Fast, Deterministic and Sparse Dimensionality Reduction.[*]

Daniel Dadush [†1], Cristóbal Guzmán [‡2], and Neil Olver [1,3]

[1]Centrum Wiskunde & Informatica, The Netherlands
[2]Pontificia Universidad Católica de Chile, Chile
[3]Vrije Universiteit Amsterdam, The Netherlands

Abstract

We provide a deterministic construction of the sparse Johnson-Lindenstrauss transform of Kane & Nelson (J.ACM 2014) which runs, under a mild restriction, in the time necessary to apply the sparse embedding matrix to the input vectors. Specifically, given a set of n vectors in $\mathbb{R}^d$ and target error ε, we give a deterministic algorithm to compute a $\{-1, 0, 1\}$ embedding matrix of rank $O((\ln n)/\varepsilon^2)$ with $O((\ln n)/\varepsilon)$ entries per column which preserves the norms of the vectors to within $1 \pm \varepsilon$. If NNZ, the number of non-zero entries in the input set of vectors, is $\Omega(d^2)$, our algorithm runs in time $O(\mathrm{NNZ} \cdot \ln n/\varepsilon)$.

One ingredient in our construction is an extremely simple proof of the Hanson-Wright inequality for subgaussian random variables, which is more amenable to derandomization. As an interesting byproduct, we are able to derive the essentially optimal form of the inequality in terms of its functional dependence on the parameters.

1 Introduction

Dimensionality reduction is an important and widely used technique in areas such as computer science, optimization and machine learning. Arguably, the most important result in dimensionality reduction is the (distributional) Johnson-Lindenstrauss (J-L) Lemma.

Lemma 1.1 ([20]). *Let d be any positive integer, and let $\varepsilon, \delta \in (0, 1/2)$. Then there exists a random $m \times d$ matrix Π, with $m = \Theta(\varepsilon^{-2}\ln(1/\delta))$, such that for any $v \in \mathbb{R}^d$,*

$$\mathbb{P}\big((1-\varepsilon)\|v\|_2^2 \le \|\Pi v\|_2^2 \le (1+\varepsilon)\|v\|_2^2\big) \ge 1-\delta.$$

In particular, a union bound implies that for any set V of n vectors in $\mathbb{R}^d$, there exists a linear map into $\mathbb{R}^m$, where $m = \Theta(\varepsilon^{-2}\ln n)$, which preserves all lengths to within $1 \pm \varepsilon$.

This result lies at the heart of many state-of-the-art algorithms for problems including approximate nearest neighbors [19], mixtures of Gaussians [12], sketching [35] and fast algorithms for numerical linear algebra [13].

It is very desirable to be able to implement the dimensionality reduction as efficiently and as conveniently as possible, and there have been a number of works in this direction. The original construction involved projecting onto a random subspace [20]. Indyk and Motwani [19] showed that a projection matrix consisting of i.i.d. Gaussian entries also works, and Achlioptas [1] showed that even i.i.d. Rademacher random variables suffice, which is particularly convenient in some applications.

All these constructions produce dense (with at best a constant fraction of zero entries), unstructured projection matrices. Ailon and Chazelle [3] showed how to construct a projection that can be applied to a vector substantially faster than an arbitrary projection matrix. The idea is to first apply a random Hadamard transform that can be very

[*] A short version of this paper was presented at the ACM-SIAM Symposium of Discrete Algorithms (SODA18)

[†] Supported by the NWO Veni grant 639.071.510.

[‡] Partially supported by FONDECYT Iniciación project 11160939. This work was partially done during a visit to Centrum Wiskunde & Informatica.

efficiently applied using the fast Fourier transform, followed by projecting by a random matrix with i.i.d. $\{-1, 0, 1\}$ entries which is very sparse. The initial Hadamard transform ensures that with high probability, the weight of the transformed vector is well spread amongst its coordinates; this guarantees that the sparse projection is effective. A number of other results in this direction have followed [4, 26, 5, 25, 17].

One disadvantage of these results is that while they speed up the embedding of dense vectors, they cannot exploit sparsity of the input vectors. This motivated work of Dasgupta et al. [11], refined by Kane and Nelson [22], on choosing projection matrices that are directly sparse. Kane and Nelson [22] achieve the same guarantees as the usual J-L Lemma, with a projection matrix containing only $O(m\varepsilon)$ nonzero entries per column. This can be multiplied by a vector u in time $O(\mathrm{NNZ}(u) \cdot m\varepsilon)$, achieving a speedup of a factor of ε compared to multiplying u by an unstructured dense matrix. For sufficiently low sparsity this can be faster than the construction of Ailon-Chazelle.

A second line of work has been on reducing the amount of randomness required to implement it, or eliminating it altogether. There are two distinct threads here.

Distributional J-L with few random bits. In one thread, the goal is a distributional J-L Lemma, of the form given in Lemma 1.1, but where the number of random bits required to describe the distribution of the random matrix Π is as small as possible. This goal may be alongside ensuring other desirable properties of Π, such as sparsity. This has implications to streaming applications.

Representative recent results include independent work of Kane and Nelson [22] and Meka [27] that achieve a construction requiring $O(\ln d + \ln(1/\varepsilon)\ln(1/\delta) + \ln(1/\delta)\ln\ln(1/\delta))$ random bits (the result of Meka is restricted to $\delta = \Omega(1/\operatorname{poly} d)$). Other works include Clarkson and Woodruff [9] and Karnin et al. [24]. It is an open question to give an explicit construction requiring $O(\ln d + \ln(1/\delta))$ bits; probabilistic arguments show that such distributions do exist.

In independent work, Cohen, Jayram and Nelson [10] provided a simplified analysis of the sparse J-L transform in [22], that avoids the expansion of the moments of $\|\Pi v\|_2^2 - 1$, as well as the intrincate combinatorics arising from such analysis. Our work is based on the same idea of directly bounding the mgf, however our proof it is arguably simpler (and provides a tight bound) since we also avoid decoupling.

Full input-dependent derandomization. In a second thread, the goal is true derandomization: given the input vectors explicitly, deterministically give an embedding matrix which has low distortion on all of the input vectors. This will be our goal. One application of these results is to derandomize approximation algorithms based on rounding of semidefinite programs, for example, MAX-CUT.

Sivakumar [33] gives a polynomial-time derandomization based on a quite general complexity-theoretic tool due to Nisan [28, 29]. He does not state an explicit bound on the running time. Engebretsen et al. [14] give a derandomization that achieves a running time of $O(mn(\ln n + 1/\varepsilon)^{O(1)})$. This matches the time required to project with an unstructured dense matrix, up to logarithmic factors. These logarithmic factors result from the somewhat brute-force nature of their approach. They derandomize the version of J-L with gaussian entries using the method of conditioned expectations. In order to do this, they fix the bits in the binary representation of the projection matrix one at a time (with a logarithmic number of bits needed per entry), and numerically approximate the integrals needed to evaluate various conditioned expectations.

Bhargava and Kosaraju [6] give a derandomization of J-L with $\{-1, 1\}$ entries, using the method of pessimistic estimators [31]. We also use pessimistic estimators, but as discussed below, our approach has a number of advantages; in particular, our algorithm is faster due to the use of sparse projection matrices.

Contributions and techniques. Our main result is a full derandomization that takes time proportional to the "embedding time", i.e., the time required to apply the projection matrix to the input vectors. Moreover, we achieve this with a sparse projection matrix: our produced matrix will have roughly $m\varepsilon$ nonzero entries per column, leading to an embedding time of $O(m\varepsilon\mathrm{NNZ}(V))$.

We will generally work with the real model of computation [7], but where in addition we allow the computation of exponentials and square roots in constant time. In order to actually implement these computations, sufficiently accurate approximations of especially exp must be used. In Section 5.3 we discuss how known results on efficient approximation of elementary functions can be used to obtain the same running time up to polylogarithmic factors in the bit model of computation.

Theorem 1.2. *Let V be a set of $n \geq d$ vectors in $\mathbb{R}^d$, with $\mathrm{NNZ}(V) = \Omega(d^2)$, $\varepsilon \in (0, 1)$ and $m \geq 12 \ln(2n)/\varepsilon^2$. Then in time $O(m\varepsilon\mathrm{NNZ}(V))$, a matrix $\Pi \in \mathbb{R}^{m\times d}$ can be found, with $\lceil m\varepsilon\rceil$ nonzero entries per column, for which*

$$(1-\varepsilon)\|v\|_2^2 \leq \|\Pi v\|_2^2 \leq (1+\varepsilon)\|v\|_2^2 \qquad \forall v \in V.$$

This matches the result of Kane and Nelson [22] for the important case $\mathrm{NNZ}(V) \geq d^2$, both with respect to the size and sparsity of the resulting projection matrix, and in terms of the time required to project the vectors. (We also match their result in full generality, except we do not obtain the desired running time.)

Broadly speaking, our approach is based on the use of pessimistic estimators. A pessimistic estimator is nothing more than a quantity that upper bounds the failure probability of an event under some distribution. In this method, one begins with a distribution for which the event occurs with nonzero probability, and iteratively specifies the random variables in such a way that the chosen pessimistic estimator never increases. If this can be done, the event must occur in the final deterministic result with positive probability, and hence probability one. This is a standard technique, but the challenge is to find an appropriate pessimistic estimator that can be updated fast enough for our purposes, and behaves well (is concave in an appropriate sense) under updates. In this respect, our approach is similar to that of Bhargava and Kosaraju [6]. The main advantage of our result is that we obtain an $O(\varepsilon)$ speedup attributable to the sparsity of the projection matrix. Moreover the pessimistic estimators that we use are rather natural, coming from moment generating functions associated with Gaussian random variables.

In order to obtain a sparse projection matrix, we adapt the arguments of Kane and Nelson [23] for obtaining a random sparse projection matrix. Their approach has two parts; the first produces a (random) sparsity pattern, describing which entries of the projection matrix will be nonzero; the second chooses the signs for these entries (in their case, i.i.d. Rademachers). We derandomize both of these steps separately, using the method of pessimistic estimators.

- For the sparsity pattern (mask) of the matrix, Kane and Nelson choose, for each column independently, precisely s nonzero entries uniformly at random, where $s = \Theta(m\varepsilon)$. Thus their sampling procedure involved dependence amongst the entries in a column, and indeed this is crucial. They argue via negative dependence to obtain the required concentration bounds.

  Although our final algorithm eschews its use in order to run as quickly as possible, our derandomization is inspired by *pipage rounding* [2, 34], a general rounding technique for rounding a fractional point in the base polytope of a matroid to a vertex. The relevant matroid in our case is the partition matroid, where we must choose precisely s entries per column; in fact, we later use a more refined partition matroid to obtain our claimed running time. We begin with an initial uniform fractional point, which corresponds to the uniform distribution over sparsity patterns with s nonzero entries per row; the final, integral, output from the pipage rounding procedure corresponds to a deterministic choice of a mask. Harvey and Olver [16] suggest how pipage rounding can be combined with the use of pessimistic estimators, as long as the pessimistic estimator is well-behaved (more precisely, concave in the directions relevant to the pipage rounding scheme). Our pessimistic estimator, which is essentially the moment generating function after an application of Hölder's inequality, satisfies these requirements.

  We have to be very efficient in order to obtain the desired running time: for example, if $\varepsilon \ll 1/\mathrm{NNZ}(V)$ then the number of entries md in the projection matrix is larger than the desired running time $m\varepsilon\mathrm{NNZ}(V)$. We adapt an approach proposed by Kane and Nelson [23] (they named it the *code construction*). For $n \geq d$, there is a "universal" sparsity pattern that satisfies the required properties and depends only on the size of the projection matrix and not on the input vectors. The requirements of this universal sparsity pattern turn out to be convenient when it comes to updating the pessimistic estimator quickly.

- In order to show that randomly signing the chosen nonzero entries works, Kane and Nelson use the Hanson-Wright inequality [15] on a matrix determined by the mask. This inequality provides concentration of a quadratic form $\boldsymbol{z}^{\mathsf{T}} A \boldsymbol{z}$, where $\boldsymbol{z}$ is a random vector with independent subgaussian entries, in terms of parameters of the matrix A.

> We thus proceed by providing a derandomized version of the Hanson-Wright inequality (to our knowledge, the first such). Existing proofs of the inequality do not seem to yield pessimistic estimators that can be quickly evaluated. We give a simpler proof, based on, but simplifying, a proof of Rudelson and Vershynin [32], that is very convenient for our purposes. The crucial idea is that the moment generating function of the quadratic form can be bounded in terms of the moment generating function of $\boldsymbol{g}^\mathsf{T} A\boldsymbol{g}$, where $\boldsymbol{g}$ is *Gaussian*, which can be handled much more easily. Further work is needed once again to get the desired running time of $O(m\varepsilon \cdot \mathrm{NNZ}(V))$ and for this we exploit some of the structure of the specific matrix obtained in the application of the Hanson-Wright inequality to dimension reduction.

If $\mathrm{NNZ}(V) = o(d^2)$, we can still obtain the indicated sparse projection matrix, but with running time $O(m(\mathrm{NNZ}(V)+d))$. The second term is comparable to the time required to even write down the projection matrix, and so is quite justified. But we do not obtain an improvement over just using a dense projection matrix. The difficulty comes in the selection of the sparsity pattern (Section 4.2); if this can be improved to time $O(s(\mathrm{NNZ}(V) + d))$ then the existing algorithm for choosing the signs is sufficiently fast.

As a further implication of our simplified proof of the Hanson-Wright inequality, we are able to substantially sharpen it, and give essentially the correct functional dependence on the parameters. We think that this is of independent interest.

The paper is structured as follows. In Section 2 we introduce our notation and some preliminary results. In Section 3 we present our new Hanson-Wright inequality, together with its proof, and the proposed pessimistic estimators for derandomizing this inequality. In Section 4 we exploit the connection between Hanson-Wright and Johnson-Lindenstrauss dimensionality reduction to obtain pessimistic estimators for the latter: this includes the signing algorithm (Section 4.1) and the masking algorithm (Section 4.2), which can be analyzed separately. In Section 5 we focus on efficient computation of the various pessimistic estimators we use; in particular, for the Hanson-Wright inequality in general (Section 5.1), and for the signing (Section 5.2) and masking (Section 5.4) components of our new fast deterministic algorithm for sparse J-L. Finally, in Section 6, we discuss further technical aspects of our sharp Hanson-Wright inequality.

2 Notation and Preliminaries

We denote the natural logarithm by $\ln$. Random variables are denoted by bold characters. For any positive integer k, $[k]$ denotes the set $\{1, 2, \dots, k\}$.

Given a matrix A, we use $\|A\|_{\mathrm{op}}$ and $\|A\|_F$ to denote the operator and Frobenius norms, respectively. Given two vectors u, v in the same vector space, $\langle u, v\rangle$ denotes the standard inner product. Given two matrices (or vectors), their Hadamard product corresponds to the component-wise multiplication, and this operation is denoted by $\odot$. The direct sum, $\bigoplus_{i=1}^k A_i$, is used to denote a block diagonal matrix with blocks $A_1, \dots, A_k$. Given a matrix A, $\mathrm{Diag}(A)$ is the diagonal matrix that coincides with A on the diagonal. Given a vector v, $\mathrm{Diag}(v)$ is a diagonal matrix with diagonal coefficients from v. We denote by $\mathrm{supp}(A)$ the support of A, where A may be a vector or matrix. Given a vector $u \in \mathbb{R}^k$ and $S \subseteq [k]$, $u_S \in \mathbb{R}^{|S|}$ denotes the restriction of u to the coordinates in S. Similarly, given $M \in \mathbb{R}^{k\times k}$ and $S, T \subseteq [k]$, $M_{S,T}$ denotes the submatrix of M indexed by rows S and columns T.

The following well-known result allows to obtain fast computation of the determinant for a rank one perturbation of a matrix.

Lemma 2.1. *Let $A \in \mathbb{R}^{k\times k}$ be an invertibe matrix, and $u, v \in \mathbb{R}^k$ vectors, then*

1. *Matrix determinant lemma:* $\det(A + uv^\mathsf{T}) = \det(A) \cdot (1 + v^\mathsf{T} A^{-1}u)$.
2. *Sherman-Morrison formula:* $(A + uv^\mathsf{T})^{-1} = A^{-1} - \dfrac{A^{-1}uv^\mathsf{T} A^{-1}}{1 + v^\mathsf{T} A^{-1}u}$.

We record the following simple inequalities regarding the logarithm.

Lemma 2.2. *For $x \geq 0$,*

1. $\ln(1+x) \le \frac{x}{2} + \frac{x}{2(1+x)}$.

2. $\ln(1+x) \ge \frac{x}{x+1}$.

3. $x\ln(1+1/x)$ *is increasing.*

Proof.

Proof of 1. Note that $\ln(1+x)$ and $\frac{x}{2} + \frac{x}{2(1+x)}$ are equal at $x=0$, so it suffices to compare derivatives. In particular, we must show that $\frac{1}{1+x} \le \frac{1}{2} + \frac{1}{2(1+x)} - \frac{x}{2(1+x)^2}$. Multiplying through by $2(1+x)^2$ and rearranging, this inequality is equivalent to $x^2 \ge 0$, which is clearly true.

Proof of 2.

$$\ln(1+x) = \ln\left(\frac{1}{1-\frac{x}{x+1}}\right) \ge \ln\left(\frac{1}{e^{-\frac{x}{x+1}}}\right) = \frac{x}{x+1} \, ,$$

where the first inequality follows from the inequality $1+y \le e^y$ applied to $y = -\frac{x}{x+1}$.

Proof of 3. Taking the derivative, we must show that $\ln(1+1/x) - \frac{1/x}{1+1/x} > 0$ for $x > 0$. Replacing x by $1/x$, this is equivalent to $\ln(1+x) > \frac{x}{1+x}$, which follows from part 2. □

We consider an ambient space $\mathbb{R}^d$ where our set of n input vectors V lie. Given $\varepsilon > 0$ and a matrix $\Pi \in \mathbb{R}^{m\times d}$ we will say it is a low-distortion projection for V if for all $v \in V$

$$(1-\varepsilon)\|v\|_2^2 \le \|\Pi v\|_2^2 \le (1+\varepsilon)\|v\|_2^2.$$

Given such a projection matrix Π we say its column sparsity is s if for all columns it has at most s nonzero coordinates.

We recall the definition of the moment generating function (mgf) of a random variable $\boldsymbol{x}$, $\psi_{\boldsymbol{x}}(\lambda) = \mathbb{E}_{\boldsymbol{x}}[e^{\lambda \boldsymbol{x}}]$. Given two real-valued random variables $\boldsymbol{y}, \boldsymbol{z}$ on the same probability space, we write $\boldsymbol{y} \preceq_{\mathrm{m}} \boldsymbol{z}$ if $\psi_{\boldsymbol{y}}(\lambda) \le \psi_{\boldsymbol{z}}(\lambda)$ for all $\lambda \in \mathbb{R}$.

A random variable $\boldsymbol{x}$ is defined to be *ν-subgaussian*, $\nu > 0$, if $\boldsymbol{x} \preceq_{\mathrm{m}} \boldsymbol{z}$ where $\boldsymbol{z} \sim \mathcal{N}(0,\nu^2)$, namely $\mathbb{E}[e^{\lambda\boldsymbol{x}}] \le e^{\lambda^2\nu^2/2}$ $\forall \lambda \in \mathbb{R}$. It is well known that up to a constant factor in the parameter, the above is equivalent to $\boldsymbol{x}$ being centered, $\mathbb{E}[\boldsymbol{x}] = 0$, and having subgaussian tails $\mathbb{P}[|\boldsymbol{x}| \ge \nu t] \le 2e^{-t^2/2}$ $\forall t \ge 0$. The former definition will however be more convenient for our purposes.

We recall that Rademacher random variables (i.e. uniform on $\{-1,1\}$) are 1-subgaussian.

Lemma 2.3. *For $\boldsymbol{x}$ Rademacher and $\lambda \in \mathbb{R}$, $\mathbb{E}[e^{\lambda\boldsymbol{x}}] \le e^{\lambda^2/2}$.*

Proof. By taking the power series expansion of e^x, we get

$$\mathbb{E}[e^{\lambda\boldsymbol{x}}] = \frac{1}{2}(e^{\lambda} + e^{-\lambda}) = \sum_{k=0}^{\infty} \frac{\lambda^{2k}}{(2k)!} \le \sum_{k=0}^{\infty} \frac{\lambda^{2k}}{2^k k!} = e^{\lambda^2/2} \, . \qquad \square$$

The following formula, which comes from direct computation, will be frequently used.

Lemma 2.4. *Let $\boldsymbol{z}$ be a standard Gaussian. Then for any $\alpha \in [0,1/2)$ and $\beta \in \mathbb{R}$*

$$\mathbb{E}_{\boldsymbol{z}}[\exp\left(\alpha\boldsymbol{z}^2 + 2\beta\boldsymbol{z}\right)] = \frac{1}{\sqrt{1-2\alpha}} \exp\left(\frac{2\beta^2}{1-2\alpha}\right).$$

Proof.

$$\begin{aligned}
\mathbb{E}_{\boldsymbol{z}}[\exp\left(\alpha \boldsymbol{z}^2 + 2\beta \boldsymbol{z}\right)] &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{\alpha x^2 + 2\beta x} e^{-x^2/2} dx \\
&= e^{\frac{2\beta^2}{1-2\alpha}} \cdot \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{-(1-2\alpha)(x - \frac{2\beta}{1-2\alpha})^2/2} dx \\
&= e^{\frac{2\beta^2}{1-2\alpha}} \cdot \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{-(1-2\alpha)x^2/2} dx \\
&= e^{\frac{2\beta^2}{1-2\alpha}} \cdot \frac{1}{\sqrt{1-2\alpha}} \cdot \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{-x^2/2} dx \\
&= e^{\frac{2\beta^2}{1-2\alpha}} \cdot \frac{1}{\sqrt{1-2\alpha}} .
\end{aligned}$$

□

3 The Hanson-Wright Inequality and Derandomization

We state below the Hanson-Wright inequality in the form described by Rudelson and Vershynin [32], restricting ourselves to the case of zeros on the diagonal, as this case suffices for dimensionality reduction. It is well known that the general case can be handled by separately analyzing the contribution of diagonal terms, which can be easily bounded as a sum of independent random variables.

Theorem 3.1 (Hanson-Wright [15, 32])**.** *Let $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_k$ be independent ν-subgaussian random variables, and let $A = (a_{ij}) \in \mathbb{R}^{k \times k}$ be a symmetric matrix with $a_{ii} = 0$ for all $i \in [k]$. Then there exists a universal constant $C_{HW} > 0$ so that for every $t > 0$,*

$$\mathbb{P}(\boldsymbol{x}^\mathsf{T} A \boldsymbol{x} > \nu^2 t) \leq \exp\left(-C_{HW} \min\left\{\tfrac{t^2}{\|A\|_F^2}, \tfrac{t}{\|A\|_{\text{op}}}\right\}\right).$$

As one of our contributions, we give a simple proof that provides an essentially sharp version of this result. The bound we achieve is stated below.

Theorem 3.2 (Sharp Hanson-Wright)**.** *Under the conditions of Theorem 3.1,*

$$\mathbb{P}(\boldsymbol{x}^\mathsf{T} A \boldsymbol{x} > \nu^2 t) \leq \exp\left(-\tfrac{1}{2}\left(\tfrac{t}{\|A\|_{\text{op}}} - \tfrac{\|A\|_F^2}{\|A\|_{\text{op}}^2} \cdot \ln(1 + t\tfrac{\|A\|_{\text{op}}}{\|A\|_F^2})\right)\right). \tag{3.1}$$

To prove this result, our main new and simple observation is that the moment generating function of the quadratic form $\boldsymbol{x}^\mathsf{T} A \boldsymbol{x}$ only gets larger when the entries of $\boldsymbol{x}$ are replaced by corresponding Gaussians. We then derive (3.1) by computing an essentially optimal bound on the moment generating function for Gaussian entries as a function of the operator and Frobenius norm. The theorem is therefore optimal with respect to the Chernoff-Cramér method, and hence we expect that it cannot in general be improved.

Comparing to the proof of Rudelson and Vershynin [32], they also reduce to the Gaussian case however with the use of an additional "decoupling step", which replaces the mgf by an average of mgfs obtained by restricting A to a randomly chosen set of rows and columns. We note that decoupling weakens the achievable bound, in particular, as far as we are aware, the best computed Hanson-Wright constant C_{HW} using their method is $1/64$ [30] whereas we achieve $\frac{1-\ln(2)}{2} \approx 0.153 > 3/20$ (see Section 6 for a full proof and discussion). Aside from this, decoupling leads to a more computationally expensive pessimistic estimator from the perspective of derandomization, crucial to the context of this paper, namely an average of many estimators (naïvely an exponential number; using pairwise independence, this can be reduced to $O(k^2)$) versus the single estimator provided by the proof of Theorem 3.2.

As an interesting consequence of Theorem 3.2, we give a slight quantitative improvement on the concentration bounds given in [1] for Rademacher Johnson-Lindenstrauss projections as a function of sparsity.

Lemma 3.3. *Let $v \in \mathbb{R}^d$ be a unit vector with s non-zero entries. For $\Pi \in \{-1, 1\}^{m \times d}$, an $m \times d$ random matrix with coefficients distributed i.i.d. Rademacher, and $\varepsilon > 0$, we have that*

$$\max\left\{\mathbb{P}(\|\Pi v\|_2^2/m \geq 1 + \varepsilon), \mathbb{P}(\|\Pi v\|_2^2/m \leq 1 - \varepsilon)\right\} \leq e^{-\frac{m}{2}\frac{s}{s-1}(\varepsilon - \ln(1+\varepsilon))}.$$

We note that this is better than the error exponent achieved by Gaussian random projections (at least for the harder upper tail estimate) by a small but explicit $s/(s-1)$ factor. The proof of the lemma follows by using sparsity dependent bounds on the operator and Frobenius norm of the relevant matrix, namely $vv^{\mathsf{T}} - \mathrm{Diag}(vv^{\mathsf{T}})$, and applying Theorem 3.2. We defer the proof of the above lemma to Section 6.

We now begin with our proof of Theorem 3.2. The crux of the proof is the following simple lemma which allows us to compare the moment generating functions of diagonal 0 quadratic forms.

Lemma 3.4. *Let $\boldsymbol{y}_1, \dots, \boldsymbol{y}_k$ and $\boldsymbol{z}_1, \dots, \boldsymbol{z}_k$ all be independent, with $\boldsymbol{y}_i \preceq_{\mathrm{m}} \boldsymbol{z}_i$ for all i. Then for a zero-diagonal matrix $A \in \mathbb{R}^{k \times k}$, $\boldsymbol{y}^{\mathsf{T}} A \boldsymbol{y} \preceq_{\mathrm{m}} \boldsymbol{z}^{\mathsf{T}} A \boldsymbol{z}$.*

Proof. It suffices to prove the lemma in the case that $\boldsymbol{y}_i = \boldsymbol{z}_i$ for all $i \neq \ell$ for some $\ell \in [k]$. The full lemma then follows by combining the inequalities

$$\mathbb{E}_{\boldsymbol{z}^{(i)}}[e^{\lambda (\boldsymbol{z}^{(i)})^{\mathsf{T}} A \boldsymbol{z}^{(i)}}] \leq \mathbb{E}_{\boldsymbol{z}^{(i+1)}}[e^{\lambda (\boldsymbol{z}^{(i+1)})^{\mathsf{T}} A \boldsymbol{z}^{(i+1)}}],$$

where $\boldsymbol{z}_j^{(i)} = \boldsymbol{z}_j$ for $i \leq j$ and $\boldsymbol{z}_j^{(i)} = \boldsymbol{y}_j$ for $j > i$. For $\lambda \in \mathbb{R}$, we have that

$$\begin{aligned}
\mathbb{E}[e^{\lambda \boldsymbol{y}^{\mathsf{T}} A \boldsymbol{y}}] &= \mathbb{E}_{\boldsymbol{y}_1, \dots, \boldsymbol{y}_k} e^{\lambda \sum_{i,j} a_{ij} \boldsymbol{y}_i \boldsymbol{y}_j} \\
&= \mathbb{E}_{\boldsymbol{y}_r : r \neq \ell} \Big[\mathbb{E}_{\boldsymbol{y}_\ell} e^{\lambda (\boldsymbol{y}_\ell \sum_{i \neq \ell} (a_{i\ell} + a_{\ell i}) \boldsymbol{y}_i + \sum_{i \neq j \neq \ell} a_{ij} \boldsymbol{y}_i \boldsymbol{y}_j)} \Big] \qquad \text{(since } a_{\ell\ell} = 0\text{)} \\
&\leq \mathbb{E}_{\boldsymbol{y}_r : r \neq \ell} \Big[\mathbb{E}_{\boldsymbol{z}_\ell} e^{\lambda (\boldsymbol{z}_\ell \sum_{i \neq \ell} (a_{i\ell} + a_{\ell i}) \boldsymbol{y}_i + \sum_{i \neq j \neq \ell} a_{ij} \boldsymbol{y}_i \boldsymbol{y}_j)} \Big] \qquad \text{(by } \boldsymbol{y}_\ell \preceq_m \boldsymbol{z}_\ell\text{)} \\
&= \mathbb{E}_{\boldsymbol{z}_r : r \neq \ell} \Big[\mathbb{E}_{\boldsymbol{z}_\ell} e^{\lambda (\boldsymbol{z}_\ell \sum_{i \neq \ell} (a_{i\ell} + a_{\ell i}) \boldsymbol{z}_i + \sum_{i \neq j \neq \ell} a_{ij} \boldsymbol{z}_i \boldsymbol{z}_j)} \Big] \\
&= \mathbb{E}_{\boldsymbol{z}_1, \dots, \boldsymbol{z}_k} [e^{\lambda \boldsymbol{z}^{\mathsf{T}} A \boldsymbol{z}}].
\end{aligned}$$

□

This result allows us to replace all subgaussian variables by Gaussians.

Corollary 3.5. *Under the assumptions of Theorem 3.1,*

$$\boldsymbol{x}^{\mathsf{T}} A \boldsymbol{x} \preceq_{\mathrm{m}} \nu^2 \boldsymbol{g}^{\mathsf{T}} A \boldsymbol{g},$$

where $\boldsymbol{g}$ is a standard k-dimensional Gaussian random vector.

Given the above, to be able to apply the Chernoff bound it suffices to get strong bounds on the moment function of $\boldsymbol{g}^{\mathsf{T}} A \boldsymbol{g}$. We recall the form of the mgf below. Let $\mu_1 \geq \cdots \geq \mu_k$ denote the eigenvalues of A. For $\lambda \in \mathbb{R}$, $|\lambda| < \frac{1}{2\|A\|_{\mathrm{op}}}$,

$$\begin{aligned}
\mathbb{E}[\exp(\lambda \boldsymbol{g}^{\mathsf{T}} A \boldsymbol{g})] &= \mathbb{E}\Big[\exp\Big(\sum_{i=1}^{k} \lambda \mu_i \boldsymbol{g}_i^2\Big)\Big] \quad \text{(by rotational symmetry of the Gaussian)} \\
&= \prod_{i=1}^{k} (1 - 2\lambda\mu_i)^{-\frac{1}{2}} \quad \text{(by Lemma 2.4)} \\
&= \det(I - 2\lambda A)^{-1/2}.
\end{aligned} \tag{3.2}$$

From here, computing an optimal mgf bound corresponds to finding the minimum value of the determinant $\det(I - 2\lambda A)$ subject to fixed upper bounds on $\|A\|_{\mathrm{op}}$ and $\|A\|_F$. In the following lemma, we give a bound on this problem which is optimal in the asymptotic regime, i.e., for $k \to \infty$, under the condition that $\|A\|_{\mathrm{op}}^2$ evenly divides $\|A\|_F^2$. We defer the proof to Section 6.

Lemma 3.6 (Determinant Lower Bound). *For $0 \leq \alpha < 1$, $\beta \geq \alpha$, we have that*

$$\begin{aligned}
&\inf\{\det(I - M) : M \in \mathbb{R}^{k \times k} \textit{ symmetric}, M_{ii} = 0 \; \forall i \in [k], k \in \mathbb{N}, \|M\|_{\mathrm{op}} \leq \alpha, \|M\|_F \leq \beta\} \\
&\qquad\qquad \geq (1 - \alpha)^{\beta^2/\alpha^2} \exp(\beta^2/\alpha),
\end{aligned} \tag{3.3}$$

where equality is attained when β^2/α^2 is an integer.

Applying the above lemma to the matrix $M = 2\lambda A$, we get the following mgf bound:

$$\mathbb{E}[\exp(\lambda \boldsymbol{g}^\mathsf{T} A \boldsymbol{g})] = \det(I - 2\lambda A)^{-1/2} \leq \exp(-\lambda \tfrac{\|A\|_F^2}{\|A\|_{\text{op}}} - \tfrac{1}{2} \tfrac{\|A\|_F^2}{\|A\|_{\text{op}}^2} \ln(1 - 2\lambda \|A\|_{\text{op}})). \tag{3.4}$$

Using the above bound with the Chernoff-Cramér method, we obtain Theorem 3.2. We state it in the following form for future reference.

Lemma 3.7. *Under the conditions of Theorem 3.1,*

$$\mathbb{P}(\boldsymbol{x}^\mathsf{T} A \boldsymbol{x} > t) \leq \inf_{\lambda > 0} e^{-\lambda t} \cdot \mathbb{E}[\exp(\lambda \boldsymbol{g}^\mathsf{T} A \boldsymbol{g})] \leq \exp\left(-\tfrac{1}{2}\left(\tfrac{t}{\|A\|_{\text{op}}} - \tfrac{\|A\|_F^2}{\|A\|_{\text{op}}^2} \cdot \ln(1 + t\tfrac{\|A\|_{\text{op}}}{\|A\|_F^2})\right)\right).$$

Proof. By homogeneity, it suffices to prove the claim when $\boldsymbol{x}$ is 1-subgaussian. We have that

$$\begin{aligned}
\mathbb{P}(\boldsymbol{x}^\mathsf{T} A \boldsymbol{x} > t) &\leq \inf_{\lambda > 0} \mathbb{P}\big(\exp(\lambda \boldsymbol{x}^\mathsf{T} A \boldsymbol{x}) > \exp(\lambda t)\big) \\
&\leq \inf_{\lambda > 0} \mathbb{E}[\exp(\lambda \boldsymbol{x}^\mathsf{T} A \boldsymbol{x})] \exp(-\lambda t) \\
&\leq \inf_{\lambda > 0} e^{-\lambda t} \cdot \mathbb{E}[\exp(\lambda \boldsymbol{g}^\mathsf{T} A \boldsymbol{g})] && (\text{ by Corollary 3.5 }) \\
&\leq \inf_{0 < \lambda < 1/(2\|A\|_{\text{op}})} e^{-\lambda t} \cdot \det(I - 2\lambda A)^{-1/2} \\
&\leq \inf_{0 < \lambda < 1/(2\|A\|_{\text{op}})} \exp\Big(- \lambda(t + \tfrac{\|A\|_F^2}{\|A\|_{\text{op}}}) - \tfrac{1}{2} \tfrac{\|A\|_F^2}{\|A\|_{\text{op}}^2} \ln(1 - 2\lambda \|A\|_{\text{op}})\Big) && (\text{ by (3.4) }) \\
&= \exp\left(-\tfrac{1}{2}\left(\tfrac{t}{\|A\|_{\text{op}}} - \tfrac{\|A\|_F^2}{\|A\|_{\text{op}}^2} \cdot \ln(1 + t\tfrac{\|A\|_{\text{op}}}{\|A\|_F^2})\right)\right).
\end{aligned}$$

The last equality follows by setting

$$\lambda = \frac{1}{2} \frac{t}{\|A\|_{\text{op}}(t + \|A\|_F^2 / \|A\|_{\text{op}})}, \tag{3.5}$$

which is a minimizer since the function in the exponent is convex and this choice sets the derivative to zero. □

3.1 Derandomization of Hanson-Wright

As opposed to other approaches for the Hanson-Wright inequality, our proof can be easily derandomized. For this we will restrict ourselves to the Rademacher case, where $\boldsymbol{x}$ is an n-dimensional vector of i.i.d. Rademacher random variables, so all components are 1-subgaussian. The derandomization is based on using the Gaussian mgf as a pessimistic estimator, and then exploiting the concavity of the mgf in order to fix the coordinates of $\boldsymbol{x}$ one by one, ensuring that the estimator does not increase.

For any $r = 0, \ldots, k$ and $\sigma \in \{-1, 1\}^r$, let $\mathcal{D}(\sigma)$ denote the distribution of a random vector in $\mathbb{R}^k$ whose i-th coordinate is deterministically σ_i for $i \leq r$, and an independent Rademacher for $i > r$; this describes a partial fixing of the variables $\boldsymbol{x}_i$. Define $\mathcal{G}(\sigma)$ similarly, except that for all $i > r$, the i-th coordinate is an independent standard Gaussian.

Our pessimistic estimator will be

$$\Phi(\sigma) := e^{-\lambda t} \cdot \mathbb{E}_{\boldsymbol{g} \sim \mathcal{G}(\sigma)} \exp(\lambda \boldsymbol{g}^\mathsf{T} A \boldsymbol{g}). \tag{3.6}$$

By Lemma 3.7,

$$\mathbb{P}_{\boldsymbol{x} \sim \mathcal{D}(\emptyset)}(\boldsymbol{x}^\mathsf{T} A \boldsymbol{x} > t) \leq \Phi(\emptyset),$$

and moreover if we choose λ optimally, $\Phi(\emptyset) < 1$ whenever the probability bound guaranteed by Theorem 3.2 is itself below 1.

All that remains is to confirm that Φ has the required concavity properties to be usable as a pessimistic estimator.

Lemma 3.8. *Let $0 \leq r < k$ and $\sigma \in \{-1, 1\}^r$. Let $\sigma^- = (\sigma, -1)$ and $\sigma^+ = (\sigma, 1)$. Then*

$$\Phi(\sigma) \geq \tfrac{1}{2}(\Phi(\sigma^-) + \Phi(\sigma^+)).$$

Proof. Let $\boldsymbol{y} \in \mathbb{R}^k$ be a random vector where $\boldsymbol{y}_{r+1}$ is an independent Rademacher, $\boldsymbol{y}_i = \sigma_i$ for $i \leq r$, and $\boldsymbol{y}_i$ is an independent standard Gaussian for $i > r+1$. Define $\boldsymbol{z} \in \mathbb{R}^k$ identically, except that $\boldsymbol{z}_{r+1}$ is an independent standard Gaussian. Notice that

$$\begin{aligned} \tfrac{1}{2}(\Phi(\sigma^-) + \Phi(\sigma^+)) &= e^{-\lambda t} \cdot \mathbb{E}_{\boldsymbol{y}_{r+1}}\big[\mathbb{E}_{\boldsymbol{y}_i : i \neq r+1}[e^{\lambda \boldsymbol{y}^\mathsf{T} A \boldsymbol{y}}]\big] \\ &\leq e^{-\lambda t} \cdot \mathbb{E}_{\boldsymbol{z}}[e^{\lambda \boldsymbol{z}^\mathsf{T} A \boldsymbol{z}}] \qquad \text{by Lemma 3.4} \\ &= \Phi(\sigma). \end{aligned}$$

□

The algorithm (modulo computation of the pessimistic estimator) is now straightforward. We sequentially choose signs $\sigma_1, \ldots, \sigma_k$, in such a way that at every step we choose the $\sigma_i \in \pm 1$ so that the pessimistic estimator does not increase. So we obtain $\sigma \in \{-1,1\}^k$ with $\Phi(\sigma) < 1$, which immediately implies that $\sigma^\mathsf{T} A \sigma \leq t$.

All that remains is to check that we can actually compute Φ; we postpone this to Section 5.1.

4 Derandomizing Sparse J-L via Pessimistic Estimators

Kane and Nelson [21, 23] demonstrated how sparse dimensionality reduction follows from the Hanson-Wright inequality. Here we will give a different proof of their randomized result, following the framework that they lay out, and show how this can be easily derandomized using pessimistic estimators derived naturally from the moment-generating functions that we use in our proof. The only aspect we will not yet discuss is the efficient computation of the pessimistic estimators, which will follow in Section 5.

We will assume from now on that $\|v\| = 1$ for all $v \in V$; this is clearly without loss of generality. Recall that $m \geq 12\ln(2n)/\varepsilon^2$. We will assume for simplicity that $m\varepsilon$ is an integer, so our goal is a projection matrix with $s := m\varepsilon$ nonzero entries per column. More precisely, we will use a J-L matrix with the following form:

$$\Pi = \frac{1}{\sqrt{s}}(\delta_{rj}\sigma_{rj})_{r\in[m], j\in[d]}.$$

Here, $\Delta := (\delta_{rj})_{r,j}$ is a masking 0-1 matrix, and its role is to sparsify the projection; and $\Sigma := (\sigma_{rj})_{r,j}$ is a ± 1 matrix, which (with proper scaling) can be seen as a standard (Rademacher) projection matrix. Notice that due to the masking most σ_{rj}'s will not be used at all in the final projection (nor in the algorithm), but keeping them makes the analysis cleaner.

The Kane-Nelson analysis, and ours as well, breaks into two parts. One is the choice of mask matrix Δ, which will be required to satisfy certain properties. The other is the choice of the sign matrix Σ once Δ has been determined. We discuss each of these parts separately.

4.1 The sign matrix

For the moment, let Δ be a given, deterministic 0-1 matrix; we will elucidate what properties we require from it. Let δ_r denote the r'th row of Δ, and for $v \in V$ let

$$A_v = \frac{1}{s}\bigoplus_{r=1}^{m}[B_{r,v} - \operatorname{Diag}(B_{r,v})],$$

with $B_{r,v} = (\delta_r \odot v)(\delta_r \odot v)^\mathsf{T}$. So each block of A_v is a symmetric rank-1 matrix less its diagonal.

Definition 4.1. We say that $\Delta \in \{0,1\}^{m\times d}$ is a **good mask matrix** if

(i) each column of Δ has precisely s nonzero entries; and

(ii) for each $v \in V$, the resulting A_v satisfies

$$\|A_v\|_{\mathrm{op}} \leq \frac{1}{s} \quad \text{and} \quad \|A_v\|_F^2 \leq \frac{2}{m}.$$

The rationale for this definition will become clear shortly. So suppose Δ is a good matrix. Let $\sigma \in \mathbb{R}^{md}$ be a vectorized version of Σ, obtained by concatenating its rows. Then for any input vector $v \in V$,

$$\begin{aligned}\|\Pi v\|_2^2 - 1 &= \frac{1}{s}\sum_{r=1}^{m}\Big(\sum_{j=1}^{d}\delta_{rj}\sigma_{rj}v_j\Big)^2 - 1 \;=\; \frac{1}{s}\sum_{r}\sum_{j\neq k}\delta_{rj}\delta_{rk}\sigma_{rj}\sigma_{rk}v_jv_k + \frac{1}{s}\sum_{r}\sum_{j}\delta_{rj}v_j^2 - 1\\ &= \sigma^{\mathsf{T}}A_v\sigma + \sum_{j}\Big(\frac{1}{s}\sum_{r}\delta_{rj}\Big)v_j^2 - 1\\ &= \sigma^{\mathsf{T}}A_v\sigma.\end{aligned}$$

The final step uses that Δ has precisely s nonzero entries per column, and that each $v \in V$ is a unit vector.

Let the entries of Σ be i.i.d. Rademachers. Then we obtain by Hanson-Wright (Theorem 3.2) and a union bound that

$$\begin{aligned}\mathbb{P}\big(\exists v \in V : \big|\|\Pi v\|_2^2 - 1\big| > \varepsilon\big) &\leq 2n \cdot \exp\Big(-\tfrac{1}{2}\Big(\tfrac{\varepsilon}{\|Av\|_{\text{op}}} - \tfrac{\|A_v\|_F^2}{\|A_v\|_{\text{op}}^2}\cdot\ln\big(1+\varepsilon\tfrac{\|A_v\|_{\text{op}}}{\|A_v\|_F^2}\big)\Big)\Big)\\ &\leq 2n\exp\big(-\tfrac{1}{2}\big(\varepsilon s - 2\varepsilon s\ln(1+1/2)\big)\big)\\ &< 2n\exp(-\varepsilon^2 m/12)\\ &\leq 1.\end{aligned} \tag{4.1}$$

Existence of a good sign matrix, being a positive probability event, is thus guaranteed. The rationale for the definition of a good mask matrix should now be apparent.

Derandomization

Since the argument for the existence of a good choice of Σ (given a mask matrix with the desired properties) was nothing more than an application of Hanson-Wright, the derandomization used in Section 3.1 applies. Our pessimistic estimator is simply

$$\Phi(\sigma) := \sum_{v\in V}\big(\Phi_v^+(\sigma) + \Phi_v^-(\sigma)\big),$$

where $\Phi_v^+(\sigma) = e^{-\lambda\varepsilon}\cdot\mathbb{E}_{\boldsymbol{g}\sim\mathcal{G}(\sigma)}\exp\{\lambda\boldsymbol{g}^{\mathsf{T}}A_v\boldsymbol{g}\}$ and $\Phi_v^-(\sigma) = e^{-\lambda\varepsilon}\cdot\mathbb{E}_{\boldsymbol{g}\sim\mathcal{G}(\sigma)}\exp\{-\lambda\boldsymbol{g}^{\mathsf{T}}A_v\boldsymbol{g}\}$, are pessimistic estimators for the events $\boldsymbol{x}^{\mathsf{T}}A_v\boldsymbol{x} > \varepsilon$ and $\boldsymbol{x}^{\mathsf{T}}A_v\boldsymbol{x} < -\varepsilon$, respectively. We set λ optimally according to (3.5), with $\|A\|_{\text{op}} = 1/s$ and $\|A\|_F^2 = 2/m$, which leads to $\lambda = s/6$. We have already seen in (4.1) that under the conditions of Theorem 1.2, this ensures that $\Phi(\emptyset) < 1$. By Lemma 3.8, for each $v \in V$

$$\Phi_v^+(\sigma) \geq \tfrac{1}{2}(\Phi_v^+(\sigma^-) + \Phi_v^+(\sigma^+))$$

(and similarly for Φ_v^-). Thus

$$\Phi(\sigma) \geq \tfrac{1}{2}(\Phi(\sigma^-) + \Phi(\sigma^+)),$$

and once again we can sequentially choose signs (of course, we only need to consider entries in the support of Δ) until we obtain some $\sigma \in \{-1,+1\}^{\text{supp}(\Delta)}$ for which $\Phi(\sigma) < 1$. This implies that

$$\|\Pi v\|_2^2 - 1 \;=\; \sigma^{\mathsf{T}}A_v\sigma \;\in\; [-\varepsilon, +\varepsilon] \qquad \forall v \in V.$$

In order to implement this algorithm, we need to be able to compare $\Phi(\sigma^+)$ with $\Phi(\sigma^-)$. We postpone the discussion of how to do this efficiently (in time comparable to applying the sparse projection) to Section 5.2.

4.2 The mask matrix

To show the existence of a good mask matrix, we follow a randomized block construction of Kane and Nelson [23]. They partition the rows into s parts of size $m/s = 1/\varepsilon$ (for simplicity, we will assume that $1/\varepsilon$ is an integer): so let $B_1 = \{1, 2, \ldots, 1/\varepsilon\}$, $B_2 = \{1/\varepsilon + 1, \ldots, 2/\varepsilon\}$, etc. For any column j and any q, we call the entries in column j and rows in B_q the entries of *block* (q, j). The distribution Kane and Nelson choose is that for each block, a single 1 entry is chosen uniformly at random, with each block being chosen independently. Clearly this yields a random matrix Δ that always has precisely s nonzero entries per column. The main difficulty in showing that Δ is a good mask matrix with some positive probability is bounding the Frobenius norm. Kane and Nelson do this by controlling its moments: the moments of

$$\|A_v\|_F^2 \;=\; \frac{1}{s^2}\sum_{j\neq k} \boldsymbol{q}_{jk} v_j^2 v_k^2, \tag{4.2}$$

where $\boldsymbol{q}_{jk} := \sum_{r\in[m]} \boldsymbol{\delta}_{rj}\boldsymbol{\delta}_{rk}$ is the number of entries that column j and k have in common. The difficulty of bounding the moments of $\|A_v\|_F^2$ comes precisely from the fact that the $\boldsymbol{q}_{jk}$'s are not independent. However, it is easy to see these variables are negatively associated, and they exploit this fact to bound the moments as if these variables were independent. We do this differently in order to provide a convenient pessimistic estimator.

Let us begin by bounding the operator norm of A_v; this requires nothing more than the block structure of A_v, and holds surely.

$$\begin{aligned}
\|A_v\|_{\text{op}} &= \frac{1}{s}\cdot \max_{r\in[m]} \Big\|(\delta_r\odot v)(\delta_r\odot v)^\top - \operatorname{Diag}((\delta_r\odot v)(\delta_r\odot v)^\top)\Big\|_{\text{op}}\\
&\le \frac{1}{s}\cdot \max_{r\in[m]}\max\Big\{\|(\delta_r\odot v)(\delta_r\odot v)^\top\|_{\text{op}}, \|\operatorname{Diag}((\delta_r\odot v)(\delta_r\odot v)^\top)\|_{\text{op}}\Big\}\\
&\le \frac{1}{s}\cdot \max_{r\in[m]}\max\Big\{\|\delta_r\odot v\|_2^2, \max_{j\in[d]} v_j^2\Big\}\\
&\le \frac{1}{s},
\end{aligned}$$

where the first inequality holds because $\|M-N\|_{\text{op}} \le \max\{\|M\|_{\text{op}}, \|N\|_{\text{op}}\}$ when M and N are both positive semidefinite.

Now we consider $\|A_v\|_F^2$. For any $\gamma > 0$ we have

$$\begin{aligned}
\mathbb{P}\big(\|A_v\|_F^2/2 > t\big) &= \mathbb{P}\Big(\frac{1}{2s^2}\sum_{r\in[m]}\sum_{j\neq k}\boldsymbol{\delta}_{rj}\boldsymbol{\delta}_{rk}v_j^2v_k^2 > t\Big)\\
&= \mathbb{P}\Big(\frac{1}{s^2}\sum_{r\in[m]}\sum_{j<k}\boldsymbol{\delta}_{rj}\boldsymbol{\delta}_{rk}v_j^2v_k^2 > t\Big)\\
&\le e^{-\gamma t}\cdot\mathbb{E}\Big[\exp\Big\{\frac{\gamma}{s^2}\sum_{j<k}\sum_{q=1}^{s}\big(\sum_{r\in B_q}\boldsymbol{\delta}_{rj}\boldsymbol{\delta}_{rk}\big)v_j^2v_k^2\Big\}\Big]\\
&= e^{-\gamma t}\cdot\mathbb{E}\Big[\prod_{q=1}^{s}\prod_{j<k}\exp\Big\{\frac{\gamma}{s^2}\big(\sum_{r\in B_q}\boldsymbol{\delta}_{rj}\boldsymbol{\delta}_{rk}\big)\Big\}^{v_j^2v_k^2}\Big]\\
&= e^{-\gamma t}\cdot\prod_{q=1}^{s}\mathbb{E}\Big[\prod_{j<k}\exp\Big\{\frac{\gamma}{s^2}\big(\sum_{r\in B_q}\boldsymbol{\delta}_{rj}\boldsymbol{\delta}_{rk}\big)\Big\}^{v_j^2v_k^2}\Big],
\end{aligned} \tag{4.3}$$

where in the last step we have used that by the choice of the distribution, the random variables are independent among blocks. Now we will use the generalized Hölder inequality, noting that

$$\sum_{j<k} v_j^2v_k^2 = \frac{1}{2}\sum_{j\neq k}v_j^2v_k^2 \le \frac{1}{2}(\sum_j v_j^2)^2 = \frac{1}{2}, \tag{4.4}$$

and hence

$$\mathbb{P}\big(\|A_v\|_F^2 > 2t\big) \le e^{-\gamma t}\cdot\prod_{q=1}^{s}\prod_{j<k}\mathbb{E}\Big[\exp\Big\{\frac{\gamma}{s^2}\big(\sum_{r\in B_q}\boldsymbol{\delta}_{rj}\boldsymbol{\delta}_{rk}\big)\Big\}\Big]^{v_j^2v_k^2}.$$

Choose $\gamma = s^2\ln 2$. Notice that for any $j<k\in[d]$, $q\in[s]$, $\sum_{r\in B_q}\boldsymbol{\delta}_{rj}\boldsymbol{\delta}_{rk}$ is distributed as a Bernoulli with parameter ε. Thus

$$\mathbb{E}\Big[2^{\sum_{r\in B_q}\boldsymbol{\delta}_{rj}\boldsymbol{\delta}_{rk}}\Big] = 1+\varepsilon.$$

Taking a union bound over all vectors $v \in V$,

$$\begin{aligned}\mathbb{P}\big(\exists v \in V : \|A_v\|_F^2 > 2 \cdot \tfrac{1}{m}\big) &\le \sum_{v \in V} e^{-(\ln 2)s^2/m} \cdot (1+\varepsilon)^{s \cdot \sum_{j<k} v_j^2 v_k^2} \\ &\le n \cdot e^{-(\ln 2)\varepsilon^2 m} \cdot (e^{\varepsilon})^{s/2} \qquad \text{by (4.4)} \\ &< n \exp\big(-\varepsilon^2 m/6\big) \qquad (4.5) \\ &< 1.\end{aligned}$$

Thus a good mask matrix does exist.

Derandomization

We will now derandomize this argument using pessimistic estimators. We will be interested in distributions over mask matrices Δ where some blocks have had the position of their single nonzero entry chosen, and some have not. We will describe such distribution with a vector $p \in \{0, \varepsilon, 1\}^{m \times d}$, satisfying

$$\sum_{r \in B_q} p_{rj} = 1 \qquad \text{for each} \quad q \in [s], j \in [d]. \qquad (4.6)$$

For any such p, let $\mathcal{M}(p)$ denote the distribution of a random matrix $(\boldsymbol{\delta}_{rj})_{r \in [m], j \in [d]}$ where $\mathbb{P}(\boldsymbol{\delta}_{rj} = 1) = p_{rj}$, and moreover each block is independent of the others and contains precisely a single 1. Let $p^{(0)}$ denote the distribution used in the existential analysis: $p^{(0)}_{rj} = \varepsilon$ for all $r \in [m]$, $j \in [d]$.

Our pessimistic estimator will be $\Psi(p) := \sum_{v \in V} \Psi_v(p)$, where

$$\Psi_v(p) := 2^{-s^2 t} \cdot \prod_{q=1}^{s} \prod_{j<k} \Big(\mathbb{E}_{\boldsymbol{\delta} \sim \mathcal{M}(p)} \big[2^{\sum_{r \in B_q} \boldsymbol{\delta}_{rj} \boldsymbol{\delta}_{rk}} \big] \Big)^{v_j^2 v_k^2}.$$

The derivation of (4.3) did not use anything about the distribution of $\boldsymbol{\delta}$ aside from the independence between blocks, and so the argument there, combined with a union bound implies that Ψ is indeed a pessimistic estimator:

$$\mathbb{P}_{\boldsymbol{\delta} \sim \mathcal{M}(p)}(\exists v \in V : \|A_v\|_F^2 > 2t) \le \Psi(p).$$

Moreover we have already seen that $\Psi(p^{(0)}) < 1$ under the conditions of Theorem 1.2.

Our algorithm will consider the blocks one-by-one in some order, and iteratively fix the position of the single nonzero entry in each block.

Lemma 4.2. *Let $p \in \{0, \varepsilon, 1\}^{m \times d}$ satisfy (4.6), and let $q' \in [s], j' \in [d]$ be indices so that $p_{rj'} = \varepsilon$ for $r \in B_{q'}$. Then there exists a $p' \in \{0, \varepsilon, 1\}^{m \times d}$, agreeing with p outside the block (q', j') but integral on the block (q', j'), for which $\Psi(p') \le \Psi(p)$.*

Proof. We first notice that for $\boldsymbol{\delta} \sim \mathcal{M}(p)$, and for any $q \in [s]$, $j \neq k$, the indicator for the event that there is a "collision" between the blocks (q, j) and (q, k), namely

$$\boldsymbol{c}_{jk,q} := \sum_{r \in B_q} \boldsymbol{\delta}_{rj} \boldsymbol{\delta}_{rk}$$

is distributed as a Bernoulli random variable with parameter $\rho_{jk,q} := \sum_{r \in B_q} p_{rj} p_{rk}$. This only holds because $p \in \{0, \varepsilon, 1\}^{m \times d}$; either both blocks are fixed, in which case the outcome is deterministic ($\rho_{jk,q} \in \{0, 1\}$); or at least one is uniform (possibly both), in which case collision occurs with probability ε, which equals $\rho_{jk,q}$. So

$$\mathbb{E}\big[2^{\boldsymbol{c}_{jk,q}}\big] = (1 - \rho_{jk,q}) + 2\rho_{jk,q} = 1 + \rho_{jk,q}\,.$$

Hence we may simplify the expression for Ψ:

$$\Psi(p) = 2^{-s^2 t} \sum_{v\in V} \prod_{q=1}^{s} \prod_{j<k} (1+\rho_{jk,q})^{v_j^2 v_k^2} . \tag{4.7}$$

We now consider $f : [0,1]^{B_{q'}} \to \mathbb{R}$ defined by

$$f(z) = 2^{-s^2 t} \sum_{v\in V} \prod_{(j,k,q)\in \mathcal{Q}} (1+\rho_{jk,q})^{v_j^2 v_k^2} \cdot \prod_{k\neq j'} \Big(1 + \sum_{r\in B_{q'}} z_r \cdot p_{rk}\Big)^{v_{j'}^2 v_k^2};$$

here, $\mathcal{Q}$ consists of all triplets (j,k,q) with $q\in[s]$, $j<k\in[d]$ except for those where $q=q'$ and $j'\in\{j,k\}$. Let $z_0 = (\varepsilon,\varepsilon,\ldots,\varepsilon) \in [0,1]^{B_{q'}}$; then $f(z_0) = \Psi(p)$.

Since

$$\sum_{k\neq j'} v_{j'}^2 v_k^2 = v_{j'}^2(1-v_{j'}^2) \leq 1/4, \tag{4.8}$$

each term of f is the composition of an affine function with a function of the form $y \to \prod_i y_i^{\alpha_i}$ with $\sum_i \alpha_i \leq 1$. Hence f is concave. Since the extreme points of the set $S = \{z\in[0,1]^{B_{q'}} : \sum_{r\in B_{q'}} z_r = 1\}$ are integral, it follows that there is a minimizer z' of f over S which is integral. Choosing p' to agree with z' on block (q',j'), we have

$$\Psi(p') = f(z') \leq f(z_0) = \Psi(p). \qquad \square$$

A universal masking

So far, our arguments have not required the condition $\mathrm{NNZ}(V) \geq d^2$ (or even the mild $n\geq d$). Unfortunately, we do not know how to compute, in general, the pessimistic estimator just discussed quickly enough to obtain the linear running-time required for Theorem 1.2: By direct application of Lemma 4.2 we are able to compute it in time $O(m\cdot \mathrm{NNZ}(V))$, rather than the $O(s\cdot\mathrm{NNZ}(V))$ time promised in Theorem 1.2.

We will now show that as long as $n\geq d$, we can construct a good mask matrix without even examining V. It will turn out that we can compute this matrix much faster, in time $O(sd^2)$. If $\mathrm{NNZ}(V) = \Omega(d^2)$, this is comparable to the time required to apply the final sparse projection matrix to V.

The idea, first proposed by Kane and Nelson in [23] as the code construction, is the following: Suppose $\Delta \in \{0,1\}^{m\times d}$ is a matrix, with precisely s nonzero entries per column, satisfying that for all $j\neq k$,

$$\frac{1}{s^2}\sum_r \delta_{rj}\delta_{rk} \leq \frac{2}{m}. \tag{4.9}$$

Then for an arbitrary input unit vector $v\in\mathbb{R}^d$

$$\|A_v\|_F^2 = \frac{1}{s^2}\sum_{j\neq k}\sum_{r\in[m]} \delta_{rj}\delta_{rk} v_j^2 v_k^2 \leq \frac{2}{m}(\|v\|_2^4 - \|v\|_4^4) \leq \frac{2}{m}.$$

Let $(e_j)_{j\in[d]}$ denote the canonical basis for $\mathbb{R}^d$. Then $\|A_{e_j+e_k}\|_F^2/2$ is exactly given by the left hand side of (4.9), and so it follows that we can reduce the task of controlling the Frobenius norms of A_v for $v\in V$ to controlling instead the Frobenius norms of A_v for all v in the particular set

$$\bar{V} := \{e_j + e_k : j<k\in[d]\}.$$

Since we only need that $\|A_v\|_F^2 \leq \frac{4}{m}$ for $v\in\bar{V}$, we set $t = 2/m$ (rather than $1/m$ as previously); otherwise our definition of Ψ is unchanged, yielding

$$\Psi(p) = \sum_{v\in\bar{V}} \Psi_v(p) = 2^{-2\varepsilon^2 m} \sum_{j<k} \prod_{q=1}^{s} (1+\rho_{jk,q}). \tag{4.10}$$

Then

$$\Psi(p^{(0)}) = 2^{-2\varepsilon^2 m} \sum_{j<k} (1+\varepsilon)^s < d^2 e^{(-2\ln 2+1)\varepsilon^2 m} < 1,$$

using that $d \le n$ and that $m \ge 12\ln(2n)/\varepsilon^2$. Furthermore, even though the vectors of $\bar{V}$ are not unit vectors, Lemma 4.2 still holds. This follows because for any $v \in \bar{V}$ and $j' \in [d]$,

$$\sum_{k \neq j'} v_{j'}^2 v_k^2 \le 1;$$

this suffices instead of (4.8) to obtain the desired concavity.

We may thus apply the derandomization via pessimistic estimators. The main advantage of this approach is that the extreme sparsity of the vectors in $\bar{V}$ make the pessimistic estimator computations easier. In Section 5.4, we show how a good universal mask matrix can be computed in time $O(s \cdot d^2)$.

5 Efficient computation of the pessimistic estimators

5.1 Computation for Hanson-Wright in general

Here we show how to compute the signing pessimistic estimator (3.6) in polynomial time. The method we describe here is not yet fast enough for use in the signing step for sparse J-L to obtain the running times claimed in Theorem 1.2.

Let us recall the setting of Section 3.1. The randomization of the algorithm is given by an i.i.d. vector of Rademacher random variables $\sigma \in \{-1,+1\}^k$. The derandomization procedure works with partial assignment vectors: Given $r = 0,\dots,k$ a partial assignment is $\sigma \in \{-1,+1\}^r$; in a derandomization step, our goal is to increase the partial assignment by choosing the next binary parameter $\sigma^+ = (\sigma,+1)$ or $\sigma^- = (\sigma,-1)$, in such a way that the pessimistic estimator (3.6) does not increase. The pessimistic estimator is given by the mgf of the quadratic form $\boldsymbol{g}^\top A \boldsymbol{g}$, over the distribution $\mathcal{G}(\sigma)$: the independent marginals of this distribution are $\boldsymbol{g}_i = \sigma_i$ (surely), for all $i \le r$, and $\boldsymbol{g}_i \sim \mathcal{N}(0,1)$, for all $i > r$.

The following lemma shows how we can evaluate the pessimistic estimator (3.6).

Lemma 5.1. *Let $t \in \{0,1,\dots,k\}$ and $\sigma \in \{-1,1\}^t$. Let M be a symmetric $k \times k$ matrix with $\|M\|_{\mathrm{op}} < 1/2$. Let $S = \{1,\dots,t\}$ and $T = \{t+1,\dots,k\}$. Then*

$$\mathbb{E}_{\boldsymbol{g}\sim\mathcal{G}(\sigma)} e^{\boldsymbol{g}^\top M \boldsymbol{g}} = \det(I - 2M_{T,T})^{-1/2} \exp\{2\sigma^\top M_{S,T}(I-2M_{T,T})^{-1} M_{T,S}\sigma + \sigma^\top M_{S,S}\sigma\}.$$

Proof. Let $\ell = k - t$. Consider the spectral decomposition $M_{T,T} = \sum_{i=1}^{\ell} \lambda_i u_i u_i^\top$, where $u_1,\dots,u_\ell$ form an orthonormal basis of $\mathbb{R}^\ell$, and the eigenvalues are in decreasing absolute value, $1/2 > |\lambda_1| \ge |\lambda_2| \ge \cdots \ge |\lambda_\ell|$. Define

$$b_i := \langle M_{T,S}\sigma, u_i\rangle, \qquad \text{for } i \in [\ell].$$

Also let $\boldsymbol{z}_i := \langle \boldsymbol{g}_T, u_i\rangle$ for $i \in [\ell]$; note that this is a standard Gaussian, and that $\boldsymbol{z}_i$ and $\boldsymbol{z}_j$ are independent for $i \neq j$.

Then

$$\begin{aligned}
\mathbb{E}\, e^{\boldsymbol{g}^\top M\boldsymbol{g}} &= \mathbb{E}\big[\exp\{\boldsymbol{g}_T^\top M_{T,T}\boldsymbol{g}_T + 2\boldsymbol{g}_T^\top M_{T,S}\sigma + \sigma^\top M_{S,S}\sigma\}\big] \\
&= \mathbb{E}\Big[\exp\Big\{ \textstyle\sum_{i=1}^{\ell} (\lambda_i \langle \boldsymbol{g}_T, u_i\rangle^2 + 2\langle \boldsymbol{g}_T, u_i\rangle\langle M_{T,S}\sigma, u_i\rangle) + \sigma^\top M_{S,S}\sigma \Big\}\Big] && \text{by orthogonality} \\
&= \prod_{i=1}^{\ell} \mathbb{E}\big[e^{\lambda_i \boldsymbol{z}_i^2 + 2\boldsymbol{z}_i b_i}\big] \cdot e^{\sigma^\top M_{S,S}\sigma} \\
&= \prod_{i=1}^{\ell} \frac{1}{\sqrt{1-2\lambda_i}} e^{2b_i^2/(1-2\lambda_i)} \cdot e^{\sigma^\top M_{S,S}\sigma} && \text{by Lemma 2.4} \\
&= \det(I-2M_{T,T})^{-1/2} \cdot \exp\{2\sigma^\top M_{S,T}(I-2M_{T,T})^{-1}M_{T,S}\sigma\} \cdot e^{\sigma^\top M_{S,S}\sigma}.
\end{aligned}$$

□

5.2 Computation for signing

The approach above is too slow to obtain the running time claimed in Theorem 1.2. In particular, the best known algorithms for computing the determinant run in matrix-vector multiplication time, which is superlinear. We will next show how to exploit the simple structure of the quadratic forms of interest in the case of dimensionality reduction to obtain a linear-time algorithm for initializing and updating the pessimistic estimator. We assume throughout that required computations, in particular the computation of exp, can be considered as fundamental operations that can be done in unit time. We delay a discussion of how these numerical computations can be done to Section 5.3.

Both the input vectors V and the mask matrix Δ are assumed to be given in a sparse representation, as a list of the locations and values of the nonzero entries. We have ample time to sort these lists if needed. So we assume that we can, for instance, iterate through $\{j : v_j \neq 0\}$ for a given $v \in V$ in time proportional to the size of this set; and the same for $\{v : v_j \neq 0\}$ for a given $j \in [d]$.

Recall that our goal is to compute some $\sigma \in \{-1, +1\}^{\mathrm{supp}(\Delta)}$ for which $\Phi(\sigma) = \sum_{v \in V}(\Phi_v^+(\sigma) + \Phi_v^-(\sigma)) < 1$, by beginning with $\Phi(\emptyset) < 1$ and sequentially appending signs to σ so that Φ does not increase. In order to obtain the desired running time, we will not recompute $\Phi(\sigma)$ from scratch each time, but rather efficiently update it each time we extend σ.

Recall that $A_v = \frac{1}{s}\bigoplus_{r\in[m]}[B_{r,v} - \mathrm{Diag}(B_{r,v})]$, with $B_{r,v} = (\delta_r \odot v)(\delta_r \odot v)^\mathsf{T}$. Given a partial fixing $\sigma \in \{-1,1\}^t$ for some $t \leq md$, let $\sigma[r]$ denote the components of σ that correspond with the r'th block $B_{r,v}$ of A_v; it exactly corresponds to the entries of σ that end up in the r'th row of the resulting matrix Σ. Thus $\sigma[r]$ has length 0 if $t \leq (r-1)d$, length d if $t \geq rd$, and length $t - (r-1)d$ otherwise.

Our goal is to efficiently compute the upper and lower tail deviation mgf's

$$\Phi_v^+(\sigma) = e^{-\lambda\varepsilon}\mathbb{E}_{\boldsymbol{g}\sim\mathcal{G}(\sigma)}e^{\lambda \boldsymbol{g}^\mathsf{T} A_v \boldsymbol{g}} = e^{-\lambda\varepsilon}\prod_{r=1}^{m}\mathbb{E}_{\boldsymbol{g}\sim\mathcal{G}(\sigma[r])}e^{\lambda \boldsymbol{g}^\mathsf{T}[B_{r,v}-\mathrm{Diag}(B_{r,v})]\boldsymbol{g}}, \tag{5.1}$$

and

$$\Phi_v^-(\sigma) = e^{-\lambda\varepsilon}\mathbb{E}_{\boldsymbol{g}\sim\mathcal{G}(\sigma)}e^{-\lambda \boldsymbol{g}^\mathsf{T} A_v \boldsymbol{g}} = e^{-\lambda\varepsilon}\prod_{r=1}^{m}\mathbb{E}_{\boldsymbol{g}\sim\mathcal{G}(\sigma[r])}e^{\lambda \boldsymbol{g}^\mathsf{T}[\mathrm{Diag}(B_{r,v})-B_{r,v}]\boldsymbol{g}},$$

respectively. As in the derandomization section, we use the optimal choice $\lambda = s/6$ above. An evaluation of each term in this product can be done with the aid of Lemma 5.1. But we may exploit the rank-1 structure of $B_{r,v}$, to speed up the calculations, as captured in the following lemma.

Lemma 5.2. *Let $u \in \mathbb{R}^k$ a vector such that $\|u\|_2^2 < 1/2$, and let $M = uu^\mathsf{T} - \mathrm{Diag}(uu^\mathsf{T})$; furthermore let $S \subseteq [k]$ and $T = [k] \setminus S$ be a partition of the coordinates, and let $\sigma \in \{-1,+1\}^S$.*

Then defining $\Theta^+ := 1 - \sum_{j\in T} 2u_j^2/(1+2u_j^2)$, $\Theta^- := 1 + \sum_{j\in T} 2u_j^2/(1-2u_j^2)$, and $\nu := \sum_{j\in S} u_j\sigma_j$, we have

$$\mathbb{E}_{\boldsymbol{g}\sim\mathcal{G}(\sigma)}\, e^{\boldsymbol{g}^\mathsf{T} M \boldsymbol{g}} = \Big(\Theta^+\prod_{j\in T}(1+2u_j^2)\Big)^{-1/2}\cdot\exp\Big\{\nu^2/\Theta^+ - \sum_{j\in S}u_j^2\Big\}, \tag{5.2}$$

$$\mathbb{E}_{\boldsymbol{g}\sim\mathcal{G}(\sigma)}\, e^{-\boldsymbol{g}^\mathsf{T} M \boldsymbol{g}} = \Big(\Theta^-\prod_{j\in T}(1-2u_j^2)\Big)^{-1/2}\cdot\exp\Big\{-\nu^2/\Theta^- + \sum_{j\in S}u_j^2\Big\}. \tag{5.3}$$

We delay the proof to the end of this section.

In what follows, and to avoid excessive repetition, we will use notation such as $\Phi_v^\pm$ to refer simultaneously to both upper and lower deviation mgf's (and similarly for other auxilliary quantities). In the pseudocode, this will allow us to describe the upper and lower deviation pessimistic estimator updates simultaneously. So for example, $\Theta_v^\pm \leftarrow \Theta_v^\pm \mp 1$ is shorthand for $\Theta_v^+ \leftarrow \Theta_v^+ - 1$ and $\Theta_v^- \leftarrow \Theta_v^- + 1$. However, in the following discussion we will focus on the upper deviation mgf.

The algorithm is divided in two phases. In the first phase we compute $\Phi_v^\pm(\emptyset)$ for each $v \in V$, and furthermore we maintain various intermediate quantities seen in the above lemma that are needed for this computation. In the second phase we will sequentially fix coordinates of σ to minimize the pessimistic estimator, updating the values of $\Phi_v^\pm$ and the various intermediate quantities as we go. The running time of each phase will be $O(s \cdot \mathrm{NNZ}(V))$.

Define, for any $v \in V$, $r \in [m]$ and partial signing σ,

$$\Theta^{\pm}_{v,r}(\sigma) := 1 \mp \tfrac{1}{3} \cdot \sum_{j \in [d] \setminus \operatorname{supp}(\sigma[r])} \delta_{rj} v_j^2 / (1 \pm \tfrac{1}{3} \delta_{rj} v_j^2),$$

$$\eta_{v,r} := \sum_{j \in \operatorname{supp}(\sigma[r])} \delta_{rj} v_j \sigma_j.$$

Then applying Lemma 5.2 (with $S = \operatorname{supp}(\sigma[r])$ and $u_j = \sqrt{\frac{\lambda}{s}} \delta_{rj} v_j = \frac{1}{\sqrt{6}} \delta_{rj} v_j$, for each $r \in [m]$) we obtain the explicit formula

$$\Phi^+_v(\sigma) = e^{-\lambda\varepsilon} \prod_{r=1}^{m} \left(\Big(\Theta^+_{v,r}(\sigma) \cdot \prod_{j \in [d] \setminus \operatorname{supp}(\sigma[r])} (1 + \tfrac{1}{3} \delta_{rj} v_j^2) \Big)^{-1/2} \cdot \exp\Big\{ \frac{1}{6} \Big(\eta^2_{v,r} / \Theta^+_{v,r}(\sigma) - \sum_{j \in \operatorname{supp}(\sigma[r])} \delta_{rj} v_j^2 \Big) \Big\} \right). \tag{5.4}$$

Computing $\Phi^{\pm}_v(\emptyset)$ for $v \in V$. Simplifying (5.4) for the case $\sigma = \emptyset$, we obtain

$$\Phi^+_v(\emptyset) = e^{-\lambda\varepsilon} \prod_{r=1}^{m} \big(\Theta^+_{v,r}(\emptyset) \big)^{-1/2} \cdot \prod_{(r,j): \delta_{rj} = 1} \big(1 + \tfrac{1}{3} v_j^2 \big)^{-1/2}. \tag{5.5}$$

Some care must be taken; since nm could be much larger than $O(s \cdot \mathrm{NNZ}(V))$, we do not necessarily have time to even write down all of the values $\Theta^+_{v,r}(\emptyset)$. However, at most $s \cdot \mathrm{NNZ}(V)$ of these values differ from one, since for each $v \in V$ and $j \in [d]$ with $v_j \neq 0$, there are only s choices of r for which $\delta_{rj} \neq 0$, and hence at most s values of $\Theta^+_{v,r}(\emptyset)$ that this entry contributes to.

So we keep track of a list $\mathcal{A}$ of "active" pairs (v, r) for which $\Theta^+_{v,r}(\emptyset)$ has been explicitly defined; undefined values are implicitly treated as being equal to 1. We also keep track of a value L_v for each $v \in V$ which indicates the last value of r for which $\Theta^+_{v,r}(\emptyset)$ was updated (or is zero if none have been updated). We will iterate through pairs (r, j) in the support of Δ in some order such that r is increasing (to be specific, let us say lexicographically increasing order). This ensures that $L_v \neq r$ precisely if $(v, r) \notin \mathcal{A}$. Altogether, this allows us to query whether a pair is active and add an active pair in constant time, and to iterate over all active pairs in time proportional to the number of pairs.

So, we consider in turn each pair (r, j) in the support of Δ, (in lexicographically increasing order). For each such, we consider each vector $v \in V$ for which $v_j \neq 0$. We subtract its contribution $\frac{1}{3} v_j^2 / [1 + \frac{1}{3} v_j^2]$ from $\Theta^+_{v,r}(\emptyset)$ (making (v, r) active if it is not already), and also multiply its contribution to the product $\prod_{(r,j):\delta_{rj}=1} \big(1 + \frac{1}{3} v_j^2\big)^{-1/2}$. This takes time $O(s \cdot \mathrm{NNZ}(V))$. Finally we compute the final values of $\Phi^+_v(\emptyset)$ for each v by multiplying the contribution of $(\Theta^+_{v,r}(\emptyset))^{-1/2}$ to $\Phi^+_v(\emptyset)$ for each active (v, r). We update Φ^-_v and associated quantities $\Theta^-_{v,r}(\emptyset)$ similarly; pseudocode is provided in Algorithm 1. (We remark that when activating a pair (v, r), we also initialize $\eta_{v,r}$ to zero; this will be updated in the next part of the algorithm.)

Updating $\Phi^{\pm}_v$ for each $v \in V$ and computing σ. Starting from $\Phi^{\pm}_v(\emptyset)$ for each $v \in V$, we will gradually update these value as we increase the length of our fixed sign vector σ. In order to do this, we will keep track of not just the values of $\Phi^{\pm}_v(\sigma)$, but also the values of $\Theta^{\pm}_{v,r}(\sigma)$, and the values of $\eta_{v,r}(\sigma) := \sum_{j \in \operatorname{supp}(\sigma[r])} \delta_{rj} v_j \sigma_j$. Again, there are too many values of $\eta_{v,r}(\sigma)$ to write down, but initially they are all zero, and all but at most $s \cdot \mathrm{NNZ}(V)$ will remain zero.

Given these values for our current choice of σ, suppose we wish to decide on the sign of σ_{rj} (where (r, j) is in the support of Δ). We need to evaluate $\Phi^+_v(\sigma^+)$ (as well as $\Phi^+_v(\sigma^-)$, which will be completely analogous). Only the term in the product (5.1) corresponding to r can be affected, and $\eta_{v,s}(\sigma^+) = \eta_{v,s}(\sigma)$ and $\Theta^+_{v,s}(\sigma^+) = \Theta^+_{v,s}(\sigma)$ for all $s \neq r$. Moreover,

$$\eta_{v,r}(\sigma^+) = \eta_{v,r}(\sigma) + \delta_{rj} v_j \tag{5.6}$$

$$\Theta^+_{v,r}(\sigma^+) = \Theta^+_{v,r}(\sigma) + \tfrac{1}{3} \delta_{rj} v_j^2 / (1 + \tfrac{1}{3} \delta_{rj} v_j^2). \tag{5.7}$$

Algorithm 1 Fast Deterministic Signing Phase 1: Initializing pessimistic estimator

```
Input
        $\Delta$    A good matrix (as in Definition 4.1), given in sparse form
        $V$    A set of sparse unit vectors in $\mathbb{R}^d$
Output
        $\sigma$    A sparse sign vector
        $\eta$    Projections of the input vectors, in sparse form

$\Phi_v^+ \leftarrow e^{-\lambda\varepsilon}$, $\Phi_v^- \leftarrow e^{-\lambda\varepsilon}$ for all $v \in V$
for $(r, j) \in \mathrm{supp}(\Delta)$, in lexicographically increasing order do
    for $v : v_j \neq 0$ do
        if $\Theta_{v,r}^+$ not defined then
            $\Theta_{v,r}^+ \leftarrow 1$, $\Theta_{v,r}^- \leftarrow 1$, $\eta_{v,r} \leftarrow 0$
        end if
        $\Theta_{v,r}^{\pm} \leftarrow \Theta_{v,r}^{\pm} \mp \frac{1}{3}v_j^2/[1 \pm \frac{1}{3}v_j^2]$
        $\Phi_v^{\pm} \leftarrow \Phi_v^{\pm} \cdot (1 \pm \frac{1}{3}v_j^2)^{-1/2}$
    end for
end for
for $(v, r) : \Theta_{v,r}^+$ defined do
    $\Phi_v^{\pm} \leftarrow \Phi_v^{\pm} \cdot (\Theta_{v,r}^{\pm})^{-1/2}$
end for

(Continued in Algorithm 2)
```

We can thus update these values, and hence Φ_v^+ via (5.4), in constant time. (The computations for updating Φ_v^- are similar.)

We proceed as follows. We consider pairs (r, j) in the support of Δ, again in lexicographically increasing order. For each, we determine what the updated value of $\Phi_v^{\pm}(\sigma^+)$ and $\Phi_v^{\pm}(\sigma^-)$ are, but only for v such that $v_j \neq 0$, since these are the only values affected. Choosing the sign that minimizes the sum of these values gives us a good choice for σ_{rj}. We make use of the same data structure described earlier (with a list of active pairs $\mathcal{A}$ along with pointers L_v to the current pair of interest for each $v \in V$) in order to look up and update the values $\Theta_{v,r}^{\pm}$ and $\eta_{v,r}$ in constant time. Pseudocode for the update phase of the algorithm can be found in Algorithm 2.

Let us see that this takes only $O(s \cdot \mathrm{NNZ}(V))$ time. Consider any $(r, j) \in \mathrm{supp}(\Delta)$. There are $T_j := |\{v \in V : v_j \neq 0\}|$ choices of v for which $v_j \neq 0$, and hence the algorithm spends $O(T_j)$ time choosing the sign σ_{rj}. The total running time is thus

$$\sum_{(r,j)\in\mathrm{supp}(\Delta)} O(T_j) = s \cdot \sum_{j\in[d]} O(T_j) = O(s\mathrm{NNZ}(V)).$$

Algorithm 2 Updating the pessimistic estimator and choosing the signs

```
    (Continued from Algorithm 1)
for $(r, j) \in \operatorname{supp}(\Delta)$, in lexicographically increasing order do
    for $v : v_j \neq 0$ do
        $\hat{\Theta}^{\pm}_{v,r} \leftarrow \Theta^{\pm}_{v,r} \pm \frac{1}{3}v_j^2/[1 \pm \frac{1}{3}v_j^2]$
        ▷ Determine new values of $\eta_{v,r}$ and $\Phi^{\pm}_v$ for both possible signing choices
        $\hat{\eta}^{(\xi)}_{v,r} \leftarrow \eta_{v,r} + \xi \cdot v_j$ for $\xi \in \{-1, +1\}$
        For $\xi \in \{-1, +1\}$, set
            $\hat{\Phi}^{\pm(\xi)}_v \leftarrow \Phi^{\pm}_v \cdot (\hat{\Theta}^{\pm}_{v,r}/\Theta^{\pm}_{v,r})^{-1/2}\Big(1 \pm \frac{1}{3}v_j^2\Big)^{1/2} \exp\Big\{\frac{1}{6}\Big(\mp\eta^2_{v,r}/\Theta^{\pm}_{v,r} \pm (\hat{\eta}^{(\xi)}_{v,r})^2/\hat{\Theta}^{\pm}_{v,r} \mp v_j^2\Big)\Big\}$    (5.8)
    end for
    Let $\sigma_{rj} \in \{-1, +1\}$ be a minimizer of $\sum_{v:v_j \neq 0}\big(\hat{\Phi}_v^{+(\sigma_{rj})} + \hat{\Phi}_v^{-(\sigma_{rj})}\big)$
    Add $\sigma_{rj}$ to the sparse vector $\sigma$
    for $v : v_j \neq 0$ do
        $\Theta^{\pm}_{v,r} \leftarrow \hat{\Theta}^{\pm}_{v,r}$,  $\eta_{v,r} \leftarrow \hat{\eta}^{(\sigma_{rj})}_{v,r}$,  $\Phi^{\pm}_v \leftarrow \hat{\Phi}^{\pm(\sigma_{rj})}_v$
    end for
end for
$\eta_v \leftarrow \eta_v/\sqrt{s}$ for all $v \in V$            ▷ Rescale $\eta$ to represent the projected vectors
return $\sigma$, $\eta$
```

Proof of Lemma 5.2. Let us first note that $\|M\|_{\text{op}} < 1/2$, as required by Lemma 5.1. Since both uu^{T} and $\operatorname{Diag}(uu^{\mathsf{T}})$ are PSD matrices, we can upper bound the operator norm by the maximum of the respective operator norms:

$$\begin{aligned}\|M\|_{\text{op}} &\leq \max\{\|uu^{\mathsf{T}}\|_{\text{op}}, \|\operatorname{Diag}(uu^{\mathsf{T}})\|_{\text{op}}\}\\ &\leq \max\{\|u\|_2^2, \max\{u_j^2 : j \in [k]\}\}\\ &< 1/2,\end{aligned}$$

as required.

Using Lemma 5.1, we have that

$$\mathbb{E}_{\boldsymbol{g} \sim \mathcal{G}(\sigma)}\, e^{\boldsymbol{g}^{\mathsf{T}} M \boldsymbol{g}} = \det(I - 2M_{T,T})^{-1/2} \exp\big\{2\sigma^{\mathsf{T}} M_{S,T}(I - 2M_{T,T})^{-1} M_{T,S}\sigma + \sigma^{\mathsf{T}} M_{S,S}\sigma\big\}. \tag{5.9}$$

Let $D := \operatorname{diag}(u_T u_T^{\mathsf{T}}) \in \mathbb{R}^{T \times T}$. By the matrix determinant lemma,

$$\begin{aligned}\det(I - 2M_{T,T}) &= \det(I + 2D) \cdot (1 - 2u_T^{\mathsf{T}}(I + 2D)^{-1}u_T)\\ &= \prod_{j \in T}(1 + 2u_j^2) \cdot \Big(1 - \sum_{j \in T}\frac{2u_j^2}{1 + 2u_j^2}\Big)\\ &= \prod_{j \in T}(1 + 2u_j^2) \cdot \Theta^{+}.\end{aligned}$$

By the Sherman-Morrison formula,

$$(I - 2M_{T,T})^{-1} = (I + 2D)^{-1} + 2\frac{(I + 2D)^{-1}u_T u_T^{\mathsf{T}}(I + 2D)^{-1}}{1 - 2u_T^{\mathsf{T}}(I + 2D)^{-1}u_T}.$$

On the other hand, since $S \cap T = \emptyset$, $M_{S,T} = u_S u_T^\mathsf{T}$. Thus,

$$\begin{aligned}
2\sigma^\mathsf{T} M_{S,T}(I - 2M_{T,T})^{-1} M_{T,S}\sigma &= 2\langle u_S, \sigma\rangle^2 \cdot u_T^\mathsf{T}(I - 2M_{T,T})^{-1} u_T \\
&= 2\langle u_S, \sigma\rangle^2 \cdot \Big(u_T^\mathsf{T}(I + 2D)^{-1} u_T + 2\frac{(u_T^\mathsf{T}(I+2D)^{-1}u_T)^2}{1 - 2u_T^\mathsf{T}(I+2D)^{-1}u_T} \Big) \\
&= 2\langle u_S, \sigma\rangle^2 \cdot \Big(\frac{u_T^\mathsf{T}(I+2D)^{-1}u_T}{1 - 2u_T^\mathsf{T}(I+2D)^{-1}u_T} \Big) \\
&= \nu^2(1 - \Theta^+)/\Theta^+.
\end{aligned}$$

Finally, we have

$$\begin{aligned}
\sigma^\mathsf{T} M_{S,S}\sigma &= \langle u_S, \sigma\rangle^2 - \|D\sigma_S\|_2^2 \\
&= \Big(\sum_{j\in S} u_j \sigma_j \Big)^2 - \sum_{j\in S} u_j^2 \\
&= \nu^2 - \sum_{j\in S} u_j^2.
\end{aligned}$$

Replacing the relevant terms in (5.9) yields (5.2). The computations for (5.3) are similar, and we omit the details. □

5.3 Numerical issues in signing

We have been working in a real model of computation where in addition, computation of square roots and exponentials takes unit time. In this section, we show that the algorithm can be implemented so that its bit complexity matches the projection time, up to logarithmic factors. It is only the computation of the signing that presents difficulties; the construction of the universal masking requires only basic arithmetic operations. For convenience, we assume in this section that $m = 12\ln(2n)/\varepsilon^2$ (rather than this value being merely a lowerbound on m).

We will need the following results due to Brent [8]. We use $M(N)$ to denote the time required to multiply two N digit integers; certainly $M(N) = O(N^2)$, with better bounds possible.

Proposition 5.1 ([8]).

- *Given a fixed interval $[a, b]$ (with $a, b > 0$), $\sqrt{x}$ can be evaluated for all $x \in [a, b]$ to a relative error of 2^{-N} in time $O(M(N))$.*
- *Given a fixed interval $[a, b]$, $\exp(x)$ can be evaluted for all $x \in [a, b]$ to a relative error of 2^{-N} in time $O(M(N)\log N)$. (For $\ln(x)$, this assumes $a > 0$.)*
- *Given a fixed interval $[a, b]$ (with $a, b > 0$), $\ln(x)$ can be evaluted for all $x \in [a, b]$ to an additive error of 2^{-N} in time $O(M(N)\log N)$.*

In all cases, only $N + O(1)$ bits of precision of the argument x are needed.

While these results require a fixed interval, standard tricks can be applied to work around this. In particular, we have the following.

Corollary 5.3. *Given R and N, $\exp(x)$ can be computed to a relative error of 2^{-N} for $x \in [-2^R, 2^R]$ in time $O(M(N + R)\log(N + R))$. The argument x need only be accurate within an additive error of $O(2^{-N})$.*

Proof. We use the identity $e^x = 2^k e^{x - k\ln 2}$, with $k \in \mathbb{Z}$ such that $|x - k\ln 2| \le 2$, to reduce the computation of exp to the fixed interval $[-2, 2]$. Given that $x \in [-R, R]$, and to obtain relative error 2^{-N}, we need $N + O(1)$ bits of precision for $x - k\ln 2$, and hence $N + R + O(1)$ bits of precision for $\ln 2$. By Proposition 5.1, this can be done in time $O(M(N + R)\log(N + R))$. □

We have been working under the assumption that our input vectors have been normalized: $\|v\|_2 = 1$ for all $v \in V$. Exact normalization is clearly too strong to deal with numerically. However, we will see that $O(\log(nd))$ bits of precision for our input vectors suffice.

First, we observe that the conditions of Theorem 1.2 ensure that the initial value $\Phi(\emptyset)$ of the pessimistic estimator is not just strictly below 1, but really bounded away from 1. More precisely, considering (4.1), we see that

$$\Phi(\emptyset) \leq (2n)^{-\frac{1}{2}(1-2\ln(\frac{3}{2}))+1/12} < 0.993.$$

We are satisfied with any final choice of $\sigma \in \{-1,+1\}^{\mathrm{supp}(\Delta)}$ for which $\Phi(\sigma) < 1$. It thus suffices if we ensure that for the σ we choose, our approximation of $\Phi(\sigma)$ is below $1 - 10^{-3}$ and has absolute error less than 10^{-3}. Now $\Phi(\sigma)$ is a sum of $2n$ terms; to maintain the required accuracy for $\Phi(\sigma)$, it suffices to approximate the individual terms $\Phi_v^{\pm}$ to within $1/(10^4 n)$.

So consider any single term $\Phi_v^+(\cdot)$ ($\Phi_v^-(\cdot)$ would be similar). We initially compute $\Phi_v^+(\emptyset)$ as a product of $O(s\mathrm{NNZ}(v))$ terms. Then each time we choose a sign σ_{rj} which impacts $\Phi_v^+(\sigma)$, which happens at most $s\mathrm{NNZ}(v)$ times, we multiply it by a constant number of terms. Since $\Phi^+(v)$ is certainly bounded by 1, it suffices that each term we multiply by is accurate to within a relative error of $c'/(nsd)$, for some appropriate c'.

We now show that the values we need to keep track of lie in a reasonable range—neither too large, nor too close to zero. First, $\Phi(\sigma)$, and hence $\Phi_v^{\pm}(\sigma)$ for each $v \in V$, is below 1. Secondly, we have for any $v \in V$ that

$$\begin{aligned}
\Phi_v^+(\sigma) &= e^{-\lambda\varepsilon} \cdot \mathbb{E}_{\boldsymbol{g}\sim\mathcal{G}(\sigma)}\big(\exp\{-\lambda \boldsymbol{g}^\top A_v \boldsymbol{g}\}\big)^{-1} \\
&\geq e^{-\lambda\varepsilon} \cdot \big(\mathbb{E}_{\boldsymbol{g}\sim\mathcal{G}(\sigma)} \exp\{-\lambda \boldsymbol{g}^\top A_v \boldsymbol{g}\}\big)^{-1} \qquad \text{(By Jensen's inequality)} \\
&= e^{-2\lambda\varepsilon}\big(\Phi_v^-(\sigma)\big)^{-1} \\
&\geq 1/(12n)^2.
\end{aligned}$$

Similarly, $\Phi_v^-(\sigma) \geq 1/(12n)^2$. This also implies that if σ and $\hat{\sigma}$ are the partial signings before and after a particular sign is chosen, the ratio between $\Phi_v^{\pm}(\sigma)$ and $\Phi_v^{\pm}(\hat{\sigma})$ is in the interval $[1/(12n)^4, (12n)^4]$. Furthermore, since the input vectors have unit norm, it can be seen from the definition of $\Theta_{v,r}^{\pm}$ that $\frac{2}{3} \leq \Theta_{v,r}^{\pm} \leq \frac{4}{3}$.

It follows that we only need to keep track of $O(\log(nsd)) = O(\log(nd))$ bits of $\Phi_v^{\pm}$ and also of all the individual terms that appear in Algorithm 1 and Algorithm 2. In particular, the absolute value of the argument of the exponential term in (5.8), for the correct signing choice, is $O(\log n)$ (and since $\eta_{v,r}^2 = O(d)$ by the Cauchy-Schwartz inequality, this argument can be computed efficiently). Proposition 5.1 and Corollary 5.3 suffice to do all calculations to sufficient precision, using only $O(\log(nd))$ bits of precision for the input vectors, and with each square root computation requiring time $O(M(\log(nd)))$ and each exponent computation requiring time $O(M(\log(nd))\log(nd))$.

Remark 5.4. It can easily be seen that the form of the projection matrix implies that it is a contraction from $\|\cdot\|_1$ to $\|\cdot\|_2$. This implies that with a small loss in the projection dimension m, only $O(\log d)$ bits of precision of the input vectors is required.

5.4 Computation for masking

Here we provide the details of the efficient computation of a universal mask matrix associated with the vectors $\bar{V} = \{e_j + e_k : j < k\}$, in time $O(s \cdot d^2)$. We provide detailed pseudocode of the procedure in Algorithm 3.

We may assume that $1/\varepsilon = O(d)$; otherwise we have a trivial construction: choose a matrix with at most one nonzero entry per row.

In the case of $\bar{V}$, our pessimistic estimator is as given in (4.10):

$$\Psi(p) = 2^{-2\varepsilon^2 m} \cdot \sum_{j<k} \prod_{q=1}^{s} (1 + \rho_{jk,q}),$$

where recall $\rho_{jk,q} := \sum_{r\in B_q} p_{rj} p_{rk}$. We consider a partial assignment vector p: namely, there exists a block (q, ℓ) such that p is integral for all blocks $(\tilde{q}, j)$ where either $j < \ell$, or $j = \ell$ and $\tilde{q} < q$; and p is uniform on all other blocks.

We are interested in derandomizing block (q, ℓ). We assume that p is stored so that for any integral block $(\tilde{q}, j)$, we can determine the choice of $r \in B_{\tilde{q}}$ such that $p_{rj} = 1$ in constant time.

Given $r \in B_q$, let $p^{(r)}$ be a vector identical to p in all blocks, except for (q, ℓ), where it has an integral assignment choosing $p_{r\ell} = 1$. Then

$$\Psi(p^{(r)}) - \Psi(p) = 2^{-2\varepsilon^2 m} \sum_{j<\ell} \prod_{\tilde{q} \neq q} (1 + \rho_{j\ell,\tilde{q}}) \Big(1 + p_{rj} - (1 + \varepsilon)\Big).$$

Now let $D_j^q := \prod_{\tilde{q}<q}(1 + \rho_{j\ell,\tilde{q}})$; notice that $D_j^q = 2^{\sum_{\tilde{q}<q} \sum_{r \in B_{\tilde{q}}} p_{rj} p_{rt}}$, so it keeps a count of the collisions among fixed variables within columns j and t. Noticing that

$$\prod_{\tilde{q} \neq q} (1 + \rho_{j\ell,\tilde{q}}) = D_j^q \cdot (1 + \varepsilon)^{s-q},$$

we obtain

$$\Psi(p^{(r)}) - \Psi(p) = 2^{-2\varepsilon^2 m} \big(1 + \varepsilon\big)^{s-q} \Big[\sum_{j<\ell} D_j^q \cdot p_{rj} - \varepsilon \sum_{j<\ell} D_j^q \Big].$$

Since the last term in this equation is independent of r, we conclude that in order to minimize the pessimistic estimator it suffices to choose $r \in B_q$ to minimize

$$\alpha^{(r)} := \sum_{j<\ell} D_j^q \cdot p_{rj}.$$

With this, the algorithm is straightforward. We proceed column by column; suppose we have fixed p for all columns $j < \ell$, but not yet for column ℓ. We then consider each block (q, ℓ) in turn, in increasing order of q. For each block, we compute $\alpha^{(r)}$ for all $r \in B_q$; since each D_j^q contributes to precisely one $\alpha^{(r)}$, this requires $O(d)$ computations. Now choose $\bar{r} \in B_q$ that minimizes $\alpha^{(r)}$ (which requires $O(1/\varepsilon) = O(d)$ computations), and set $p_{\bar{r}\ell} = 1$.

In order to move onto the next block, we need the values of D_j^{q+1} for all $j < \ell$. We can compute these values immediately from the values of D_j^q, as follows. If $p_{\bar{r}j} = 1$, meaning that $\rho_{j\ell,q} = 1$, then $D_j^{q+1} = 2D_j^q$; otherwise, $\rho_{j\ell,q} = 0$, and so $D_j^{q+1} = D_j^q$.

In conclusion, to compute all entries of Δ in column ℓ, we consider s blocks and require $O(d)$ time per block. The total running time of the algorithm is thus $O(s \cdot d^2)$ (which is $O(s \cdot \mathrm{NNZ}(V))$ given the assumption that $\mathrm{NNZ}(V) = \Omega(d^2)$).

Algorithm 3 Universal Masking Algorithm

```
Input
        $d$    Input dimension
Output
        $\Delta$    Good matrix for $\bar{V}$, given in sparse form

Initialize $\Delta = (\delta_{rj})_{r\in[m],j\in[d]}$ to be empty (a sparse matrix with all entries zero)
for $\ell = 1, \ldots, d$ do
    $D_j^1 \leftarrow 1$ for all $j < \ell$
    for $q = 1, 2, \ldots, s$ do
        $\alpha^{(r)} \leftarrow 0$ for all $r \in B_q$
        for $j < \ell$ do
            Choose $r \in B_q$ s.t. $\delta_{rj} = 1$
            $\alpha^{(r)} \leftarrow \alpha^{(r)} + D_j^q$
        end for
        $\bar{r} = \arg\min_{r\in B_q} \alpha^{(r)}$
        $\delta_{\bar{r},\ell} \leftarrow 1$

        ▷ Compute $D_j^{q+1}$
        for $j < \ell$ do
            if $\delta_{\bar{r}j} = 1$ then
                $D_j^{q+1} \leftarrow 2D_j^q$
            else
                $D_j^{q+1} \leftarrow D_j^q$
            end if
        end for
    end for
end for
return $\Delta = (\delta_{rj})_{r\in[m],j\in[d]}$, in sparse form
```

6 A Sharp Hanson-Wright Inequality

In the following sections, we prove the technical estimate from Lemma 3.6 and give applications of our functional form of the Hanson-Wright inequality to get an essentially optimal Hanson-Wright constant and to get improved concentration bounds for Rademacher projections.

6.1 The Determinant Lower Bound

In this section, we give the proof of determinant lower bound in Lemma 3.6. To begin, we show that this problem is in fact an eigenvalue optimization problem.

Lemma 6.1. *For $k \in \mathbb{N}$, $0 \le \alpha < 1$, $\beta \ge 1$,*

$$\min\Big\{\det(I - M) : M \in \mathbb{R}^{k\times k} \textit{ symmetric}, M_{ii} = 0\ \forall i \in [k], \|M\|_{\mathrm{op}} \le \alpha, \|M\|_F \le \beta\Big\}$$

$$= \min\Big\{\prod_{i=1}^{k}(1 - \lambda_i) : |\lambda_i| \le \alpha, \forall i \in [k], \sum_{i=1}^{k}\lambda_i^2 \le \beta^2, \sum_{i=1}^{k}\lambda_i = 0\Big\}.$$

Proof. Firstly, we note that the minimum is clearly achieved for both problems since the feasible regions are compact and the objective functions are continuous.

We now show that a solution for one problem can be converted to a solution for the other of same value. Starting with a $k \times k$ diagonal 0 matrix M as above, we claim that the eigenvalues $\lambda_1, \dots, \lambda_k$ form a solution for the other side. This follows since $\sum_{i=1}^k \lambda_i = \mathrm{trace}(M) = \sum_{i=1}^k M_{ii} = 0$, $\sum_{i=1}^k \lambda_i^2 = \|M\|_F^2 \leq \beta$, $\max_{i \in [k]} |\lambda_i| = \|M\|_{\mathrm{op}}$ and $\det(I - M) = \prod_{i=1}^k (1 - \lambda_i)$.

For the other direction, starting from $\lambda_1, \dots, \lambda_k$ and given the above identities, we must only show that there exists a symmetric diagonal 0 matrix $M \in \mathbb{R}^{k \times k}$ having $\lambda_1, \dots, \lambda_k$ as its eigenvalues. By the Schur-Horn theorem [18, Theorem 5], the set of achievable diagonals for the set of symmetric matrices with eigenvalues $\lambda_1, \dots, \lambda_k$ corresponds exactly to the permutahedron

$$\text{convex.hull}((\lambda_{\pi[1]}, \dots, \lambda_{\pi[k]}) : \pi : [k] \to [k] \text{ a permutation}).$$

Since $\sum_{i=1}^k \lambda_i = 0$, by averaging over all permutations, we see that the vector $(0, \dots, 0)$ is an achievable diagonal, as needed. □

Though the most natural approach would be to find the exact minimizer for any fixed dimension k, we are currently only able to get a sharp bound for the minimum value when we allow $k \to \infty$. This allows us to use a natural continuous relaxation, which is sharp when α^2 evenly divides β^2, where we can show that the optimal solution is supported on at most 2 distinct eigenvalues. This reduces the problem to one on two variables, which we can solve directly.

Proof of Lemma 3.6. By Lemma 6.1 and taking logs, the problem reduces to showing that

$$\begin{aligned} \inf\Big\{\sum_{i=1}^k \ln(1-\lambda_i) :\ k \in \mathbb{N}, |x_i| \leq \alpha \ \ \forall i \in [k], \sum_{i=1}^k \lambda_i = 0, \sum_{i=1}^k \lambda_i^2 \leq \beta^2 \ \ \forall i \in [k]\Big\} \\ \geq (\beta^2/\alpha^2)\ln(1-\alpha) + \beta^2/\alpha. \end{aligned} \tag{6.1}$$

Since we can clearly drop any zero λ_is without affecting the constraints or objective, we may assume that all the λ_is are non-zero. By modelling multiplicities and allowing fractionality, we see that program (6.1) has value at least

$$\begin{aligned} \inf\Big\{\sum_{i=1}^k l_i \ln(1-\lambda_i) :& k \in \mathbb{N}, l_i \geq 0 \ \forall i \in [k], \lambda_1, \dots, \lambda_k \text{ distinct non-zero }, \\ & |\lambda_i| \leq \alpha \ \ \forall i \in [k], \ \sum_{i=1}^k l_i\lambda_i = 0, \ \sum_{i=1}^k l_i \lambda_i^2 \leq \beta^2\Big\}. \end{aligned} \tag{6.2}$$

For the above relaxed formulation, we claim that we can assume that $k = 2$. Since the eigenvalue sum is 0, we need at least one positive and one negative eigenvalue, so we must clearly have $k \geq 2$. To show $k \leq 2$, fix any distinct non-zero $(\lambda_1, \dots, \lambda_k)$ which can be extended to a feasible solution. Note that this reduces the problem to a linear program over $l_1, \dots, l_k$. Given that the LP is feasible by assumption and bounded (since all the $\lambda_i \neq 0$), the optimal value can be obtained at a basic feasible solution. Since there are only 2 non-trivial constraints on $l_1, \dots, l_k$ apart from non-negativity, any basic feasible solution can have support on at most two of $l_1, \dots, l_k$, as needed.

Given the above, program (6.2) reduces to

$$\begin{aligned} \inf\{l_1 \ln(1-\lambda_1) + l_2 \ln(1-\lambda_2) :& \lambda_1 \in (0,\alpha), \ \lambda_2 \in (-\alpha, 0), \ l_1, \ l_2 \geq 0, \\ & l_1\lambda_1 + l_2\lambda_2 = 0, \ l_1\lambda_1^2 + l_2\lambda_2^2 \leq \beta^2\}. \end{aligned} \tag{6.3}$$

Since $l_1 \ln(1-\lambda_1) + l_2 \ln(1-\lambda_2) \leq -l_1\lambda_1 - l_2\lambda_2 = 0$, we can always achieve $l_1\lambda_1^2 + l_2\lambda_2^2 = \beta^2$ by scaling without increasing the objective value. Given this, for feasible λ_1, λ_2 above, it is easy to check that setting

$$l_1 = \frac{\beta^2}{\lambda_1(\lambda_1 - \lambda_2)} \quad \text{and} \quad l_2 = \frac{\beta^2}{\lambda_2(\lambda_2 - \lambda_1)}$$

is both feasible and optimal. Thus, the program (6.3) reduces further to

$$\beta^2 \inf\Big\{\frac{1}{\lambda_1 - \lambda_2}\Big(\frac{\ln(1-\lambda_1)}{\lambda_1} - \frac{\ln(1-\lambda_2)}{\lambda_2}\Big) : \lambda_1 \in (0,\alpha), \lambda_2 \in (-\alpha, 0)\Big\}. \tag{6.4}$$

The following claim shows that the objective function is monotonically decreasing in λ_1 and λ_2 over its domain.

Claim 6.2. *The function*

$$g(x,y) := \frac{1}{x-y}\left(\frac{\ln(1-x)}{x} - \frac{\ln(1-y)}{y}\right)$$

is decreasing in x and y on the domain $(0,1)\times(-1,0)$.

Proof. We will show the integral description

$$g(x,y) = -\int_0^1 z \cdot \frac{1}{1-zx} \cdot \frac{1}{1-zy} dz.$$

The claim is then immediate, since for any fixed $z \in [0,1]$, the integrand is clearly increasing on the domain.

Expanding $\ln(1-x)$ and $\ln(1-y)$, we have

$$\begin{aligned} g(x,y) &= \frac{1}{x-y}\sum_{k\geq 0}\frac{(-x^k+y^k)}{k+1} \\ &= -\sum_{k\geq 0}\frac{1}{k+1}\sum_{i=0}^{k-1} x^i y^{k-i-1} \\ &= -\sum_{i\geq 0}\sum_{k\geq i+1}\frac{1}{k+1}x^i y^{k-i-1} \\ &= -\sum_{i\geq 0}\sum_{j\geq 0}\frac{1}{i+j+2}x^i y^j. \end{aligned}$$

Now we use the identity $\frac{1}{i+j+2} = \int_0^1 z^{i+j+1}dz$. We obtain (using the dominated convergence theorem to justify the exchange of integral and summation)

$$\begin{aligned} g(x,y) &= -\int_0^1 z \cdot \sum_{i\geq 0}\sum_{j\geq 0}(zx)^i (zy)^j dz \\ &= -\int_0^1 z \cdot \frac{1}{1-zx}\frac{1}{1-zy}dz, \end{aligned}$$

as required. □

Given the claim, the infimum of (6.4) is thus obtained by the limit

$$\lim_{\lambda_1\to\alpha,\lambda_2\to 0}\frac{\beta^2}{\lambda_1-\lambda_2}\Big(\frac{\ln(1-\lambda_1)}{\lambda_1} - \frac{\ln(1-\lambda_2)}{\lambda_2}\Big) = \frac{\beta^2}{\alpha^2}\ln(1-\alpha) + \beta^2/\alpha, \tag{6.5}$$

as needed.

We now show that this bound is sharp for the original program (6.1) as long as α^2 divides β^2. In particular, it suffices to construct a sequence of solutions $\lambda_{1,t}, \lambda_{2,t}$ to (6.4) such that $\lambda_{1,t} \to \alpha$ and $\lambda_{2,t} \to 0$ as $t \to \infty$, for which the corresponding multiplicities $l_{1,t} = \frac{\beta^2}{\lambda_{1,t}(\lambda_{1,t}-\lambda_{2,t})}$ and $l_{2,t} = \frac{\beta^2}{\lambda_{2,t}(\lambda_{2,t}-\lambda_{1,t})}$ are positive integers. For this purpose, take the sequence $\lambda_{1,t} = \alpha\sqrt{\frac{t}{t+1}}$ and $\lambda_{2,t} = -\alpha\sqrt{\frac{t}{t+1}}\frac{1}{t}$. A direct computation reveals that $l_{1,t} = \frac{\beta^2}{\alpha^2}$ and $l_{2,t} = \frac{\beta^2}{\alpha^2}t$, which are integral for $t \in \mathbb{N}$ by the assumption that α^2 divides β^2. Since this sequence clearly converges to the desired limit, this proves the claim.

□

6.2 The Hanson-Wright Constant

The following lemma uses the functional form of Theorem 3.2 to derive optimal tradeoffs between the linear and quadratic behavior of the Hanson Wright inequality.

Lemma 6.3. *Under the conditions of Theorem 3.1, for any $\gamma > 0$, we have that*

$$\mathbb{P}(\boldsymbol{x}^\mathsf{T} A \boldsymbol{x} > t\nu^2) \le \exp\left(-\tfrac{1-\ln(1+\gamma)/\gamma}{2\gamma} \min\left\{\frac{t^2}{\|A\|_F^2}, \frac{\gamma t}{\|A\|_{\mathrm{op}}}\right\}\right). \tag{6.6}$$

Proof. Letting $\alpha := \|A\|_{\mathrm{op}}$, $\beta := \|A\|_F$, by Theorem 3.2 we recall that

$$\mathbb{P}(\boldsymbol{x}^\mathsf{T} A \boldsymbol{x} > t\nu^2) \le e^{-\frac{1}{2}f(t)} \text{ where } f(t) = \frac{t}{\alpha} - \frac{\beta^2}{\alpha^2}\ln(1 + t\frac{\alpha}{\beta^2}).$$

To prove the lemma, it thus suffices to show that

$$f(t) \ge \frac{1 - \ln(1+\gamma)/\gamma}{\gamma} \min\left\{\frac{t^2}{\beta^2}, \frac{\gamma t}{\alpha}\right\}. \tag{6.7}$$

For $t = \gamma\beta^2/\alpha$, equality is attained between $\frac{t^2}{\beta^2}$ and $\frac{\gamma t}{\alpha}$, and it can be verified that they both equal $f(t)$. All that remains is to show that $\frac{\alpha}{\gamma t}f(t)$ is increasing and $\frac{\beta^2}{t^2}f(t)$ is decreasing. This can be done by direct computation of derivatives. □

Using the above lemma, we derive the optimal Hanson-Wright constant C_{HW} for Theorem 3.1. In particular, this corresponds to setting $\gamma = 1$ above, which gives $C_{\mathrm{HW}} = \frac{1-\ln 2}{2} \approx 0.153 > 3/20$.

As for the different tradeoffs, letting $\gamma \to \infty$ improves the bound for small t (the quadratic regime), where the constant in front the $t^2/\|A\|_F^2$ terms converges to $1/4$. Letting $\gamma \to 0$ improves the bound for large t (the linear regime), where the constant in front of $t/\|A\|_{\mathrm{op}}$ converges to $1/2$.

6.3 Tighter Concentration for Rademacher Projections

In this section, we prove Lemma 3.3, which gives sparsity dependent concentration bound for Rademacher projections.

We begin with the following lemma, which shows that our functional form of Hanson-Wright has good monotonicity properties.

Lemma 6.4. *The probability bound in Theorem 3.2 is monotone increasing in $\|A\|_{\mathrm{op}}$ and $\|A\|_F$.*

Proof. Letting $\alpha := \|A\|_{\mathrm{op}}$ and $\beta := \|A\|_F$, we must show that

$$\exp\left(-\tfrac{1}{2}\left(\tfrac{t}{\alpha} - \tfrac{\beta^2}{\alpha^2} \cdot \ln(1 + t\tfrac{\alpha}{\beta^2})\right)\right)$$

is monotone increasing in α and β. Equivalently, it suffices to show that the function $\frac{t}{\alpha} - \frac{\beta^2}{\alpha^2}\ln(1 + t\frac{\alpha}{\beta^2})$ in the exponent is monotone decreasing in α and β. To prove this for α, replacing t by $t\beta^2$, this is equivalent to showing that $\frac{t}{\alpha} - \frac{1}{\alpha^2}\ln(1 + t\alpha)$ is decreasing. Taking a derivative with respect to α, we require that

$$-\frac{t}{\alpha^2} + 2\frac{1}{\alpha^3}\ln(1 + t\alpha) - \frac{t}{\alpha^2(1 + t\alpha)} \le 0,$$

or equivalently,

$$\ln(1 + t\alpha) \le \frac{t\alpha}{2} + \frac{t\alpha}{2(1 + t\alpha)},$$

which follows from Lemma 2.2 part 1 applied to $x = t\alpha$. To prove it for β, it suffices to show that $\frac{\beta^2}{\alpha}\ln(1 + t\frac{\alpha}{\beta^2})$ is increasing in β. After replacing β^2 by $\beta t\alpha$, this is equivalent to showing that $\beta\ln(1 + 1/\beta)$ is increasing, which follows from Lemma 2.2 part 3. □

The following technical lemma bounds the matrix norms required to analyze Rademacher projections.

Lemma 6.5. *Let $v \in \mathbb{R}^d$ be a unit vector with s non-zero entries. Then for $H = vv^\mathsf{T} - \mathrm{Diag}(vv^\mathsf{T})$, we have that $\|H\|_{\mathrm{op}} \leq 1 - 1/s$ and $\|H\|_F^2 = \|v\|_2^4 - \|v\|_4^4 \leq 1 - 1/s$.*

Proof. We may assume that $s = d$ without loss of generality, by restricting to the support of v. The claim is trivial for $d = 1$, so assume $d \geq 2$.

The bound on $\|H\|_F^2$ is immediate from $\|v\|_2 = 1$ and $\|v\|_4^4 \geq \frac{1}{d}\|v\|_2^4$.

To prove that $\|H\|_{\mathrm{op}} \leq 1 - 1/d$, it suffices to show that for all $w \in \mathbb{R}^d$,

$$-\tfrac{1}{2} \cdot \|w\|_2^2 \cdot \|v\|_2^2 \leq w^\mathsf{T}\Big(vv^\mathsf{T} - \mathrm{Diag}(vv^\mathsf{T})\Big)w \leq (1 - 1/d) \cdot \|w\|_2^2 \cdot \|v\|_2^2,$$

where we have introduced a term $\|v\|_2^2$ where appropriate to make the inequalities insensitive to the magnitude of v.

We begin with the bound on the maximum eigenvalue. Rewriting our desired inequality, we wish to show that

$$U := (d-1)\sum_{i=1}^{d}\sum_{j=1}^{d} v_i^2 w_j^2 - d\Big(\sum_{i=1}^{d} v_i w_i\Big)^2 + d\sum_{i=1}^{d} v_i^2 w_i^2.$$

is nonnegative. This follows by noting that U may be rewritten as a sum of squares:

$$U = \frac{1}{2}\sum_{i=1}^{d}\sum_{j=1}^{d}\Big((d-1)(v_i w_j - v_j w_i)^2 + (v_i w_i - v_j w_j)^2\Big).$$

For the lower bound on the minimum eigenvalue, we wish to show that

$$L := 2\Big(\sum_{i=1}^{d} v_i w_i\Big)^2 - 2\sum_{i=1}^{d} v_i^2 w_i^2 + \sum_{i=1}^{d}\sum_{j=1}^{d} v_i^2 w_j^2$$

is nonnegative. Again, we rewrite this as a sum of squares:

$$\begin{aligned} L &= \Big(\sum_{i=1}^{d} v_i w_i\Big)^2 + \sum_{i \neq j} v_i v_j w_i w_j + \sum_{i \neq j} v_i^2 w_j^2 \\ &= \Big(\sum_{i=1}^{d} v_i w_i\Big)^2 + \frac{1}{2}\sum_{i \neq j}(v_i w_j + v_j w_i)^2. \end{aligned}$$

□

We now prove the claimed concentration bound.

Proof of Lemma 3.3. Let H be as in Lemma 6.5 and let M denote the $(md) \times (md)$ block diagonal symmetric matrix with m blocks corresponding to H/m. Index the entries of M by $M_{(ki),(lj)}$, $k, l \in [m]$, $i, j \in [d]$, where $M_{(ki),(lj)} = 0$ if $k \neq l$ and equal to H_{ij}/m otherwise. By Lemma 6.5, we have that $\|M\|_{\mathrm{op}} = \|H\|_{\mathrm{op}}/m \leq (1 - 1/s)/m$ and $\|M\|_F^2 = \|H\|_F^2/m \leq (1 - 1/s)/m$.

For the Rademacher projection matrix $\Pi \in \{-1, 1\}^{m \times d}$, a direct computation reveals that

$$\|\Pi v\|_2^2/m - 1 = \sum_{r \in [m]}\sum_{i,j \in [d]} M_{(r,i),(r,j)}\Pi_{r,i}\Pi_{r,j}\,.$$

Thus applying Theorem 3.2 to right hand side quadratic form, we get that

$$\begin{aligned} \max\Big\{\mathbb{P}(\|\Pi v\|_2^2/m \geq 1 + \varepsilon), \mathbb{P}(\|\Pi v\|_2^2/m \leq 1 - \varepsilon)\Big\} &\leq \exp\Big(-\frac{1}{2}\Big(\frac{\varepsilon}{\|M\|_{\mathrm{op}}} - \frac{\|M\|_F^2}{\|M\|_{\mathrm{op}}^2}\ln\big(1 + \frac{\varepsilon\|M\|_{\mathrm{op}}}{\|M\|_F^2}\big)\Big)\Big) \\ &\leq \exp\big(-\frac{m}{2}\frac{s}{s-1}(\varepsilon - \ln(1+\varepsilon))\big), \end{aligned}$$

where the last inequality follows by setting $\|M\|_{\mathrm{op}}$ and $\|M\|_F$ to their upper bounds together with Lemma 6.4. □

References

[1] D. Achlioptas. Database-friendly random projections: Johnson-Lindenstrauss with binary coins. *J. Comput. Syst. Sci.*, 66(4):671–687, June 2003.

[2] A. A. Ageev and M. Sviridenko. Pipage rounding: A new method of constructing algorithms with proven performance guarantee. *Journal of Combinatorial Optimization*, 8(3):307–328, 2004.

[3] N. Ailon and B. Chazelle. The fast Johnson-Lindenstrauss transform and approximate nearest neighbors. *SIAM Journal on Computing*, 39(1):302–322, 2009.

[4] N. Ailon and E. Liberty. Fast dimension reduction using Rademacher series on dual BCH codes. *Discrete & Computational Geometry*, 42(4):615, 2009.

[5] N. Ailon and E. Liberty. An almost optimal unrestricted fast Johnson-Lindenstrauss transform. *ACM Transactions on Algorithms*, 9(3):21, 2013.

[6] A. Bhargava and S. R. Kosaraju. Derandomization of dimensionality reduction and SDP based algorithms. In F. Dehne, A. López-Ortiz, and J.-R. Sack, editors, *Proceedings of the 9th International Workshop on Algorithms and Data Structures (WADS)*, pages 396–408, 2005.

[7] L. Blum, M. Shub, and S. Smale. On a theory of computation and complexity over the real numbers: NP-completeness, recursive functions and universal machines. *Bull. Am. Math. Soc., New Ser.*, 21(1):1–46, 1989.

[8] R. P. Brent. Fast multiple-precision evaluation of elementary functions. *J. ACM*, 23(2):242–251, Apr. 1976.

[9] K. L. Clarkson and D. P. Woodruff. Numerical linear algebra in the streaming model. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC)*, pages 205–214, 2009.

[10] M. B. Cohen, T. Jayram, and J. Nelson. Simple Analyses of the Sparse Johnson-Lindenstrauss Transform. In R. Seidel, editor, *1st Symposium on Simplicity in Algorithms (SOSA 2018)*, volume 61 of *OpenAccess Series in Informatics (OASIcs)*, pages 15:1–15:9, Dagstuhl, Germany, 2018. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

[11] A. Dasgupta, R. Kumar, and T. Sarlós. A sparse Johnson-Lindenstrauss transform. In *Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC)*, pages 341–350. ACM, 2010.

[12] S. Dasgupta. Learning mixtures of gaussians. In *Proceedings of the 40th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 634–644, 1999.

[13] P. Drineas and M. W. Mahoney. RandNLA: Randomized numerical linear algebra. *Commun. ACM*, 59(6):80–90, May 2016.

[14] L. Engebretsen, P. Indyk, and R. O'Donnell. Derandomized dimensionality reduction with applications. In *Proceedings of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 705–712, 2002.

[15] D. L. Hanson and F. T. Wright. A bound on tail probabilities for quadratic forms in independent random variables. *Ann. Math. Statist.*, 42(3):1079–1083, 06 1971.

[16] N. J. A. Harvey and N. Olver. Pipage rounding, pessimistic estimators and matrix concentration. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 926–945, 2014.

[17] A. Hinrichs and J. Vybíral. Johnson-Lindenstrauss lemma for circulant matrices. *Random Structures & Algorithms*, 39(3):391–398, 2011.

[18] A. Horn. Doubly stochastic matrices and the diagonal of a rotation matrix. *Amer. J. Math.*, 76:620–630, 1954.

[19] P. Indyk and R. Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In *Proceedings of the 30th Annual ACM Symposium on Theory of Computing (STOC)*, pages 604–613, 1998.

[20] W. B. Johnson and J. Lindenstrauss. Extensions of Lipschitz mappings into a Hilbert space. *Contemporary mathematics*, 26(189-206):1, 1984.

[21] D. M. Kane and J. Nelson. A derandomized sparse Johnson-Lindenstrauss transform. *Electronic Colloquium on Computational Complexity (ECCC)*, 17:98, 2010.

[22] D. M. Kane and J. Nelson. Sparser Johnson-Lindenstrauss transforms. In *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1195–1206, 2012.

[23] D. M. Kane and J. Nelson. Sparser Johnson-Lindenstrauss transforms. *J. ACM*, 61(1):4:1–4:23, 2014.

[24] Z. S. Karnin, Y. Rabani, and A. Shpilka. Explicit dimension reduction and its applications. In *Proceedings of the 26th Annual IEEE Conference on Computational Complexity (CCC)*, pages 262–272, 2011.

[25] F. Krahmer and R. Ward. New and improved Johnson-Lindenstrauss embeddings via the restricted isometry property. *SIAM Journal on Mathematical Analysis*, 43(3):1269–1281, 2011.

[26] E. Liberty, N. Ailon, and A. Singer. Dense fast random projections and lean Walsh transforms. *Lecture Notes in Computer Science*, 5171:512–522, 2008.

[27] R. Meka. Almost optimal explicit Johnson-Lindenstrauss transformations. *arXiv preprint arXiv:1011.6397*, 2010.

[28] N. Nisan. Psuedorandom generators for space-bounded computation. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing (STOC)*, pages 204–212, 1990.

[29] N. Nisan. RL $\subseteq$ SC. In *Proceedings of the 24th Annual ACM Symposium on Theory of Computing (STOC)*, pages 619–623, 1992.

[30] D. Pollard. Lecture notes on empirical processes. Lecture Notes, 2015.

[31] P. Raghavan. Probabilistic construction of deterministic algorithms: Approximating packing integer programs. *Journal of Computer and System Sciences*, 37, 1988.

[32] M. Rudelson and R. Vershynin. Hanson-Wright inequality and sub-gaussian concentration. *Electron. Commun. Probab.*, 18:9 pages., 2013.

[33] D. Sivakumar. Algorithmic derandomization via complexity theory. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing (STOC)*, pages 619–626. ACM, 2002.

[34] A. Srinivasan. Distributions on level-sets with applications to approximation algorithms. In *Proceedings of the 42nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 588–597, 2001.

[35] D. P. Woodruff. Sketching as a tool for numerical linear algebra. *Foundations and Trends in Theoretical Computer Science*, 10:1–157, 2014.