper in 2007 [14], two groups of researchers have demonstrated that numerical implementations of variants of the algorithms we introduce in this paper can perform very well in practice.

- In 2008, Rokhlin and Tygert [22] describe a variant of our random projection algorithm, and they demonstrate that their algorithm runs in time

$$O(\ln(\ell) + \kappa \ln(1/\epsilon)nd + d^2\ell),$$

where ℓ is an “oversampling” parameter and κ is a condition number. Importantly (at least for very high-precision applications of this random sampling methodology), they reduce the dependence on ϵ from $1/\epsilon$ to $\ln(1/\epsilon)$. Moreover, by choosing $\ell \geq 4d^2$, they demonstrate that $\kappa \leq 3$. Although this bound is inferior to ours, they also consider a class of matrices for which choosing $\ell = 4d$ empirically produced a condition number $\kappa < 3$, which means that for this class of matrices their running time is

$$O(\ln(d) + \kappa \ln(1/\epsilon)nd + d^3).$$

Their numerical experiments on this class of matrices clearly indicate that their implementations of variants of our algorithms perform well for certain matrices as small as thousands of rows by hundreds of columns.

- In 2009, Avron et al. [4,5] introduced a randomized least-squares solver based directly on our algorithms. They call it Blendepik, and by considering a much broader class of matrices, they demonstrate that their solver “beats LAPACK’s direct dense least-squares solver by a large margin on essentially any dense tall matrix.” Beyond providing additional theoretical analysis, including backward error analysis bounds for our algorithm, they consider five (and numerically implement three) random projection strategies (i.e., discrete Fourier transform, discrete Cosine transform, discrete Hartely transform, Walsh-Hadamard transform, and a Kac random walk), and they evaluate their algorithms on a wide range of matrices of various sizes and various “localization” or “coherence” properties. Based on these results that empirically show the superior performance of randomized algorithms such as those we introduce and analyze in this paper on a wide class of matrices, they go so far as to “suggest that random-projection algorithms should be incorporated into future versions of LAPACK.”

1.4 Outline

After a brief review of relevant background in Sect. 2, Sect. 3 presents a structural result outlining conditions on preconditioner matrices that are sufficient for relative-error approximation. Then, we present our main sampling-based algorithm for approximating least squares approximation in Sect. 4 and in Sect. 5 we present a second projection-based algorithm for the same problem. Preliminary versions of parts of

this paper have appeared as conference proceedings in the 17th ACM-SIAM Symposium on Discrete Algorithms [12] and in the 47th IEEE Symposium on Foundations of Computer Science [24]; and the original Technical Report version of this journal paper has appeared on the arXiv [14]. In particular, the core of our analysis in this paper was introduced in [12], where an expensive-to-compute probability distribution was used to construct a relative-error approximation sampling algorithm for the least squares approximation problem. Then, after the development of the Fast Johnson-Lindenstrauss transform [1,24] proved that similar ideas could be used to improve the running time of randomized algorithms for the least squares approximation problem. In this paper, we have combined these ideas, treated the two algorithms in a manner to highlight their similarities and differences, and considerably simplified the analysis.

2 Preliminaries

2.1 Notation

We let $[n]$ denote the set $\{1, 2, \dots, n\}$; $\ln x$ denotes the natural logarithm of x and $\log_2 x$ denotes the base two logarithm of x . For any matrix $A \in \mathbb{R}^{n \times d}$, $A_{(i)}$, $i \in [n]$ denotes the i -th row of A as a row vector and $A^{(j)}$, $j \in [d]$ denotes the j -th column of A as a column vector. Also, given a random variable X , we let $\mathbf{E}[X]$ denote its expectation and $\mathbf{Var}[X]$ denote its variance.

We will make frequent use of matrix and vector norms. More specifically, we let

$$\|A\|_F^2 = \sum_{i=1}^n \sum_{j=1}^d A_{ij}^2$$

denote the square of the Frobenius norm of A , and we let

$$\|A\|_2 = \sup_{x \in \mathbb{R}^d, \|x\|_2=1} \|Ax\|_2$$

denote the spectral norm of A . For any vector $x \in \mathbb{R}^n$, its ℓ_2 -norm (or Euclidean norm) is equal to the square root of the sum of the squares of its elements, while its ℓ_∞ norm is defined as $\|x\|_\infty = \max_{i \in [n]} |x_i|$.

2.2 Linear Algebra background

We now review relevant definitions and facts from linear algebra; for more details, see [7,6,16,25]. Let the rank of $A \in \mathbb{R}^{n \times d}$ be $\rho \leq \min\{n, d\}$. The Singular Value Decomposition (SVD) of A is denoted by $A = U_A \Sigma_A V_A^T$, where $U_A \in \mathbb{R}^{n \times \rho}$ is the matrix of left singular vectors, $\Sigma_A \in \mathbb{R}^{\rho \times \rho}$ is the diagonal matrix of non-zero singular values, and $V_A \in \mathbb{R}^{d \times \rho}$ is the matrix of right singular vectors. Let $\sigma_i(A)$, $i \in [\rho]$, denote the i th non-zero singular value of A , and $\sigma_{\max}(A)$ and $\sigma_{\min}(A)$ denote the maximum and minimum singular value of A . The condition number of A is

$\kappa(A) = \sigma_{\max}(A)/\sigma_{\min}(A)$. The Moore-Penrose generalized inverse, or pseudoinverse, of A may be expressed in terms of the SVD as $A^\dagger = V_A \Sigma_A^{-1} U_A^T$ [6]. Finally, for any orthogonal matrix $U \in \mathbb{R}^{n \times \ell}$, let $U^\perp \in \mathbb{R}^{n \times (n-\ell)}$ denote an orthogonal matrix whose columns are an orthonormal basis spanning the subspace of \mathbb{R}^n that is orthogonal to the column space of U . In terms of U_A^\perp , the optimal value of the least squares residual of Eq. (1) is

$$\mathcal{Z} = \min_{x \in \mathbb{R}^d} \|Ax - b\|_2 = \left\| U_A^\perp U_A^{\perp T} b \right\|_2.$$

2.3 Markov’s inequality and the union bound

We will make frequent use of the following fundamental result from probability theory, known as Markov’s inequality [19]. Let X be a random variable assuming non-negative values with expectation $\mathbf{E}[X]$. Then, for all $t > 0$,

$$X \leq t \cdot \mathbf{E}[X]$$

with probability at least $1 - t^{-1}$.

We will also need the so-called union bound. Given a set of random events $\mathcal{E}_1, \mathcal{E}_2, \dots, \mathcal{E}_n$ holding with respective probabilities p_1, p_2, \dots, p_n , the probability that all events hold (i.e., the probability of the union of those events) is upper bounded by $\sum_{i=1}^n p_i$.

2.4 The randomized Hadamard transform

The randomized Hadamard transform was introduced in [1] as one step in the development of a fast version of the Johnson-Lindenstrauss lemma [1, 18]. Recall that the (non-normalized) $n \times n$ matrix of the Hadamard transform H_n may be defined recursively as follows:

$$H_n = \begin{bmatrix} H_{n/2} & H_{n/2} \\ H_{n/2} & -H_{n/2} \end{bmatrix}, \quad \text{with} \quad H_2 = \begin{bmatrix} +1 & +1 \\ +1 & -1 \end{bmatrix}.$$

The $n \times n$ normalized matrix of the Hadamard transform is equal to $\frac{1}{\sqrt{n}} H_n$; hereafter, we will denote this normalized matrix by H . Now consider a diagonal matrix $D \in \mathbb{R}^{n \times n}$ such that the diagonal entries D_{ii} are set to $+1$ with probability $1/2$ and to -1 with probability $1/2$ in n independent trials. The product HD is the Randomized Hadamard Transform and has two useful properties. First, when applied to a vector, it “spreads out” its energy, in the sense of providing a bound for its infinity norm (see Sect. 4.2). Second, computing the product HDx for any vector $x \in \mathbb{R}^n$ takes $O(n \log_2 n)$ time. Even better, if we only need to access, say, r elements in the transformed vector, then those r elements can be computed in $O(n \log_2 r)$ time [2]. We will expand on the latter observation in the proofs of Theorems 2 and 3.

3 Our algorithms as preconditioners

Both of our algorithms may be viewed as preconditioning the input matrix A and the target vector b with a carefully-constructed data-independent random matrix X . For our random sampling algorithm, we let $X = S^T H D$, where S is a matrix that represents the sampling operation and $H D$ is the Randomized Hadamard Transform, while for our random projection algorithm, we let $X = T H D$, where T is a random projection matrix. Thus, we replace the least squares approximation problem of Eq. (1) with the least squares approximation problem

$$\tilde{\mathcal{Z}} = \min_{x \in \mathbb{R}^d} \|X(Ax - b)\|_2. \tag{6}$$

We explicitly compute the solution to the above problem using a traditional deterministic algorithm [16], e.g., by computing the vector

$$\tilde{x}_{\text{opt}} = (XA)^\dagger Xb. \tag{7}$$

Alternatively, one could use standard iterative methods such as the the Conjugate Gradient Normal Residual method (CGNR, see [16] for details), which can produce an ϵ -approximation to the optimal solution of Eq. (6) in $O(\kappa(XA)rd \ln(1/\epsilon))$ time, where $\kappa(XA)$ is the condition number of XA and r is the number of rows of XA .

3.1 A structural result sufficient for relative-error approximation

In this subsection, we will state and prove a lemma that establishes sufficient conditions on any matrix X such that the solution vector \tilde{x}_{opt} to the least squares problem of Eq. (6) will satisfy relative-error bounds of the form (4) and (5). Recall that the SVD of A is $A = U_A \Sigma_A V_A^T$. In addition, for notational simplicity, we let $b^\perp = U_A^\perp U_A^{\perp T} b$ denote the part of the right hand side vector b lying outside of the column space of A .

The two conditions that we will require of the matrix X are:

$$\sigma_{\min}^2(XU_A) \geq 1/\sqrt{2}; \quad \text{and} \tag{8}$$

$$\left\| U_A^T X^T X b^\perp \right\|_2^2 \leq \epsilon \mathcal{Z}^2/2, \tag{9}$$

for some $\epsilon \in (0, 1)$. Several things should be noted about these conditions. First, although condition (9) depends on the right hand side vector b , Algorithms 1 and 2 will satisfy it without using any information from b . Second, although condition (8) only states that $\sigma_i^2(XU_A) \geq 1/\sqrt{2}$, for all $i \in [d]$, for both of our randomized algorithms we will show that $|1 - \sigma_i^2(XU_A)| \leq 1 - 2^{-1/2}$, for all $i \in [d]$. Thus, one should think of XU_A as an approximate isometry. Third, condition (9) simply states that $Xb^\perp = XU_A^\perp U_A^{\perp T} b$ remains approximately orthogonal to XU_A . Finally, note that the following lemma is a deterministic statement, since it makes no explicit reference to either of our randomized algorithms. Failure probabilities will enter later when we show that our randomized algorithms satisfy conditions (8) and (9).

Lemma 1 Consider the overconstrained least squares approximation problem of Eq. (1) and let the matrix $U_A \in \mathbb{R}^{n \times d}$ contain the top d left singular vectors of A . Assume that the matrix X satisfies conditions (8) and (9) above, for some $\epsilon \in (0, 1)$. Then, the solution vector \tilde{x}_{opt} to the least squares approximation problem (6) satisfies:

$$\|A\tilde{x}_{\text{opt}} - b\|_2 \leq (1 + \epsilon)\mathcal{Z}, \quad \text{and} \tag{10}$$

$$\|x_{\text{opt}} - \tilde{x}_{\text{opt}}\|_2 \leq \frac{1}{\sigma_{\min}(A)}\sqrt{\epsilon}\mathcal{Z}. \tag{11}$$

Proof Let us first rewrite the down-scaled regression problem induced by X as

$$\min_{x \in \mathbb{R}^d} \|Xb - XAx\|_2^2 = \min_{y \in \mathbb{R}^d} \|X(Ax_{\text{opt}} + b^\perp) - XA(x_{\text{opt}} + y)\|_2^2 \tag{12}$$

$$\begin{aligned} &= \min_{y \in \mathbb{R}^d} \|Xb^\perp - XAy\|_2^2 \\ &= \min_{z \in \mathbb{R}^d} \|Xb^\perp - XU_Az\|_2^2. \end{aligned} \tag{13}$$

(12) follows since $b = Ax_{\text{opt}} + b^\perp$ and (13) follows since the columns of the matrix A span the same subspace as the columns of U_A . Now, let $z_{\text{opt}} \in \mathbb{R}^d$ be such that $U_Az_{\text{opt}} = A(x_{\text{opt}} - \tilde{x}_{\text{opt}})$, and note that z_{opt} minimizes Eq. (13). The latter fact follows since

$$\begin{aligned} \|Xb^\perp - XA(x_{\text{opt}} - \tilde{x}_{\text{opt}})\|_2^2 &= \|Xb^\perp - X(b - b^\perp) + XA\tilde{x}_{\text{opt}}\|_2^2 \\ &= \|XA\tilde{x}_{\text{opt}} - Xb\|_2^2. \end{aligned}$$

Thus, by the normal Eq. (3), we have that

$$(XU_A)^T XU_Az_{\text{opt}} = (XU_A)^T Xb^\perp.$$

Taking the norm of both sides and observing that under condition (8) we have $\sigma_i((XU_A)^T XU_A) = \sigma_i^2(XU_A) \geq 1/\sqrt{2}$, for all i , it follows that

$$\|z_{\text{opt}}\|_2^2/2 \leq \|(XU_A)^T XU_Az_{\text{opt}}\|_2^2 = \|(XU_A)^T Xb^\perp\|_2^2. \tag{14}$$

Using condition (9) we observe that

$$\|z_{\text{opt}}\|_2^2 \leq \epsilon\mathcal{Z}^2. \tag{15}$$

To establish the first claim of the lemma, let us rewrite the norm of the residual vector as

$$\begin{aligned} \|b - A\tilde{x}_{\text{opt}}\|_2^2 &= \|b - Ax_{\text{opt}} + Ax_{\text{opt}} - A\tilde{x}_{\text{opt}}\|_2^2 \\ &= \|b - Ax_{\text{opt}}\|_2^2 + \|Ax_{\text{opt}} - A\tilde{x}_{\text{opt}}\|_2^2 \end{aligned} \tag{16}$$

$$= \mathcal{Z}^2 + \|U_A z_{\text{opt}}\|_2^2 \tag{17}$$

$$\leq \mathcal{Z}^2 + \epsilon \mathcal{Z}^2, \tag{18}$$

where (16) follows by Pythagoras, since $b - Ax_{\text{opt}} = b^\perp$, which is orthogonal to A , and consequently to $A(x_{\text{opt}} - \tilde{x}_{\text{opt}})$; (17) follows by the definition of z_{opt} and \mathcal{Z} ; and (18) follows by (15) and the orthogonality of U_A . The first claim of the lemma follows since $\sqrt{1 + \epsilon} \leq 1 + \epsilon$.

To establish the second claim of the lemma, recall that $A(x_{\text{opt}} - \tilde{x}_{\text{opt}}) = U_A z_{\text{opt}}$. If we take the norm of both sides of this expression, we have that

$$\|x_{\text{opt}} - \tilde{x}_{\text{opt}}\|_2^2 \leq \frac{\|U_A z_{\text{opt}}\|_2^2}{\sigma_{\min}^2(A)} \tag{19}$$

$$\leq \frac{\epsilon \mathcal{Z}^2}{\sigma_{\min}^2(A)}, \tag{20}$$

where (19) follows since $\sigma_{\min}(A)$ is the smallest singular value of A and since the rank of A is d ; and (20) follows by (15) and the orthogonality of U_A . Taking the square root, the second claim of the lemma follows. \square

If we make no assumption on b , then (11) from Lemma 1 may provide a weak bound in terms of $\|x_{\text{opt}}\|_2$. If, on the other hand, we make the additional assumption that a constant fraction of the norm of b lies in the subspace spanned by the columns of A , then (11) can be strengthened. Such an assumption is reasonable, since most least-squares problems are practically interesting if at least some part of b lies in the subspace spanned by the columns of A .

Lemma 2 *Using the notation of Lemma 1 and assuming that $\|U_A U_A^T b\|_2 \geq \gamma \|b\|_2$, for some fixed $\gamma \in (0, 1]$ it follows that*

$$\|x_{\text{opt}} - \tilde{x}_{\text{opt}}\|_2 \leq \sqrt{\epsilon} \left(\kappa(A) \sqrt{\gamma^{-2} - 1} \right) \|x_{\text{opt}}\|_2. \tag{21}$$

Proof Since $\|U_A U_A^T b\|_2 \geq \gamma \|b\|_2$, it follows that

$$\begin{aligned} \mathcal{Z}^2 &= \|b\|_2^2 - \|U_A U_A^T b\|_2^2 \\ &\leq (\gamma^{-2} - 1) \|U_A U_A^T b\|_2^2 \\ &\leq \sigma_{\max}^2(A) (\gamma^{-2} - 1) \|x_{\text{opt}}\|_2^2. \end{aligned}$$

This last inequality follows from $U_A U_A^T b = Ax_{\text{opt}}$, which implies

$$\|U_A U_A^T b\|_2 = \|Ax_{\text{opt}}\|_2 \leq \|A\|_2 \|x_{\text{opt}}\|_2 = \sigma_{\max}(A) \|x_{\text{opt}}\|_2.$$

By combining this with Eq. (11) of Lemma 1, the lemma follows. □

4 A sampling-based randomized algorithm

In this section, we present our randomized sampling algorithm for the least squares approximation problem of Eq. (1). We also state and prove an associated quality-of-approximation theorem.

4.1 The main algorithm and main theorem

Algorithm 1 takes as input a matrix $A \in \mathbb{R}^{n \times d}$, a vector $b \in \mathbb{R}^n$, and an error parameter $\epsilon \in (0, 1)$. This algorithm starts by preprocessing the matrix A and the vector b with the Randomized Hadamard Transform. It then constructs a smaller problem by sampling uniformly at random a small number of constraints from the preprocessed problem. Our main quality-of-approximation theorem (Theorem 2 below) states that with constant probability over the random choices made by the algorithm, the vector \tilde{x}_{opt} returned by this algorithm will satisfy the relative-error bounds of Eqs. (4) and (5) and will be computed quickly.

Input: $A \in \mathbb{R}^{n \times d}$, $b \in \mathbb{R}^n$, and an error parameter $\epsilon \in (0, 1)$.

Output: $\tilde{x}_{\text{opt}} \in \mathbb{R}^d$.

1. Let r assume the value of Eq. (22).
2. Let S be an empty matrix.
3. **For** $t = 1, \dots, r$ (i.i.d. trials with replacement) **select uniformly at random** an integer from $\{1, 2, \dots, n\}$.
 - **If** i is selected, **then** append the column vector $(\sqrt{n/r}) e_i$ to S , where $e_i \in \mathbb{R}^n$ is an all-zeros vector except for its i -th entry which is set to one.
4. Let $H \in \mathbb{R}^{n \times n}$ be the normalized Hadamard transform matrix.
5. Let $D \in \mathbb{R}^{n \times n}$ be a diagonal matrix with

$$D_{ii} = \begin{cases} +1, & \text{with probability } 1/2 \\ -1, & \text{with probability } 1/2 \end{cases}$$
6. Compute and return $\tilde{x}_{\text{opt}} = (S^T H D A)^\dagger S^T H D b$.

Algorithm 1: A fast random sampling algorithm for least squares approximation

In more detail, after preprocessing with the Randomized Hadamard Transform of Sect. 2.4, Algorithm 1 samples exactly r constraints from the preprocessed least

squares problem, rescales each sampled constraint by $\sqrt{n/r}$, and solves the least squares problem induced on just those sampled and rescaled constraints. (Note that the algorithm explicitly computes only those rows of HDA and only those elements of HDb that need to be accessed.) More formally, we will let $S \in \mathbb{R}^{n \times r}$ denote a sampling matrix specifying which of the n constraints are to be sampled and how they are to be rescaled. This matrix is initially empty and is constructed as described in Algorithm 1. Then, we can consider the problem

$$\tilde{Z} = \min_{x \in \mathbb{R}^d} \left\| S^T HDAx - S^T HDb \right\|_2,$$

which is just a least squares approximation problem involving the r constraints sampled from the matrix A after the preprocessing with the Randomized Hadamard Transform. The minimum ℓ_2 -norm vector $\tilde{x}_{\text{opt}} \in \mathbb{R}^d$ among those that achieve the minimum value \tilde{Z} in this problem is

$$\tilde{x}_{\text{opt}} = \left(S^T HDA \right)^\dagger S^T HDb,$$

which is the output of Algorithm 1.

Theorem 2 Suppose $A \in \mathbb{R}^{n \times d}$, $b \in \mathbb{R}^n$, and let $\epsilon \in (0, 1)$. Run Algorithm 1 with

$$r = \max \left\{ 48^2 d \ln(40nd) \ln \left(100^2 d \ln(40nd) \right), 40d \ln(40nd)/\epsilon \right\} \quad (22)$$

and return \tilde{x}_{opt} . Then, with probability at least 0.8, the following two claims hold: first, \tilde{x}_{opt} satisfies

$$\|A\tilde{x}_{\text{opt}} - b\|_2 \leq (1 + \epsilon)Z;$$

and, second, if we assume that $\|U_A U_A^T b\|_2 \geq \gamma \|b\|_2$ for some $\gamma \in (0, 1]$, then \tilde{x}_{opt} satisfies

$$\|x_{\text{opt}} - \tilde{x}_{\text{opt}}\|_2 \leq \sqrt{\epsilon} \left(\kappa(A) \sqrt{\gamma^{-2} - 1} \right) \|x_{\text{opt}}\|_2.$$

Finally,

$$n(d + 1) + 2n(d + 1) \log_2(r + 1) + O\left(r d^2\right)$$

time suffices to compute the solution \tilde{x}_{opt} .

Remark Assuming that $d \leq n \leq e^d$, and using $\max\{a_1, a_2\} \leq a_1 + a_2$, we get that

$$r = O\left(d(\ln d)(\ln n) + \frac{d \ln n}{\epsilon}\right).$$

Thus, the running time of Algorithm 1 becomes

$$O\left(nd \ln \frac{d}{\epsilon} + d^3(\ln d)(\ln n) + \frac{d^3 \ln n}{\epsilon}\right).$$

Assuming that $\frac{n}{\ln n} = \Omega(d^2)$, the above running time reduces to

$$O\left(nd \ln \frac{d}{\epsilon} + \frac{nd \ln d}{\epsilon}\right).$$

It is worth noting that improvements over the standard $O(nd^2)$ time could be derived with weaker assumptions on n and d . However, for the sake of clarity of presentation, we only focus on the above setting.

Remark The assumptions in our theorem have a natural geometric interpretation.¹ In particular, they imply that our approximation becomes worse as the angle between the vector b and the column space of A increases. To see this, let $\mathcal{Z} = \|Ax_{\text{opt}} - b\|_2$, and note that $\|b\|_2^2 = \|U_A U_A^T b\|_2^2 + \mathcal{Z}^2$. Hence the assumption $\|U_A U_k^T b\|_2 \geq \gamma \|b\|_2$ can be simply stated as

$$\mathcal{Z} \leq \sqrt{1 - \gamma^2} \|b\|_2.$$

The fraction $\mathcal{Z}/\|b\|_2$ is the sine of the angle between b and the column space of A ; see page 242 of [16]. Thus, $\sqrt{\gamma^{-2} - 1}$ is a bound on the tangent between b and the column space of A ; see page 244 of [16]. This means that the bound for $\|x_{\text{opt}} - \tilde{x}_{\text{opt}}\|_2$ is proportional to this tangent.

4.2 The effect of the Randomized Hadamard Transform

In this subsection, we state a lemma that quantifies the manner in which HD approximately “uniformizes” information in the left singular subspace of the matrix A . We state the lemma for a general $n \times d$ orthogonal matrix U such that $U^T U = I_d$, although we will be interested in the case when $n \gg d$ and U consists of the top d left singular vectors of the matrix A .

Lemma 3 *Let U be an $n \times d$ orthogonal matrix and let the product HD be the $n \times n$ Randomized Hadamard Transform of Sect. 2.4. Then, with probability at least 0.95,*

$$\|(HDU)_{(i)}\|_2^2 \leq \frac{2d \ln(40nd)}{n}, \quad \text{for all } i \in [n]. \tag{23}$$

Proof We follow the proof of Lemma 2.1 in [1]. In that lemma, the authors essentially prove that the Randomized Hadamard Transform HD “spreads out” input vectors.

¹ We would like to thank Ilse Ipsen for pointing out to us this geometric interpretation.

More specifically, since the columns of the matrix U (denoted by $U^{(j)}$ for all $j \in [d]$) are unit vectors, they prove that for fixed $j \in [d]$ and fixed $i \in [n]$,

$$\Pr \left[\left| \left(HDU^{(j)} \right)_i \right| \geq s \right] \leq 2e^{-s^2n/2}.$$

(Note that we consider d vectors in \mathbb{R}^n whereas [1] considered n vectors in \mathbb{R}^d and thus the roles of n and d are inverted in our proof.) Let $s = \sqrt{2n^{-1} \ln(40nd)}$ to get

$$\Pr \left[\left| \left(HDU^{(j)} \right)_i \right| \geq \sqrt{2n^{-1} \ln(40nd)} \right] \leq \frac{1}{20nd}.$$

From a standard union bound, this immediately implies that with probability at least $1 - 1/20$,

$$\left| \left(HDU^{(j)} \right)_i \right| \leq \sqrt{2n^{-1} \ln(40nd)} \tag{24}$$

holds for all $i \in [n]$ and $j \in [d]$. Using

$$\| (HDU)_{(i)} \|_2^2 = \sum_{j=1}^d \left(HDU^{(j)} \right)_i^2 \leq \frac{2d \ln(40nd)}{n} \tag{25}$$

for all $i \in [n]$, we conclude the proof of the lemma. □

4.3 Satisfying condition (8)

We now establish the following lemma which states that all the singular values of $S^T HDU_A$ are close to one. The proof of Lemma 4 depends on a bound for approximating the product of a matrix times its transpose by sampling (and rescaling) a small number of columns of the matrix. This bound appears as Theorem 4 in the Appendix and is an improvement over prior work of ours in [13].

Lemma 4 *Assume that Eq. (23) holds. If*

$$r \geq 48^2 d \ln(40nd) \ln \left(100^2 d \ln(40nd) \right) \tag{26}$$

then, with probability at least 0.95,

$$\left| 1 - \sigma_i^2 \left(S^T HDU_A \right) \right| \leq 1 - \frac{1}{\sqrt{2}},$$

holds for all $i \in [d]$.

Proof Note that for all $i \in [d]$

$$\begin{aligned} \left| 1 - \sigma_i^2 \left(S^T H D U_A \right) \right| &= \left| \sigma_i \left(U_A^T D H^T H D U_A \right) - \sigma_i \left(U_A^T D H^T S S^T H D U_A \right) \right| \\ &\leq \left\| U_A^T D H^T H D U_A - U_A^T D H^T S S^T H D U_A \right\|_2. \end{aligned} \tag{27}$$

In the above, we used the fact that $U_A^T D H^T H D U_A = I_d$. We now can view $U_A^T D S S^T H^T H D U_A$ as an approximation to the product of two matrices $U_A^T D H^T = (H D U_A)^T$ and $H D U_A$ by randomly sampling and rescaling columns of $(H D U_A)^T$. Thus, we can leverage Theorem 4 from the Appendix. More specifically, consider the matrix $(H D U_A)^T$. Obviously, since $H, D,$ and U_A are orthogonal matrices, $\|H D U_A\|_2 = 1$ and $\|H D U_A\|_F = \|U_A\|_F = \sqrt{d}$. Let $\beta = (2 \ln(40nd))^{-1}$; since we assumed that Eq. (23) holds, we note that the columns of $(H D U_A)^T$, which correspond to the rows of $H D U_A$, satisfy

$$\frac{1}{n} \geq \beta \frac{\|(H D U_A)_{(i)}\|_2^2}{\|H D U_A\|_F^2}, \quad \text{for all } i \in [n]. \tag{28}$$

Thus, applying Theorem 4 with β as above, $\epsilon = 1 - (1/\sqrt{2})$, and $\delta = 1/20$ implies that

$$\left\| U_A^T D H^T H U_A - U_A^T D H^T S S^T H D U_A \right\|_2 \leq 1 - \frac{1}{\sqrt{2}}$$

holds with probability at least $1 - 1/20 = 0.95$. For the above bound to hold, we need r to assume the value of Eq. (26). Finally, we note that since $\|H D U_A\|_F^2 = d \geq 1$, the assumption of Theorem 4 on the Frobenius norm of the input matrix is always satisfied. Combining the above with inequality (27) concludes the proof of the lemma. \square

4.4 Satisfying condition (9)

We next prove the following lemma, from which it will follow that condition (9) is satisfied by Algorithm 1. The proof of this lemma depends on bounds for randomized matrix multiplication algorithms that appeared in [11].

Lemma 5 *If Eq. (23) holds and $r \geq 40d \ln(40nd)/\epsilon$, then with probability at least 0.9,*

$$\left\| \left(S^T H D U_A \right)^T S^T H D b^\perp \right\|_2^2 \leq \epsilon \mathcal{Z}^2 / 2.$$

Proof Recall that $b^\perp = U_A^\perp U_A^{\perp T} b$ and that $\mathcal{Z} = \|b^\perp\|_2$. We start by noting that since $\|U_A^T D H^T H D b^\perp\|_2^2 = \|U_A^T b^\perp\|_2^2 = 0$ it follows that

$$\left\| \left(S^T H D U_A \right)^T S^T H D b^\perp \right\|_2^2 = \left\| U_A^T D H^T S S^T H D b^\perp - U_A^T D H^T H D b^\perp \right\|_2^2.$$

Thus, we can view $(S^T H D U_A)^T S^T H D b^\perp$ as approximating the product of two matrices $(H D U_A)^T$ and $H D b^\perp$ by randomly sampling columns from $(H D U_A)^T$ and rows/elements from $H D b^\perp$. Note that the sampling probabilities are uniform and do not depend on the norms of the columns of $(H D U_A)^T$ or the rows of $H D b^\perp$. However, we can still apply the results of Table 1 (second row) in page 150 of [11]. More specifically, since we condition on Eq. (23) holding, the rows of $H D U_A$ (which of course correspond to columns of $(H D U_A)^T$) satisfy

$$\frac{1}{n} \geq \beta \frac{\| (H D U_A)_{(i)} \|_2^2}{\| H D U_A \|_F^2}, \quad \text{for all } i \in [n], \tag{29}$$

for $\beta = (2 \ln(40nd))^{-1}$. Applying the result of Table 1 (second row) of [11] we get

$$\mathbf{E} \left[\left\| \left(S^T H D U_A \right)^T S^T H D b^\perp \right\|_2^2 \right] \leq \frac{1}{\beta r} \| H D U_A \|_F^2 \| H D b^\perp \|_2^2 = \frac{d \mathcal{Z}^2}{\beta r}.$$

In the above we used $\| H D U_A \|_F^2 = d$. Markov’s inequality now implies that with probability at least 0.9,

$$\left\| \left(S^T H D U_A \right)^T S^T H D b^\perp \right\|_2^2 \leq \frac{10d \mathcal{Z}^2}{\beta r}.$$

Setting $r \geq 20\beta^{-1}d/\epsilon$ and using the value of β specified above concludes the proof of the lemma. □

4.5 Completing the proof of Theorem 2

We now complete the proof of Theorem 2. First, let $\mathcal{E}_{(23)}$ denote the event that Eq. (23) holds; clearly, $\Pr [\mathcal{E}_{(23)}] \geq 0.95$. Second, let $\mathcal{E}_{4,5(23)}$ denote the event that both Lemmas 4 and 5 hold conditioned on $\mathcal{E}_{(23)}$ holding. Then,

$$\begin{aligned} \mathcal{E}_{4,5(23)} &= 1 - \overline{\mathcal{E}_{4,5(23)}} \\ &= 1 - \Pr [(\text{Lemma 4 does not hold} \mid \mathcal{E}_{(23)}) \text{ OR } (\text{Lemma 5 does not hold} \mid \mathcal{E}_{(23)})] \\ &\geq 1 - \Pr [\text{Lemma 4 does not hold} \mid \mathcal{E}_{(23)}] \\ &\quad - \Pr [\text{Lemma 5 does not hold} \mid \mathcal{E}_{(23)}] \\ &\geq 1 - 0.05 - 0.1 = 0.85. \end{aligned}$$

In the above, $\bar{\mathcal{E}}$ denotes the complement of event \mathcal{E} . In the first inequality we used the union bound and in the second inequality we leveraged the bounds for the failure probabilities of Lemmas 4 and 5 given that Eq. (23) holds. We now let \mathcal{E} denote the event that both Lemmas 4 and 5 hold, without any a priori conditioning on event $\mathcal{E}(23)$; we will bound $\Pr[\mathcal{E}]$ as follows:

$$\begin{aligned} \Pr[\mathcal{E}] &= \Pr[\mathcal{E}|\mathcal{E}(23)] \cdot \Pr[\mathcal{E}(23)] + \Pr[\mathcal{E}|\bar{\mathcal{E}}(23)] \cdot \Pr[\bar{\mathcal{E}}(23)] \\ &\geq \Pr[\mathcal{E}|\mathcal{E}(23)] \cdot \Pr[\mathcal{E}(23)] \\ &= \Pr[\mathcal{E}_{4,5}|\mathcal{E}(23)] \cdot \Pr[\mathcal{E}(23)] \\ &\geq 0.85 \cdot 0.95 \geq 0.8. \end{aligned}$$

In the first inequality we used the fact that all probabilities are positive. The above derivation immediately bounds the success probability of Theorem 2. Combining Lemmas 4 and 5 with the structural results of Lemma 1 and setting r as in Eq. (22) concludes the proof of the accuracy guarantees of Theorem 2.

We now discuss the running time of Algorithm 1. First of all, by the construction of S , the number of non-zero entries in S is r . In Step 6 we need to compute the products $S^T H D A$ and $S^T H D b$. Recall that A has d columns and thus the running time of computing both products is equal to the time needed to apply $S^T H D$ on $(d + 1)$ vectors. First, note that in order to apply D on $(d + 1)$ vectors in \mathbb{R}^n , $n(d + 1)$ operations suffice. In order to estimate how many operations are needed to apply $S^T H$ on $(d + 1)$ vectors, we use the results of Theorem 2.1 (see also Sect. 7) of Ailon and Liberty [2], which state that at most $2n(d + 1) \log_2(|S| + 1)$ operations are needed for this operation. Here $|S|$ denotes the number of non-zero elements in the matrix S , which is at most r . After this preprocessing, Algorithm 1 must compute the pseudoinverse of an $r \times d$ matrix, or, equivalently, solve a least-squares problem on r constraints and d variables. This operation can be performed in $O(rd^2)$ time since $r \geq d$. Thus, the entire algorithm runs in time

$$n(d + 1) + 2n(d + 1) \log_2(r + 1) + O(rd^2).$$

5 A projection-based randomized algorithm

In this section, we present a projection-based randomized algorithm for the least squares approximation problem of Eq. (1). We also state and prove an associated quality-of-approximation theorem.

5.1 The main algorithm and main theorem

Algorithm 2 takes as input a matrix $A \in \mathbb{R}^{n \times d}$, a vector $b \in \mathbb{R}^n$, and an error parameter $\epsilon \in (0, 1/2)$. This algorithm also starts by preprocessing the matrix A and right hand side vector b with the Randomized Hadamard Transform. It then constructs a smaller problem by performing a “sparse projection” on the preprocessed problem.

Our main quality-of-approximation theorem (Theorem 3 below) will state that with constant probability (over the random choices made by the algorithm) the vector \tilde{x}_{opt} returned by this algorithm will satisfy the relative-error bounds of Eqs. (4) and (5) and will be computed quickly.

Input: $A \in \mathbb{R}^{n \times d}$, $b \in \mathbb{R}^n$, and an error parameter $\epsilon \in (0, 1/2)$.

Output: $\tilde{x}_{\text{opt}} \in \mathbb{R}^d$.

1. Let q and k assume the values of Eqs. (30) and (31).
2. Let $T \in \mathbb{R}^{k \times n}$ be a random matrix with

$$T_{ij} = \begin{cases} +\sqrt{\frac{1}{kq}}, & \text{with probability } q/2 \\ -\sqrt{\frac{1}{kq}}, & \text{with probability } q/2 \\ 0, & \text{with probability } 1 - q, \end{cases}$$
 for all i, j independently.
3. Let $H \in \mathbb{R}^{n \times n}$ be the normalized Hadamard transform matrix.
4. Let $D \in \mathbb{R}^{n \times n}$ be a diagonal matrix with

$$D_{ii} = \begin{cases} +1, & \text{with probability } 1/2 \\ -1, & \text{with probability } 1/2 \end{cases}$$
5. Compute and return $\tilde{x}_{\text{opt}} = (THDA)^\dagger THDb$.

Algorithm 2: A fast random projection algorithm for least squares approximation

In more detail, Algorithm 2 begins by preprocessing the matrix A and right hand vector b with the Randomized Hadamard Transform HD of Sect. 2.4. This algorithm explicitly computes only those rows of HDA and those elements of HDb that need to be accessed to perform the sparse projection. After this initial preprocessing, Algorithm 2 will perform a “sparse projection” by multiplying HDA and HDb by the sparse matrix T (described in more detail in Sect. 5.2). Then, we can consider the problem

$$\tilde{\mathcal{Z}} = \min_{x \in \mathbb{R}^d} \left\| THDAx - THDb \right\|_2,$$

which is just a least squares approximation problem involving the matrix $THDA \in \mathbb{R}^{k \times d}$ and the vector $THDb \in \mathbb{R}^k$. The minimum ℓ_2 -norm vector $\tilde{x}_{\text{opt}} \in \mathbb{R}^d$ among those that achieve the minimum value $\tilde{\mathcal{Z}}$ in this problem is

$$\tilde{x}_{\text{opt}} = (THDA)^\dagger THDb,$$

which is the output of Algorithm 2.

Theorem 3 Suppose $A \in \mathbb{R}^{n \times d}$, $b \in \mathbb{R}^n$, and let $\epsilon \in (0, 1/2)$. Run Algorithm 2 with²

$$q \geq \frac{C_q d \ln(40nd)}{n} (2 \ln n + 16d + 16) \tag{30}$$

$$k \geq \max \left\{ C_k \left(118^2 d + 98^2 \right), \frac{60d}{\epsilon} \right\} \tag{31}$$

and return \tilde{x}_{opt} . Then, with probability at least 0.8, the following two claims hold: first, \tilde{x}_{opt} satisfies

$$\|A\tilde{x}_{\text{opt}} - b\|_2 \leq (1 + \epsilon)\mathcal{Z};$$

and, second, if we assume that $\|U_A U_A^T b\|_2 \geq \gamma \|b\|_2$ for some $\gamma \in (0, 1]$ then \tilde{x}_{opt} satisfies

$$\|x_{\text{opt}} - \tilde{x}_{\text{opt}}\|_2 \leq \sqrt{\epsilon} \left(\kappa(A) \sqrt{\gamma^{-2} - 1} \right) \|x_{\text{opt}}\|_2.$$

Finally, the expected running time of the algorithm is (at most)

$$n(d + 1) + 2n(d + 1) \log_2(nkq + 1) + O(kd^2).$$

Remark Assuming that $d \leq n \leq e^d$ we get that

$$q = O\left(\frac{d^2 \ln n}{n}\right) \quad \text{and} \quad k = O\left(\frac{d}{\epsilon}\right).$$

Thus, the expected running time of Algorithm 2 becomes

$$O\left(nd \ln \frac{d}{\epsilon} + \frac{d^3}{\epsilon}\right).$$

Finally, assuming $n = \Omega(d^2)$, the above running time reduces to

$$O\left(nd \ln \frac{d}{\epsilon} + \frac{nd}{\epsilon}\right).$$

It is worth noting that improvements over the standard $O(nd^2)$ time could be derived with weaker assumptions on n and d .

² C_q and C_k are the unspecified constants of Lemma 6.

5.2 Sparse projection matrices

In this subsection, we state a lemma about the action of a sparse random matrix operating on a vector. Recall that given any set of n points in Euclidean space, the Johnson–Lindenstrauss lemma states that those points can be mapped via a linear function to $k = O(\epsilon^{-2} \ln n)$ dimensions such that the distances between all pairs of points are preserved to within a multiplicative factor of $1 \pm \epsilon$; see [18] and references therein for details.

Formally, let $\epsilon \in (0, 1/2)$ be an error parameter, $\delta \in (0, 1)$ be a failure probability, and $\alpha \in [1/\sqrt{n}, 1]$ be a “uniformity” parameter. In addition, let q be a “sparsity” parameter defining the expected number of nonzero elements per row, and let k be the number of rows in our matrix. Then, define the $k \times n$ random matrix T as in Algorithm 2. Matoušek proved the following lemma, as the key step in his version of the Ailon–Chazelle result [1, 18].

Lemma 6 *Let T be the sparse random matrix of Algorithm 2, where $q = C_q \alpha^2 \ln(\frac{n}{\epsilon \delta})$ for some sufficiently large constant C_q (but still such that $q \leq 1$), and $k = C_k \epsilon^{-2} \ln(\frac{4}{\delta})$ for some sufficiently large constant C_k (but such that k is integral). Then for every vector $x \in \mathbb{R}^n$ such that $\|x\|_\infty / \|x\|_2 \leq \alpha$, we have that with probability at least $1 - \delta$*

$$| \|Tx\|_2 - \|x\|_2 | \leq \epsilon \|x\|_2 .$$

Remark In order to achieve sufficient concentration for all vectors $x \in \mathbb{R}^n$, the linear mapping defining the Johnson–Lindenstrauss transform is typically “dense,” in the sense that almost all the elements in each of the k rows of the matrix defining the mapping are nonzero. In this case, implementing the mapping on d vectors (in, e.g., a matrix A) via a matrix multiplication requires $O(ndk)$ time. This is not faster than the $O(nd^2)$ time required to compute an exact solution to the problem of Eq. (1) if k is at least d . The Ailon–Chazelle result [1, 18] states that the mapping can be “sparse,” in the sense that only a few of the elements in each of the k rows need to be nonzero, provided that the vector x is “well-spread,” in the sense that $\|x\|_\infty / \|x\|_2$ is close to $1/\sqrt{n}$. This is exactly what the preprocessing with the randomized Hadamard transform guarantees.

5.3 Proof of Theorem 3

In this subsection, we provide a proof of Theorem 3. Recall that by the results of Sect. 3.1, in order to prove Theorem 3, we must show that the matrix THD constructed by Algorithm 2 satisfies conditions (8) and (9) with probability at least .5. The next two subsections focus on proving that these conditions hold; the last subsection discusses the running time of Algorithm 2.

5.3.1 Satisfying condition (8)

In order to prove that all the singular values of $THDU_A$ are close to one, we start with the following lemma which provides a means to bound the spectral norm of a matrix. This lemma is an instantiation of lemmas that appeared in [3, 15].

Lemma 7 *Let M be a $d \times d$ symmetric matrix and define the grid*

$$\Omega = \left\{ x : x \in \frac{1}{2\sqrt{d}}\mathbb{Z}^d, \|x\|_2 \leq 1 \right\}. \tag{32}$$

In words, Ω includes all d -dimensional vectors x whose coordinates are integer multiples of $(2\sqrt{d})^{-1}$ and satisfy $\|x\|_2 \leq 1$. Then, the cardinality of Ω is at most e^{4d} . In addition, if for every $x, y \in \Omega$ we have that $|x^T My| \leq \epsilon'$, then for every unit vector x we have that $|x^T Mx| \leq 4\epsilon'$.

We next establish Lemma 8, which states that all the singular values of $THDU_A$ are close to one with constant probability. The proof of this lemma depends on the bound provided by Lemma 7 and it immediately shows that condition (8) is satisfied by Algorithm 2.

Lemma 8 *Assume that Lemma 3 holds. If q and k satisfy:*

$$q \geq \frac{C_q d \ln(40nd)}{n} (2 \ln n + 16d + 16) \tag{33}$$

$$k \geq C_k (118^2 d + 98^2), \tag{34}$$

then, with probability at least 0.95,

$$\left| 1 - \sigma_i^2(THDU_A) \right| \leq 1 - (1/\sqrt{2})$$

holds for all $i \in [d]$. Here C_q and C_k are the unspecified constants of Lemma 6.

Proof Define the symmetric matrix $M = U_A^T D H^T T^T T H D U_A - I_d \in \mathbb{R}^{d \times d}$, recall that $I_d = U_A^T D H^T H D U_A$, and note that

$$\left| 1 - \sigma_i^2(THDU_A) \right| \leq \|M\|_2 \tag{35}$$

holds for all $i \in [d]$. Consider the grid Ω of Eq. (32) and note that there are no more than e^{8d} pairs $(x, y) \in \Omega \times \Omega$, since $|\Omega| \leq e^{4d}$ by Lemma 7. Since $\|M\|_2 = \sup_{\|x\|_2=1} |x^T Mx|$, in order to show that $\|M\|_2 \leq 1 - 2^{-1/2}$, it suffices by Lemma 7 to show that $|x^T My| \leq (1 - 2^{-1/2})/4$, for all $x, y \in \Omega$. To do so, first, consider a single x, y pair.

Let

$$\begin{aligned} \Delta_1 &= \|THDU_A(x + y)\|_2^2 - \|H DU_A(x + y)\|_2^2 \\ \Delta_2 &= \|THDU_Ax\|_2^2 - \|H DU_Ax\|_2^2 \\ \Delta_3 &= \|THDU_Ay\|_2^2 - \|H DU_Ay\|_2^2, \end{aligned}$$

and note that

$$\Delta_1 = (x + y)^T U_A^T D H^T T^T T H D U_A (x + y) - (x + y)^T (x + y).$$

By multiplying out the right hand side of the above equation and rearranging terms, it follows that

$$x^T M y = x^T U_A^T D H^T T^T T H D U_A y - x^T y = \frac{1}{2} (\Delta_1 + \Delta_2 + \Delta_3). \tag{36}$$

In order to use Lemma 6 to bound the quantities $\Delta_1, \Delta_2,$ and $\Delta_3,$ we need a bound on the uniformity ratio $\|H D U_A x\|_\infty / \|H D U_A x\|_2.$ To do so, note that

$$\begin{aligned} \frac{\|H D U_A x\|_\infty}{\|H D U_A x\|_2} &= \frac{\max_{i \in [n]} |(H D U_A)_{(i)} x|}{\|H D U_A x\|_2} \leq \frac{\max_{i \in [n]} \|(H D U_A)_{(i)}\|_2 \|x\|_2}{\|x\|_2} \\ &\leq \sqrt{\frac{2d \ln(40nd)}{n}}. \end{aligned}$$

The above inequalities follow by $\|H D U_A x\|_2 = \|x\|_2$ and Lemma 3. This holds for both our chosen points x and y and in fact for all $x \in \Omega.$ Let $\epsilon_1 = 3/125$ and let $\delta = 1/(60e^{8d})$ (these choices will be explained shortly). Then, it follows from Lemma 6 that by setting $\alpha = \sqrt{2d \ln(40nd)}/n$ and our choices for k and $q,$ each of the following three statements holds with probability at least $1 - \delta:$

$$\begin{aligned} |\Delta_1| &\leq \epsilon_1 \|H D U_A(x + y)\|_2^2 = \epsilon_1 \|x + y\|_2^2 \leq 4\epsilon_1 \\ |\Delta_2| &\leq \epsilon_1 \|H D U_A x\|_2^2 = \epsilon_1 \|x\|_2^2 \leq \epsilon_1 \\ |\Delta_3| &\leq \epsilon_1 \|H D U_A y\|_2^2 = \epsilon_1 \|y\|_2^2 \leq \epsilon_1. \end{aligned}$$

Thus, combining the above with Eq. (36), for this single pair of vectors $(x, y) \in \Omega \times \Omega,$

$$\left| x^T M y \right| = \left| x^T U_A^T D H^T T^T T H D U_A y - x^T y \right| \leq \frac{1}{2} 6\epsilon_1 = 3\epsilon_1 \tag{37}$$

holds with probability at least $1 - 3\delta.$ Next, recall that there are no more than e^{8d} pairs of vectors $(x, y) \in \Omega \times \Omega,$ and we need Eq. (37) to hold for all of them. Since we set $\delta = 1/(60e^{8d})$ then it follows by a union bound that Eq. (37) holds for all pairs of vectors $(x, y) \in \Omega \times \Omega$ with probability at least 0.95. Additionally, let us set $\epsilon_1 = 3/125,$ which implies that $|x^T M y| \leq 9/125 \leq (1 - 2^{-1/2})/4$ thus concluding the proof of the lemma.

Finally, we discuss the values of the parameters q and k . Since $\delta = 1/(60e^{8d})$, $\epsilon_1 = 3/125$, and $\alpha = \sqrt{2d \ln(40nd)}/n$, the appropriate values for q and k emerge after elementary manipulations from Lemma 6. \square

5.3.2 Satisfying condition (9)

In order to prove that condition (9) is satisfied, we start with Lemma 9. In words, this lemma states that given vectors x and y we can use the random sparse projection matrix T to approximate $|x^T y|$ by $|x^T T^T T y|$, provided that $\|x\|_\infty$ (or $\|y\|_\infty$, but not necessarily both) is bounded. The proof of this lemma is elementary but tedious and is deferred to Sect. 6.2 of the Appendix.

Lemma 9 *Let x, y be vectors in \mathbb{R}^n such that $\|x\|_\infty \leq \alpha$. Let T be the $k \times n$ sparse projection matrix of Sect. 5.2, with sparsity parameter q . If $q \geq \alpha^2$, then*

$$\mathbf{E} \left[\left| x^T T^T T y - x^T y \right|^2 \right] \leq \frac{2}{k} \|x\|_2^2 \|y\|_2^2 + \frac{1}{k} \|y\|_2^2.$$

The following lemma proves that condition (9) is satisfied by Algorithm 2. The proof of this lemma depends on the bound provided by Lemma 9. Recall that $b^\perp = U_A^\perp U_A^{\perp T} b$ and thus $\|b^\perp\|_2 = \|U_A^\perp U_A^{\perp T} b\|_2 = \mathcal{Z}$.

Lemma 10 *Assume that Eq. (23) holds. If $k \geq 60d/\epsilon$ and $q \geq 2n^{-1} \ln(40nd)$, then, with probability at least 0.9,*

$$\left\| (T H D U_A)^T T H D b^\perp \right\|_2^2 \leq \epsilon \mathcal{Z}^2 / 2.$$

Proof We first note that since $U_A^T b^\perp = 0$, it follows that $U_A^{(j)T} b^\perp = U_A^{(j)T} D H^T H D b^\perp = 0$, for all $j \in [d]$. Thus, we have that

$$\left\| U_A^T D H^T T^T T H D b^\perp \right\|_2^2 = \sum_{j=1}^d \left(\left((H D U_A)^{(j)} \right)^T T^T T H D b^\perp - U_A^{(j)T} D H^T H D b^\perp \right)^2. \tag{38}$$

We now bound the expectation of the left hand side of Eq. (38) by using Lemma 9 to bound each term on the right hand side of Eq. (38). Using Eq. (24) of Lemma 3 we get that

$$\left\| (H D U_A)^{(j)} \right\|_\infty \leq \sqrt{2n^{-1} \ln(40nd)}$$

holds for all $j \in [d]$. By our choice of the sparsity parameter q the conditions of Lemma 9 are satisfied. It follows from Lemma 9 that

$$\begin{aligned} & \mathbf{E} \left[\left\| U_A^T D H^T T^T T H D b^\perp \right\|_2^2 \right] \\ &= \sum_{j=1}^d \mathbf{E} \left[\left(\left((H D U_A)^{(j)} \right)^T T^T T H D b^\perp - U_A^{(j)T} D H^T H D b^\perp \right)^2 \right] \\ &\leq \sum_{j=1}^d \left(\frac{2}{k} \left\| (H D U_A)^{(j)} \right\|_2^2 \left\| H D b^\perp \right\|_2^2 + \frac{1}{k} \left\| H D b^\perp \right\|_2^2 \right) \\ &= \frac{3d}{k} \left\| H D b^\perp \right\|_2^2 = \frac{3d}{k} \mathcal{Z}^2. \end{aligned}$$

The last line follows since $\left\| (H D U_A)^{(j)} \right\|_2 = 1$, for all $j \in [d]$. Using Markov’s inequality, we get that with probability at least 0.9,

$$\left\| U_A^T D H^T T^T T H D b^\perp \right\|_2^2 \leq \frac{30d}{k} \mathcal{Z}^2.$$

The proof of the lemma is concluded by using the assumed value of k . □

5.3.3 Proving Theorem 3

By our choices of k and q as in Eqs. (31) and (30), it follows that both conditions (8) and (9) are satisfied. Combining with Lemma 1 we immediately get the accuracy guarantees of Theorem 3. The failure probability of Algorithm 2 can be bounded using an argument similar to the one used in Sect. 4.5.

In order to complete the proof we discuss the running time of Algorithm 2. First of all, by the construction of T , the expected number of non-zero entries in T is kqn . In Step 5 we need to compute the products $T H D A$ and $T H D b$. Recall that A has d columns and thus the running time of computing both products is equal to the time needed to apply $T H D$ on $(d + 1)$ vectors. First, note that in order to apply D on $(d + 1)$ vectors in \mathbb{R}^n , $n(d + 1)$ operations suffice. In order to estimate how many operations are needed to apply $T H$ on $(d + 1)$ vectors, we use the results of Theorem 2.1 (see also Sect. 7) of Ailon and Liberty [2], which state that at most $2n(d + 1) \log_2 (|T| + 1)$ operations are needed for this operation. Here $|T|$ denotes the number of non-zero elements in the matrix T , which – in expectation – is nkq . After this preprocessing, Algorithm 2 must compute the pseudoinverse of a $k \times d$ matrix, or, equivalently, solve a least-squares problem on k constraints and d variables. This operation can be performed in $O(kd^2)$ time since $k \geq d$. Thus, the entire algorithm runs in expected time

$$\begin{aligned} & n(d + 1) + 2n(d + 1) \mathbf{E} [\log_2 (|T| + 1)] + O(kd^2) \leq n(d + 1) \\ & + 2n(d + 1) \log_2 (nkq + 1) + O(kd^2). \end{aligned}$$

6 Appendix

6.1 Approximating matrix multiplication

Let $A \in \mathbb{R}^{m \times n}$ be any matrix. Consider the following algorithm (which is essentially the algorithm in page 876 of [13]) that constructs a matrix $C \in \mathbb{R}^{m \times c}$ consisting of c rescaled columns of A . We will seek a bound on the approximation error $\|AA^T - CC^T\|_2$, which we will provide in Theorem 4. A variant of this theorem appeared as Theorem 7 in [13]; this version modifies and supersedes Eq. (47) of Theorem 7 in the following manner: first, we will assume that the spectral norm of A is bounded and is at most one (this is a minor normalization assumption). Second, and most importantly, we will need to set c to be at least the value of Eq. (40) for the theorem to hold. This second assumption was omitted from the statement of Eq. (47) in Theorem 7 of [13].

Data : $A \in \mathbb{R}^{m \times n}$, $p_i \geq 0, i \in [n]$ s.t. $\sum_{i \in [n]} p_i = 1$, positive integer $c \leq n$.
Result : $C \in \mathbb{R}^{m \times c}$
 Initialize $S \in \mathbb{R}^{m \times c}$ to be an all-zero matrix.
for $t = 1, \dots, c$ **do**
 Pick $i_t \in [n]$, where $\Pr(i_t = i) = p_i$;
 $S_{i_t t} = 1/\sqrt{cp_{i_t}}$;
end
 Return $C = AS$;

Algorithm 3: The EXACTLY(c) algorithm.

Theorem 4 *Let $A \in \mathbb{R}^{m \times n}$ with $\|A\|_2 \leq 1$. Construct C using the EXACTLY(c) algorithm and let the sampling probabilities p_i satisfy*

$$p_i \geq \beta \frac{\|A^{(i)}\|_2^2}{\|A\|_F^2} \tag{39}$$

for all $i \in [n]$ for some constant $\beta \in (0, 1]$. Let $\epsilon \in (0, 1)$ be an accuracy parameter and assume $\|A\|_F^2 \geq 1/24$. If

$$c \geq \frac{96 \|A\|_F^2}{\beta \epsilon^2} \ln \left(\frac{96 \|A\|_F^2}{\beta \epsilon^2 \sqrt{\delta}} \right) \tag{40}$$

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

$$\|AA^T - CC^T\|_2 \leq \epsilon.$$

Proof Consider the EXACTLY(c) algorithm. Then

$$AA^T = \sum_{i=1}^n A^{(i)} A^{(i)T}.$$

Similar to [23] we shall view the matrix AA^T as the true mean of a bounded operator valued random variable, whereas $CC^T = AS(AS)^T = ASS^T A^T$ will be its empirical mean. Then, we will apply Lemma 1 of [21]. To this end, define a random vector $y \in \mathbb{R}^m$ as

$$\Pr \left[y = \frac{1}{\sqrt{p_i}} A^{(i)} \right] = p_i$$

for $i \in [n]$. The matrix $C = AS$ has columns $\frac{1}{\sqrt{c}}y^1, \frac{1}{\sqrt{c}}y^2, \dots, \frac{1}{\sqrt{c}}y^c$, where y^1, y^2, \dots, y^c are c independent copies of y . Using this notation, it follows that

$$\mathbf{E} \left[yy^T \right] = AA^T \tag{41}$$

and

$$CC^T = ASS^T A^T = \frac{1}{c} \sum_{t=1}^c y^t y^{tT}.$$

Finally, let

$$M = \|y\|_2 = \frac{1}{\sqrt{p_i}} \|A^{(i)}\|_2. \tag{42}$$

We can now apply Lemma 1, p. 3 of [21]. Notice that from Eq. (41) and our assumption on the spectral norm of A , we immediately get that

$$\left\| \mathbf{E} \left[yy^T \right] \right\|_2 = \|AA^T\|_2 \leq \|A\|_2 \|A^T\|_2 \leq 1.$$

Then, Lemma 1 of [21] implies that

$$\|CC^T - AA^T\|_2 < \epsilon, \tag{43}$$

with probability at least $1 - (2c)^2 \exp \left(-\frac{c\epsilon^2}{16M^2 + 8M^2\epsilon} \right)$. Let δ be the failure probability of Theorem 4; we seek an appropriate value of c in order to guarantee $(2c)^2$

$\exp\left(-\frac{c\epsilon^2}{16M^2+8M^2\epsilon}\right) \leq \delta$. Equivalently, we need to satisfy

$$\frac{c}{\ln(2c/\sqrt{\delta})} \geq \frac{2}{\epsilon^2} (16M^2 + 8M^2\epsilon).$$

Recall that $\epsilon < 1$, and combine Eqs. (42) and (39) to get $M^2 \leq \|A\|_F^2 / \beta$. Combining with the above equation, it suffices to choose a value of c such that

$$\frac{c}{\ln(2c/\sqrt{\delta})} \geq \frac{48}{\beta\epsilon^2} \|A\|_F^2,$$

or, equivalently,

$$\frac{2c/\sqrt{\delta}}{\ln(2c/\sqrt{\delta})} \geq \frac{96}{\beta\epsilon^2\sqrt{\delta}} \|A\|_F^2.$$

We now use the fact that for any $\eta \geq 4$, if $x \geq 2\eta \ln \eta$ then $\frac{x}{\ln x} \geq \eta$. Let $x = 2c/\sqrt{\delta}$, let $\eta = 96 \|A\|_F^2 / (\beta\epsilon^2\sqrt{\delta})$, and note that $\eta \geq 4$ if $\|A\|_F^2 \geq 1/24$, since β, ϵ , and δ are at most one. Thus, it suffices to set

$$\frac{2c}{\sqrt{\delta}} \geq 2 \frac{96 \|A\|_F^2}{\beta\epsilon^2\sqrt{\delta}} \ln\left(\frac{96 \|A\|_F^2}{\beta\epsilon^2\sqrt{\delta}}\right),$$

which concludes the proof of the theorem. □

6.2 The proof of Lemma 9

Let $T \in \mathbb{R}^{k \times n}$ be the sparse projection matrix constructed via Algorithm 2 (see Sect. 5.1), with sparsity parameter q . In addition, given $x, y \in \mathbb{R}^n$, let $\Delta = x^T T^T T y - x^T y$. We will derive a bound for

$$\mathbf{E}[\Delta^2] = \mathbf{E}\left[\left(x^T T^T T y - x^T y\right)^2\right].$$

Let $t_{(i)}$ be the i -th row of T as a row vector, for $i \in [k]$, in which case

$$\Delta = \sum_{i=1}^k \left(x^T t_{(i)}^T t_{(i)} y - \frac{1}{k} x^T y\right).$$

Rather than computing $\mathbf{E}[\Delta^2]$ directly, we will instead use that $\mathbf{E}[\Delta^2] = (\mathbf{E}[\Delta])^2 + \mathbf{Var}[\Delta]$. We first claim that $\mathbf{E}[\Delta] = 0$. By linearity of expectation,

$$\mathbf{E}[\Delta] = \sum_{i=1}^k \left[\mathbf{E} \left[x^T t_{(i)}^T t_{(i)} y \right] - \frac{1}{k} x^T y \right]. \tag{44}$$

We first analyze $t_{(i)} = t$ for some fixed i (w.l.o.g. $i = 1$). Let t_i denote the i -th element of the vector t and recall that $\mathbf{E}[t_i] = 0$, $\mathbf{E}[t_i t_j] = 0$ for $i \neq j$, and also that $\mathbf{E}[t_i^2] = 1/k$. Thus,

$$\begin{aligned} \mathbf{E} \left[x^T t^T t y \right] &= \mathbf{E} \left[\sum_{i=1}^n \sum_{j=1}^n x_i t_i t_j y_j \right] = \sum_{i=1}^n \sum_{j=1}^n x_i \mathbf{E} [t_i t_j] y_j = \sum_{i=1}^n x_i \mathbf{E} [t_i^2] y_i \\ &= \frac{1}{k} x^T y. \end{aligned}$$

By combining the above with Eq. (44), it follows that $\mathbf{E}[\Delta] = 0$, and thus that $\mathbf{E}[\Delta^2] = \mathbf{Var}[\Delta]$. In order to provide a bound for $\mathbf{Var}[\Delta]$, note that

$$\mathbf{Var}[\Delta] = \sum_{i=1}^k \mathbf{Var} \left[x^T t_{(i)}^T t_{(i)} y - \frac{1}{k} x^T y \right] \tag{45}$$

$$= \sum_{i=1}^k \mathbf{Var} \left[x^T t_{(i)}^T t_{(i)} y \right]. \tag{46}$$

Equation (45) follows since the k random variables $x^T t_{(i)}^T t_{(i)} y - \frac{1}{k} x^T y$ are independent (since the elements of T are independent) and Eq. (46) follows since $\frac{1}{k} x^T y$ is constant. In order to bound Eq. (46), we first analyze $t_{(i)} = t$ for some i (w.l.o.g. $i = 1$). Then,

$$\begin{aligned} \mathbf{Var} \left[x^T t^T t y \right] &= \mathbf{E} \left[(x^T t^T t y)^2 \right] - \left(\mathbf{E} \left[x^T t^T t y \right] \right)^2 \\ &= \mathbf{E} \left[(x^T t^T t y)^2 \right] - \frac{1}{k^2} (x^T y)^2. \end{aligned} \tag{47}$$

We will bound the $\mathbf{E} \left[(x^T t^T t y)^2 \right]$ term directly:

$$\begin{aligned} \mathbf{E} \left[\left(\sum_{i=1}^n \sum_{j=1}^n x_i t_i t_j y_j \right)^2 \right] &= \mathbf{E} \left[\sum_{i_1=1}^n \sum_{i_2=1}^n \sum_{j_1=1}^n \sum_{j_2=1}^n x_{i_1} x_{i_2} t_{i_1} t_{i_2} t_{j_1} t_{j_2} y_{j_1} y_{j_2} \right] \\ &= \sum_{i_1=1}^n \sum_{i_2=1}^n \sum_{j_1=1}^n \sum_{j_2=1}^n x_{i_1} x_{i_2} \mathbf{E} [t_{i_1} t_{i_2} t_{j_1} t_{j_2}] y_{j_1} y_{j_2}. \end{aligned} \tag{48}$$

Notice that if any of the four indices i_1, i_2, j_1, j_2 appears only once, then the expectation $\mathbf{E} [t_{i_1} t_{i_2} t_{j_1} t_{j_2}]$ corresponding to those indices equals zero. This expectation is non-zero if the four indices are paired in couples or if all four are equal. That is, non-zero expectation happens if

- (A) : $i_1 = i_2 \neq j_1 = j_2$ ($n^2 - n$ terms)
- (B) : $i_1 = j_1 \neq i_2 = j_2$ ($n^2 - n$ terms)
- (C) : $i_1 = j_2 \neq i_2 = j_1$ ($n^2 - n$ terms)
- (D) : $i_1 = i_2 = j_1 = j_2$ (n terms).

For case (A), let $i_1 = i_2 = \ell$ and let $j_1 = j_2 = p$, in which case the corresponding terms in Eq. (48) become:

$$\begin{aligned} \sum_{\ell=1}^n \sum_{p=1:p \neq \ell}^n x_\ell^2 \mathbf{E} [t_\ell^2 t_p^2] y_p^2 &= \sum_{\ell=1}^n \sum_{p=1:p \neq \ell}^n x_\ell^2 \mathbf{E} [t_\ell^2] \mathbf{E} [t_p^2] y_p^2 \\ &= \frac{1}{k^2} \sum_{\ell=1}^n \sum_{p=1:p \neq \ell}^n x_\ell^2 y_p^2 \\ &= \frac{1}{k^2} \sum_{\ell=1}^n \sum_{p=1:p \neq \ell}^n x_\ell^2 y_p^2 + \frac{1}{k^2} \sum_{p=1}^n x_p^2 y_p^2 - \frac{1}{k^2} \sum_{p=1}^n x_p^2 y_p^2 \\ &= \frac{1}{k^2} \|x\|_2^2 \|y\|_2^2 - \frac{1}{k^2} \sum_{p=1}^n x_p^2 y_p^2. \end{aligned}$$

Similarly, cases (B) and (C) give:

$$\begin{aligned} \sum_{\ell=1}^n \sum_{p=1:p \neq \ell}^n x_\ell x_p \mathbf{E} [t_\ell^2 t_p^2] y_\ell y_p &= \frac{1}{k^2} (x^T y)^2 - \frac{1}{k^2} \sum_{p=1}^n x_p^2 y_p^2 \\ &\quad \text{(where } i_1 = j_1 = \ell \text{ and } i_2 = j_2 = p), \text{ and} \\ \sum_{\ell=1}^n \sum_{p=1:p \neq \ell}^n x_\ell x_p \mathbf{E} [t_\ell^2 t_p^2] y_\ell y_p &= \frac{1}{k^2} (x^T y)^2 - \frac{1}{k^2} \sum_{p=1}^n x_p^2 y_p^2 \\ &\quad \text{(where } i_1 = j_2 = \ell \text{ and } i_2 = j_1 = p). \end{aligned}$$

Finally, for case (D), let $i_1 = i_2 = j_1 = j_2 = \ell$, in which case:

$$\sum_{\ell=1}^n x_\ell^2 \mathbf{E} [t_\ell^4] y_\ell^2 = \frac{1}{k^2 q} \sum_{\ell=1}^n x_\ell^2 y_\ell^2,$$

where we have used that $\mathbf{E}[t_\ell^4] = 1/(k^2q)$. By combining these four terms for each of the k terms in the sum, it follows from Eqs. (46) and (47) that

$$\begin{aligned} \mathbf{E}[\Delta^2] &= k \left(\frac{1}{k^2} \|x\|_2^2 \|y\|_2^2 + \frac{2}{k^2} (x^T y)^2 - \frac{3}{k^2} \sum_{p=1}^n x_p^2 y_p^2 + \frac{1}{k^2 q} \sum_{p=1}^n x_p^2 y_p^2 - \frac{1}{k^2} (x^T y)^2 \right) \\ &\leq \frac{2}{k} \|x\|_2^2 \|y\|_2^2 + \frac{1}{kq} \sum_{p=1}^n x_p^2 y_p^2. \end{aligned} \quad (49)$$

In the above we used $(x^T y)^2 \leq \|x\|_2^2 \|y\|_2^2$. Since we assumed that $\|x\|_\infty \leq \alpha$, the second term on the right hand side of Eq. (49) is bounded by $\frac{\alpha^2}{kq} \|y\|_2^2$ and the lemma follows since we have assumed that $q \geq \alpha^2$.

References

1. Ailon, N., Chazelle, B.: Approximate nearest neighbors and the fast Johnson–Lindenstrauss transform. In: Proceedings of the 38th Annual ACM Symposium on Theory of Computing, pp. 557–563 (2006)
2. Ailon, N., Liberty, E.: Fast dimension reduction using Rademacher series on dual BCH codes. In Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms, pages 1–9, (2008)
3. Arora, S., Hazan, E., Kale, S.: A fast random sampling algorithm for sparsifying matrices. In: Proceedings of the 10th International Workshop on Randomization and Computation, pp. 272–279 (2006)
4. Avron, H., Maymounkov, P., Toledo, S.: Blendenpik: Supercharging LAPACK’s least-squares solver. Manuscript (2009)
5. Avron, H., Maymounkov, P., Toledo, S.: Blendenpik: Supercharging LAPACK’s least-squares solver. *SIAM J. Scientific Comput.* (2010, to appear)
6. Ben-Israel, A., Greville, T.N.E.: *Generalized Inverses: Theory and Applications*. Springer, New York (2003)
7. Bhatia, R.: *Matrix Analysis*. Springer-Verlag, New York (1997)
8. Clarkson, K.L., Woodruff, D.P.: Numerical linear algebra in the streaming model. In: Proceedings of the 41st Annual ACM Symposium on Theory of Computing, pp. 205–214 (2009)
9. Coppersmith, D., Winograd, S.: Matrix multiplication via arithmetic progressions. *J. Symbolic Comput.* **9**(3), 251–280 (1990)
10. Dahlquist, G., Sjöberg, B., Svensson, P.: Comparison of the method of averages with the method of least squares. *Math. Comput.* **22**(104), 833–845 (1968)
11. Drineas, P., Kannan, R., Mahoney, M.W.: Fast Monte Carlo algorithms for matrices I: approximating matrix multiplication. *SIAM J. Comput.* **36**, 132–157 (2006)
12. Drineas, P., Mahoney, M.W., Muthukrishnan, S.: Sampling algorithms for ℓ_2 regression and applications. In: Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 1127–1136 (2006)
13. Drineas, P., Mahoney, M.W., Muthukrishnan, S.: Relative-error CUR matrix decompositions. *SIAM J. Matrix Anal. Appl.* **30**, 844–881 (2008)
14. Drineas, P., Mahoney, M.W., Muthukrishnan, S., Sarlós, T.: Faster least squares approximation. Technical report. Preprint: arXiv:0710.1435 (2007)
15. Feige, U., Ofek, E.: Spectral techniques applied to sparse random graphs. *Random Struct. Algorithms* **27**(2), 251–275 (2005)
16. Golub, G.H., Van Loan, C.F.: *Matrix Computations*. Johns Hopkins University Press, Baltimore (1996)
17. Ibarra, O.H., Moran, S., Hui, R.: A generalization of the fast LUP matrix decomposition algorithm and applications. *J. Algorithms* **3**, 45–56 (1982)
18. Matoušek, J.: On variants of the Johnson–Lindenstrauss lemma. *Random Struct. Algorithms* **33**(2), 142–156 (2008)

19. Motwani, R., Raghavan, P.: *Randomized Algorithms*. Cambridge University Press, New York (1995)
20. Nguyen, N.H., Do, T.T., Tran, T.D.: A fast and efficient algorithm for low-rank approximation of a matrix. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing*, pp. 215–224 (2009)
21. Oliveira, R.I.: Sums of random Hermitian matrices and an inequality by Rudelson. arXiv:1004.3821v1 (2010)
22. Rokhlin, V., Tygert, M.: A fast randomized algorithm for overdetermined linear least-squares regression. *Proc. Natl. Acad. Sci. USA* **105**(36), 13212–13217 (2008)
23. Rudelson, M., Vershynin, R.: Sampling from large matrices: an approach through geometric functional analysis. *J. ACM*, **54**(4):Article 21 (2007)
24. Sarlós, T.: Improved approximation algorithms for large matrices via random projections. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science*, pp. 143–152 (2006)
25. Stewart, G.W., Sun, J.G.: *Matrix Perturbation Theory*. Academic Press, New York (1990)
26. Stigler, S.M.: *The History of Statistics: The Measurement of Uncertainty before 1900*. Harvard University Press, Cambridge (1986)