

Subspace Learning from Extremely Compressed Measurements

Akshay Krishnamurthy ^{*1}, Martin Azizyan ^{†2}, and Aarti Singh ^{‡2}

¹Computer Science Department, Carnegie Mellon University

²Machine Learning Department, Carnegie Mellon University

April 4, 2014

Abstract

We consider learning the principal subspace of a large set of vectors from an extremely small number of compressive measurements of each vector. Our theoretical results show that even a constant number of measurements per column suffices to approximate the principal subspace to arbitrary precision, provided that the number of vectors is large. This result is achieved by a simple algorithm that computes the eigenvectors of an estimate of the covariance matrix. The main insight is to exploit an averaging effect that arises from applying a different random projection to each vector. We provide a number of simulations confirming our theoretical results.

1 Introduction

Low rank approximation is a fundamental preprocessing task in a wide variety of machine learning and signal processing applications. Given a set of vectors $x_1, \dots, x_n \in \mathbb{R}^d$, low rank approximation, popularly referred to as principal component analysis, aims to find a low dimensional subspace of \mathbb{R}^d that captures a large amount of the energy present in the vectors $\{x_t\}_{t=1}^n$. It is well known that if we concatenate the vectors as columns of a $d \times n$ matrix X , then the span of the left singular vectors corresponding to the k largest singular values of X forms the best k -dimensional subspace in many senses. For example, it minimizes the squared Euclidean reconstruction error, $\|X - \Pi X\|_F^2$, over all k -dimensional projections Π .

There are many situations where obtaining this best subspace is computationally challenging, motivating several different theoretical and empirical studies. A number of clever approaches, ranging from sparsification [1] and column sampling [2], to sketching and streaming techniques [9, 4] have been used to reduce the computational burden associated with computing the principal subspace when n and d are large.

An equally pressing concern is the cost of data acquisition, motivating lines of work on low rank approximation from missing data and compressive measurements. The matrix completion literature focuses on finding approximations from randomly or adaptively sampled entries of the data matrix [12, 11, 8, 6].

*akshaykr@cs.cmu.edu

†mazizyan@cs.cmu.edu

‡aarti@cs.cmu.edu

With compressive measurements, only a few linear combinations of each column of X are acquired, which mathematically translates to observing $Y = RX$ where R is a $m \times d$ random matrix with $m \ll d$. Usually, the columns of Y are projected back into d -dimensions and the principal subspace of these vectors is used as an estimator [13, 7, 10]. In this paper, we deviate from this strategy, and propose an algorithm that uses a different sensing operation per column to approximate the principal subspace of the vectors $\{x_t\}_{t=1}^n$.

This difference, on one hand, adds some complexity to the algorithm, as a new random projection must be applied to each column. On the other hand, resampling the projection leads to an averaging phenomenon that allows us to undersample each column significantly more than existing approaches. In particular, we show that our algorithm only needs two compressive measurements per column provided there are enough columns, while existing approaches require $O(k/\epsilon)$ measurements to achieve an ϵ -approximation with target rank k [7].

Before proceeding, let us briefly mention two motivating applications. In time-series analysis, typically we have extremely long data sequences and would like to extract some signal from the sequence to assist in some inferential task. In many of these problems, observing any time-point of the data sequence is expensive; it may require performing some scientific experiment. Our results show that one can significantly undersample each time point of the series, which can tremendously reduce measurement overhead, while still extracting the principal subspace from the data sequence.

A related application is in signal processing over a distributed sensor network with high measurement and communication cost. Suppose each of the n sensors records a d -dimensional vector of measurements, and our goal is to make inferences about the $d \times n$ matrix of measurements. Since measurement costs are high, one might be interested in compressive sampling at each sensor, but if the compression operator is shared across the sensors, an expensive synchronization step must be performed before data acquisition. Our method avoids the need for synchronization, since each sensor uses its own compression operator, while guaranteeing a good approximation to the matrix of measurements.

2 Results

Before proceeding to our main results, let us set up some notation used throughout the manuscript. We are interested in recovering the principal k -dimensional subspace Π of the set of vectors $x_1, \dots, x_n \in \mathbb{R}^d$ which we will concatenate into a matrix $X \in \mathbb{R}^{d \times n}$. We assume that each vector x_t has bounded norm, that is $\max_{t \in [n]} \|x_t\|_2^2 \leq \mu^1$. Let $\Sigma = \frac{1}{n} \sum_{t=1}^n x_t x_t^T$ be the covariance matrix and notice that the optimal k -dimensional subspace Π , in terms of minimizing squared Euclidean reconstruction error, is spanned by the top- k eigenvectors of Σ . A key parameter that governs the performance of our algorithm is the **eigengap** γ_k , defined as the difference between the k th largest eigenvalue and the $k + 1$ st largest eigenvalue of the covariance matrix Σ .

We measure the error of an estimate $\hat{\Pi}$ for the principal subspace in terms of the spectral norm $\|\hat{\Pi} - \Pi\|_2$, where $\|M\|_2$ is the largest singular value of M . It is not too hard to show that when $\hat{\Pi}$ and Π are projection matrices, $\|\hat{\Pi} - \Pi\|_2$ corresponds to the sine of the largest principal angle between the two subspaces, and therefore it is an appropriate measure of error in our setting.

We are interested in the compressive sampling framework, in which for some $m \geq 1$, we observe each vector x_t through $2m$ compressive measurements $y_{it} = a_{it}^T x_t, i \in [2m]$ for vectors a_i drawn independently from the d -dimensional unit sphere. An equivalent characterization of this framework is to pick a $2m$ -dimensional random projection $\Phi_t \in \mathbb{R}^{2m \times d}$ uniformly at random and observe $y_t = \Phi_t x_t$. In our algorithm, we will write this as two m -dimensional random projections Φ_t and Ψ_t per column with observations $y_t =$

¹We use $[n]$ to denote $\{1, \dots, n\}$.

Algorithm 1 Compressive Subspace Learning

Input: Compression parameter m , target rank k .

$$\hat{\Sigma} = 0 \in \mathbb{R}^{d \times d}$$

for $x_t \in \mathbb{R}^d$ in the data stream **do**

 Let Φ_t, Ψ_t be m -dimensional random projections.

 Acquire $y_t = \Phi_t x_t$ and $z_t = \Psi_t x_t$.

 Update $\hat{\Sigma} = \hat{\Sigma} + \frac{1}{2}(y_t z_t^T + z_t y_t^T)$

end for

Output: $\hat{\Pi} = \text{span}(u_1, \dots, u_k)$ the top k eigenvectors of $\hat{\Sigma}$.

$\Phi_t x_t$ and $z_t = \Psi_t x_t$. To connect the two representations, let A_t denote the $d \times m$ matrix whose columns are the vectors $a_{it}, i \in [m]$ and define $\Phi_t = P_{A_t}$, the the projection onto the span of the columns A_t . Define B_t similarly to A_t , but with the second m vectors and let $\Psi_t = P_{B_t}$.

Our algorithm, which we call Compressive Subspace Learning (CSL), is conceptually very simple (See Algorithm 1). For each vector x_t , we observe y_t, z_t and form an estimate for $C_t = x_t x_t^T$ with:

$$\hat{C}_t = \frac{1}{2}(y_t z_t^T + z_t y_t^T) \quad (1)$$

We estimate the covariance Σ with $\hat{\Sigma} = \sum_{t=1}^n \hat{C}_t$ and use the span of top k eigenvectors of $\hat{\Sigma}$ as an estimate for the principal subspace Π . Let $\hat{\Pi}$ denote the output of our algorithm, namely the span of the top k eigenvectors of $\hat{\Sigma}$. Note that $\hat{\Sigma}$ needs to be appropriately rescaled before it provides a reasonable estimate of Σ , but since we are only interested in the principal subspace, this normalization is only necessary for analysis.

In addition to the statistical guarantee in Theorem 1 below, we also mention some practical considerations. The algorithm can be implemented in the streaming model with each column vector x_t being streamed through memory. In this model, the algorithm requires $O(d^2)$ time to process each vector and $O(d^2)$ space in total, to store $\hat{\Sigma}$.

The distributed sensor network model we described earlier is quite appealing for CSL. Recall that we had a network of n sensors, each observing a d -dimensional vector x_t , with high measurement and communication cost. A naïve procedure that observes and transmits the vectors in full has $O(nd)$ communication overhead, but also $O(nd)$ measurement cost, which can be prohibitively expensive. The traditional compressive sampling approach of using a single random projection needs only $O(mn)$ measurements but synchronizing the projection before data acquisition requires $O(nmd)$ communication cost. This is non-negligible overhead as one typically requires $m \asymp k/\epsilon$ to achieve error ϵ with target rank k [7].

Using CSL, each sensor can generate its own random projection, take $2m$ measurements of its signal x_t , and send only two d -dimensional vectors (namely y_t and z_t) over the network. Thus CSL makes $O(mn)$ measurements while suffering only $O(nd)$ communication overhead, achieving the best properties of both other approaches. Moreover, provided n is large, CSL can succeed even with m constant, which is significantly smaller than other methods.

Our main result is the following statistical guarantee on the performance of the algorithm:

Theorem 1. *With probability $\geq 1 - \delta$:*

$$\|\hat{\Pi} - \Pi\|_2 \leq \frac{1}{\gamma_k} \left(\sqrt{\frac{44\mu^2}{nm} \log(d/\delta)} + \frac{2}{3} \frac{\mu d^2}{m^2 n} \log(d/\delta) \right) \quad (2)$$




Figure 1: Comparison of our algorithm with that of Halko *et. al* [7] on a low rank matrix $d = 20, k = 2$. Our algorithm achieves low error for fixed m , provided that n is large, while their algorithm does not.

So that one can achieve spectral norm error $\leq \epsilon$ provided that:

$$n \geq \max \left\{ \frac{176\mu^2 \log(d/\delta)}{m\gamma_k^2 \epsilon^2}, \frac{4}{3} \frac{\mu d^2}{\gamma_k \epsilon m^2} \log(d/\delta) \right\} \quad (3)$$

Before turning to the proof, some remarks are in order:

We are interested in the setting where m is small compared to d , in which case the second term in Equation 2 is active. Thus in interpreting the theorem, one should focus on the second terms in both bounds to see the dependence on d and m . On the other hand, the relationship between n and ϵ is dictated by the first terms.

Since the theorem holds for any $m \geq 1$, it is clear that even two compressive measurements per column suffice to approximate the principal subspace, provided that $n \gg d$, which is common in a number of signal processing and machine learning applications. This is in sharp contrast with a number of other results on matrix approximation from compressive measurements, where the same random projection is used on each column, in which case such a bound is not possible unless $m \geq k/\epsilon$ [7]. The justification for this is that using different random projections allows our algorithm to exploit averaging across the columns to capture the principal directions of the matrix. On the other hand, if the same projection is used, the measurements across columns become highly correlated and one does not experience a law-of-large-numbers phenomenon.

To make this concrete, consider the setting where X is a rank one matrix with identical columns. If the same projection is used across the columns, then there is no hope of approximating the principal direction with small m , irrespective of n . On the other hand, our bound says that by leveraging additional randomness in the data acquisition process, we can achieve an ϵ -approximation provided that n is large. In Figure 1 we demonstrate this phenomenon.

Results of this flavor are usually stated in terms of the number of measurements per column that suffice to achieve error ϵ . Stated this way, our bound says that:

$$m \asymp \frac{d}{\sqrt{n}} \sqrt{\frac{\mu}{\epsilon \gamma_k} \log(d/\delta)}$$

is sufficient. In comparison, Halko *et. al* only require $m \asymp k/\epsilon$ to achieve error ϵ with target rank k [7] while the state-of-the-art results on matrix completion require observing $k \log(d/\delta)$ coordinates per column [12, 11]. Both sets of results are better than ours when k is small compared to d/\sqrt{n} . We chose to state our




Figure 2: Left: Geometry of the distribution of y_t . Right: Distribution of y_t (green dots) given x_t (blue line).

bound as in Theorem 1, as it highlights the averaging effect mentioned above. However, some improvement in our result may be possible with a more careful analysis.

Our result is similar in spirit to a recent analysis by Gonen *et. al* who study the subspace learning problem when only a small number of entries of each column are observed [6]. They also show that one can approximate the principal subspace of a matrix X using only a few observations per column. Their analysis is somewhat simpler than ours, due to the fact that the distribution induced by the random projection operator is analytically more challenging than the subsampling operator. Interestingly, they also show that one *cannot* learn the subspace when observing only one coordinate per column, which may be related to why we require two compressive measurements, although of course the sampling paradigm is quite different.

It may seem curious that the target rank k is absent from the bound, but notice that there is an interaction between k , the bound on the column norms μ , and the eigengap. The target rank indirectly has some influence on the sample complexity n via both the eigengap γ_k and the length bound μ .

Proof. The main component of the proof is an exact characterization of the distribution of the observations y_t, z_t . Equipped with this distribution, the result is an application of two fairly well known results in the matrix perturbation literature. The first is the Matrix-Bernstein inequality, which analogously to the Bernstein inequality, formalizes the concentration of measure phenomenon. The second is the Davis-Kahan sine theorem, which characterizes the deviation of eigenvectors under a random perturbation of the matrix. We state both theorems here without proof:

Theorem 2 (Matrix Bernstein Inequality [14]). *Let X_1, \dots, X_n be a sequence of independent, random, self-adjoint matrices with dimension d . Assume that each random matrix satisfies:*

$$\mathbb{E}X_k = 0, \quad \|X_k\|_2 \leq R \text{ a.s.}$$

Then, for all $t \geq 0$,

$$\mathbb{P} \left(\left\| \sum_k X_k \right\|_2 \geq t \right) \leq d \exp \left\{ \frac{-t^2/2}{\sigma^2 + Rt/3} \right\} \quad (4)$$

where $\sigma^2 = \left\| \sum_k \mathbb{E}(X_k^2) \right\|_2$.




Figure 3: Top row: Rates of convergence alongside rescaled rates $\sqrt{n} \times \|\hat{\Pi} - \Pi\|_2$ across a number of problem settings. Left two plots show exactly low rank case while the right two plots show low rank approximation of high rank matrices. Bottom row: Influence of other problem parameters on error. From left to right: number of measurements per column (m), problem dimension (d), target rank (k), and eigengap (γ_k).

Theorem 3 (Davis-Kahan Theorem [5, 3]). *Let A and M be $d \times d$ self-adjoint matrices and denote the eigenvalues of A as $\lambda_1 \geq \dots \geq \lambda_d$. Define $\Pi_k(A)$ (resp. $\Pi_k(M)$) to be the projection onto the top- k eigenvectors of A . Then:*

$$\|\Pi_k(A) - \Pi_k(M)\|_2 \leq \frac{\|A - M\|_2}{\gamma_k(A)} \quad (5)$$

Where $\gamma_k(A) = \min_{j>k} |\lambda_k - \lambda_j|$ is the eigengap for A .

Before diving into details of our proof, let us first capture some intuition. Informally, the Davis-Kahan theorem says that the principal subspace of $\hat{\Sigma}$ is close to the principal subspace of Σ provided that $\|\hat{\Sigma} - \Sigma\|_2$, once $\hat{\Sigma}$ is appropriately normalized, is small. At the same time, the Matrix Bernstein inequality shows that if we can write $\hat{\Sigma} - \Sigma$ as a sum of centered random matrices, then we can control the spectral norm deviation. Since $\Sigma = \frac{1}{n} \sum_{t=1}^n x_t x_t^T$ and by defining $C_t = x_t x_t^T$, it suffices to show that \hat{C}_t , defined in Equation 1, is close to C_t . To see why this is true, we examine the vectors y_t and z_t .

Let us focus on y_t . An equivalent way to analyze y_t is to consider Φ_t fixed and draw the direction of x_t uniformly. By rotational invariance, we can think of Φ_t as projecting onto the first m standard basis elements, in which case it is easy to see that that $\|y_t\|_2^2$ has the same distribution as $\omega_t \|x_t\|_2^2$ where $\omega_t \sim \text{Beta}(\frac{m}{2}, \frac{d-m}{2})^2$.

To capture the distribution of y_t , we need a more geometric argument (See the left panel of Figure 2). The angle θ between x_t and Φ_t is $\cos^{-1}(\sqrt{\omega_t})$, which means that the magnitude of y_t in the direction of x_t is $\|y_t\|_2 \cos(\theta) = \omega_t \|x_t\|_2$. Moreover, the magnitude of y_t in the orthogonal direction is $\|x_t\|_2 \sqrt{\omega_t(1-\omega_t)}$ and this direction is chosen uniformly at random (subject to being orthogonal to x_t). Thus y_t is distributed as:

$$\omega_t x_t + \|x_t\|_2 \sqrt{\omega_t(1-\omega_t)} W_t \alpha_t$$

² ω_t can also be written as the ratio $\frac{a}{a+b}$ where $a \sim \chi_m^2, b \sim \chi_{d-m}^2$.

where $W_t \in \mathbb{R}^{d \times d-1}$ is a basis for the set of vectors orthogonal to x_t and α_t is drawn uniformly from the unit sphere in \mathbb{R}^{d-1} and independently from ω_t .

Analogously z_t is distributed as:

$$\eta_t x_t + \|x_t\|_2 \sqrt{\eta_t(1-\eta_t)} W_t \beta_t$$

Where the random variables η_t, β_t are independent from each other and from ω_t, α_t .

Since α_t, β_t are drawn uniformly on the sphere, it is clear that $\mathbb{E}\alpha_t = \mathbb{E}\beta_t = 0$, so that $\mathbb{E}y_t = \mathbb{E}z_t = \frac{m}{d}x_t$. Thus $\frac{d^2}{m^2}\hat{C}_t$ is an unbiased estimator for C_t . We are interested in approximating $\Sigma = \frac{1}{n}\sum_t C_t$, so the appropriate random variables to control are $X_t = \frac{d^2}{nm^2}\hat{C}_t - \frac{1}{n}C_t$.

To apply Theorem 2, we just need to compute the variance σ^2 and verify that $\|X_t\|_2$ is bounded. The latter is immediate; since Φ_t, Ψ_t are projection operators, it must be the case that $\|y_t\|_2, \|z_t\|_2 \leq \|x_t\|_2$, so that we can set $R = 2\frac{d^2}{nm^2}\mu$ (assuming $m \leq d$).

Some careful calculations, detailed in the appendix show that the variance can be bounded by:

$$\sigma^2 \leq 22\frac{\mu^2}{mn}.$$

We can now apply the inequality, which, once inverted, says that with probability $\geq 1 - \delta$:

$$\left\| \frac{d^2}{nm^2}\hat{\Sigma} - \Sigma \right\|_2 \leq \sqrt{\frac{44\mu^2}{nm} \log(d/\delta)} + \frac{2}{3} \frac{\mu d^2}{nm^2} \log(d/\delta). \quad (6)$$

Plugging this into the bound in Theorem 3 establishes Equation 2. Equation 3 is just a rearrangement of Equation 2. \square

3 Experiments

We complement our theoretical study of the compressive subspace learning algorithm with a number of simulations. Recall that our main goal is to capture the dependence between the error ϵ and the number of columns n . We do not believe that our results are tight in their dependence on the other parameters d, m, γ_k, μ . However we do verify qualitatively the influence of the other parameters on the error.

Our first simulation is a comparison between CSL and the compressive singular value decomposition algorithm of Halko *et. al* [7]. Their algorithm simply projects the columns of X onto a fixed m -dimensional subspace of \mathbb{R}^d and uses principal subspace of $Y = \Phi X$ as the estimate for Π . The main difference between their algorithm and ours is that we use a different m -dimensional projection per column, which leads to significant improvement in performance as demonstrated in Figure 1. In that figure, we plot the error $\|\hat{\Pi} - \Pi\|_2$ as a function of the number of columns n , for two different choices of m . The matrix rank and the target rank k is 2 and $d = 20$.

As predicted by Theorem 1, the error for our algorithm converges quickly to zero with n , with better convergence for larger m . On the other hand, the HMT algorithm does not enjoy any performance improvements as n increases, which is also predicted by their theoretical results. Thus we empirically see the averaging effect that is formalized by our theory.

We also verify that the *rate* of convergence in Theorem 1 is substantiated by empirical simulations. To this end, in Figure 3, we plot both the error alongside the rescaled error $\sqrt{n}\|\hat{\Pi} - \Pi\|_2$ as a function of n across several problem settings. In the left two plots we consider the exactly low rank case, while on the right two plots we consider matrices that have some non-zero eigengap γ_k , but are not exactly low rank.

The first thing to notice is that in all simulations, $\sqrt{n}\|\hat{\Pi} - \Pi\|_2$ does appear to level out to some constant value, which demonstrates that our algorithm does converge at $n^{-1/2}$ rate. Moreover, in comparing between the trials, we see that increasing the dimensionality degrades the performance of the algorithm. On the other hand, increasing the target rank k seems to have less influence.

We study the effect of d, m, k and γ_k in the bottom row of Figure 3. Our interpretation of the results is somewhat more qualitative, as we do not expect our theory to precisely capture the dependence between the error and these parameters. In the first figure, we see that increasing the number of measurements per column m significantly improves the performance of the algorithm, as one might expect. It does appear that the dependence is $\epsilon \asymp 1/m$, rather than the inverse-quadratic dependence in Theorem 1. In terms of problem dimension d , plotted in the second figure, the dependence appears to be linear rather than quadratic as predicted by Theorem 1.

As we mentioned earlier, there is no explicit dependence on the target rank k in Theorem 1, although it does play some role indirectly through γ_k and μ . However, in the third plot from the left of the bottom row of Figure 3, we clearly see a linear dependence between k and ϵ . We suspect that a more careful analysis of our algorithm can explicitly introduce the dependence on k while possibly removing one factor of d , leading to a bound that is more reminiscent of existing matrix completion and compression results [7, 12, 11].

Lastly we plot the effect of the eigengap γ_k on the error. We comment on two main effects: (1) increasing γ_k improves the performance of the algorithm, and (2) for fixed eigengap, the performance improves with n . This qualitatively justifies γ_k in the denominator of Equation 2.

4 Conclusion

In this manuscript, we demonstrate how one can approximate the principal subspace of a data matrix from very few compressive measurements per column. The main insight is that by using an independent random compression operator on each column, we can effectively take averages across columns, which preserves the signal but diminishes the noise stemming from the compression.

These results are practically and theoretically appealing yet several challenges still remain. First, our simulations suggest that the error should scale linearly with d and k and inversely linearly with m , yet our error bound has worse dependences. We would like to improve these dependences so that the theoretical results are more predictive of our experimental findings. Secondly, as we mentioned before, in theory, our algorithm requires at least two compressive measurements per column, although we did not observe such a requirement in simulation. While Gonen *et. al* proved that such a requirement is necessary in the missing data setting, their justification does not immediately carry over to our setting [6]. It would be worthwhile to understand the differences between these sampling paradigms and to show either that two compressive measurements per column are necessary or that one suffices.

Lastly, it is not clear what the fundamental limits are for the compressive subspace learning problem. Given the compressed measurements y_t, z_t for each column x_t , is there a lower bound on the error achievable by *any* algorithm? Is there an algorithm that achieves, or nearly achieves, this bound?

We hope to address these questions in future work.

Acknowledgements

This research is supported in part by NSF awards IIS-1116458, IIS-1247658 and CAREER IIS-1252412 and an AFOSR YIP award. AK is supported in part by a NSF Graduate Research Fellowship.

References

- [1] Dimitris Achlioptas and Frank Mcsherry. Fast computation of low-rank matrix approximations. *Journal of the ACM*, April 2007.
- [2] Christos Boutsidis, Petros Drineas, and Malik Magdon-Ismael. Near Optimal Column-Based Matrix Reconstruction. In *IEEE Symposium on Foundations of Computer Science*. IEEE, October 2011.
- [3] Chandler Davis and W. M. Kahan. The Rotation of Eigenvectors by a Perturbation. III. *SIAM Journal on Numerical Analysis*, March 1970.
- [4] Petros Drineas and Ravi Kannan. Pass efficient algorithms for approximating large matrices. In *ACM-SIAM Symposium on Discrete Algorithms*. Society for Industrial and Applied Mathematics, January 2003.
- [5] G H Golub and C F Van Loan. *Matrix Computations*. Johns Hopkins University Press, 1996.
- [6] Alon Gonen, Dan Rosenbaum, Yonina Eldar, and Shai Shalev-Shwartz. The Sample Complexity of Subspace Learning with Partial Information. *arXiv:1402.4844*, 2014.
- [7] N. Halko, P. G. Martinsson, and J. A. Tropp. Finding Structure with Randomness: Probabilistic Algorithms for Constructing Approximate Matrix Decompositions. *SIAM Review*, January 2011.
- [8] A Krishnamurthy and A Singh. Low-Rank Matrix and Tensor Completion via Adaptive Sampling. *Advances in Neural Information Processing Systems*, 2013.
- [9] Edo Liberty. Simple and deterministic matrix sketching. In *ACM International Conference on Knowledge Discovery and Data Mining*, New York, New York, USA, August 2013. ACM Press.
- [10] Edo Liberty, Franco Woolfe, Per-Gunnar Martinsson, Vladimir Rokhlin, and Mark Tygert. Randomized algorithms for the low-rank approximation of matrices. *Proceedings of the National Academy of Sciences*, December 2007.
- [11] Sahand Negahban and Martin J. Wainwright. Restricted strong convexity and weighted matrix completion: optimal bounds with noise. *The Journal of Machine Learning Research*, March 2012.
- [12] Benjamin Recht. A Simpler Approach to Matrix Completion. *The Journal of Machine Learning Research*, February 2011.
- [13] Tamas Sarlos. Improved Approximation Algorithms for Large Matrices via Random Projections. In *IEEE Symposium on Foundations of Computer Science*. IEEE, 2006.
- [14] Joel A. Tropp. User-Friendly Tail Bounds for Sums of Random Matrices. *Foundations of Computational Mathematics*, August 2011.

A Bounding the Variance

Here we bound the variance:

$$\sigma^2 = \left\| \sum_{t=1}^n \mathbb{E} X_t^2 \right\|, \quad X_t = \frac{d^2}{2m^2n} (y_t z_t^T + z_t y_t^T) - \frac{1}{n} x_t x_t^T.$$

Recall that:

$$\begin{aligned} y_t &= \omega_t x_t + \|x_t\| \sqrt{\omega_t(1-\omega_t)} W_t \alpha_t \\ z_t &= \eta_t x_t + \|x_t\| \sqrt{\eta_t(1-\eta_t)} W_t \beta_t \end{aligned}$$

where $\omega_t, \eta_t \sim \text{Beta}(\frac{m}{2}, \frac{d-m}{2})$, $W_t \in \mathbb{R}^{d \times d-1}$ is an basis for the set of vectors orthogonal to x_t and α_t, β_t are drawn uniformly from the unit sphere in \mathbb{R}^{d-1} . Everything is independent.

Notice that $\mathbb{E}\omega_t^2 = \mathbb{E}\eta_t^2 = \frac{m(m+2)}{d(d+2)}$, so that:

$$\begin{aligned} \mathbb{E} y_t z_t^T y_t z_t^T &= \mathbb{E} \left(\omega_t \eta_t \|x_t\|^2 + \|x_t\|^2 \sqrt{\omega_t(1-\omega_t)} \sqrt{\eta_t(1-\eta_t)} \alpha_t^T \beta_t \right) \times \\ &\quad \times \left(\omega_t \eta_t x_t x_t^T + \omega_t \|x_t\| \sqrt{\eta_t(1-\eta_t)} x_t \beta_t^T W_t^T + \eta_t \|x_t\| \sqrt{\omega_t(1-\omega_t)} W_t \alpha_t^T x_t^T \right. \\ &\quad \left. + \|x_t\|^2 \sqrt{\omega_t(1-\omega_t)} \sqrt{\eta_t(1-\eta_t)} W_t \alpha_t \beta_t^T W_t^T \right) \\ &= \|x_t\|^2 \left(\frac{m(m+2)}{d(d+2)} \right)^2 x_t x_t^T + \|x_t\|^4 \left(\frac{m(d-m)}{d(d+2)(d-1)} \right)^2 W_t W_t^T. \end{aligned}$$

This follows since the last two terms in the second line make no contribution to the sum.

We also have:

$$\begin{aligned} &\mathbb{E} y_t z_t^T z_t y_t^T \\ &= \mathbb{E} \left(\eta_t^2 \|x_t\|^2 + \|x_t\|^2 (\eta_t - \eta_t^2) \|\beta_t\|^2 \right) \times \\ &\quad \times \left(\omega_t^2 x_t x_t^T + \omega_t \|x_t\| \sqrt{\omega_t(1-\omega_t)} (x_t \alpha_t^T W_t^T + W_t \alpha_t x_t^T) + \|x_t\|^2 (\omega_t - \omega_t^2) W_t \alpha_t \alpha_t^T W_t^T \right) \\ &= \|x_t\|^2 \left(\frac{m(m+2)}{d(d+2)} \right)^2 x_t x_t^T + \|x_t\|^4 \frac{m^2(m+2)(d-m)}{d^2(d+2)^2(d-1)} W_t W_t^T \\ &\quad + \|x_t\|^2 \frac{m^2(m+2)(d-m)}{d^2(d+2)^2(d-1)} x_t x_t^T + \|x_t\|^4 \left(\frac{m(d-m)}{d(d+2)(d-1)} \right)^2 W_t W_t^T \end{aligned}$$

Thus we can bound the variance by:

$$\begin{aligned}
\sigma^2 &= \frac{d^4}{m^4 n^2} \left\| \sum_{t=1}^n \frac{1}{2} \mathbb{E} y_t z_t^T y_t z_t^T + \frac{1}{2} \mathbb{E} y_t z_t^T z_t y_t^T - \frac{m^4}{d^4} x_t x_t^T x_t x_t^T \right\| \\
&\leq \frac{d^4}{m^4 n^2} \sum_{t=1}^n \left\| \frac{1}{2} \mathbb{E} y_t z_t^T y_t z_t^T + \frac{1}{2} \mathbb{E} y_t z_t^T z_t y_t^T - \frac{m^4}{d^4} x_t x_t^T x_t x_t^T \right\| \\
&\leq \frac{d^4}{m^4 n^2} \sum_{t=1}^n \|x_t\|^4 \left(\frac{m(d-m)}{d(d+2)(d-1)} \right)^2 \|W_t W_t^T\| + \frac{1}{2} \|x_t\|^4 \left(\frac{m^2(m+2)(d-m)}{d^2(d+2)^2(d-1)} \right) \|W_t W_t^T\| \\
&\quad + \frac{1}{2} \|x_t\|^2 \left(\frac{m^2(m+2)(d-m)}{d^2(d+2)^2(d-1)} \right) \|x_t x_t^T\| + \|x_t\|^2 \left(\frac{m^2(m+2)^2}{d^2(d+2)^2} - \frac{m^4}{d^4} \right) \|x_t x_t^T\| \\
&\leq \frac{d^4 \mu^2}{m^4 n} \left[\left(\frac{m(d-m)}{d(d+2)(d-1)} \right)^2 + \left(\frac{m^2(m+2)(d-m)}{d^2(d+2)^2(d-1)} \right) + \left(\frac{m^2(m+2)^2}{d^2(d+2)^2} - \frac{m^4}{d^4} \right) \right]
\end{aligned}$$

The main point is that this expression is actually $O(\frac{\mu^2}{mn})$, which gives us the variance bound used in our theorem. To conclude we need to control the rational functions. The last one can be bounded as:

$$\begin{aligned}
&\left(\frac{m^2(m+2)^2}{d^2(d+2)^2} - \frac{m^4}{d^4} \right) \\
&= \frac{m^2(m^2 + 4m + 2)d^2 - m^2(d^2 + 4d + 2)}{d^2 d^2 (d+2)^2} \\
&= \frac{m^2(m^2 d^2 + 4m d^2 + 2d^2 - m^2 d^2 - 4m^2 d - 2m^2)}{d^2 d^2 (d+2)^2} \\
&\leq \frac{m^2(4md^2 + 2d^2)}{d^2 d^2 (d+2)^2} \leq \frac{6m^3}{d^4}
\end{aligned}$$

The first two terms can together be bounded by:

$$\begin{aligned}
&\left(\frac{m(d-m)}{d(d+2)(d-1)} \right)^2 + \left(\frac{m^2(m+2)(d-m)}{d^2(d+2)^2(d-1)} \right) \\
&= \frac{m^2}{d^2} \left[\frac{(d-m)^2 + (m+2)(d-m)(d-1)}{(d+2)^2(d-1)^2} \right] \\
&= \frac{m^2(d-m)}{d^2} \frac{d-m+md+2d-m-2}{(d+2)^2(d-1)^2} \\
&\leq \frac{m^2(d-m)}{d^2} \frac{3d+md}{(d+2)^2(d-1)^2} \\
&\leq \frac{m^2(3+m)}{(d+2)^2(d-1)^2} \leq \frac{16m^3}{d^4}
\end{aligned}$$

Here in the last step we bounded $3 \leq 3m$ and $d-1 \geq d/2$, which is true provided $d \geq 2$.

Plugging these two bounds into our bound above shows that the variance is:

$$\sigma^2 \leq \frac{22\mu^2}{mn}.$$