

Restricted Isometry Property under High Correlations

Shiva Prasad Kasiviswanathan*

Mark Rudelson†

Abstract

Matrices satisfying the Restricted Isometry Property (RIP) play an important role in the areas of compressed sensing and statistical learning. RIP matrices with optimal parameters are mainly obtained via probabilistic arguments, as explicit constructions seem hard. In this paper, we try to bridge this gap between random and deterministic designs by introducing a new model for restricted isometry designs that incorporates a fixed matrix into the construction. Our construction starts with a fixed (deterministic) matrix X satisfying some simple stable rank condition, and we show that the matrix XR , where R is a random matrix drawn from various popular probabilistic models (including, subgaussian, sparse, low-randomness, satisfying convex concentration property), satisfies the RIP with high probability. These theorems have various applications in signal recovery, deep learning, random matrix theory, dimensionality reduction, etc. Additionally, motivated by an application for understanding the effectiveness of word vector embeddings popular in natural language processing and machine learning applications, we investigate the RIP of the matrix $XR^{(\ell)}$ where $R^{(\ell)}$ is formed by taking all possible (disregarding order) ℓ -way entrywise products of the columns of a random matrix R .

1 Introduction

A vector $\mathbf{x} \in \mathbb{R}^d$ is said to be k -sparse if it has at most k nonzero coordinates. Sparsity is a structure of wide applicability, with a broad literature dedicated to its study in various scientific fields (see, e.g., [Foucart and Rauhut, 2017, Eldar and Kutyniok, 2012]). Given an $\epsilon \in (0, 1)$, an $n \times d$ matrix M (typically with $n \ll d$) is said to satisfy the (k, ϵ) -Restricted Isometry Property (RIP) [Candes and Tao, 2005] if it approximately preserves the Euclidean norm in the following sense: for every k -sparse vector, we have

$$(1 - \epsilon)\|\mathbf{x}\| \leq \|M\mathbf{x}\| \leq (1 + \epsilon)\|\mathbf{x}\|.$$

RIP is a fundamental property of a matrix that enables recovery of a sparse high-dimensional signal from its compressed measurement. Given this, matrices satisfying the restricted isometry property have found many interesting applications in high-dimensional statistics, machine learning, and compressed sensing [Foucart and Rauhut, 2017, Eldar and Kutyniok, 2012, Wainwright, 2019]. Restricted isometry is also closely related to other matrix properties such as restricted nullspace, restricted eigenvalue, and pairwise incoherence [Wainwright, 2019].

Various probabilistic models are known to generate random matrices that satisfy the restricted isometry property with a value of k which is (almost) linear n . For example, generating entries of M i.i.d. from a common distribution (like satisfying subgaussianity) and then normalizing the columns of M to unit norm, guarantees RIP with high probability provided $n = \Omega((k/\epsilon^2) \log(d/k))$ [Baraniuk et al., 2008]. We refer the reader to [Vershynin, 2010] for additional references to the probabilistic RIP literature. The restricted isometry property also holds for a rich class

*Amazon, USA. kasivisw@gmail.com.

†University of Michigan, Ann Arbor, MI, USA. rudelson@umich.edu. Partially supported by NSF grant, DMS-1464514.

of structured random matrices, where usually the best known bounds for n have additional log factors in d [Foucart and Rauhut, 2017]. The use of randomness still remains pivotal for near-optimal results.

At the same time, verifying whether a given matrix satisfies the restricted isometry property is tricky, as the problem of certifying RIP of a matrix in the worst case is NP-hard [Bandeira et al., 2013a]. Furthermore, even determining the RIP value ϵ up to a certain approximation factor is hard in the average-case sense, as shown by [Wang et al., 2016] using a reduction from the Planted Clique Problem. While there are constructions of RIP matrices when $k = O(\sqrt{n})$, most methods however break down when k is at least some constant times \sqrt{n} , see [Bandeira et al., 2013b] for a survey of deterministic methods. The best unconditional explicit construction to date is due to Bourgain et al. [2011] which gives a RIP guarantee for $k = \Theta(n^{1/2+\delta})$ for some unspecified small constant $\delta > 0$, in the regime $n = \Omega(d^{1-\epsilon})$ and with matrix containing complex valued entries. Gamarnik [2018] recently showed why explicit RIP matrix construction is a “hard” challenge, by connecting it to a question in the field of extremal combinatorics which has been open for decades.

So on the one hand while it is easy to generate matrices satisfying the restricted isometry property through random designs, designing natural families of deterministic matrices satisfying RIP seem to run into hard barriers. The main conceptual contribution of this paper is a step towards bridging this gap. We construct families of practically motivated not-completely random matrices that satisfy RIP. We focus on a natural way of incorporating a fixed (deterministic) matrix into RIP constructions, raising the question of whether there are other interesting models that lie in the intersection of random and deterministic designs.

Our Contributions. We establish the restricted isometry property for a wide class of matrices, which can be factorized through (possibly non-i.i.d.) random matrices. In particular, we will be interested in the class of matrices which have a XR -factorization, where X is a fixed (deterministic) $n \times p$ matrix and R is a $p \times d$ random matrix. The XR -model (product of a deterministic and random matrix) in the context of RIP captures a variety of applications some of which we discuss later. The main challenge in establishing RIP comes from the fact that the entries in the matrix XR could be highly correlated, even if the entries in R are independent.¹

Our main result is that if we start with any deterministic matrix X , satisfying a very mild easy to check condition, the matrix XR satisfies RIP (with high probability) for a R constructed from a variety of popular probabilistic models. All we need is that the stable rank (or numerical rank) of X is not “too small”. Stable rank of a matrix X (denoted by $\text{sr}(X)$), defined as the squared ratio of Frobenius and spectral norms of X is a commonly used robust surrogate to usual matrix rank in linear algebra. Stable rank of a matrix is at most its usual rank. Computing the stable rank of a matrix is a polynomial time operation, meaning that given the factorization XR one could easily verify whether the required conditions on X is satisfied. We investigate many common constructions of the random matrix R .²

- *Columns of R are independent subgaussian random vectors:* In this setting, we obtain RIP on the XR matrix, if the stable rank of X satisfies $\text{sr}(X) = \Omega(k \log(d/k))$ (see Theorem 3.1). Note that this setting includes the (standard random matrix) case where R is an i.i.d. subgaussian random matrix. The dependence on d, k in this stable rank condition cannot be improved in general.
- *Generated by l -wise independent distributions:* We ask: can we reduce the amount of randomness in R ? We answer in affirmative by showing that one can achieve restricted isometry with the same condition on $\text{sr}(X)$ as above, by only requiring that R be generated from a $2 \text{sr}(X)$ -wise independent distribution (see Theorem C.3).
- *Sparse-subgaussian matrix:* Sparsity is a desirable property in R as it leads to computational speedup when working with sparse signals. We use the standard Bernoulli-Subgaussian process for generating a sparse matrix R , where R is defined as the Hadamard matrix product of an i.i.d. Bernoulli(β) random matrix and an i.i.d.

¹Note that w.r.t. RIP condition, the XR -model behaves quite differently from the RX -model as we explain in Appendix B.

²All the results are high probability statements, and for simplicity we omit the dependence on certain parameters such as ϵ, ψ_2 -norm of subgaussian vectors, etc.

subgaussian random matrix. In this case, we get RIP if the stable rank of X satisfies, $\text{sr}(X) = \Omega(k \log(d/k)/\beta)$, where β is the Bernoulli parameter (see Theorem C.9).

- *Columns of R are independent vectors satisfying convex concentration:* Our result holds for distributions that satisfy the so-called convex concentration property. The convex concentration property of a random vector was first observed by Talagrand who first proved it for the uniform measure on the discrete cube and for general product measures with bounded support [Talagrand, 1988, 1995]. Vectors satisfying convex concentration are regularly used in statistical analysis, as it includes random vectors drawn from a centered multivariate Gaussian distribution with arbitrary covariance matrix, random vectors uniformly distributed on the sphere, random vectors satisfying the *logarithmic Sobolev* inequality, among many others [Ledoux, 2001]. Ignoring the dependence on the convex concentration constant, again we get RIP on XR matrix if the stable rank of X satisfies, $\text{sr}(X) = \Omega(k \log(d/k))$ (see Theorem C.12).
- *ℓ -way column Hadamard-product construction:* Motivated by an application in understanding the effectiveness of word vector embeddings (also referred to as just word embeddings) in linear classification tasks, we investigate a correlated random matrix, formed by taking all possible (disregarding order) ℓ -way entrywise products³ of the columns of an i.i.d. centered bounded random matrix R (see Definition 4). Let $R^{(\ell)} \in \mathbb{R}^{p \times \binom{d}{\ell}}$ denote this constructed matrix starting from R (with $R^{(1)} = R$). We establish RIP on $XR^{(\ell)}$ for $\ell = 2$ if $\text{sr}(X) = \Omega(k^2 \log(d^2/k))$ (see Theorem 4.1), and for $\ell \geq 3$ if $\text{sr}(X) = \Omega(k^3 \log(d^\ell/k))$ (see Theorem 4.3). Notice that the dependence on $\text{sr}(X)$ on the sparsity parameter k is worse than in the previous cases. However, the value of k used in the motivating application is typically small.

Our proofs rely on the Hanson-Wright inequality which provides a large deviation bound for quadratic forms of i.i.d. subgaussian random variables [Rudelson and Vershynin, 2013], along with its recent extensions [Zhou, 2015, Adamczak et al., 2015]. While the proof is simple when the columns of R are independent subgaussian random vectors, various challenges arise when dealing with other models of the random matrix. One general idea is get to a concentration bound on $\|XZ\mathbf{u}\|$, where X is a fixed matrix, Z is some random matrix, and \mathbf{u} is a fixed k -sparse vector, and then use a net argument over sparse vectors on the sphere. In the case of column Hadamard-product, getting this concentration bound is tricky as it involves analyzing some high order homogeneous chaos in terms of the random variables, and we use a different idea based on bounding the ψ_2 -norm.

Throughout the proofs the challenge comes in dealing with dependences that arise both in the random matrix, and in the product matrix. This is especially true for the Hadamard-product construction, where *without* X (i.e., X is the identity matrix), the problem is much simpler and the matrix $R^{(\ell)}$ satisfies RIP under a milder condition that $p = \Omega(k \log(d^\ell/k))$ [Rudelson and Vershynin, 2008, Arora et al., 2018a]. Intuitively in the $R^{(\ell)}$ case one has only to ensure that $\|R^{(\ell)}\mathbf{u}\|$ is close to a constant with high probability for a fixed sparse unit vector \mathbf{u} . In the case of $XR^{(\ell)}$, in addition to it, one has to guarantee that the direction of the vector $R^{(\ell)}\mathbf{u}$ is more or less uniformly distributed over the sphere, making the problem significantly harder to analyze. So while it is tempting to conjecture that the right dependence of k in $\text{sr}(X)$ in the Hadamard-product case should also be linear, it appears challenging to obtain this linear bound, and it is plausible that this superlinear dependence on k might in fact be unavoidable.

Applications. The restricted isometry property is widely utilized in compressed sensing and statistical learning literature. Here we mention a few interesting applications of our restricted isometry results.

- (1) **Effectiveness of Word/Sequence Embeddings.** Consider a vocabulary set \mathcal{W} (say, all words in a particular language). Word embeddings, which associates with each word \mathbf{w} in a vocabulary set \mathcal{W} a vector representation $\mathbf{v}_w \in \mathbb{R}^p$, is a basic building block in Natural Language Processing pipelines and algorithms. Word embeddings were recently popularized via embedding schemes such as *word2vec* [Mikolov et al., 2013] and *GloVe* [Pennington et al., 2014]. Word embeddings pretrained on large sets of raw text have demonstrated remarkable success when

³The entrywise product of ℓ vectors $\mathbf{v}_1, \dots, \mathbf{v}_\ell \in \mathbb{R}^p$ is the vector $\hat{\mathbf{v}} \in \mathbb{R}^p$ with j entry equaling $\hat{v}_j = \prod_{i=1}^{\ell} v_{ij}$.

used as features to a supervised learner in various applications ranging from question-answering [Zhou et al., 2015] to sentiment classification of text documents [Maas et al., 2011]. The general intuition behind word embeddings is that they are known to capture the “similarities” between words. There has also been recent work on creating representations for *word sequences* such as *phrases* or *sentences* with methods ranging from simple composition of the word vectors to sophisticated architectures as *recurrent neural networks* (see e.g., [Arora et al., 2017] and references therein)

Understanding the theoretical properties of these word embeddings is an area of active interest. In a recent result, Arora et al. [2018a] introduced the scheme of *Distributed Cooccurrence* (DisC) embedding (Definition 6) for a word sequence that produces a compressed representation of a Bag-of- L -cooccurrences vector.⁴ They showed that if one uses i.i.d. ± 1 -random vectors as word embeddings, then a linear classifier trained on these compressed DisC embeddings performs “as good as” as a similar classifier trained on the original Bag-of- L -cooccurrences vectors. This was the first result that provided provable quantification of the power of any text embedding.⁵ They achieve this result by connecting this problem with the theory of compressed sensing, an idea that we build upon here. Let $V \in \mathbb{R}^{p \times d}$ be a matrix whose columns are the embeddings for all the words in \mathcal{W} . Arora et al. [2018a] result relies on establishing the restricted isometry property on the matrix $V^{(\ell)} \in \mathbb{R}^{p \times \binom{d}{\ell}}$ where V is an i.i.d. ± 1 -random matrix (i.e., random vectors are used as word embeddings).

Linear transformations are regularly used to transfer between different embeddings or to adapt to a new domain [Bollegala et al., 2017, Arora et al., 2018b]. The linear transformation can encode *contextual information*, an idea utilized recently by [Khodak et al., 2018] who applied a linear transformation on the DisC embedding scheme to construct a new embedding scheme (referred to as *à la carte* embedding), and empirically showed that it outperforms many other popular word sequence embedding schemes. Now akin to [Arora et al., 2018a] result on DisC embeddings, our results shed some theoretical insights into the performance on linear transformations of DisC embeddings. In particular, our RIP results on $XV^{(\ell)}$, where $V \in \mathbb{R}^{p \times d}$ is an i.i.d. centered bounded random matrix, provides provable performance guarantees on linear transformation (defined by X) of DisC embeddings in a linear classification setting, under a stable rank assumption on X (see Corollary 4.5). We expand on this application in Section 4.1.

- (2) **Deep Linear Networks.** A deep linear network $\mathcal{D} : \mathbb{R}^d \rightarrow \mathbb{R}^n$ is a neural network that has multiple hidden layers but have no nonlinearities between layers (typically $n \ll d$). That is, for a given datapoint \mathbf{x} , the output $\mathcal{D}(\mathbf{x})$ is computed via a series $\mathcal{D}(x) = W_t W_{t-1} \dots W_1 \mathbf{x}$ of matrix multiplications. Such models are commonly used as a tool for theoretically understanding deep neural networks [Arora et al., 2019, Laurent and Brecht, 2018]. A simple question that arises is whether low-dimensional output of a deep linear network can be used as a surrogate representation for classification, i.e., what happens if instead of $(\mathbf{x}_1, v_1), \dots, (\mathbf{x}_b, v_b)$ where v_i is the label on \mathbf{x}_i , $(\mathcal{D}(\mathbf{x}_1), v_1), \dots, (\mathcal{D}(\mathbf{x}_b), v_b)$ is used for classification? Our RIP results help answer this question. For example, if W_1 is a matrix from one of the above random matrix families⁶, and the matrix $X = W_t W_{t-1} \dots W_2$ satisfies the stable rank condition, then we get that the matrix $W_t W_{t-1} \dots W_1$ satisfies RIP with high probability. Then the known results in compressed learning (see for example Theorem D.2) can be used to relate the performance of a linear classifier trained on the compressed representation $(\mathcal{D}(\mathbf{x}_1), v_1), \dots, (\mathcal{D}(\mathbf{x}_b), v_b)$ to a similar classifier trained on the original representation $(\mathbf{x}_1, v_1), \dots, (\mathbf{x}_b, v_b)$.

⁴An ℓ -cooccurrence is a set of ℓ words. Bag-of- L -cooccurrences for a word sequence counts the number of times any possible ℓ -cooccurrence (from the vocabulary set \mathcal{W}) for $\ell \in [L]$ appears in the word sequence. See Definition 5. Linear classification models are empirically known to perform well over these simple representations [Wang and Manning, 2012, Arora et al., 2018a].

⁵Arora et al. [2018a] have additional results on the powerfulness of low-memory LSTM (Long Short Term Memory) network embeddings, by showing that the LSTMs under certain initialization can simulate these DisC embeddings. Our more general RIP results can be easily used to generalize these results too (details omitted).

⁶Randomness is commonly used in deep networks for initializing weights [Arora et al., 2019] before training, and sometimes random weights are used without training [Saxe et al., 2011].

- (3) Linear Transformation and Johnson-Lindenstrauss Embedding.** *Johnson-Lindenstrauss (JL) embedding lemma* states that any set of m points in high dimensional Euclidean space \mathbb{R}^d can be embedded into $p = O(\log(m)/\epsilon^2)$ dimensions, without distorting the distance between any two points by more than a factor between $1 - \epsilon$ and $1 + \epsilon$ [Johnson and Lindenstrauss, 1984]. JL lemma has become a valuable tool for dimensionality reduction [Woodruff, 2014]. Typically, the embedding is constructed through a random linear map referred to as an ϵ -JL matrix. Let $R \in \mathbb{R}^{p \times d}$ be an ϵ -JL matrix. Consider a set of m points $\mathbf{a}_1, \dots, \mathbf{a}_m \in \mathbb{R}^d$, and let $R\mathbf{a}_1, \dots, R\mathbf{a}_m$ be their low-dimensional embedding. A natural question to ask is: under what fixed linear transformations of $R\mathbf{a}_1, \dots, R\mathbf{a}_m$ does this distance preservation property still hold? More concretely, let $X \in \mathbb{R}^{n \times p}$ be a linear transformation. Does the JL distance preservation property still hold for $XR\mathbf{a}_1, \dots, XR\mathbf{a}_m$? Our results answers this question because of the close connection between JL-embedding and the restricted isometry property. Krahmer and Ward [2011] showed that, when the columns of a matrix satisfying RIP are multiplied by independent random signs, any $(O(\log m), O(\epsilon))$ -RIP matrix becomes an ϵ -JL matrix for a fixed set of m vectors with probability at least $1 - m^{-\Omega(1)}$ ⁷. This means, one can now use our RIP results discussed above. For example, if the entries of R are drawn i.i.d. from a centered symmetric subgaussian distribution, then for a X (scaled to have unit Frobenius norm), if $\text{sr}(X) = \tilde{\Omega}(\log(m)/\epsilon^2)$ ⁸, we have that $XR\mathbf{a}_1, \dots, XR\mathbf{a}_m$ satisfies the JL-embedding property with high probability (as multiplying with random signs does not change the distribution R).
- (4) Compressed Sensing under Linear Transformations.** Compressed sensing algorithms are designed to recover approximately sparse signals, and a popular sufficient condition for a matrix to succeed for the purposes of compressed sensing is given by the restricted isometry property (see, e.g., Theorem B.3). Another question that can be addressed using our results is how does the recovery guarantee hold if we apply a linear transformation on the compressed signals. Our RIP results establishes conditions on X , for the class of random matrices R described earlier, under which given $\mathbf{y} = XR\mathbf{x} + \mathbf{e}$ where \mathbf{e} is a noise vector, one can recover an approximately sparse vector \mathbf{x} .
- (5) Source Separation.** Separation of underdetermined mixtures is an important problem in signal processing. Sparsity has often been exploited for source separation. The general goal with source separation is that given a signal matrix $S \in \mathbb{R}^{n_1 \times c}$, a mixing matrix $M \in \mathbb{R}^{n_1 \times a}$, and a dictionary $\Phi \in \mathbb{R}^{b \times c}$, is to find a matrix of coefficients $C \in \mathbb{R}^{a \times b}$ such that $S \approx MC\Phi$ and C is as sparse as possible. The dictionary is generally overcomplete, i.e., $b > c$. The connection between source separation and compressed sensing was first noted by [Blumensath and Davies, 2007].⁹ It is easy to recast (see Appendix A for details) the source separation problem as a compressed sensing problem of the form: $\mathbf{s} = XRC\mathbf{c}$, where the goal is to estimate the sparse \mathbf{c} (entries of the matrix C). Here, $\mathbf{s} \in \mathbb{R}^n$, $X \in \mathbb{R}^{n \times p}$, $R \in \mathbb{R}^{p \times d}$, and $\mathbf{c} \in \mathbb{R}^d$, with $n = n_1c$, $d = ab$, $p = ac$. Hence, our RIP results establishes conditions on the mixing matrix M , for the classes of random matrices (dictionaries) described earlier, under which we can get recovery guarantees on \mathbf{c} (entries of the matrix C).
- (6) Singular Values of Correlated Random Matrices.** Next application is a simple consequence of the restricted isometry property. Understanding the singular values of random matrices is an important problem, with lots of applications in machine learning and in the field of non-asymptotic theory of random matrices [Vershynin, 2016]. Restricted isometry property of $M = XR$ can be interpreted in terms of the extreme singular values of submatrices of M . Indeed, the restricted isometry property (assuming for simplicity that X has unit Frobenius norm¹⁰) equivalently states that the inequality

$$\sqrt{1 - \epsilon} \leq \lambda_{\min}(M_I^\top M_I) \leq \lambda_{\max}(M_I^\top M_I) \leq \sqrt{1 + \epsilon}. \quad (1)$$

⁷The connection in other direction going from JL-embedding to RIP is also well-known [Baraniuk et al., 2008].

⁸The $\tilde{\Omega}$ notation hides polylog factors in $d, 1/\epsilon$.

⁹In a variant of this problem, called the *blind source separation*, even the mixing matrix M is assumed to be unknown.

¹⁰Otherwise results can be appropriately scaled.

holds for all $n \times k$ submatrices $M_I = XR_I$, those formed by the columns of M indexed by sets I of size k . If the columns of R are independent random vectors drawn from distribution \mathcal{R} over \mathbb{R}^p , we can think of our restricted isometry results as bounds on the singular values of matrices formed by linear transformation of random vectors (or equivalently, the eigenvalues of the sample covariance matrix drawn from the distribution $X\mathcal{R}$). For example, let $Q \in \mathbb{R}^{p \times k}$ be a matrix whose columns are independent random vectors from a subgaussian distribution or satisfying convex concentration property, then our RIP results through (1) bounds *all* the singular values on XQ (with high probability) under the condition that $\text{sr}(X) = \Omega(k \log(1/\epsilon)/\epsilon^2)$. Previously only a spectral norm bound on XQ for the cases where Q has independent entries drawn from a subgaussian or heavy-tailed distributions was known [Rudelson and Vershynin, 2013, Vershynin, 2011].

2 Preliminaries

Notation. We use $[n]$ to denote the set $\{1, \dots, n\}$. We denote the Hadamard (elementwise) product by \odot . For a set $I \subseteq [d]$, I^c denotes its complement set.

Vectors used in the paper are by default column vectors and are denoted by boldface letters. For a vector \mathbf{v} , \mathbf{v}^\top denotes its transpose, $\|\mathbf{v}\|$ and $\|\mathbf{v}\|_1$ denote the L_2 and L_1 -norm respectively, and $\text{supp}(\mathbf{v})$ its support. $\mathbf{e}_i \in \mathbb{R}^d$ denote the standard basis with i th coordinate 1. For a matrix M , $\|M\|$ denotes its spectral norm, $\|M\|_F$ denotes its Frobenius norm, and M_{ij} denotes its (i, j) th entry, $\text{diag}(M)$ denotes the diagonal of M , and $\text{offdiag}(M)$ denotes the off-diagonal of M . \mathbb{I}_n represents the identity matrix in dimension n . For a vector \mathbf{x} and set of indices S , let \mathbf{x}_S be the vector formed by the entries in \mathbf{x} whose indices are in S , and similarly, M_S is the matrix formed by columns of M whose indices are in S .

The Euclidean sphere in \mathbb{R}^d centered at origin is denoted by S^{d-1} . We call a vector $\mathbf{a} \in \mathbb{R}^d$, k -sparse, if it has at most k non-zero entries. Denote by Σ_k the set of all vectors $\mathbf{a} \in S^{d-1}$ with support size at most k : $\Sigma_k = \{\mathbf{a} \in S^{d-1} : |\text{supp}(\mathbf{a})| \leq k\}$. Stable rank (denoted by $\text{sr}()$) of a matrix M is defined as:

$$\text{sr}(M) = \|M\|_F^2 / \|M\|^2.$$

Stable rank cannot exceed the usual rank. The stable rank is a more robust notion than the usual rank because it is largely unaffected by tiny singular values.

Throughout this paper C, c, C' , also with subscripts, denote positive absolute constants, whose value may change from line to line. In Appendix B we discuss additional preliminaries about subgaussian/subexponential random variables, sparse recovery, and Hanson-Wright inequality.

Restricted Isometry. Candes and Tao [2005] introduced the following isometry condition on matrices M . It is perhaps one of the most popular property of a matrix which makes it “good” for compressed sensing.

Definition 1. Let M be an $n \times d$ matrix with real entries. Let $\epsilon \in (0, 1)$ and let $k < d$ be an integer. We say that M satisfies (k, ϵ) -RIP if for every k -sparse vector $\mathbf{u} \in S^{d-1}$ (i.e., $\mathbf{u} \in \Sigma_k$)

$$(1 - \epsilon) \leq \|M\mathbf{u}\| \leq (1 + \epsilon).$$

Thus, M acts almost as an isometry when we restrict attention to k -sparse vectors. Geometrically, the restricted isometry property guarantees that the geometry of k -sparse vectors \mathbf{u} is well preserved by the measurement matrix M . It turns out that in this case, given a (noisy) compressed measurement $M\mathbf{x}$, for approximately sparse \mathbf{x} , one can recover \mathbf{x} using a simple convex program. Theorem B.3 provides a bound on the worst-case recovery performance for uniformly bounded noise. Similar recovery results using RIP on the measurement matrix M are also well-known under other interesting settings (we refer the reader to the survey on this topic in [Eldar and Kutyniok, 2012]).

We investigate the restricted isometry property of matrices $M \in \mathbb{R}^{n \times d}$ that can be factored as $M = XR$, where R will be some random matrix and X is a fixed matrix. Now RIP is not invariant under scaling, i.e., given a RIP

matrix M , changing M to some aM for some $a \in \mathbb{R}$ scales both the left and right hand side of the inequality in the Definition 1 by a . So while working in the XR -model we have to adjust for the scaling introduced by X , and we use a generalization of Definition 1 appropriate for this setting.

Definition 2 (RIP Condition). *A matrix $M = XR$ satisfies (k, ϵ) -RIP if for every k -sparse vector $\mathbf{u} \in S^{d-1}$ (i.e., $\mathbf{u} \in \Sigma_k$)*

$$(1 - \epsilon)\|X\|_F \leq \|XR\mathbf{u}\| \leq (1 + \epsilon)\|X\|_F.$$

This scaling by $\|X\|_F$ is unavoidable, because even if R is a matrix with centered uncorrelated entries of unit variance, then $\mathbb{E}[\|XR\mathbf{u}\|^2] = \|X\|_F^2$ for any \mathbf{u} with $\|\mathbf{u}\| = 1$. Note that it is easy to reconstruct standard RIP scenarios by choosing X appropriately. For example, if R contains i.i.d. subgaussian random variables with zero mean and unit variance, then setting $X = \mathbb{I}_n/\sqrt{n}$, leads to standard setting of $XR = R/\sqrt{n}$ and $\|X\|_F = 1$.

Hanson-Wright Inequality. An important concentration tool used in this paper is the Hanson-Wright inequality (Theorem 2.1) that investigates concentration of a quadratic form of independent centered subgaussian random variables, and its recent extensions [Zhou, 2015, Adamczak et al., 2015]. A slightly weaker version of this inequality was first proved in [Hanson and Wright, 1971].

Theorem 2.1 (Hanson-Wright Inequality [Rudelson and Vershynin, 2013]). *Let $\mathbf{x} = (x_1, \dots, x_n) \in \mathbb{R}^n$ be a random vector with independent components x_i which satisfy $\mathbb{E}[x_i] = 0$ and $\|x_i\|_{\psi_2} \leq K$. Let M be an $n \times n$ matrix. Then for every $t \geq 0$,*

$$\Pr \left[\left| \mathbf{x}^\top M \mathbf{x} - \mathbb{E}[\mathbf{x}^\top M \mathbf{x}] \right| > t \right] \leq 2 \exp \left(-c \min \left\{ \frac{t^2}{K^4 \|M\|_F^2}, \frac{t}{K^2 \|M\|} \right\} \right).$$

3 Restricted Isometry of XR with ‘‘Random’’ R

In this section, we investigate the restricted isometry property for the class of XR matrices, for various classes of random matrices $R \in \mathbb{R}^{p \times d}$. Missing details from this section are collected in Appendix C.

As a warmup, we start with the simplest case where R is a centered i.i.d. subgaussian random matrix, and build on this result, where we consider various other general families of R such as those constructed using low-randomness, with sparsity structure, or satisfying a convex concentration condition.

Theorem 3.1 presents the result in case R is an i.i.d. subgaussian random matrix. The proof idea here is quite simple, but provides a framework that will be helpful later. Under the stable rank condition, a net argument, along with Hanson-Wright inequality implies, $\Pr[\exists I \subset [d], |I| = k, \|XR_I\| \geq C_0 \epsilon \|X\|_F] \leq \exp(-c_0 \epsilon^2 \text{sr}(X)/K^4)$, where R_I is the $p \times k$ submatrix of R with columns from the set I . Also, by the same Hanson-Wright, for any $\mathbf{u} \in S^{d-1}$, $\Pr[\|XR\mathbf{u}\| - \|X\|_F \geq \epsilon \|X\|_F] \leq \exp(-c_1 \epsilon^2 \text{sr}(X)/K^4)$. Using a bound on the net size for sparse vectors gives all the required ingredients for the following theorem.

Theorem 3.1. *Let X be an $n \times p$ matrix. Let $R = (R_{ij})$ be a $p \times d$ matrix whose entries are with independent entries such that $\mathbb{E}[R_{ij}] = 0$, $\mathbb{E}[R_{ij}^2] = 1$, and $\|R_{ij}\|_{\psi_2} \leq K$. Let $\epsilon \in (0, 1)$, and let $k \in \mathbb{N}$ be a number satisfying $\text{sr}(X) \geq CK^4 \frac{k}{\epsilon^2} \log\left(\frac{d}{k}\right)$. Then with probability at least $1 - \exp(-c\epsilon^2 \text{sr}(X)/K^4)$, the matrix XR satisfies (k, ϵ) -RIP, i.e.,*

$$\forall \mathbf{u} \in \Sigma_k, (1 - \epsilon)\|X\|_F \leq \|XR\mathbf{u}\| \leq (1 + \epsilon)\|X\|_F.$$

Note that under the assumption the stable rank, the probability $1 - \exp(-c\epsilon^2 \text{sr}(X)/K^4)$ is at least $1 - \exp(-c'k)$.

Remark 3.2. *Notice that there is no direct condition on n , except that comes through the stable rank assumption on X , as $\text{sr}(X) \leq \text{rank}(X) \leq \min\{n, p\}$. Indeed one should expect this to happen. For example, if we take a matrix*

X and add a bunch of zero rows to it this would increase n , but should not change the recovery properties. This suggests that we need a notion of a “true” dimension of the range of X , and the stable rank is the one.

This assumption on the stable rank is optimal up to constant factors. For example, when $X = \mathbb{I}_n$ is the identity matrix and R is a standard Gaussian matrix, then $\text{sr}(X) = \text{rank}(X) = n$, and therefore the stable rank condition just becomes $n = \Omega((k/\epsilon^2) \log(d/k))$. We know that up to constant, this dependence of n on d and in k is optimal [Foucart et al., 2010]. This shows that the lower bound on the stable rank in Theorem 3.1 cannot be improved in general.

In the following, we investigate other popular families of random matrices, extending the result in Theorem 3.1 in various directions.

(a) Independent Subgaussian Columns: Consider the case where the columns of the matrix R are independently drawn isotropic subgaussian random vectors. In this case the proof proceeds as in Theorem 3.1, by applying in this case an extension of Hanson-Wright inequality (Theorem 2.1) to subgaussian random vectors [Vershynin, 2016], and by noting that for any fixed vector $\mathbf{u} \in S^{d-1}$, $R\mathbf{u}$ will be an isotropic subgaussian random vector as it is a linear combination of columns of R which are all independent.

(b) Low Randomness: Optimal use of randomness is an important consideration when designing the matrix R . In the dimensionality reduction literature much attention has been given in obtaining explicit constructions of R minimizing the number of random bits used (see, e.g., [Kane et al., 2011] and references therein). For a fixed X , we show that as long as R satisfies $2 \text{sr}(X)$ -wise independence, then XR satisfies RIP under the same (up to constant) stable rank condition on X as in Theorem 3.1. So in effect, one can reduce the number of random bits from $O(pd)$ (in Theorem 3.1) to $O(\text{sr}(X) \log(pd))$.

Again we start as in Theorem 3.1 with the Hanson-Wright inequality, but in this case the proof has to deal with the lack of independence in R . Our general strategy in this proof will be to rely on higher moments where we can treat certain variables as independent. The main technical step is to derive a concentration bound for $\|XR\mathbf{u}\|$ for a fixed $\mathbf{u} \in S^{d-1}$, which in this case is obtained through analyzing certain higher moments of $\|XR\mathbf{u}\|^2 - \mathbb{E}[\|XR\mathbf{u}\|^2]$ using moment generating functions. For $2 \text{sr}(X)$ -wise independent R , we establish (k, ϵ) -RIP on XR if $\text{sr}(X) = \Omega((k/\epsilon^2) \log(d/k))$. See Theorem C.3 for a formal statement.

(c) Sparse-Subgaussian: Sparsity is a desirable property in the measurement matrix because it leads to faster computation matrix-vector product. For example, if R is drawn from a distribution over matrices having at most s non-zeroes per column, then $R\mathbf{x}$ can be computed in time $s\|\mathbf{x}\|_0$. We use the sparse-subgaussian model of random matrices.

Our result here (Theorem C.9) utilizes the recently introduced sparse Hanson-Wright inequality [Zhou, 2015]. Similar to Hanson-Wright inequality, sparse Hanson-Wright inequality provides a large deviation bound for a quadratic form. However, in this case, the quadratic form is sparse, and is of the form $(\mathbf{x} \odot \zeta)^\top M (\mathbf{x} \odot \zeta)$ where \mathbf{x} is a random vector with independent centered subgaussian components (ψ_2 -norm bounded by K) and ζ contains independent Bernoulli(β) random variables. We use the sparse Hanson-Wright inequality to obtain the necessary subgaussian concentration bounds (Lemmas C.7, C.8) in a proof framework similar to Theorem 3.1. For a Bernoulli-subgaussian matrix R , we establish (k, ϵ) -RIP on XR if $\text{sr}(X) = \Omega((K^4 k / (\beta \epsilon^2)) \log(d/k))$. See Theorem C.9 for a formal statement.

(d) Under Convex Concentration: Convex concentration property is a generalization of standard concentration property (such as Gaussian concentration) by requiring concentration to hold only for 1-Lipschitz convex functions.

Definition 3 (Convex Concentration Property). *Let \mathbf{x} be a random vector in \mathbb{R}^d . We will say that \mathbf{x} has the convex concentration property (c.c.p) with constant K if for every 1-Lipschitz convex function $\phi : \mathbb{R}^d \rightarrow \mathbb{R}$, we have $\mathbb{E}[|\phi(\mathbf{x})|] < \infty$ and for every $t > 0$,*

$$\Pr[|\phi(\mathbf{x}) - \mathbb{E}[\phi(\mathbf{x})]| \geq t] \leq 2 \exp(-t^2/K^2).$$

The class of distributions satisfying c.c.p is extremely broad [Ledoux, 2001]. Some examples include: i) Gaussian random vectors drawn from $N(0, \Sigma)$ have c.c.p with $K^2 = 2\|\Sigma\|$, ii) random vectors that are uniformly distributed on the sphere $\sqrt{d}S^{d-1}$ have c.c.p with constant $K = 2$, iii) subclass of logarithmically concave random vectors, iv) random vectors with possibly dependent entries which satisfy a *Dobrushin* type condition, and v) random vectors satisfying the *logarithmic Sobolev* inequality.

The main technical part here is to show that a linear combination of independent vectors having a convex concentration property would have this property as well (Lemma C.11). To this end, we start with a fixed $\mathbf{u} = (u_1, \dots, u_d) \in S^{d-1}$ and construct a martingale with variables $\mathbb{E}[\phi(\sum_{i=1}^d \mathbf{r}_i u_i) | \mathbf{r}_1, \dots, \mathbf{r}_j]$ where $\phi : \mathbb{R}^p \rightarrow \mathbb{R}$ is a 1-Lipschitz convex function and \mathbf{r}_i is the i th column in R , and then apply Azuma's inequality to it. Once we have this established, verifying RIP consists of applying the Hanson-Wright inequality of [Adamczak et al., 2015] for isotropic vectors having convex concentration property in a proof framework similar to Theorem 3.1. For a matrix R with independent columns satisfying c.c.p with constant K , we establish (k, ϵ) -RIP on XR if $\text{sr}(X) = \Omega((K^4 k / \epsilon^2) \log(d/k))$. See Theorem C.12 for a formal statement.

4 Restricted Isometry under ℓ -way Column Hadamard-product

In this section, we investigate the restricted isometry property for a class of correlated random matrices motivated by theoretically understanding the effectiveness of word vector embeddings. Missing details from this section are collected in Appendix D.

To introduce this setting, let us start with the definition of a matrix product operation introduced by [Arora et al., 2018a] to construct their distributed cooccurrence (DisC) word embeddings.

Definition 4 (*ℓ -way Column Hadamard-product Operation*). *Let M be an $p \times d$ matrix, and let $\ell \in \mathbb{N}$. The ℓ -way column Hadamard-product operation constructs a $p \times \binom{d}{\ell}$ matrix $M^{(\ell)}$ whose columns indexed by a sequence $1 \leq i_1 < i_2 < \dots < i_\ell \leq d$ is the elementwise product of the i_1, i_2, \dots, i_ℓ -th columns of M , i.e., $(i_1, i_2, \dots, i_\ell)$ -th column in $M^{(\ell)}$ is $\mathbf{m}_{i_1} \odot \dots \odot \mathbf{m}_{i_\ell}$, where \mathbf{m}_j for $j \in [d]$ is the j th column in M .*

Arora et al. [2018a], based on an application of a result of [Foucart and Rauhut, 2017] on RIP for *bounded orthonormal systems*, showed that if $R \in \mathbb{R}^{n \times d}$ is an i.i.d. random sign matrix, and if $n = \Omega((k/\epsilon^2) \log(d^\ell/\gamma))$, then with probability at least $1 - \gamma$, $R^{(\ell)}$ satisfies restricted isometry property. In this paper, we investigate RIP on $XR^{(\ell)}$. It is not hard to see that achieving RIP for $R^{(\ell)}$ is *simpler* than achieving RIP for $XR^{(\ell)}$. Intuitively, in the $R^{(\ell)}$ case one has only to ensure that $\|R^{(\ell)}\mathbf{u}\|$ is close to a constant with high probability for a fixed sparse unit vector \mathbf{u} . In the case of $XR^{(\ell)}$, in addition to it, one has to guarantee that the direction of the vector $R^{(\ell)}\mathbf{u}$ is more or less uniformly distributed over the sphere.

For $\ell = 1$, Theorem 3.1 holds, and we get RIP under the condition $\text{sr}(X) = \Omega((k/\epsilon^2) \log(d/k))$. Below, we first analyze the case of $\ell = 2$, and then for larger ℓ 's. In the motivating application discussed in Section 4.1, ℓ is generally a small constant.

Analysis for $\ell = 2$. Let R be a random matrix with centered τ -bounded entries.¹¹ We prove that the matrix $XR^{(2)}$ satisfies RIP with high probability provided that $\text{sr}(X)$ is sufficiently large. Note that we do not require the entries in R to be identically distributed.

Theorem 4.1. *Let X be an $n \times p$ matrix, and let R be a $p \times d$ random matrix with independent entries R_{ij} such that $\mathbb{E}[R_{ij}] = 0, \mathbb{E}[R_{ij}^2] = 1$, and $|R_{ij}| \leq \tau$ almost surely. Let $\epsilon \in (0, 1)$, and let $k \in \mathcal{N}$ be a number satisfying*

¹¹The centered assumption is necessary because to hope for RIP with high probability, we must have it in average. The boundedness assumption is an artifact of our proof approach. The product of subgaussians is not subgaussian. However, the class of bounded random variables is closed under product.

$\text{sr}(X) \geq \frac{C\tau^8 k^2}{\epsilon^2} \log\left(\frac{d^2}{k}\right)$. Then with probability at least $1 - \exp(-c\epsilon^2 \text{sr}(X)/(k\tau^8))$, the matrix $XR^{(2)}$ satisfies the (k, ϵ) -RIP property, i.e., for any $\mathbf{u} \in \mathcal{S}^{\binom{d}{2}-1}$ with $|\text{supp}(\mathbf{u})| \leq k$,

$$(1 - \epsilon) \|\mathbf{u}\|_{\text{F}} \leq \left\| XR^{(2)} \mathbf{u} \right\| \leq (1 + \epsilon) \|\mathbf{u}\|_{\text{F}}.$$

Proof. Let $\mathbf{u} \in \mathbb{R}^{\binom{d}{2}}$ be a vector with $|\text{supp}(\mathbf{u})| \leq k$. Let u_{ij} be the (i, j) th element in \mathbf{u} with $(i, j) \in \binom{[d]}{2}$. The random variable $\left\| XR^{(2)} \mathbf{u} \right\|^2$ is order 4 homogenous chaos in terms of the random variables R_{ij} . Establishing concentration for this chaos can be a difficult task, so we will approach the problem from a different angle. Let $l \in [p]$, and define

$$y_l = \sum_{(i,j) \in \binom{[d]}{2}} R_{li} R_{lj} u_{ij}.$$

Note that the random variables y_l , $l \in [p]$ are independent. We will estimate the ψ_2 -norm of y_l and use the Hanson-Wright inequality (Theorem 2.1) to establish the concentration for the norm of $XR^{(2)} \mathbf{u} = X\mathbf{y}$ (where $\mathbf{y} = (y_1, \dots, y_p)$). Note that the support of \mathbf{u} contains the pair (i, j) with $u_{ij} \neq 0$. Directly from the triangle inequality, one can get

$$\left\| \sum_{(i,j) \in \binom{[d]}{2}} R_{li} R_{lj} u_{ij} \right\|_{\psi_2} \leq \sum_{(i,j) \in \text{supp}(\mathbf{u})} \|R_{li} R_{lj} u_{ij}\|_{\psi_2} = O(\tau^2 \|\mathbf{u}\|_1) = O(\tau^2 \sqrt{k} \|\mathbf{u}\|), \quad (2)$$

since $|\text{supp}(\mathbf{u})| \leq k$. This estimate is, however, too wasteful. We will prove a more precise one using a special decomposition of the vector \mathbf{u} . It is based on a novel induction procedure. For simplicity, we ignore the subscript l , and denote R_{li} by r_i and R_{lj} by r_j . We explain the induction idea here, deferring the entire proof to Appendix D.

The support of vector \mathbf{u} can be viewed as an $I \times I^{co}$ matrix with at most k non-zero entries. If each row of this matrix contains at most one non-zero entry, we can condition on r_j , $j \in I^{co}$ and get a bound on the ψ_2 -norm which does not depend on k using Hoeffding's inequality. The same is true if each column of the matrix contains at most one non-zero entry. This suggests that intuitively, the worst case scenario occurs when the non-zero entries of \mathbf{u} form a $\sqrt{k} \times \sqrt{k}$ submatrix. This submatrix can be split into the sum of \sqrt{k} rows, which in combination with Cauchy-Schwarz inequality allows to bound the ψ_2 -norm by $k^{1/4}$. But it is not clear how to extend this splitting to the case of \mathbf{u} with an arbitrary support. However, the $\sqrt{k} \times \sqrt{k}$ matrix can be split in a slightly different way. Namely, we separate the first row of the $\sqrt{k} \times \sqrt{k}$ matrix of support of \mathbf{u} , so that the remaining part is a $(\sqrt{k} - 1) \times \sqrt{k}$ matrix. After that, we separate the first column of the remaining matrix which leaves a $(\sqrt{k} - 1) \times (\sqrt{k} - 1)$ matrix. Alternating between rows and columns, we get a $(\sqrt{k} - t) \times (\sqrt{k} - t)$ matrix after $2t$ steps. This process ends in $2\sqrt{k}$ steps, which yields the same $k^{1/4}$ bound for the ψ_2 -norm. We show below that this method can be extended to a vector \mathbf{u} with an arbitrary support.

Lemma 4.2. *Let $\mathbf{u} \in \mathbb{R}^{\binom{d}{2}}$ be a vector with $|\text{supp}(\mathbf{u})| \leq k$. Let r_i , $i \in [d]$ be independent random variables with $\mathbb{E}[r_i] = 0$, $\mathbb{E}[r_i^2] = 1$, and $|r_i| \leq \tau$ almost surely. Then*

$$\left\| \sum_{(i,j) \in \binom{[d]}{2}} r_i r_j u_{ij} \right\|_{\psi_2} \leq 16\tau^2 k^{1/4} \|\mathbf{u}\|.$$

From Lemma 4.2, we finish the proof of the theorem by combining the Hanson-Wright inequality (Theorem 2.1) and a net argument.

Recall that for any k -sparse vector $\mathbf{u} \in \mathbb{R}^{\binom{d}{2}}$, $XR^{(2)}\mathbf{u} = X\mathbf{y}$, where $\mathbf{y} = (y_1, \dots, y_p)$ is a random vector with independent coordinates such that for all $l \in [p]$

$$\mathbb{E}[y_l] = 0, \quad \mathbb{E}[y_l^2] = \|\mathbf{u}\|^2, \quad \text{and } \|y_l\|_{\psi_2} \leq C\tau^2 k^{1/4} \|\mathbf{u}\| \quad (\text{from Lemma 4.2}).$$

By the volumetric estimate we can choose an $(1/2C_2)$ -net \mathcal{N} in the set of all k -sparse vectors in $S^{\binom{d}{2}-1}$ such that

$$|\mathcal{N}| \leq \binom{\binom{d}{2}}{k} (6C_2)^k \leq \exp\left(k \log\left(\frac{C_0 d^2}{k}\right)\right).$$

Using Corollary B.1, for any $\mathbf{u} \in S^{\binom{d}{2}-1}$ with $|\text{supp}(\mathbf{u})| \leq k$ (and $\mathbf{y} = R^{(2)}\mathbf{u}$),

$$\Pr\left[\left|\|X\mathbf{y}\| - \|X\|_{\text{F}}\right| > \epsilon \|X\|_{\text{F}}\right] \leq 2 \exp\left(-\frac{C\epsilon^2}{\max_l \|y_l\|_{\psi_2}^4} \text{sr}(X)\right) \leq 2 \exp\left(-\frac{C_1\epsilon^2}{\tau^8 k} \text{sr}(X)\right).$$

Combining this with the union bound over $\mathbf{u} \in \mathcal{N}$ and using the assumption on $\text{sr}(X)$, we get

$$\Pr\left[\exists \mathbf{u} \in \mathcal{N}, \left|\|XR^{(2)}\mathbf{u}\| - \|X\|_{\text{F}}\right| > \epsilon \|X\|_{\text{F}}\right] \leq \exp\left(k \log\left(\frac{C_0 d^2}{k}\right)\right) \cdot 2 \exp\left(-\frac{C_1\epsilon^2}{\tau^8 k} \text{sr}(X)\right). \quad (3)$$

Similar from Hanson-Wright (Corollary B.2), we can derive that

$$\Pr\left[\exists I \in \binom{[d]}{2}, |I| = k, \|XR_I^{(2)}\| > C_1\epsilon \|X\|_{\text{F}}\right] \leq \exp\left(-\frac{c_1\epsilon^2}{\tau^8 k} \text{sr}(X)\right). \quad (4)$$

Now using (3) and (4) and an approximation idea (as in Theorem 3.1) completes the proof of Theorem 4.1. \square

Analysis for $\ell \geq 3$. While extending the stronger ψ_2 -norm estimate given by Lemma 4.2 to these larger ℓ 's seems tricky, the looser bound obtained through a triangle inequality argument (as in (2)) still holds. This leads to the following theorem.

Theorem 4.3. *Let X be an $n \times p$ matrix, and let R be a $p \times d$ random matrix with independent entries R_{ij} such that $\mathbb{E}[R_{ij}] = 0$, $\mathbb{E}[R_{ij}^2] = 1$, and $|R_{ij}| \leq \tau$ almost surely. Let $\ell \geq 3$ be a constant. Let $\epsilon \in (0, 1)$, and let $k \in \mathcal{N}$ be a number satisfying $\text{sr}(X) \geq \frac{C\tau^{4\ell} k^3}{\epsilon^2} \log\left(\frac{d^\ell}{k}\right)$. Then with probability at least $1 - \exp(-c\epsilon^2 \text{sr}(X)/(k^2\tau^{4\ell}))$, the matrix $XR^{(\ell)}$ satisfies the (k, ϵ) -RIP property, i.e., for any $\mathbf{u} \in S^{\binom{d}{\ell}-1}$ with $|\text{supp}(\mathbf{u})| \leq k$,*

$$(1 - \epsilon) \|X\|_{\text{F}} \leq \left\|XR^{(\ell)}\mathbf{u}\right\| \leq (1 + \epsilon) \|X\|_{\text{F}}.$$

Remark 4.4. *It is tempting to conjecture that the dependence on k in the $\text{sr}(X)$ condition in both Theorems 4.1 and 4.3 should be linear, but as mentioned above analyzing these matrix families appears challenging, and it is plausible that a superlinear dependence on k might in fact be unavoidable.*

4.1 Application to Understanding Effectiveness of Word Embeddings

Word embeddings which represent the ‘‘meaning’’ of each word via a low-dimensional vector, have been widely utilized for many natural language processing and machine learning applications. We refer the reader to [Mikolov et al., 2013, Pennington et al., 2014] for additional background about word vector embeddings. Individual word embeddings can be extended to embed word sequences (such as a phrase or sentence) in multiple ways. There has been a recent effort to better understand the effectiveness of these embeddings, in terms of the information they

encode and how this relates to performance on downstream tasks [Arora et al., 2017, 2018a, Khodak et al., 2018]. In this paper, we work with a recently introduced word sequence embedding scheme, called *distributed cooccurrence* (DisC) embedding, that has been shown to be empirically effective for downstream classification tasks and also supports some theoretical justification [Arora et al., 2018a].

We in fact investigate linear transformations of these DisC embedding vectors (Definition 7). Linear transformations are commonly applied over existing embeddings to construct new embeddings in the context of *domain adaptation*, *transfer learning*, etc. For example, recently [Khodak et al., 2018] applied a (learnt) linear transformation on the DisC embedding vectors to construct a new embedding scheme, referred to as *à la carte* embedding, which they empirically show regularly outperforms the DisC embedding. Our results show, under some conditions, these linearly transformed DisC embeddings have provable performance guarantees for linear classification. Before stating our result formally, we need some definitions.

Let \mathcal{W} denote the vocabulary set (collection of some words) with $|\mathcal{W}| = d$. We assume each word $\mathbf{w} \in \mathcal{W}$ has a vector representation $\mathbf{v}_w \in \mathbb{R}^p$. Let $V \in \mathbb{R}^{p \times d}$ denote the matrix whose columns are these word embeddings. A ℓ -gram is a *contiguous* sequence of ℓ words in a longer word sequence. Here, $\ell = 2$ is referred commonly to as bigram, $\ell = 3$ is referred to as trigram, etc.¹² The Bag-of- L -grams representation of a word sequence is a vector that counts the number of times any possible ℓ -gram (from the vocabulary set \mathcal{W}) for some $\ell \in [L]$ appears in the sequence. In practice, L is set to a small value, typically ≤ 4 . Tweaked versions of this simple representation is known to perform well for many downstream classification tasks [Wang and Manning, 2012]. It is common to ignore the ordering of words in an ℓ -gram and to define ℓ -gram as an unordered collection of ℓ words (also referred to as ℓ -cooccurrence). Also as in [Arora et al., 2018a] for simplicity, we assume that each ℓ -cooccurrence contains a word at most once (as noted by [Arora et al., 2018a] this can be ensured by merging words during a preprocessing step).

Definition 5 (Bag-of- L -cooccurrences). *Given a k -word sequence $S = (w_1, \dots, w_k)$ and $L \in \mathbb{N}$, we define its Bag-of- L -cooccurrences vector \mathbf{c}_S as the concatenation of vectors $\mathbf{c}^{(1)}, \dots, \mathbf{c}^{(L)}$ where*

$$\forall \ell \in [L], \mathbf{c}^{(\ell)} \in \mathbb{R}^{\binom{d}{\ell}} \text{ defined as } \mathbf{c}^{(\ell)} = \sum_{t=1}^{k-\ell+1} \mathbf{e}_{\{w_t, \dots, w_{t+\ell-1}\}},$$

Here, \mathbf{c}_S is a $\sum_{\ell=1}^L \binom{d}{\ell}$ dimensional vector. Here, $\binom{d}{\ell}$ is number of possible ℓ -cooccurrences in a d word vocabulary set.

In practice, one would also consider only small word sequences, and therefore k will be a small number. The embedding of an ℓ -cooccurrence is defined as the elementwise product of the embeddings of its constituent words.

Definition 6 (DisC Embedding [Arora et al., 2018a]). *Given a k -word sequence $S = (w_1, \dots, w_k)$ and $L \in \mathbb{N}$, we define its L -DisC embedding \mathbf{v}_{DisC} as the pL -dimensional vector formed by concatenation of vectors $\mathbf{v}^{(1)}, \dots, \mathbf{v}^{(L)} \in \mathbb{R}^p$ where*

$$\forall \ell \in [L], \mathbf{v}^{(\ell)} \in \mathbb{R}^p \text{ defined as } \mathbf{v}^{(\ell)} = \sum_{t=1}^{k-\ell+1} \mathbf{v}_{w_t} \odot \dots \odot \mathbf{v}_{w_{t+\ell-1}},$$

i.e., $\mathbf{v}^{(\ell)}$ is the sum of the ℓ -cooccurrence embeddings of all ℓ -cooccurrences in the document.

¹²For example, if the word sequence equals (“the”, “cow”, “jumps”, “over”, “the”, “moon”), then it has 5 different words. The collection of 2-grams (bigrams) for this sequence would be “the cow”, “cow jumps”, ..., “the moon”. Similarly, the collection of 3-grams (trigrams) would be “the cow jumps”, “cow jumps over”, ..., “over the moon”. In this case, the collection of 2-cooccurrences and 3-cooccurrences are same as the collection of 2-grams and 3-grams respectively.

From Definitions 4, 5, and 6, it can be observed that if the columns of $V \in \mathbb{R}^{p \times d}$ are the words embeddings of the words in vocabulary set \mathcal{W} , then

$$\mathbf{v}^{(\ell)} = V^{(\ell)} \mathbf{c}^{(\ell)},$$

where $V^{(\ell)}$ is the ℓ -way column Hadamard-product constructed out of V . The columns of $V^{(\ell)}$ contain the DisC embedding of all possible ℓ -cooccurrences in the vocabulary set \mathcal{W} (and thus $V^{(1)} = V$).

Definition 7 (Linearly Transformed DisC Embedding). *Given a matrix of word vectors $V \in \mathbb{R}^{p \times d}$ and a set of matrices $X^{(1)}, \dots, X^{(L)} \in \mathbb{R}^{n \times p}$. A linearly transformed L -DisC embedding matrix $\mathbb{V}^{(L)}$ of V is a block-diagonal matrix with blocks $X^{(\ell)} V^{(\ell)}, \ell \in [L]$*

$$\mathbb{V}^{(L)} = \begin{bmatrix} X^{(1)} V^{(1)} & \mathbf{0}_{n \times \binom{d}{2}} & \dots & \mathbf{0}_{n \times \binom{d}{L}} \\ \mathbf{0}_{n \times d} & X^{(2)} V^{(2)} & \dots & \mathbf{0}_{n \times \binom{d}{L}} \\ \dots & \dots & \dots & \dots \\ \mathbf{0}_{n \times d} & \mathbf{0}_{n \times \binom{d}{2}} & \dots & X^{(L)} V^{(L)} \end{bmatrix} \quad (5)$$

For a k -word sequence $S = (w_1, \dots, w_k)$, the linearly transformed DisC embedding is then defined as $\mathbf{v}_{\widetilde{\text{DisC}}} = \mathbb{V}^{(L)} \mathbf{c}_S$, where \mathbf{c}_S is the Bag-of- L -cooccurrences vector.

Note that $X^{(1)}, \dots, X^{(L)}$ are fixed matrices, and can be used to encode any contextual or prior information. If they are all identity matrices, then $\mathbf{v}_{\widetilde{\text{DisC}}} = \mathbf{v}_{\text{DisC}}$.

Compressed Learning Problem. Notice that the dimensionality of \mathbf{c}_S is $\sum_{\ell=1}^L \binom{d}{\ell}$, whereas that of $\mathbf{v}_{\widetilde{\text{DisC}}}$ is only nL , i.e., $\mathbf{v}_{\widetilde{\text{DisC}}}$ is a compressed representation of \mathbf{c}_S . Therefore, there is inherent computational advantage of working with $\mathbf{v}_{\widetilde{\text{DisC}}}$ than \mathbf{c}_S , and $\mathbf{v}_{\widetilde{\text{DisC}}}$ can also be computed efficiently in practice [Khodak et al., 2018]. The question then arises is: How well does a linear classifier trained on this compressed representation perform say compared to a similar classifier trained on the uncompressed (original) representations?¹³ For theoretically answering this question, we use random vectors as word embeddings, instead of pretrained vectors, as also considered by Arora et al. [2018a] in their analysis of DisC embeddings.¹⁴ Our analysis uses a result from [Arora et al., 2018a] (Theorem D.2), that bounds the loss of a linear classifier trained on the compressed domain compared to a similar classifier trained on the original domain, along with the RIP result established in Theorems 3.1 ($\ell = 1$), 4.1 ($\ell = 2$), and 4.3 ($\ell \geq 3$). To formally state the results, we need some more notation. L is assumed to be a constant below. Some additional background about this problem of compressed learning is provided in Appendix D.1.

We consider the standard binary classification task using labeled data. Using the notation introduced earlier, let \mathcal{X} denote the set $\{(\mathbf{c}, v) : \mathbf{c} \text{ is a Bag-of-}L\text{-cooccurrences vector of a word sequence of length at most } k, \|\mathbf{c}\| \leq \alpha, v \in \{-1, 1\}\}$, with v indicating the label on \mathbf{c} . Note that \mathbf{c} is $h = \sum_{\ell=1}^L \binom{d}{\ell}$ dimensional vector. Let \mathcal{D} be a distribution over \mathcal{X} . Consider a labeled dataset $(\mathbf{c}_1, v_1), \dots, (\mathbf{c}_b, v_b)$ drawn i.i.d. from \mathcal{D} . Consider a linear loss function, $f(\langle \mathbf{c}, \theta \rangle; v)$ for $\theta \in \mathbb{R}^h$, where $f : \mathbb{R} \times \{-1, 1\} \rightarrow \mathbb{R}$ is assumed to be convex and λ -Lipschitz in the first parameter. Define, $f_{\mathcal{D}}(\theta) = \mathbb{E}_{(\mathbf{c}, v) \sim \mathcal{D}} [f(\langle \mathbf{c}, \theta \rangle; v)]$ as the generalization (distributional) loss on the original Bag-of- L -cooccurrences domain for parameter θ . Let $\theta^* \in \text{argmin}_{\theta \in \mathbb{R}^h} f_{\mathcal{D}}(\theta)$. Under this notation, the following corollary follows from combining Theorems 3.1, 4.1, 4.3, and D.2. It shows (ignoring Lipschitz and scale parameter) that the difference in the generalization loss by operating on the compressed domain vs. operating on the original domain is $\approx O(\sqrt{\epsilon})$. Note that if a word sequence has length at most k , then all its ℓ -cooccurrences will have cardinality at most k .

Corollary 4.5. *Let $V \in \mathbb{R}^{p \times d}$ be an i.i.d. centered τ -bounded random matrix. Let $X^{(1)}, \dots, X^{(L)} \in \mathbb{R}^{n \times p}$ be a set of matrices (picked independent of V) with $\text{sr}(X^{(1)}) = \Omega((k\tau^4/\epsilon^2) \log(d/k))$, $\text{sr}(X^{(2)}) = \Omega((k^2\tau^8/\epsilon^2) \log(d^2/k))$,*

¹³As mentioned earlier, linear classifiers trained on Bag-of- L -grams or Bag-of- L -cooccurrences representations perform well in practice.

¹⁴Arora et al. [2018a] also provide some empirical evidence on the performance of using random signs for word embeddings.

$\mathbf{sr}(X^{(\ell)}) = \Omega((k^3\tau^{4\ell}/\epsilon^2) \log(d^\ell/k))$ for $3 \leq \ell \leq L$. Let $\hat{\vartheta}$ be the minimizer of a classifier trained on the linearly transformed DisC embeddings $(\mathbb{V}^{(L)}\mathbf{c}_1, v_1), \dots, (\mathbb{V}^{(L)}\mathbf{c}_b, v_b)$ defined as $\hat{\vartheta} \in \operatorname{argmin}_{\vartheta \in \mathbb{R}^{nL}} \frac{1}{b} \sum_i f(\langle \mathbb{V}^{(L)}\mathbf{c}_i, \vartheta \rangle; v_i) + \frac{1}{2C} \|\vartheta\|^2$ (for an appropriate choice of C), then with probability at least $(1 - L \exp(-k))(1 - \delta)$, we have

$$f_{\mathcal{D}}(\hat{\vartheta}) \leq f_{\mathcal{D}}(\theta^*) + O\left(\lambda\alpha\|\theta^*\|\sqrt{\epsilon + \frac{1}{b} \log \frac{1}{\delta}}\right),$$

where $f_{\mathcal{D}}(\vartheta) = \mathbb{E}_{(\mathbf{c}, v) \sim \mathcal{D}}[f(\langle \mathbb{V}^{(L)}\mathbf{c}, \vartheta \rangle; y)]$ is the generalization loss on the compressed domain for parameter ϑ .

We end this discussion by noting that we started with random vectors for word embeddings but added the flexibility of incorporating contextual information through some fixed linear transform. While this is a good starting theoretical model for understanding the effectiveness of these embeddings, actual pretrained embeddings used in practice would not satisfy the restricted isometry property (as they will have high coherence). Bridging this gap is an interesting research direction.

References

- Radoslaw Adamczak et al. A note on the hanson-wright inequality for random vectors with dependencies. *Electronic Communications in Probability*, 20, 2015.
- Sanjeev Arora, Yingyu Liang, and Tengyu Ma. A simple but tough-to-beat baseline for sentence embeddings. *ICLR*, 2017.
- Sanjeev Arora, Mikhail Khodak, Nikunj Saunshi, and Kiran Vodrahalli. A compressed sensing view of unsupervised text embeddings, bag-of-n-grams, and lstms. *ICLR*, 2018a.
- Sanjeev Arora, Yuanzhi Li, Yingyu Liang, Tengyu Ma, and Andrej Risteski. Linear algebraic structure of word senses, with applications to polysemy. *Transactions of the Association of Computational Linguistics*, 6:483–495, 2018b.
- Sanjeev Arora, Nadav Cohen, Noah Golowich, and Wei Hu. A convergence analysis of gradient descent for deep linear neural networks. In *ICLR*, 2019.
- Afonso S Bandeira, Edgar Dobriban, Dustin G Mixon, and William F Sawin. Certifying the restricted isometry property is hard. *IEEE transactions on information theory*, 59(6):3448–3450, 2013a.
- Afonso S Bandeira, Matthew Fickus, Dustin G Mixon, and Percy Wong. The road to deterministic matrices with the restricted isometry property. *Journal of Fourier Analysis and Applications*, 19(6):1123–1149, 2013b.
- Richard Baraniuk, Mark Davenport, Ronald DeVore, and Michael Wakin. A simple proof of the restricted isometry property for random matrices. *Constructive Approximation*, 28(3):253–263, 2008.
- Thomas Blumensath and Mike Davies. Compressed sensing and source separation. In *International Conference on Independent Component Analysis and Signal Separation*, pages 341–348. Springer, 2007.
- Danushka Bollegala, Kohei Hayashi, and Ken-ichi Kawarabayashi. Learning linear transformations between counting-based and prediction-based word embeddings. *PLoS one*, 12(9):e0184544, 2017.
- Jean Bourgain, Stephen Dilworth, Kevin Ford, Sergei Konyagin, Denka Kutzarova, et al. Explicit constructions of rip matrices and related problems. *Duke Mathematical Journal*, 159(1):145–185, 2011.
- Robert Calderbank, Sina Jafarpour, and Robert Schapire. Compressed learning: Universal sparse dimensionality reduction and learning in the measurement domain. *preprint*, 2009.
- Emmanuel J Candès. The restricted isometry property and its implications for compressed sensing. *Comptes Rendus Mathématique*, 346(9):589–592, 2008.

- Emmanuel J Candes and Terence Tao. Decoding by linear programming. *IEEE transactions on information theory*, 51(12):4203–4215, 2005.
- Yonina C Eldar and Gitta Kutyniok. *Compressed sensing: theory and applications*. Cambridge University Press, 2012.
- Simon Foucart and Holger Rauhut. A mathematical introduction to compressive sensing. *Bull. Amer. Math. Soc*, 54:151–165, 2017.
- Simon Foucart, Alain Pajor, Holger Rauhut, and Tino Ullrich. The gelfand widths of ℓ_p -balls for $0 < p \leq 1$. *Journal of Complexity*, 26(6):629–640, 2010.
- David Gamarnik. Explicit construction of rip matrices is ramsey-hard. *arXiv preprint arXiv:1805.11238*, 2018.
- David Lee Hanson and Farroll Tim Wright. A bound on tail probabilities for quadratic forms in independent random variables. *The Annals of Mathematical Statistics*, 42(3):1079–1083, 1971.
- William B Johnson and Joram Lindenstrauss. Extensions of lipschitz mappings into a hilbert space. *Contemporary mathematics*, 26(189-206):1, 1984.
- Daniel Kane, Raghu Meka, and Jelani Nelson. Almost optimal explicit johnson-lindenstrauss families. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 628–639. Springer, 2011.
- Mikhail Khodak, Nikunj Saunshi, Yingyu Liang, Tengyu Ma, Brandon Stewart, and Sanjeev Arora. A la carte embedding: Cheap but effective induction of semantic feature vectors. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, volume 1, pages 12–22, 2018.
- Felix Krahmer and Rachel Ward. New and improved johnson-lindenstrauss embeddings via the restricted isometry property. *SIAM Journal on Mathematical Analysis*, 43(3):1269–1281, 2011.
- Thomas Laurent and James Brecht. Deep linear networks with arbitrary loss: All local minima are global. In *International Conference on Machine Learning*, pages 2908–2913, 2018.
- Michel Ledoux. *The concentration of measure phenomenon*. American Mathematical Soc., 2001.
- Kyle Luh and Van Vu. Dictionary learning with few samples and matrix concentration. *IEEE Transactions on Information Theory*, 62(3):1516–1527, 2016.
- Andrew L Maas, Raymond E Daly, Peter T Pham, Dan Huang, Andrew Y Ng, and Christopher Potts. Learning word vectors for sentiment analysis. In *Proceedings of the 49th annual meeting of the association for computational linguistics: Human language technologies-volume 1*, pages 142–150. Association for Computational Linguistics, 2011.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119, 2013.
- Rajeev Motwani and Prabhakar Raghavan. *Randomized algorithms*. Cambridge university press, 1995.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543, 2014.
- Mark Rudelson and Roman Vershynin. On sparse reconstruction from fourier and gaussian measurements. *Communications on Pure and Applied Mathematics*, 61(8):1025–1045, 2008.
- Mark Rudelson and Roman Vershynin. Hanson-wright inequality and sub-gaussian concentration. *Electronic Communications in Probability*, 18, 2013.
- Andrew M Saxe, Pang Wei Koh, Zhenghao Chen, Maneesh Bhand, Bipin Suresh, and Andrew Y Ng. On random weights and unsupervised feature learning. In *ICML*, 2011.

- Daniel A Spielman, Huan Wang, and John Wright. Exact recovery of sparsely-used dictionaries. In *Conference on Learning Theory*, pages 37–1, 2012.
- Michel Talagrand. An isoperimetric theorem on the cube and the kintchine-kahane inequalities. *Proceedings of the American Mathematical Society*, 104(3):905–909, 1988.
- Michel Talagrand. Concentration of measure and isoperimetric inequalities in product spaces. *Publications Mathématiques de l’Institut des Hautes Etudes Scientifiques*, 81(1):73–205, 1995.
- Roman Vershynin. Introduction to the Non-asymptotic Analysis of Random Matrices. *arXiv preprint arXiv:1011.3027*, 2010.
- Roman Vershynin. Spectral norm of products of random and deterministic matrices. *Probability theory and related fields*, 150(3-4):471–509, 2011.
- Roman Vershynin. High dimensional probability. *An Introduction with Applications*, 2016.
- Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.
- Liming Wang and Yuejie Chi. Blind deconvolution from multiple sparse inputs. *IEEE Signal Processing Letters*, 23(10):1384–1388, 2016.
- Sida Wang and Christopher D Manning. Baselines and bigrams: Simple, good sentiment and topic classification. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Short Papers-Volume 2*, pages 90–94. Association for Computational Linguistics, 2012.
- Tengyao Wang, Quentin Berthet, and Yaniv Plan. Average-case hardness of rip certification. In *Advances in Neural Information Processing Systems*, pages 3819–3827, 2016.
- David P. Woodruff. Sketching as a tool for numerical linear algebra. *FoT-TCS*, 10(1–2):1–157, 2014. URL <http://dx.doi.org/10.1561/04000000060>.
- Guangyou Zhou, Tingting He, Jun Zhao, and Po Hu. Learning continuous word embedding with metadata for question retrieval in community question answering. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, volume 1, pages 250–259, 2015.
- Shuheng Zhou. Sparse hanson-wright inequalities for subgaussian quadratic forms. *arXiv preprint arXiv:1510.05517*, 2015.

A Source Separation Construction

In this section, we explain how the source separation problem can be recast into a setting where we have the product of fixed and random matrices. Let $S = MC\Phi$, where $M \in \mathbb{R}^{n_1 \times a}$, $C \in \mathbb{R}^{a \times b}$, and $\Phi \in \mathbb{R}^{b \times c}$. Let $\mathbf{s}_1^\top, \dots, \mathbf{s}_{n_1}^\top$ be the rows of the matrix S . Let $\mathbf{c}_1^\top, \dots, \mathbf{c}_a^\top$ be the rows of the matrix C . Let $M = (M_{ij})$. Let $A \in \mathbb{R}^{n_1 \times ac}$ be a matrix such that M_{ij} equals $A_{i,(j-1)c+1}$ for all $i \in [n_1]$ and $j \in [a]$, and all other entries of A equals 0. It follows from this construction that:

$$\begin{bmatrix} \mathbf{s}_1 \\ \vdots \\ \mathbf{s}_{n_1} \end{bmatrix}_{n_1 c} = \begin{bmatrix} A \\ \vdots \\ A \end{bmatrix}_{n_1 c \times ac} \begin{bmatrix} \Phi^\top & & 0 \\ & \ddots & \\ 0 & & \Phi^\top \end{bmatrix}_{ac \times ab} \begin{bmatrix} \mathbf{c}_1 \\ \vdots \\ \mathbf{c}_a \end{bmatrix}_{ab}.$$

Let us further define matrices $X_i \in \mathbb{R}^{n_1 \times c}$ as

$$\begin{bmatrix} A \\ \vdots \\ A \end{bmatrix} = [X_1 | X_2 | \dots | X_a].$$

Now it is easy to observe that

$$\begin{bmatrix} A \\ \vdots \\ A \end{bmatrix} \begin{bmatrix} \Phi^\top & & 0 \\ & \ddots & \\ 0 & & \Phi^\top \end{bmatrix} = \underbrace{\begin{bmatrix} X_1 & X_2 & \dots & X_a \end{bmatrix}}_X \underbrace{\begin{bmatrix} \Phi^\top & & 0 \\ & \ddots & \\ 0 & & \Phi^\top \end{bmatrix}}_R = [X_1\Phi^\top | X_2\Phi^\top | \dots | X_a\Phi^\top].$$

Notice that if the dictionary matrix Φ is a random matrix, then each $X_i\Phi^\top$ is a product of a fixed and random matrix. Therefore, each of them fit in the product model considered in this paper. Also it follows, that if each $X_i\Phi^\top$ satisfies (k, ϵ) -RIP with probability $1 - \gamma$, then $[X_1\Phi^\top | X_2\Phi^\top | \dots | X_a\Phi^\top]$ satisfies (k, ϵ) -RIP with probability $1 - a\gamma$.

B Missing Preliminaries

Subgaussian and Subexponential Random Variables. The class of subgaussian random variables is natural and quite wide. Let us start by formally defining subgaussian random variables and vectors. A number of equivalent definitions are used in the literature.

Definition 8 (Subgaussian Random Variable and Vector). *We call a random variable $x \in \mathbb{R}$ subgaussian if there exists a constant $C > 0$ if $\Pr[|x| > t] \leq 2 \exp(-t^2/C^2)$ for all $t \geq 0$. We say that a random vector $\mathbf{x} \in \mathbb{R}^d$ is subgaussian if the one-dimensional marginals $\langle \mathbf{x}, \mathbf{y} \rangle$ are subgaussian random variables for all $\mathbf{y} \in \mathbb{R}^d$.*

The class of subgaussian random variables includes many random variables that arise naturally in data analysis, such as standard normal, Bernoulli, spherical, bounded (where the random variable x satisfies $|x| \leq M$ almost surely for some fixed M). The natural generalizations of these random variables to higher dimension are all subgaussian random vectors. For many *isotropic convex sets*¹⁵ \mathcal{K} (such as the hypercube), a random vector \mathbf{x} uniformly distributed in \mathcal{K} is subgaussian.

Definition 9 (ψ_2 -norm of a Subgaussian Random Variable and Vector). *The ψ_2 -norm of a subgaussian random variable $x \in \mathbb{R}$, denoted by $\|x\|_{\psi_2}$ is:*

$$\|x\|_{\psi_2} = \sup_{a \geq 1} a^{-1/2} (\mathbb{E}[|x|^a])^{1/a}.$$

The ψ_2 -norm of a subgaussian random vector $\mathbf{x} \in \mathbb{R}^d$ is:

$$\|\mathbf{x}\|_{\psi_2} = \sup_{\mathbf{y} \in S^{d-1}} \|\langle \mathbf{x}, \mathbf{y} \rangle\|_{\psi_2}.$$

We also will work with a class of subexponential random variables, those with at least an exponential tail.

Definition 10 (Subexponential Random Variable and ψ_1 -norm). *A random variable x that satisfies $\Pr[|x| > t] \leq 2 \exp(-t/C)$ for all $t \geq 0$ is called a subexponential random variable. The subexponential norm of x , denoted $\|x\|_{\psi_1}$, is defined as $\|x\|_{\psi_1} = \sup_{a \geq 1} a^{-1} (\mathbb{E}[|x|^a])^{1/a}$.*

An immediate consequence from the above definitions of subgaussian and subexponential random variables is that,

$$\text{For subgaussian } x : (E[|x|^a])^{1/a} \leq \|x\|_{\psi_2} \sqrt{a}, \quad \forall a \geq 1$$

$$\text{For subexponential } x : (E[|x|^a])^{1/a} \leq \|x\|_{\psi_1} a, \quad \forall a \geq 1.$$

¹⁵A convex set \mathcal{K} in \mathbb{R}^d is called isotropic if a random vector chosen uniformly from \mathcal{K} according to the volume is isotropic. A random vector $\mathbf{x} \in \mathbb{R}^d$ is isotropic if for all $\mathbf{y} \in \mathbb{R}^d$, $\mathbb{E}[\langle \mathbf{x}, \mathbf{y} \rangle^2] = \|\mathbf{y}\|^2$.

Corollaries from Hanson-Wright Inequality (Theorem 2.1). A simple corollary of Hanson-Wright inequality from Theorem 2.1 is a concentration inequality for random vectors with independent subgaussian components.

Corollary B.1 (Subgaussian Concentration [Rudelson and Vershynin, 2013]). *Let M be a fixed $n \times d$ matrix. Let $\mathbf{x} = (x_1, \dots, x_n) \in \mathbb{R}^n$ be a random vector with independent components x_i which satisfies $\mathbb{E}[x_i] = 0$, $\mathbb{E}[x_i^2] = 1$ and $\|x_i\|_{\psi_2} \leq K$. Then for every $t \geq 0$,*

$$\Pr[|\|M\mathbf{x}\| - \|M\|_F| > t] \leq 2 \exp\left(\frac{-ct^2}{K^4\|M\|^2}\right).$$

Another corollary of this inequality is a bound on the spectral norm of product of deterministic and random matrices.

Corollary B.2 (Spectral Norm of the Product [Rudelson and Vershynin, 2013]). *Let B be a fixed $n \times p$ matrix, and let $G = (G_{ij})$ be a $p \times d$ random matrix with independent entries that satisfy: $\mathbb{E}[G_{ij}] = 0$, $\mathbb{E}[G_{ij}^2] = 1$, and $\|G_{ij}\|_{\psi_2} \leq K$. Then for any $a, b > 0$,*

$$\Pr\left[\|BG\| > CK^2(a\|B\|_F + b\sqrt{d}\|B\|)\right] \leq 2 \exp(-a^2 \mathbf{sr}(B) - b^2 d)$$

RIP and Compressed Sensing. Common random matrices such as Gaussian, Bernoulli, Fourier all satisfy the restricted isometry property. For example, if the entries of $M \in \mathbb{R}^{n \times d}$ are independent and identically distributed subgaussian random variables with zero mean and unit variance. Assume that $n \geq Ck \log(2d/k)$ where C depends only on ϵ, γ , and the subgaussian moment. Then with probability at least $1 - \gamma$, the matrix $\tilde{M} = M/\sqrt{n}$ satisfies the restricted isometry property.

The following classic result, shows that one can recover from compressed signal (even in presence of noise) if the measurement matrix satisfies RIP.

Theorem B.3 ([Theorem 1.2, [Candès, 2008]). *Suppose that M satisfies the $(2k, \epsilon)$ -RIP with $\epsilon < \sqrt{2} - 1$ and let $\mathbf{y} = M\mathbf{x} + \mathbf{e}$ where $\|\mathbf{e}\| \leq \delta$. Let $\hat{\mathbf{x}} = \operatorname{argmin}_{\mathbf{z}} \|\mathbf{z}\|_1$ subject to $\|M\mathbf{z} - \mathbf{y}\| \leq \delta$. Then $\hat{\mathbf{x}}$ obeys*

$$\|\hat{\mathbf{x}} - \mathbf{x}\| \leq C_0 \frac{\sigma(\mathbf{x})_1}{\sqrt{k}} + C_2 \delta \quad \text{where,}$$

$$C_0 = 2 \frac{1 - (1 - \sqrt{2})\epsilon}{1 - (1 + \sqrt{2})\epsilon}, \quad C_2 = 4 \frac{\sqrt{1 + \epsilon}}{1 - (1 + \sqrt{2})\epsilon}, \quad \text{and } \sigma(\mathbf{x})_1 = \min_{k\text{-sparse } \mathbf{x}'} \|\mathbf{x} - \mathbf{x}'\|_1.$$

Restricted Isometry under the XR - vs. the RX -model. Our results establish the restricted isometry condition for a matrix that can be factorized as XR , for a deterministic X and a random R , under some conditions on X and R . One could ask what happens if we change the model and work in a setting where the factorization is of the form RX (again for random R , and deterministic X). Unfortunately, in this case one could not hope to get a mild assumption on X . In fact, a necessary and sufficient condition on X for a matrix RX to satisfy RIP is that X satisfies RIP. The necessary condition is trivial. Indeed, if $\|X\mathbf{u}\|$ is essentially different from $\|X\mathbf{v}\|$ for two sparse unit vectors \mathbf{u}, \mathbf{v} , then $\mathbb{E}[\|RX\mathbf{u}\|]$ will be essentially different from $\mathbb{E}[\|RX\mathbf{v}\|]$ (unless, of course, X satisfies RIP). The proof for the sufficiency is not much harder, and follows from standard techniques in the literature.

C Missing Details from Section 3

We provide missing details from Section 3. We start with the proof of Theorem 3.1.

Theorem C.1 (Theorem 3.1 Restated). *Let X be an $n \times p$ matrix. Let $R = (R_{ij})$ be a $p \times d$ matrix whose entries are with independent entries such that $\mathbb{E}[R_{ij}] = 0$, $\mathbb{E}[R_{ij}^2] = 1$, and $\|R_{ij}\|_{\psi_2} \leq K$. Let $\epsilon \in (0, 1)$, and let $k \in \mathbb{N}$ be a number satisfying $\text{sr}(X) \geq CK^4 \frac{k}{\epsilon^2} \log\left(\frac{d}{k}\right)$. Then with probability at least $1 - \exp(-c\epsilon^2 \text{sr}(X)/K^4)$, the matrix XR satisfies (k, ϵ) -RIP, i.e.,*

$$\forall \mathbf{u} \in \Sigma_k, (1 - \epsilon) \|X\|_F \leq \|XR\mathbf{u}\| \leq (1 + \epsilon) \|X\|_F.$$

Proof. First, we prove a spectral norm bound for every $n \times k$ submatrix of XR . Fix a subset $I \subset [d]$ with $|I| = k$. Let R_I be the matrix R restricted to the columns indexed by I . Using Corollary B.2 (with $a = \epsilon/K^2$, $b = \epsilon/K^2$), $\|XR_I\|$ satisfies

$$\Pr \left[\|XR_I\| > C_1 \epsilon (\|X\|_F + \sqrt{k} \|X\|) \right] \leq 2 \exp(-\text{sr}(X) \epsilon^2 / K^4 - k \epsilon^2 / K^4).$$

By assumption, $\text{sr}(X) \geq k \log(d/k)$, therefore $\|X\|_F \geq \sqrt{k \log(d/k)} \|X\| \geq \sqrt{k} \|X\|$. Then

$$\Pr [\|XR_I\| > 2C_1 \epsilon \|X\|_F] \leq 2 \exp(-\text{sr}(X) \epsilon^2 / K^4).$$

Using a union bound, and the fact that $\binom{d}{k} \leq \left(\frac{de}{k}\right)^k$,

$$\begin{aligned} \Pr [\exists I \in [d], |I| = k, \|XR_I\| > 2C_1 \epsilon \|X\|_F] &\leq \binom{d}{k} \cdot 2 \exp(\text{sr}(X) \epsilon^2 / K^4) \\ &\leq 2 \exp(-\text{sr}(X) \epsilon^2 / K^4 + k \log(ed/k)) \leq 2 \exp(-c_0 \text{sr}(X) \epsilon^2 / K^4), \end{aligned} \quad (6)$$

for appropriate choice of constants.

Let us now fix an $\mathbf{u} \in S^{d-1}$. A simple consequence of the Hanson-Wright inequality (see Theorem 3.2, [Rudelson and Vershynin, 2013]) is that for any $t \geq 0$,

$$\Pr [|\|XR\mathbf{u}\| - \|X\|_F| \geq t] \leq 2 \exp\left(\frac{-c_1 t^2}{K^4 \|X\|^2}\right).$$

Setting $t = \epsilon \|X\|_F$, we get

$$\Pr [|\|XR\mathbf{u}\| - \|X\|_F| \geq \epsilon \|X\|_F] \leq 2 \exp\left(\frac{-c_1 \epsilon^2 \text{sr}(X)}{K^4}\right).$$

Let Σ_k be the set of all k -sparse vectors in S^{d-1} . By a volumetric estimate, this set has an $(1/2C_2)$ -net \mathcal{N} of cardinality smaller than

$$\binom{d}{k} \cdot (6C_2)^k \leq \exp\left(k \log \frac{C_0 d}{k}\right).$$

Taking the union bound over this net, we obtain

$$\Pr [\forall \mathbf{u} \in \mathcal{N}, |\|XR\mathbf{u}\| - \|X\|_F| \leq \epsilon \|X\|_F] \geq 1 - \exp\left(\frac{-c_1 \epsilon^2 \text{sr}(X)}{K^4} + k \log \frac{C_0 d}{k}\right). \quad (7)$$

We now use a simple approximation idea for extending the above argument from the net to all k -sparse vectors. Let us first assume that events described in (7) and (6) happen. We can write any $\mathbf{u} \in \Sigma_k$ as $\mathbf{u} = \mathbf{a} + \mathbf{b}$, where \mathbf{a} in \mathcal{N} , and \mathbf{b} is such that $|\text{supp}(\mathbf{b})| \leq k$ and $\|\mathbf{b}\| \leq 1/(2C_2)$. Let $I_{\mathbf{b}} = \text{supp}(\mathbf{b}) \subset [d]$. Let $\tilde{\mathbf{b}}$ be \mathbf{b} restricted to $\text{supp}(\mathbf{b})$.

$$\begin{aligned} \|XR\mathbf{u}\| &= \|XR\mathbf{a} + XR\mathbf{b}\| \leq \|XR\mathbf{a}\| + \|XR\mathbf{b}\| = \|XR\mathbf{a}\| + \|XR_{I_{\mathbf{b}}}\tilde{\mathbf{b}}\| \\ &\leq \|XR\mathbf{a}\| + \|XR_{I_{\mathbf{b}}}\|\tilde{\mathbf{b}}\| \leq (1 + \epsilon) \|X\|_F + \frac{1}{2C_2} \|XR_{I_{\mathbf{b}}}\| \\ &\leq (1 + \epsilon_1) \|X\|_F, \end{aligned}$$

where for the spectral norm bound for $\|XR_{I_b}\|$ we use (6) and the bound on $\|XR_{I_a}\|$ follows from (7). Similarly,

$$\|XR_{\mathbf{u}}\| \geq (1 - \epsilon)\|X\|_F - \frac{\epsilon}{2C_2}\|XR_{I_b}\| \geq (1 - \epsilon_2)\|X\|_F.$$

Adjusting the constants, and removing the conditioning completes the proof of the theorem. \square

C.1 Restricted Isometry of XR with Low Randomness R

In this section, we operate under a weaker randomness assumption on R . In particular, we will use the notion of l -wise independence to capture low randomness to construct R . When truly random bits are costly to generate or supplying them in advance requires too much space, the standard idea is to use l -wise independence which allows one to maintain a succinct data structure for storing the random bits.

Definition 11. *A sequence of random variables x_1, \dots, x_m is called l -wise independent if every l of them are independent. More formally, x_1, \dots, x_m drawn from some distribution \mathcal{D} over a range Υ if for all i_1, i_2, \dots, i_l (all unique) and $t_1, \dots, t_l \in \Upsilon$,*

$$\Pr_{x_1, \dots, x_m \sim \mathcal{D}}[x_{i_1} = t_1, \dots, x_{i_l} = t_l] = \Pr[x_{i_1} = t_1] \cdots \Pr[x_{i_l} = t_l].$$

For simplicity, we will work with Rademacher random variables (random signs). Constructing n l -wise independent random signs from $O(l \log(n))$ truly independent random signs using simple families of hash functions is a well-known idea [Motwani and Raghavan, 1995].

Let R satisfy $2 \text{sr}(X)$ -wise independence. Our general strategy in this proof will be to rely on higher moments where we can treat certain variables as independent. We start with $Q \in \mathbb{R}^{p \times k}$ with i.i.d. ± 1 entries, and establish an bound on $\mathbb{E}[\|XQ\|^{2 \text{sr}(X)}]$. Using Markov's inequality for higher moments and a union bound gives our first result, $\Pr[\exists I \subset [d], |I| = k, \|XR_I\| > C_0 \epsilon \|X\|_F] \leq \exp(-c_0 \epsilon^2 \text{sr}(X))$ (Lemma C.2). We then derive a concentration bound for $\|XR_{\mathbf{u}}\|$ for a fixed $\mathbf{u} \in S^{d-1}$. For this, we investigate certain higher moments of $\|XR_{\mathbf{u}}\|^2 - \mathbb{E}[\|XR_{\mathbf{u}}\|^2]$ using moment generating functions. The result then follows using a net argument over the set of sparse vectors on the sphere.

Lemma C.2. *Let X be an $n \times p$ matrix and $\epsilon \in (0, 1)$. Let R be a $p \times d$ matrix whose entries are $2 \text{sr}(X)$ -wise independent ± 1 -random variables. Let $k \in \mathbb{N}$ be a number satisfying $\text{sr}(X) \geq C \frac{k}{\epsilon^2} \log\left(\frac{d}{k}\right)$. Then*

$$\Pr[\forall I \subset [d], |I| = k, \|XR_I\| \geq C_0 \epsilon \|X\|_F] \leq \exp(-c_0 \epsilon^2 \text{sr}(X)),$$

where R_I the $p \times k$ submatrix of R with columns from the set I .

Proof. Denote for shortness $m = \text{sr}(X)$. First let us fix $I \subset [d], |I| = k$ and introduce a $p \times k$ matrix Q with i.i.d. ± 1 entries. Then expanding the trace of the power of a matrix and using the $2m$ -wise independence,

$$\mathbb{E}[\|XR_I\|^{2m}] \leq \mathbb{E}[\|XR_I\|_F^{2m}] \leq \mathbb{E}[\text{tr}(R_I^\top X^\top XR_I)^m] = \mathbb{E}[\text{tr}(Q^\top X^\top XQ)^m] = \mathbb{E}[\|XQ\|_F^{2m}] \leq k \mathbb{E}[\|XQ\|^{2m}].$$

Here, we use the fact that since expectation and trace are both linear, they commute, and also the fact the Frobenius norm is at most the rank times the spectral norm.

Since the matrix Q is subgaussian, the last expectation can be estimated using Corollary B.2 (with $a = b = s$) as

$$\Pr\left[\|XQ\| \geq C_1 s (\|X\|_F + \sqrt{k} \|X\|)\right] \leq 2 \exp(-s^2 \text{sr}(X) - s^2 k).$$

By the stable rank assumption on X , $\|X\|_F \geq \sqrt{k} \|X\|$. Setting $s = a/(2C_1 \|X\|_F)$,

$$\Pr[\|XQ\| \geq a] \leq 2 \exp(-ca^2/\|X\|^2).$$

Using this we get,

$$\begin{aligned}
\mathbb{E}[\|XQ\|^{2m}] &\leq (C_1 \|X\|_F)^{2m} + 2 \int_{C_1 \|X\|_F}^{\infty} ma^{2m-1} \Pr[\|XQ\| \geq a] da \\
&\leq (C_1 \|X\|_F)^{2m} + 2 \int_0^{\infty} ma^{2m-1} \exp\left(-c \frac{a^2}{\|X\|^2}\right) da \\
&\leq (C_1 \|X\|_F)^{2m} + C_2^{2m} m^m \|X\|^{2m} \\
&= (C_1 \|X\|_F)^{2m} + C_2^{2m} \left(\frac{\|X\|_F}{\|X\|}\right)^{2m} \|X\|^{2m} \\
&\leq (C'_1 \|X\|_F)^{2m}, \tag{8}
\end{aligned}$$

where the last inequality follows from the assumption on $\text{sr}(X)$. Using Markov's inequality and the union bound, we obtain

$$\begin{aligned}
\Pr[\exists I \subset [d], |I| = k, \|XR_I\| > \epsilon C'_1 \|X\|_F] &\leq \binom{d}{k} \cdot \max_{|I|=k} \frac{\mathbb{E}[\|XR_I\|^{2m}]}{(\epsilon C'_1 \|X\|_F)^{2m}} \\
&\leq k \exp\left(k \log\left(\frac{ed}{k}\right) - 2\epsilon^2 m\right) \leq \exp(-c_0 \epsilon^2 m),
\end{aligned}$$

where we used the uniform bound on $\mathbb{E}[\|XR_I\|^{2m}]$ as established in (8). \square

Theorem C.3. *Let X be an $n \times p$ matrix. Let R be a $p \times d$ matrix whose entries are $2 \text{sr}(X)$ -wise independent ± 1 -random variables. Let $\epsilon \in (0, 1)$, and let $k \in \mathbb{N}$ be a number satisfying $\text{sr}(X) \geq C \frac{k}{\epsilon^2} \log\left(\frac{d}{k}\right)$. Then with probability at least $1 - \exp(-c\epsilon^2 \text{sr}(X))$, the matrix XR satisfies (k, ϵ) -RIP, i.e.,*

$$\forall \mathbf{u} \in \Sigma_k, (1 - \epsilon) \|X\|_F \leq \|XR\mathbf{u}\| \leq (1 + \epsilon) \|X\|_F.$$

Proof. We derive a small ball probability estimate. Fix $I \subset [d]$, $|I| = k$, and let $\mathbf{u} \in S^{d-1}$ be a vector with $\text{supp}(\mathbf{u}) \subset I$. Let $l = c\epsilon^2 \text{sr}(X)$ and let Q be a $p \times k$ matrix Q with i.i.d. ± 1 entries (as in Lemma C.2). Then

$$\mathbb{E}[\mathbf{u}^\top R^\top X^\top XR\mathbf{u}] = \mathbb{E}[\mathbf{u}^\top Q^\top X^\top XQ\mathbf{u}] = \|X\|_F^2.$$

Also the $2 \text{sr}(X)$ -wise independence implies that

$$\mathbb{E}\left[(\mathbf{u}^\top R^\top X^\top XR\mathbf{u} - \mathbb{E}[\mathbf{u}^\top R^\top X^\top XR\mathbf{u}])^{2l}\right] = \mathbb{E}\left[(\mathbf{u}^\top Q^\top X^\top XQ\mathbf{u} - \mathbb{E}[\mathbf{u}^\top Q^\top X^\top XQ\mathbf{u}])^{2l}\right]. \tag{9}$$

Here $Q\mathbf{u}$ is a vector with independent centered subgaussian coordinates. Let $\mathbf{z} = Q\mathbf{u}$, and $A = X^\top X$. Let $A = (A_{ij})$. We can rewrite $\mathbf{u}^\top Q^\top X^\top XQ\mathbf{u} - \mathbb{E}[\mathbf{u}^\top Q^\top X^\top XQ\mathbf{u}]$ as $\mathbf{z}^\top A\mathbf{z} - \mathbb{E}[\mathbf{z}^\top A\mathbf{z}]$. By centering and independence of the entries in \mathbf{z} , we can represent, with $\mathbf{z} = (z_1, \dots, z_p)$.

$$\mathbf{z}^\top A\mathbf{z} - \mathbb{E}[\mathbf{z}^\top A\mathbf{z}] = \sum_{ij} A_{ij} z_i z_j - \sum_i A_{ii} \mathbb{E}[z_i^2] = \sum_i A_{ii} (z_i^2 - \mathbb{E}[z_i^2]) + \sum_{i,j,i \neq j} A_{ij} z_i z_j.$$

We are interested in bounding $\mathbb{E}[(\mathbf{z}^\top A\mathbf{z} - \mathbb{E}[\mathbf{z}^\top A\mathbf{z}])^{2l}]$. Here,

$$\mathbb{E}\left[\left(\sum_i A_{ii} (z_i^2 - \mathbb{E}[z_i^2]) + \sum_{i,j,i \neq j} A_{ij} z_i z_j\right)^{2l}\right] \leq 2^{2l} \left(\mathbb{E}\left[\left(\sum_i A_{ii} (z_i^2 - \mathbb{E}[z_i^2])\right)^{2l}\right] + \mathbb{E}\left[\left(\sum_{i,j,i \neq j} A_{ij} z_i z_j\right)^{2l}\right]\right). \tag{10}$$

1. Let us first focus on bounding, $\mathbb{E}[\sum_i (A_{ii}(z_i^2 - \mathbb{E}[z_i^2]))^{2l}]$. Since, $z_i^2 - \mathbb{E}[z_i^2]$ are all independent mean-zero subexponential random variables¹⁶, and

$$\|z_i^2 - \mathbb{E}[z_i^2]\|_{\Psi_1} \leq 2\|z_i^2\|_{\Psi_1} \leq 4\|z_i\|_{\Psi_2}^2 = 4,$$

as for ± 1 -random variables the Ψ_2 -norm is 1. Let $S_0 = \sum_i A_{ii}(z_i^2 - \mathbb{E}[z_i^2])$. The moment generating function of S_0 can be bound using standard techniques (see [Vershynin, 2010, Proposition 5.16]). We get that $|\lambda| \leq c_1/\max_i |A_{ii}|$,

$$\mathbb{E}[\exp(\lambda S_0)] \leq \exp(c_2 \lambda^2 \sum_{i=1}^p A_{ii}^2).$$

Now using subexponential tail estimates yields,

$$\mathbb{E} \left[\left(\sum_i A_{ii}(z_i^2 - \mathbb{E}[z_i^2]) \right)^{2l} \right] \leq (C'_3 \|A\| \cdot 2l)^{2l} = (C_3 \|X\|^{2l})^{2l}. \quad (11)$$

2. We now bound $\mathbb{E} \left[\left(\sum_{i,j,i \neq j} A_{ij} z_i z_j \right)^{2l} \right]$. Let $S = \sum_{i,j,i \neq j} A_{ij} z_i z_j$. For bounding $\mathbb{E}[S^{2l}]$, we use the following result implicit in [Rudelson and Vershynin, 2013]. The proof is based on decoupling and reduction to normal random variables arguments and is omitted here.

Claim C.4 (From Theorem 1.1, [Rudelson and Vershynin, 2013]). *Let $A = (A_{ij})$ be a $p \times p$ matrix. Let x_1, \dots, x_p be a sequence of independent random variables such that $\mathbb{E}[x_i] = 0$, $\mathbb{E}[x_i^2] = 1$, and $\|x_i\|_{\psi_2}$ is bounded. Let $S = \sum_{i,j,i \neq j} A_{ij} x_i x_j$. Then the moment generating function $\mathbb{E}[\exp(\lambda S)] \leq \exp(C\lambda^2 \|A\|_F^2)$ for all $\lambda \leq c/\|A\|$.*

Claim C.4 along with subgaussian tail estimates yields,

$$\mathbb{E} \left[\left(\sum_{i,j,i \neq j} A_{ij} z_i z_j \right)^{2l} \right] = \mathbb{E}[S^{2l}] = (C'_4 \|A\|_F \cdot \sqrt{2l})^{2l} \leq (C_4 \|X\| \|X\|_F \sqrt{l})^{2l}. \quad (12)$$

Here for the last inequality we used that $\|A\|_F = \|X^\top X\|_F \leq \|X\| \|X\|_F$.

Plugging in the bounds from (11) and (12) in (10) yields,

$$\mathbb{E} \left[\left(\sum_i A_{ii}(z_i^2 - \mathbb{E}[z_i^2]) + \sum_{i,j,i \neq j} A_{ij} z_i z_j \right)^{2l} \right] \leq (C_5 \sqrt{l} \|X\| \|X\|_F)^{2l},$$

as $l < \text{sr}(X)$. In other words (from (9)),

$$\mathbb{E} \left[\left(\mathbf{u}^\top R^\top X^\top X R \mathbf{u} - \mathbb{E}[\mathbf{u}^\top R^\top X^\top X R \mathbf{u}] \right)^{2l} \right] = \mathbb{E} \left[\left(\mathbf{u}^\top R^\top X^\top X R \mathbf{u} - \|X\|_F^2 \right)^{2l} \right] \leq (C_5 \sqrt{l} \|X\| \|X\|_F)^{2l}.$$

Combining this with Markov's inequality,

$$\begin{aligned} \Pr \left[\left| \|\mathbf{u}^\top R \mathbf{u}\| - \|X\|_F \right| \geq \epsilon' \|X\|_F \right] &\leq \Pr \left[\left| \|\mathbf{u}^\top R \mathbf{u}\|^2 - \|X\|_F^2 \right| \geq \epsilon'^2 \|X\|_F^2 \right] \\ &= \Pr \left[\left| \mathbf{u}^\top R^\top X^\top X R \mathbf{u} - \|X\|_F^2 \right| \geq \epsilon \|X\|_F^2 \right] \\ &\leq \frac{(C_5 \sqrt{l} \|X\|_F \|X\|)^{2l}}{(\epsilon \|X\|_F^2)^{2l}} \leq \exp(-2l), \end{aligned}$$

¹⁶A random variable x is subgaussian if and only if x^2 is subexponential, therefore $\|x\|_{\Psi_2}^2 \leq \|x^2\|_{\Psi_1} \leq 2\|x\|_{\Psi_2}^2$.

if the constant c in the definition of l is chosen sufficiently small, and setting $\epsilon'^2 = \epsilon$. For the first inequality, we used the fact that for $a_1, a_2, a_3 \geq 0$, $\Pr[|a_1^2 - a_2^2| \geq a_3^2] \geq \Pr[|a_1 - a_2|^2 \geq a_3^2] = \Pr[|a_1 - a_2| \geq a_3]$.

We now finish the proof using a net argument as in Theorem 3.1. Let \mathcal{N} be an $(1/2C_2)$ -net \mathcal{N} over Σ_k . As mentioned in Theorem 3.1,

$$|\mathcal{N}| \leq \exp\left(k \log \frac{C_0 d}{k}\right).$$

Taking the union bound over this net, we obtain

$$\Pr[\forall \mathbf{u} \in \mathcal{N}, \quad |\|XR\mathbf{u}\| - \|X\|_{\text{F}}| \leq \epsilon' \|X\|_{\text{F}}] \geq 1 - \exp\left(-2l + k \log \frac{C_0 d}{k}\right) \geq 1 - \exp(-l),$$

where we used our choice of l in the last inequality.

Reinitializing ϵ and using the spectral norm bound from Lemma C.2 in conjunction with the approximation idea used in Theorem 3.1 yields that with probability at least $1 - \exp(-c\epsilon^2 \text{sr}(X))$, $|\|XR\mathbf{u}\| - \|X\|_{\text{F}}| \leq \epsilon \|X\|_{\text{F}}$ for all $\mathbf{u} \in \Sigma_k$. The theorem is proved. \square

Comparing Theorems C.3 and 3.1. While the assumption on the stable rank does not change (by more than a constant) between these two theorems, we have a drastic reduction in the number of random bits from $O(pd)$ (in Theorem 3.1) to $O(\text{sr}(X) \log(pd))$ (in Theorem C.3). Storing the XR matrix when R is an i.i.d. random matrix takes $O(np + pd)$ words of memory, while if R is $2 \text{sr}(X)$ -wise independent storing XR only requires $O(np + \text{sr}(X) \log(pd))$ words of memory.

C.2 Restricted Isometry of XR with Sparse Random R

In this section, we investigate the restricted isometry property when R is a sparse random matrix. We use the following popular probabilistic model for our sparse random matrices [Spielman et al., 2012, Luh and Vu, 2016, Wang and Chi, 2016].

Definition 12 (Bernoulli-Subgaussian Sparse Random Matrix). *We say that $R \in \mathbb{R}^{p \times d}$ satisfies the Bernoulli-Subgaussian model with parameter $\beta \in (0, 1)$ if $R = \Omega \odot \Gamma$, where $\Omega \in \mathbb{R}^{p \times d}$ is an i.i.d. Bernoulli matrix where each entry is 1 independently with probability β , and $\Gamma = (\Gamma_{ij})$ is a random matrix with independent entries that satisfy: $\mathbb{E}[\Gamma_{ij}] = 0$, $\mathbb{E}[\Gamma_{ij}^2] = 1$, and $\|\Gamma_{ij}\|_{\psi_2} \leq K$, and \odot denotes the Hadamard product.*

We note that the sparsity of XR is manipulated by the Bernoulli distribution, and the non-zero entries of XR obey the subgaussian distribution, thereby facilitating a very general model of the sparse matrix.

Our results in this section will rely on the sparse Hanson-Wright inequality from Zhou [2015] (Theorem C.5). Similar to Hanson-Wright inequality (Theorem 2.1), sparse Hanson-Wright inequality provides a large deviation bound for a quadratic form. However, in this case, the quadratic form is sparse, and is of the form $(\mathbf{x} \odot \zeta)^\top M (\mathbf{x} \odot \zeta)$ where \mathbf{x} is a random vector with independent centered subgaussian components and ζ contains independent Bernoulli random variables.

Theorem C.5 (Restated from Theorem 1.1 [Zhou, 2015]). *Let $\mathbf{x} = (x_1, \dots, x_d) \in \mathbb{R}^d$ be a random vector with independent components x_i which satisfy $\mathbb{E}[x_i] = 0$ and $\|x_i\|_{\psi_2} \leq K$ is bounded. Let $\zeta = (\zeta_1, \dots, \zeta_d) \in \{0, 1\}^d$ be a random vector independent of \mathbf{x} , with independent Bernoulli random variables ζ_i such that $\mathbb{E}[\zeta_i] = \beta$. Let*

$M = (M_{ij})$ be an $d \times d$ matrix. Then for every $t \geq 0$,¹⁷

$$\Pr \left[\left| (\mathbf{x} \odot \zeta)^\top M (\mathbf{x} \odot \zeta) - \mathbb{E}[(\mathbf{x} \odot \zeta)^\top M (\mathbf{x} \odot \zeta)] \right| > t \right] \leq 2 \exp \left(-c \min \left\{ \frac{t^2}{K^4(\beta \|\text{diag}(M)\|_F^2 + \beta^2 \|\text{offdiag}(M)\|_F^2)}, \frac{t}{K^2 \|M\|} \right\} \right) \leq 2 \exp \left(-c \min \left\{ \frac{t^2}{\beta K^4 \|M\|_F^2}, \frac{t}{K^2 \|M\|} \right\} \right)$$

We now show that using the above sparse Hanson-Wright inequality (Theorem C.5), one can obtain a concentration inequality for sparse random vectors. The following simple lemma follows a proof strategy as that used in Corollary B.1.

Lemma C.6. *Let M be a fixed $n \times d$ matrix. Let \mathbf{x} and ζ be random vectors as in Theorem C.5. Then for any $\epsilon \geq 0$,*

$$\Pr \left[\left| \|M(\mathbf{x} \odot \zeta)\| - \sqrt{\beta} \|M\|_F \right| > \epsilon \sqrt{\beta} \|M\|_F \right] \leq 2 \exp \left(\frac{-c\beta\epsilon^2 \text{sr}(X)}{K^4} \right).$$

Proof. Let $\mathbf{y} = \mathbf{x} \odot \zeta$. Let $Q = M^\top M$. We apply Theorem C.5 for random vector \mathbf{y} and matrix Q . We have, $\mathbf{y}^\top Q \mathbf{y} = \|M\mathbf{y}\|^2$ and $\mathbb{E}[\mathbf{y}^\top Q \mathbf{y}] = \beta \|M\|_F^2$. Thus, we obtain for any $t \geq 0$,

$$\Pr \left[\left| \|M\mathbf{y}\|^2 - \beta \|M\|_F^2 \right| > t \right] \leq 2 \exp \left(-c \min \left\{ \frac{t^2}{\beta K^4 \|M^\top M\|_F^2}, \frac{t}{K^2 \|M\|^2} \right\} \right).$$

We now can use the fact that, $\|M^\top M\|_F \leq \|M^\top\| \|M\|_F = \|M\| \|M\|_F$. Setting $t = \rho\beta \|M\|_F^2$. It follows that

$$\Pr \left[\left| \|M\mathbf{y}\|^2 - \beta \|M\|_F^2 \right| > \rho\beta \|M\|_F^2 \right] \leq 2 \exp \left(-c\beta \min\{\rho^2, \rho\} \frac{\|M\|_F^2}{K^4 \|M\|^2} \right).$$

Consider the event $\| \|M\mathbf{y}\|^2 - \beta \|M\|_F^2 \leq \rho\beta \|M\|_F^2$. Dividing both sides by $\beta \|M\|_F^2$ gives,

$$\left| \frac{\|M\mathbf{y}\|^2}{\beta \|M\|_F^2} - 1 \right| \leq \rho. \quad (13)$$

Now let $\rho = \max\{\epsilon, \epsilon^2\}$. Note that for any $r \geq 0$, $\max\{|r-1|, |r-1|^2\} \leq |r^2-1|$. Let $r^2 = \frac{\|M\mathbf{y}\|^2}{\beta \|M\|_F^2}$. If $\rho = \epsilon$, then (13) implies

$$\left| \frac{\|M\mathbf{y}\|}{\sqrt{\beta} \|M\|_F} - 1 \right| \leq \left| \frac{\|M\mathbf{y}\|^2}{\beta \|M\|_F^2} - 1 \right| \leq \epsilon.$$

If $\rho = \epsilon^2$, then (13) implies

$$\left| \frac{\|M\mathbf{y}\|}{\sqrt{\beta} \|M\|_F} - 1 \right|^2 \leq \left| \frac{\|M\mathbf{y}\|^2}{\beta \|M\|_F^2} - 1 \right| \leq \epsilon^2 \implies \left| \frac{\|M\mathbf{y}\|}{\sqrt{\beta} \|M\|_F} - 1 \right| \leq \epsilon.$$

Putting together the above two inequalities shows that under the event $\| \|M\mathbf{y}\|^2 - \beta \|M\|_F^2 \leq \rho\beta \|M\|_F^2$, implies $\| \|M\mathbf{y}\| - \sqrt{\beta} \|M\|_F \leq \epsilon \sqrt{\beta} \|M\|_F$. Using this along with the observation that $\min\{\rho^2, \rho\} = \epsilon^2$, we get the claimed result,

$$\Pr \left[\left| \|M\mathbf{y}\| - \sqrt{\beta} \|M\|_F \right| > \epsilon \sqrt{\beta} \|M\|_F \right] \leq 2 \exp \left(-c\beta\epsilon^2 \frac{\|M\|_F^2}{K^4 \|M\|^2} \right).$$

□

¹⁷The expectation $\mathbb{E}[\cdot]$ is now over both the randomness in \mathbf{x} and ζ .

Lemma C.7. Let X be a fixed $n \times p$ matrix, and let $R = \Omega \odot \Gamma$ be a sparse random matrix as in Definition 12. Let $\epsilon > 0$. Then for any fixed $\mathbf{u} \in S^{d-1}$,

$$\Pr \left[\left| \|XR\mathbf{u}\| - \sqrt{\beta}\|X\|_F \right| > \epsilon\sqrt{\beta}\|X\|_F \right] \leq 2 \exp \left(\frac{-c\beta\epsilon^2 \mathbf{sr}(X)}{K^4} \right).$$

Proof. Let us first fix $\mathbf{u} \in S^{d-1}$. By concatenating the rows of R , we can view R as a long vector in \mathbb{R}^{pd} . Consider the linear operator $T : L_2^{pd} \rightarrow L_2^n$ defined as $T(R) = XR\mathbf{u}$. Note that here \mathbf{u} is a fixed unit vector. Let us apply Lemma C.6 to the linear operator $T(R)$. The Frobenius norm of T equals $\|X\|$, and the operator norm of T can be bound as $\|XR\mathbf{u}\| \leq \|X\| \|R\|_F \|\mathbf{u}\| \leq \|X\|$ (as $\|R\|_F$ is the Euclidean norm of R as a vector in L_2^{pd}). From Lemma C.6, for any $\epsilon > 0$, we get the claimed bound,

$$\Pr \left[\left| \|XR\mathbf{u}\| - \sqrt{\beta}\|X\|_F \right| > \epsilon\sqrt{\beta}\|X\|_F \right] \leq 2 \exp \left(\frac{-c\beta\epsilon^2 \mathbf{sr}(X)}{K^4} \right).$$

□

Lemma C.8. Let X be a fixed $n \times p$ matrix, and let $R = \Omega \odot \Gamma$ be a sparse random matrix as in Definition 12. Let $\epsilon > 0$. Let $k \in \mathbb{N}$ be a number satisfying $\mathbf{sr}(X) \geq C \frac{K^4}{\epsilon^2} \frac{k}{\beta} \log \left(\frac{d}{k} \right)$. Then

$$\Pr \left[\exists I \in [d], |I| = k, \|XR_I\| > C_1(1 + \epsilon)\sqrt{\beta}\|X\|_F \right] \leq \exp \left(\frac{-c\epsilon^2\beta \mathbf{sr}(X)}{K^4} \right).$$

Proof. Let Q be a $p \times k$ matrix drawn from the same distribution as R . From Lemma C.7 for a fixed $\mathbf{v} \in S^{k-1}$

$$\Pr \left[\left| \|XQ\mathbf{v}\| - \sqrt{\beta}\|X\|_F \right| > (1 + \epsilon)\sqrt{\beta}\|X\|_F \right] \leq 2 \exp \left(\frac{-c_1\beta\epsilon^2 \mathbf{sr}(X)}{K^4} \right).$$

Let \mathcal{M} be a $1/2$ -net on S^{k-1} in the Euclidean metric. Standard volumetric argument gives that $|\mathcal{M}| \leq 5^k$ [Vershynin, 2010]. By a union bound, with probability at least

$$5^k \cdot 2 \exp \left(\frac{-c_1\epsilon^2\beta \mathbf{sr}(X)}{K^4} \right),$$

every $\mathbf{v} \in \mathcal{M}$ satisfies $\|XQ\mathbf{v}\| \leq (1 + \epsilon)\sqrt{\beta}\|X\|_F$. By our choice of $\mathbf{sr}(X)$, $5^k \cdot 2 \exp(-c_1\epsilon^2\beta \mathbf{sr}(X)/K^4) \leq \exp(-c_2\epsilon^2\beta \mathbf{sr}(X)/K^4)$. Let us condition on this probability event happening. Since every $\mathbf{v} \in S^{k-1}$ can be written as $\mathbf{v} = \mathbf{a} + \mathbf{b}$, where $\mathbf{a} \in \mathcal{M}$ and $\|\mathbf{b}\| \leq 1/2$, we get

$$\|XQ\| = \max_{\mathbf{v} \in S^{k-1}} \|XQ\mathbf{v}\| \leq \max_{\mathbf{a} \in \mathcal{M}} \|XQ\mathbf{a}\| + \max_{\mathbf{b}: \|\mathbf{b}\| \leq 1/2} \|XQ\mathbf{b}\| \leq (1 + \epsilon)\sqrt{\beta}\|X\|_F + \frac{1}{2}\|XQ\| \implies \|XQ\| \leq 2(1 + \epsilon)\sqrt{\beta}\|X\|_F.$$

Using this along with a union bound over all $p \times k$ submatrices of R , we obtain

$$\Pr \left[\exists I \subset [d], |I| = k, \|XR_I\| > C_1(1 + \epsilon)\sqrt{\beta}\|X\|_F \right] \leq \binom{d}{k} \cdot \exp \left(\frac{-c_2\beta\epsilon^2 \mathbf{sr}(X)}{K^4} \right) \leq \exp \left(\frac{-c_3\beta\epsilon^2 \mathbf{sr}(X)}{K^4} \right),$$

by our assumption on $\mathbf{sr}(X)$. □

Theorem C.9. Let X be an $n \times p$ matrix. Let $\hat{R} = \Omega \odot \Gamma$ be a $p \times d$ matrix, where $\Omega \in \mathbb{R}^{p \times d}$ is an i.i.d. Bernoulli matrix where each entry is 1 independently with probability β , and $\Gamma = (\Gamma_{ij})$ is a random matrix with independent entries that satisfy: $\mathbb{E}[\Gamma_{ij}] = 0$, $\mathbb{E}[\Gamma_{ij}^2] = 1$, and $\|\Gamma_{ij}\|_{\psi_2} \leq K$. Let $R = \hat{R}/(c_0\sqrt{\beta})$. Let $\epsilon \in (0, 1)$, and let $k \in \mathbb{N}$ be a number satisfying $\mathbf{sr}(X) \geq CK^4 \frac{k}{\beta\epsilon^2} \log \left(\frac{d}{k} \right)$. Then with probability at least $1 - \exp(-c\beta\epsilon^2 \mathbf{sr}(X)/K^4)$, the matrix XR satisfies (k, ϵ) -RIP, i.e.,

$$\forall \mathbf{u} \in \Sigma_k, (1 - \epsilon)\|X\|_F \leq \|XR\mathbf{u}\| \leq (1 + \epsilon)\|X\|_F.$$

Proof. From Lemma C.7,

$$\Pr \left[\left| \|X \hat{R} \mathbf{u}\| - \sqrt{\beta} \|X\|_F \right| > \epsilon \sqrt{\beta} \|X\|_F \right] \leq 2 \exp \left(\frac{-c_1 \beta \epsilon^2 \text{sr}(X)}{K^4} \right).$$

We now use a standard net argument. Let \mathcal{N} be an $(1/2C_2)$ -net \mathcal{N} over Σ_k as in Theorem 3.1. Taking the union bound over this net, we obtain

$$\begin{aligned} \Pr \left[\forall \mathbf{u} \in \mathcal{N}, \left| \|X \hat{R} \mathbf{u}\| - \sqrt{\beta} \|X\|_F \right| \leq \epsilon \sqrt{\beta} \|X\|_F \right] &\geq 1 - \exp \left(\frac{-c_1 \beta \epsilon^2 \text{sr}(X)}{K^4} + k \log \frac{C_0 d}{k} \right) \\ &\geq 1 - \exp \left(\frac{-c_2 \beta \epsilon^2 \text{sr}(X)}{K^4} \right), \end{aligned}$$

where we used the assumption on $\text{sr}(X)$. Now using the spectral norm bound from Lemma C.8 in conjunction with the approximation idea used in Theorem 3.1 yields that with probability at least $1 - \exp(-c\beta\epsilon^2 \text{sr}(X)/K^4)$, $c_0(1-\epsilon)\sqrt{\beta} \|X\|_F \leq \|X \hat{R} \mathbf{u}\| \leq c_0(1+\epsilon)\sqrt{\beta} \|X\|_F$ for all $\mathbf{u} \in \Sigma_k$. We finish by noting that $R = \hat{R}/(c_0\sqrt{\beta})$. \square

C.3 Restricted Isometry of XR under Convex Concentration Property on R

We investigate the case where $R \in \mathbb{R}^{p \times d}$ is composed of independent columns satisfying the convex concentration property. In this case, we utilize the recent result of [Adamczak et al., 2015], who proved the Hanson-Wright inequality for isotropic random vectors having convex concentrations property.

Theorem C.10 (Theorem 2.3 [Adamczak et al., 2015]). *Let $\mathbf{x} = (x_1, \dots, x_d) \in \mathbb{R}^d$ be a random vector with independent components x_i which satisfy $\mathbb{E}[x_i] = 0$. Let \mathbf{x} satisfy the convex concentration property with constant K . Let M be a $d \times d$ matrix. Then for every $t \geq 0$,*

$$\Pr[|\mathbf{x}^\top M \mathbf{x} - \mathbb{E}[\mathbf{x}^\top M \mathbf{x}]| > t] \leq 2 \exp \left(-c \min \left\{ \frac{t^2}{2K^4 \|M\|_F^2}, \frac{t}{K^2 \|M\|} \right\} \right).$$

Our first result here is to show that a linear combination of independent vectors having a convex concentration property would have this property as well (Lemma C.11). To this end, we start with a fixed $\mathbf{u} = (u_1, \dots, u_d) \in S^{d-1}$ and construct a martingale with variables $\mathbb{E}[\phi(\sum_{i=1}^d \mathbf{r}_i u_i) \mid \mathbf{r}_1, \dots, \mathbf{r}_j]$ where $\phi : \mathbb{R}^p \rightarrow \mathbb{R}$ is a 1-Lipschitz convex function and \mathbf{r}_i is the i th column in R , and then apply Azuma's inequality to it. Once we have this established, verifying RIP consists of applying the Hanson-Wright inequality of [Adamczak et al., 2015] in a proof framework similar to Theorem 3.1.

Lemma C.11. *Let R be a $p \times d$ random matrix with independent columns satisfying the convex concentration property with constant K . Let $\mathbf{u} \in S^{d-1}$. Then $R\mathbf{u}$ is a random vector satisfying the convex concentration property with constant $O(K)$.*

Proof. Let $\mathbf{u} = (u_1, \dots, u_d)$. Let $\phi : \mathbb{R}^p \rightarrow \mathbb{R}$ be a 1-Lipschitz convex function. Let $\mathbf{r}_1, \dots, \mathbf{r}_d \in \mathbb{R}^p$ be the columns of R . For any $j \in [d]$, the function $V_j : \mathbb{R}^p \rightarrow \mathbb{R}$ defined by

$$V_j(\mathbf{a}) = \phi \left(\sum_{i \neq j} \mathbf{r}_i u_i + \mathbf{a} u_j \right) - \mathbb{E} \left[\phi \left(\sum_{i \neq j} \mathbf{r}_i u_i + \mathbf{a} u_j \right) \right],$$

is also convex and Lipschitz with the Lipschitz constant $|u_j|$. The convex concentration property implies that the random variable

$$\Delta_j = \mathbb{E} \left[\phi \left(\sum_{i=1}^d \mathbf{r}_i u_i \right) \mid \mathbf{r}_1, \dots, \mathbf{r}_j \right] - \mathbb{E} \left[\phi \left(\sum_{i=1}^d \mathbf{r}_i u_i \right) \mid \mathbf{r}_1, \dots, \mathbf{r}_{j-1} \right]$$

is subgaussian with ψ_2 -norm $\|\Delta\|_{\psi_2} \leq K|u_j|$. To verify it, we can condition on $\mathbf{r}_i, i \neq j$, and then integrate over $\mathbf{r}_1, \dots, \mathbf{r}_{j-1}$ using Jensen's inequality for the ψ_2 -norm. The random variables

$$m_j = \mathbb{E} \left[\phi \left(\sum_{i=1}^d \mathbf{r}_i u_i \right) \mid \mathbf{r}_1, \dots, \mathbf{r}_j \right], \quad j \in \{0, \dots, d\},$$

form a martingale with the martingale differences Δ_j . Hence, Azuma's inequality implies that

$$m_d - m_0 = \phi \left(\sum_{i=1}^d \mathbf{r}_i u_i \right) - \mathbb{E} \left[\phi \left(\sum_{i=1}^d \mathbf{r}_i u_i \right) \right]$$

is subgaussian random variable with $\|m_d - m_0\|_{\psi_2} \leq CK\|\mathbf{u}\| = O(K)$. This completes the proof. \square

Theorem C.12. *Let X be an $n \times p$ matrix. Let R be a $p \times d$ matrix with mean zero independent columns satisfying the convex concentration property (Definition 3) with constant K . Let $\epsilon \in (0, 1)$, and let $k \in \mathbb{N}$ be a number satisfying $\mathbf{sr}(X) \geq CK^4 \frac{k}{\epsilon^2} \log\left(\frac{d}{k}\right)$. Then with probability at least $1 - \exp(-c\epsilon^2 \mathbf{sr}(X)/K^4)$, the matrix XR satisfies (k, ϵ) -RIP, i.e.,*

$$\forall \mathbf{u} \in \Sigma_k, (1 - \epsilon) \|X\|_F \leq \|XR\mathbf{u}\| \leq (1 + \epsilon) \|X\|_F.$$

Proof. Fix any $\mathbf{u} \in S^{d-1}$. From Lemma C.11, we get that $R\mathbf{u}$ satisfies the c.c.p with constant $O(K)$. Using $R\mathbf{u}$ as the random vector, a simple consequence of the Hanson-Wright inequality from Theorem C.10 is for all $t \geq 0$,

$$\Pr[|\|XR\mathbf{u}\| - \|X\|_F| \geq t] \leq 2 \exp\left(\frac{-c_1 t^2}{K^4 \|X\|^2}\right).$$

Setting $t = \epsilon \|X\|_F$, we get

$$\Pr[|\|XR\mathbf{u}\| - \|X\|_F| \geq \epsilon \|X\|_F] \leq 2 \exp\left(\frac{-c_1 \epsilon^2 \mathbf{sr}(X)}{K^4}\right). \quad (14)$$

The remainder of the proof follows the framework of Theorem 3.1. From the Hanson-Wright inequality from Theorem C.10, we obtain that

$$\Pr[\exists I \in [d], |I| = k, \|XR_I\| > C_1 \epsilon \|X\|_F] \leq 2 \exp(-c_0 \epsilon^2 \mathbf{sr}(X)/K^4).$$

Verifying RIP consists now of applying (14) to a net in the set of sparse vectors on the unit sphere as in Theorem 3.1 and the details are omitted here. \square

D Missing Details from Section 4

Proof of Lemma 4.2. We start with the decoupling argument as in [Rudelson and Vershynin, 2013]. Let $\delta_1, \dots, \delta_d$ be i.i.d. Bernoulli(1/2) random variables and denote by I the random set $I = \{i \in [d] : \delta_i = 1\}$. Let $a > 0$, and assume that for any realization of I ,

$$\mathbb{E} \left[\exp \left(\frac{1}{2a} \sum_{i \in I, j \notin I} r_i r_j u_{ij} \right) \right] \leq 2.$$

Then by Jensen's inequality,

$$\begin{aligned}
\mathbb{E} \left[\exp \left(\frac{1}{2a} \sum_{(i,j) \in \binom{[d]}{2}} r_i r_j u_{ij} \right) \right] &\leq \mathbb{E} \left[\exp \left(\frac{1}{2a} \sum_{(i,j) \in \binom{[d]}{2}} 2 \mathbb{E}_\delta [\delta_i (1 - \delta_j) r_i r_j u_{ij}] \right) \right] \\
&\leq \mathbb{E} \left[\mathbb{E}_\delta \left[\exp \left(\frac{1}{a} \sum_{(i,j) \in \binom{[d]}{2}} \delta_i (1 - \delta_j) r_i r_j u_{ij} \right) \right] \right] \\
&\leq \max_{I \subset [d]} \mathbb{E} \left[\exp \left(\frac{1}{2a} \sum_{i \in I, j \notin I} r_i r_j u_{ij} \right) \right] \leq 2.
\end{aligned}$$

This calculation means that

$$\left\| \sum_{(i,j) \in \binom{[d]}{2}} r_i r_j u_{ij} \right\|_{\psi_2} \leq \max_{I \subset [d]} \left\| \sum_{i \in I, j \notin I} r_i r_j u_{ij} \right\|_{\psi_2}. \quad (15)$$

From now on, we fix $I \subset [d]$ and focus on bounding $\left\| \sum_{i \in I, j \notin I} r_i r_j u_{ij} \right\|_{\psi_2}$.

We will decompose the random variable $\sum_{i \in I, j \notin I} r_i r_j u_{ij}$ into the sum of at most $4\sqrt{k}$ subgaussian random variables corresponding to disjoint parts of $\text{supp}(\mathbf{u})$ and bound the ψ_2 -norms of these variables separately. To this end, we introduce the following induction procedure. Let P_1, P_2 be coordinate projections of $I \times I^c$ on the first and the second coordinates respectively. Set $F_1 = \text{supp}(\mathbf{u})$ and let $I_1 = P_1(F_1)$ and $J_0 = P_2(F_1)$. Choose a set $U_1 \subset F_1$ such that $P_1(U_1) = I_1$ and for any $i \in I_1$, $|P_1^{-1}(i) \cap U_1| = 1$. This means that we lift the projection I_1 in the set F_1 choosing a single point in F_1 for any point in the projection. The set U_1 can and will be chosen to contain the set $F_1 \cap (I \times \{j\})$ for some $j \in J_0$. If $F_1 = U_1$, we set $W_1 = \emptyset$ and stop. Otherwise, set $E_1 = F_1 \setminus U_1$, and $J_1 = P_2(E_1)$. By construction, $J_1 \subset J_0$, and $|J_1| \leq |J_0| - 1$. Similarly to what we have done before, choose a set $W_1 \subset E_1$ such that $P_2(W_1) = J_1$ and for any $j \in J_1$, $|P_2^{-1}(j) \cap W_1| = 1$. We choose the set W_1 so that it contains the set $E_1 \cap (\{i\} \times I^c)$ for some $i \in I_1$. If $E_1 = W_1$, we stop. Otherwise, set $F_2 = E_1 \setminus W_1$, and $I_2 = P_1(F_2)$. Again by construction, $F_2 \subset F_1$, $I_2 \subset I_1$, and $|I_2| \leq |I_1| - 1$. This finishes the first step of the induction. After that, we repeat this process using the set F_2 in place of F_1 . At each step of the process, we create disjoint sets U_1, U_2, \dots and W_1, W_2, \dots contained in $\text{supp}(\mathbf{u})$. Since these sets are non-empty, this process eventually stops. After it stops, we obtain a decomposition

$$\text{supp}(\mathbf{u}) \cap (I \times I^c) = \bigcup_{t=1}^s U_t \cup \bigcup_{t=1}^s W_t,$$

where all terms are disjoint. Moreover, the construction yields $I_1 = P_1(U_1) \supset I_2 = P_1(U_2) \supset \dots \supset P_1(U_s)$ with $|U_1| > |U_2| > \dots > |U_s|$. Similarly, $J_1 = P_2(W_1) \supset J_2 = P_2(W_2) \supset \dots \supset P_2(W_s)$ with $|W_1| > |W_2| > \dots > |W_s|$. In view of the decomposition above, one of the quantities $\sum_{j=1}^s |U_j|$ or $\sum_{j=1}^s |W_j|$ is at most $|\text{supp}(\mathbf{u}) \cap (I \times I^c)|/2$. Assume that this is the former one. Then

$$\frac{k}{2} \geq \frac{|\text{supp}(\mathbf{u}) \cap (I \times I^c)|}{2} \geq \sum_{t=1}^s |U_t| \geq \sum_{t=1}^s (|U_s| + s - t) \geq \frac{s(s-1)}{2},$$

which implies $s \leq \sqrt{k} + 1$. The estimate in the other case is exactly the same.

Let us use the decomposition we constructed to estimate $\left\| \sum_{i \in I, j \notin I} r_i r_j u_{ij} \right\|_{\psi_2}$. For $t \in [s]$, denote

$$v_t = \sum_{(i,j) \in U_t} r_i r_j u_{ij} \quad \text{and} \quad w_t = \sum_{(i,j) \in W_t} r_i r_j u_{ij}.$$

The random variables r_i and r_j are independent for all $i \in I$ and $j \in I^c$. Hence, by Hoeffding's inequality,

$$\|v_t\|_{\psi_2} \leq \sup_{|a_1|, \dots, |a_s| \leq K} \left\| \sum_{(i,j) \in U_t} r_i a_j u_{ij} \right\|_{\psi_2} \leq 4\tau \sup_{|a_1|, \dots, |a_s| \leq K} \left(\sum_{(i,j) \in U_t} a_j^2 u_{ij}^2 \right)^{1/2} \leq 4\tau^2 \left(\sum_{(i,j) \in U_t} u_{ij}^2 \right)^{1/2},$$

and $\|w_t\|_{\psi_2}$ satisfies the same estimate. Finally,

$$\begin{aligned} \left\| \sum_{i \in I, j \notin I} r_i r_j u_{ij} \right\|_{\psi_2} &\leq \sum_{t=1}^s \|v_t\|_{\psi_2} + \sum_{t=1}^s \|w_t\|_{\psi_2} \leq 4\tau^2 \sum_{t=1}^s \left(\sum_{(i,j) \in U_t} u_{ij}^2 \right)^{1/2} + 4\tau^2 \sum_{t=1}^s \left(\sum_{(i,j) \in W_t} u_{ij}^2 \right)^{1/2} \\ &\leq 8\tau^2 \sqrt{s} \left(\sum_{i \in I, j \notin I} u_{ij}^2 \right)^{1/2} \leq 8\tau^2 \sqrt{s} \|\mathbf{u}\|. \end{aligned}$$

Combining this with (15) and taking into account that $s \leq \sqrt{k} + 1$, we complete the proof of the lemma. \square

Theorem D.1 (Theorem 4.3 Restated). *Let X be an $n \times p$ matrix, and let R be a $p \times d$ random matrix with independent entries R_{ij} such that $\mathbb{E}[R_{ij}] = 0$, $\mathbb{E}[R_{ij}^2] = 1$, and $|R_{ij}| \leq \tau$ almost surely. Let $\ell \geq 3$ be a constant. Let $\epsilon \in (0, 1)$, and let $k \in \mathcal{N}$ be a number satisfying $\mathbf{sr}(X) \geq \frac{C\tau^{4\ell} k^3}{\epsilon^2} \log\left(\frac{d^\ell}{k}\right)$. Then with probability at least $1 - \exp(-c\epsilon^2 \mathbf{sr}(X)/(k^2 \tau^{4\ell}))$, the matrix $XR^{(\ell)}$ satisfies the (k, ϵ) -RIP property, i.e., for any $\mathbf{u} \in S^{(d)}_{k-1}$ with $|\text{supp}(\mathbf{u})| \leq k$,*

$$(1 - \epsilon) \|X\|_{\text{F}} \leq \left\| XR^{(\ell)} \mathbf{u} \right\| \leq (1 + \epsilon) \|X\|_{\text{F}}.$$

Proof. Let $\mathbf{u} \in \mathbb{R}^{(d)}$ be a vector with $|\text{supp}(\mathbf{u})| \leq k$. Let $u_{i_1 \dots i_\ell}$ be the (i_1, \dots, i_ℓ) th element in \mathbf{u} (with $1 \leq i_1 < i_2 < \dots < i_\ell \leq d$). Let $l \in [p]$, and define

$$y_l = \sum_{(i_1, \dots, i_\ell) \in \binom{[d]}{\ell}} R_{li_1} R_{li_2} \dots R_{li_\ell} u_{i_1 \dots i_\ell}.$$

The random variables y_l , $l \in [p]$ are independent. From triangle inequality,

$$\begin{aligned} \|y_l\|_{\psi_2} &= \left\| \sum_{(i_1, \dots, i_\ell) \in \binom{[d]}{\ell}} R_{li_1} R_{li_2} \dots R_{li_\ell} u_{i_1 \dots i_\ell} \right\|_{\psi_2} \leq \sum_{(i_1, \dots, i_\ell) \in \text{supp}(\mathbf{u})} \|R_{li_1} R_{li_2} \dots R_{li_\ell} u_{i_1 \dots i_\ell}\|_{\psi_2} = O(\tau^\ell \|\mathbf{u}\|_1) \\ &= O(\tau^\ell \sqrt{k} \|\mathbf{u}\|). \end{aligned} \tag{16}$$

The rest of the proof proceeds as in Theorem 4.1, by using the $\|y_l\|_{\psi_2}$ bound from (16). In this case, we use a net \mathcal{N} on the set of k -sparse vectors in $S^{(d)}_{k-1}$. \square

D.1 Compressed Learning Background

We discuss some background on compressed learning. Consider the standard binary classification task with labeled data from $\mathcal{X} = \{(\mathbf{x}, v) : \mathbf{x} \in \mathbb{R}^d, \mathbf{x} \text{ is } k\text{-sparse}, \|\mathbf{x}\| \leq \alpha, v \in \{-1, 1\}\}$, with v indicating the label on \mathbf{x} . We focus on linear loss functions having the form, $f(\langle \mathbf{x}, \theta \rangle; v)$ for $\theta \in \mathbb{R}^d$, where $f : \mathbb{R} \times \{-1, 1\} \rightarrow \mathbb{R}$ is assumed to be convex and Lipschitz in the first parameter. This type of program captures a variety of important learning problems, e.g., the linear regression is captured by setting $f(\langle \mathbf{x}, \theta \rangle; v) = (v - \langle \mathbf{x}, \theta \rangle)^2$, logistic regression is captured by setting $f(\langle \mathbf{x}, \theta \rangle; v) = \ln(1 + \exp(-v\langle \mathbf{x}, \theta \rangle))$, support vector machine is captured by setting $f(\langle \mathbf{x}, \theta \rangle; v) = \text{hinge}(v\langle \mathbf{x}, \theta \rangle)$, where $\text{hinge}(a) = 1 - a$ if $a \leq 1$ and 0 otherwise.

For a distribution \mathcal{D} over \mathcal{X} , define, $f_{\mathcal{D}}(\theta) = \mathbb{E}_{(\mathbf{x}, v) \sim \mathcal{D}}[f(\langle \mathbf{x}, \theta \rangle; v)]$ as the generalization loss on the original domain for parameter θ . Let $\theta^* \in \text{argmin}_{\theta \in \mathbb{R}^d} f_{\mathcal{D}}(\theta)$. The following theorem from [Arora et al. \[2018a\]](#), building upon a result by [Calderbank et al., 2009](#), relates the generalization loss obtained by solving the classification problem on the compressed domain to the loss on the original domain.

Theorem D.2 (Restated from Theorem 4.2 [\[Arora et al., 2018a\]](#)). *Let $M \in \mathbb{R}^{n \times d}$ be a (k, ϵ) -RIP matrix. Let $(\mathbf{x}_1, v_1), \dots, (\mathbf{x}_b, v_b)$ be a set of labeled data drawn i.i.d. from a distribution \mathcal{D} over \mathcal{X} . Let f be a λ -Lipschitz convex linear loss function and $\theta^* \in \text{argmin}_{\theta \in \mathbb{R}^d} f_{\mathcal{D}}(\theta)$ be its minimizer over \mathcal{D} . Let $\hat{\vartheta} \in \mathbb{R}^n$ be the minimizer of a classifier trained on $(M\mathbf{x}_1, v_1), \dots, (M\mathbf{x}_b, v_b)$ minimizing the L_2 -regularized empirical loss function $\frac{1}{b} \sum_i f(\langle M\mathbf{x}_i, \vartheta \rangle; v_i) + \frac{1}{2C} \|\vartheta\|^2$, i.e., $\hat{\vartheta} \in \text{argmin}_{\vartheta \in \mathbb{R}^n} \frac{1}{b} \sum_i f(\langle M\mathbf{x}_i, \vartheta \rangle; v_i) + \frac{1}{2C} \|\vartheta\|^2$ (for appropriate constant C), then with probability at least $1 - \delta$, we have*

$$f_{\mathcal{D}}(\hat{\vartheta}) \leq f_{\mathcal{D}}(\theta^*) + O\left(\lambda\alpha\|\theta^*\|\sqrt{\epsilon + \frac{1}{b} \log \frac{1}{\delta}}\right),$$

where $f_{\mathcal{D}}(\vartheta) = \mathbb{E}_{(\mathbf{x}, v) \sim \mathcal{D}}[f(\langle M\mathbf{x}, \vartheta \rangle; v)]$ is the generalization loss on the compressed domain for parameter ϑ .