Applications

High Dim Probability & Linear Algebra for ML and Sig Proc

Instructor: Namrata Vaswani
Professor of Electrical and Computer Engineering

Iowa State University

Non-asymptotic Random Matrix Theory: what is it?

Weak and Strong Laws of Large Numbers provide asymptotic results: WLLN says that for i.i.d. random variables, under simple assumptions on the "niceness" of distributions, the sample mean converges to the population mean in probability as the number of samples, m, tends to infinity.

But it does not tell us what m to pick to ensure a certain error bound with at least a certain probability.

Similarly, asymptotic results exist that show that the max, min singular values of a random Gaussian matrix A of size $N \times n$ with $N > n$, converge to $\sqrt{N} \pm \sqrt{n}$ as $N, n \to \infty$ with $N/n = C$.

This course provides results that help obtain finite sample (non-asymptotic) high probability bounds on the minimum and maximum singular values of large matrices

- with either independent entries
- or matrices of the form $\mathbf{Z}\mathbf{Z}'$ with $\mathbf{Z}$ containing independent subG entries
- or on sums of independent random matrices that have bounded norms
- or...

There is more

Applications I

- Compressed Sensing / Sparse Recovery: Given $y := Ax$ recover x from y when y is shorter than x. Use sparsity of x.
- Low-rank Matrix Completion: Given a subset of entries of a low-rank matrix M, complete the matrix
  - given $y = \mathcal{P}_\Omega(M)$, find M. Ω: set of indices of the observed entries
- Matrix Sensing: given a set of n linear functions of M, find M using the fact that M is low-rank
  - given $y = \mathcal{A}(M)$ where $\mathcal{A}(.)$ is a linear operator, find M. This can be written as $y_i =< A_i, M >$ where $< A, B >= trace(A'B)$ is the usual inner product.
- Robust PCA: given $Y := X + L$, find X and L
  - $L =$ unknown low rank matrix.
  - $X =$ sparse matrix (corresponds to outliers)
- Phase retrieval: compute vector x from $y := |Ax|^2$. Here $|.|$ means element-wise magnitude of the vector. More specifically $y_i = |A^i x|^2$ (here A^i is the i-th row of A).
  - the term phase retrieval comes from Fourier imaging where A is the DFT matrix; but now it's used more generally for any matrix A
- Sparse PR
- Low Rank PR

Applications II

- Bounding the degree of dense or sparse random graphs
- Community detection in networks/graphs
  - A random n/w with n nodes, node connections independent, two communities, nodes within a community are connected w.p. p, nodes in different communities are connected with probability $q < p$. If q sufficiently smaller than p, it should be possible to detect the communities given a realization of the graph
- Ranking and individualized ranking estimation
- CS: projection imaging - MRI, CT, single-pixel camera, radar, ...
- MC: recommendation system design, e.g., Netflix problem
- Matrix sensing: one special case is phase retrieval. Notice that we can rewrite $y_i = A^i x x' A^{i'} = < xx', A^i A^{i'} >$
- RPCA: recommendation system design in the presence of outliers, Video analytics, Survey data analysis,
- Phase retrieval: astronomy, X-ray crystallography,...
- Graphs: model computer or social networks. Degree of graph: expected number of neighbors of a given node
  Community detection: detecting communities in social n/w

Applications: some details I

1. Compressive Sensing / Sparse Recovery
   Recover an s-sparse n-length vector $\mathbf{x}$ from $\mathbf{y} := \mathbf{A}\mathbf{x}$ with $\mathbf{A}$ being an $m \times n$ matrix with $m < n$. Suppose $\mathbf{A}$ is random Gaussian. Then how large does m need to be so that $\mathbf{x}$ can be exactly recovered?
   Answer: if $s \ll n$, $m \ll n$ suffices, with high probability, to exactly recover x using an "efficient algorithm". In particular, m of order $s \log(n/s)$ suffices.
   "efficient algorithm": an algorithm that is guaranteed to run in time that is at most polynomial in the signal dimension, n.
   "intractable algorithm": algo with complexity exponential in the signal dimension or in the sparsity: $O(e^n)$ or $O(e^s)$ etc.
   App: This problem with $\mathbf{A}$ being random Fourier occurs in medical imaging: MRI and CT.
2. PCA
   Given a symmetric p.s.d. matrix $\mathbf{M}$, find the column span of its top r eigenvectors, or, sometimes, find the "optimal" rank r approximation. Here "optimal" means $\min_{\mathbf{L}} ||\mathbf{M} - \mathbf{L}||_2$
   App: find the r directions with the largest variance for a dataset, $\mathbf{D}$: here $\mathbf{M} = (\mathbf{D} - \mu)(\mathbf{D} - \mu)^T$, with μ the mean (expected value) of the data vectors
   If an $n \times m$ matrix $\mathbf{D}$ is approx rank r, it can be expressed as $\mathbf{D} = \mathbf{L} + \mathbf{W}$ where $\mathbf{L}$ is rank r and $\mathbf{W}$ is the residual error. If columns of $\mathbf{W}$ are mutually independent and either bounded or sub-Gaussian, how large should m be in order to guarantee accurate subpsace recovery?

Applications: some details II

3. Low rank matrix sensing and completion
   Sensing: recover $\mathbf{M}$ from $\mathbf{y}_i := \langle \mathbf{A}_i, \mathbf{M} \rangle, \ i = 1, 2, \dots, m$.
   Completion: recover $\mathbf{M}$ from a subset of its entries.
4. Phase Retrieval (PR)
5. Sparse PR
6. Low Rank PR
7. Community detection in networks/graphs
8. Bounding the degree of dense or sparse random graphs
9. ....

Alternating Min

- Goal: compute $\min_{x,y} f(x,y)$ when $f(.)$ is non-convex
- Clearly $\min_{x,y} f(x,y) = \min_x(\min_y f(x,y))$ but of course in most cases, RHS is also hard to compute.
- Consider the class of problems where the min is easy when one variable is fixed, i.e., $\min_y f(x_0,y)$ is easy for a given x_0 and $\min_x f(x,y_0)$ is easy for a given y_0.
  - A common solution: Alt-Min
  - Start with an initial guess x_0.
  - Compute $y_1 \in \arg\min_y f(x_0,y)$
  - Compute $x_1 \in \arg\min_y f(x,y_1)$
  - Repeat above until a stopping criterion is met.
- Guarantees? Till very recently none. Recent work:
  - If initialized carefully, Alt-Min gets to within a small error of the true solution in a finite number of iterations. Possible to bound this number also.
  - A common approach to initialization: "spectral method" - compute the top eigenvector of an appropriately defined matrix
  - Guarantees exist for Matrix Completion and for Phase Retrieval

Gradient descent based approaches for non-convex problems

- With a suitable initialization, it is possible to get a guarantee

- Truncated gradient descent idea of "truncated Wirtinger flow" paper: the gradient turns out to be a weighted average of certain vectors; discard those weights that are too large and compute a truncated gradient estimate
-

Compressive Sensing

1. Goal is to find an n-length sparse vector x (a vector with at most s nonzero entries with $s \ll n$) from only m measurements $\mathbf{y} = \mathbf{A}x$ when $m < n$. Here A is an $m \times n$ matrix.
2. It has been shown by Candes et al that if A satisfies the s-RIP (restricted isometry property), then exact recovery of an s-sparse x can be guaranteed.
3. s-RIP: A matrix A satisfies this if for any subset $T \subset \{1, 2, \dots, n\}$ of size s, the $m \times |T|$ sub-matrix A_T is an approximate isometry. This means that, for all vectors $z \in \Re^s$,

$$(1 - \delta_s)\|z\|^2 \leq \|A_T z\|^2 \leq (1 + \delta_s)\|z\|^2$$

for a $\delta_s < 1$. δ_s is called the Restricted Isometry Constant. The smaller it is the better is the approx isometry.
This requirement is equivalent to requiring that

$$(1 - \delta_s) \leq s_{\min}(A_T) \leq s_{\max}(A_T) \leq 1 + \delta_s$$

for all subsets T of size s. Here $s_{\min}, s_{\max}$ denote the min and max singular values.
4. For matrices A with independent, zero mean sub-Gaussian rows, and sub-G norm a constant, and a fixed subset T of size s, we can use the subG rows' Theorem to show that $\sigma_{\max}(A_T), \sigma_{\min}(A_T)$ are upper/lower bounded by $\sqrt{m} \pm C(\sqrt{s} + t)$ w.p. at least $1 - \exp(-ct^2)$.

5. By using another union bound over all sets T of size s (there are $\binom{n}{s} \le e^{s \log n/s}$ such subsets), the above bound holds w.p. at least $1 - \exp(s \log(n/s) - ct^2)$. To ensure that this probability is large enough (is at least $1 - \exp(-s \log(n/s))$), we need to set $t = C\sqrt{s \log n/s}$ with a C large enough. With this choice of t, to ensure that the min singular value is at least $\sqrt{m}(1 - \delta_s)$, we need $m \ge Cs \log(n/s)/\delta_s^2$.

PCA:

1. Given a dataset with covariance matrix Σ, PCA finds the directions of largest variability in the dataset by computing the eigenvectors of Σ with the r largest eigenvalues, and projecting each data vector onto the subspace spanned by these r eigenvectors. One can argue that, for a fixed r, this projection minimizes the expected squared reconstruction error. For more details on PCA, see `https://www.ece.iastate.edu/~namrata/MachLearn_SigProc/Summary_Notes.pdf` (see latest dropbox file actually, to be updated).
2. The true cov matrix Σ is unknown. Assuming the data is zero mean, we can compute its empirical estimate as

$$\Sigma_m = \sum_{i=1}^{m} X_i X_i'$$

3. Clearly, $\mathbb{E}[\Sigma_m] = \Sigma$.

4. Assume that the X_i's are sub-G. Can use the subG rows' Theorem to argue the following: suppose $m \geq n$; then,
    1. $\mathbb{E}[\|\Sigma_m - \Sigma\|] \leq CK^2\sqrt{n/m}\|\Sigma\|$
    2. $\Pr(\|\Sigma_m - \Sigma\| > CK^2(\sqrt{n} + t)/\sqrt{m}) \leq \exp(-t^2)$
    3. Thus, to get an estimate Σ_m that is within $\epsilon\|\Sigma\|$ error of the true Σ, we need $m \geq C\frac{K^2}{\|\Sigma\|}n$.
5. Davis-Kahan sin theta theorem: bounds the perturbation of eigen-vectors or of subspaces of eigenvectors when a symmetric matrix is perturbed. Subspaces version:

Theorem (Davis-Kahan for principal subspaces)

Let $\mathbf{D}, \hat{\mathbf{D}}$ *be Hermitian matrices of size* $n \times n$ *with top* r *eigenvectors denoted by the* $n \times r$ *matrices with ortho cols* $\mathbf{U}, \hat{\mathbf{U}}$. *Then*

$$SD(\hat{\mathbf{U}}, \mathbf{U}) := \|(\mathbf{I} - \hat{\mathbf{U}}\hat{\mathbf{U}}')\mathbf{U}\| \leq \frac{\|(\hat{\mathbf{D}} - \mathbf{D})\mathbf{U}\|}{\lambda_r(\mathbf{D}) - \lambda_{r+1}(\hat{\mathbf{D}})} \leq \frac{\|(\hat{\mathbf{D}} - \mathbf{D}\|}{\lambda_r(\mathbf{D}) - \lambda_{r+1}(\mathbf{D}) - \|\hat{\mathbf{D}} - \mathbf{D}\|}$$

This shows that if Σ_m is close to Σ in 2-norm, then their corresponding top eigenvectors also span subspaces that are close to each other in the subspace distance defined by SD. This measures the sine of the largest principal angle between the two subspaces.

Applications Details IV

6. The goal of PCA is to find the span of $\mathbf{U}$ which is the matrix of top r eigenvectors of Σ.
7. By using the above result followed by the Theorem given earlier to bound $\|\Sigma_m - \Sigma\|$ we can decide on the required sample complexity m (required number of samples/data-points m to guarantee accurate principal subspace recovery. For a general matrix Σ, we will need $m \geq CK^2 n/\epsilon^2$ to get an ϵ accurate estimate of Σ
8. Can get a modified result with a lower sample complexity for the setting where Σ is approximately low rank: the effective rank, $trace(\Sigma)/\|\Sigma\|$, is much smaller than n.

Low Rank Matrix Recovery

1. LR Matrix Completion , LR Matrix Sensing, Compressive PCA (LR Compressive Sensing), LR Phase Retrieval
2. Non-convex solutions to all these problems consist of a spectral initialization step, that provides the initial estimate of the column span of the unknown LR matrix, followed by either an alternating minimization algorithm or a gradient descent (GD) method.
3. Spectral init: compute top r left singular vectors of a carefully defined matrix (usually a sum of independent matrices).
4. To show that the spectral init output indeed is a good approx to the true column span of the LR matrix, the typical approach involves use of the Davis-Kahan sin theta theorem, followed by use of one of the results from earlier to obtain high probability bounds on each of the terms in the Davis-Kahan bound.
5. In some cases, Davis-Kahan can be replaced by Wedin's sin theta theorem.

6. For analyzing the iterations, one again arrives at a subspace error bound either using one of the above two results, or directly using other ideas. The terms of this bound are then bounded with high probability using one of the results from above.
7. Typically: either one of the two matrix results given above is used or ideas similar to the proof of the sub-Gaussian result are used to bound the terms directly.

Community detection in graphs:

- A random n/w with n nodes, node connections independent, two communities, nodes within a community are connected w.p. p, nodes in different communities are connected with probability $q < p$. If q sufficiently smaller than p, it should be possible to detect the communities given a realization of the graph
  Spectral clustering algorithm.
  (i) Compute adjacency matrix of the graph ($A_{ij} = 1$ if i connected to j, zero otherwise). A is symmetric.
  (ii) Compute the eigenvector of A corresponding to second largest eigenvalue. Denote by $v_2(A)$
  (iii) Partition the nodes based on the signs of entries of $v_2(A)$
  Idea this works: Can show that $\mathbb{E}[A]$ is a rank-2 matrix. Its first eigenvector is the all-ones vector scaled by $1/\sqrt{n}$; the second eigenvector is a vector of $+1/\sqrt{n}, -1/\sqrt{n}$'s with the sign indicating which community it belongs to. Write $A = \mathbb{E}[A] + H$. If H is small, one can argue that the eigenvectors of A will be close to those of $\mathbb{E}[A]$ (Davis-Kahan sin theta

theorem) as long as the eigen-gaps are large enough. Here this means $\min((p-q)/2, q) \geq \mu > 0$.
To bound $\|H\| = \|A - \mathbb{E}[A]\|$, we use Theorem 4.4.5 of book
Since the entries of A are independent and each is either $Bern(p)$ or $Bern(q)$, the matrix contains sub-Gaussian independent entries. The rows may not be isotropic so our previous result does not apply, but an easier result that only provides an upper bound on $\|A\|$ (Theorem 4.4.5 of book) applies. This result does not use isotropy. Proved using $\|A\| = \max_{x,y \ unit \ norm} x'Ay$; for fixed x, y, $x'Ay = \sum_{i,j} x_i y_j A_{ij}$; can show this is a sub-G with sub-G norm CK^2; then use subG tail bound; followed by epsilon-net argument and union bound.

A detailed set of slides on Low Rank Phase Retrieval and its linear version, Compressive PCA, follows.

Old slides on sparse recovery

The sparse recovery / compressed sensing problem

- Given $y := Ax$ where A is a fat matrix, find x.
  - underdetermined system, without any other info, has infinite solutions
- Key applications where this occurs: Computed Tomography (CT) or MRI
  - CT: acquire radon transform of cross-section of interest
  - typical set up: obtain line integrals of the cross-section along a set of parallel lines at a given angle, and repeated for a number of angles from 0 to π), common set up: 22 angles, 256 parallel lines per angle
  - by Fourier slice theorem, can use radon transform to compute the DFT along radial lines in the 2D-DFT plane
  - Projection MRI is similar, directly acquire DFT samples along radial lines
  - parallel lines is most common type of CT, other geometries also used.
- Given 22x256 data points of 2D-DFT of the image, need to compute the 256x256 image

Limitation of zero-filling

- A traditional solution: zero filling + I-DFT
  - set the unknown DFT coeff's to zero, take I-DFT
  - not good: leads to spatial aliasing
- Zero-filling is the minimum energy (2-norm) solution, i.e. it solves $\min_x \|x\|_2 \ s.t. \ y = Ax$. Reason
  - clearly, min energy solution in DFT domain is to set all unknown coefficients to zero, i.e. zero-fill
  - (energy in signal) = (energy in DFT)*2π, so min energy solution in DFT domain is also the min energy solution
- The min energy solution will not be sparse because 2-norm is not sparsity promoting
  - In fact it will not be sparse in any other ortho basis either because $\|x\|_2 = \|\Phi x\|_2$ for any orthonormal Φ. Thus min energy solution is also min energy solution in Φ basis and thus is not sparse in Φ basis either
- But most natural images, including medical images, are approximately sparse (or are sparse in some basis)

Sparsity in natural signals/images

- Most natural images, including medical images, are approximately sparse (or are sparse in some basis)
  - e.g. angiograms are sparse
  - brain images are well-approx by piecewise constant functions (gradient is sparse): sparse in TV norm
  - brain, cardiac, larynx images are approx. piecewise smooth: wavelet sparse
- Sparsity is what lossy data compression relies on: JPEG-2000 uses wavelet sparsity, JPEG uses DCT sparsity
- But first acquire all the data, then compress (throw away data)
- In MRI or CT, we are just acquiring less data to begin with - can we still achieve exact/accurate reconstruction?

Use sparsity as a regularizer

- Min energy solution $\min_x \|x\|_2 \ s.t. \ y = Ax$ is not sparse, but is easy to compute $\hat{x} = A'(AA')^{-1}y$
- Can we try to find the min sparsity solution, i.e. find $\min_x \|x\|_0 \ s.t. \ y = Ax$
- Claim: If true signal, x_0, is exactly S-sparse, this will have a unique solution that is EXACTLY equal to x_0 if $spark(A) > 2S$
  - $spark(A) =$ smallest number of columns of A that are linearly dependent.
  - in other words, any set of (spark-1) columns are always linearly independent
- proof in class
- Even when x is approx-sparse this will give a good solution
- But finding the solution requires a combinatorial search: $O(\sum_{k=1}^{S} choosemk) = O(m^S)$

Practical solutions [Chen,Donoho'95] [Mallat,Zhang'93]

- Basis Pursuit: replace ℓ_0 norm by ℓ_1 norm: closest norm to ℓ_0 that is convex

$$\min_x \|x\|_1 \ s.t. \ y = Ax$$

- Greedy algorithms: Matching Pursuit, Orthogonal MP
- Key idea: all these methods "work" if columns of A are sufficiently "incoherent"
- "work": give exact reconstruction for exactly sparse signals and zero noise, give small error recon for approx. sparse (compressible) signals or noisy measurements

Compressive Sensing

- name: instead of capturing entire signal/image and then compressing, can we just acquire less data?
- i.e. can we compressively sense?
- MRI (or CT): data acquired one line of Fourier projections at a time (or random transform samples at one angle at a time)
- if need less data: faster scan time
- new technologies that use CS idea:
  - single-pixel camera,
  - A-to-D: take random samples in time: works when signal is Fourier sparse
  - imaging by random convolution
  - decoding "sparse" channel transmission errors.
- Main contribution of CS: theoretical results

General form of Compressive Sensing

- Assume that an N-length signal, z, is S-sparse in the basis Φ, i.e. $z = \Phi x$ and x is S-sparse.
- We sense

$$y := \Psi z = \underbrace{\Psi \Phi}_{} Ax$$

- It is assumed that Ψ is "incoherent w.r.t. Φ"
  - or that $A := \Psi \Phi$ is "incoherent"
- Find x, and hence $z = \Phi x$, by solving

$$\min_x \|x\|_1 \ s.t. \ y = Ax$$

- A random Gaussian matrix, Ψ, is "incoherent" w.h.p for S-sparse signals if it contains $O(S \log N)$ rows
- And it is also incoherent w.r.t. any orthogonal basis, Φ w.h.p. This is because if Ψ is r-G, then $\Psi\Phi$ is also r-G (ϕ any orthonormal matrix).
- Same property for random Bernoulli.

Quantifying "incoherence"

- Rows of A need to be "dense", i.e. need to be computing a "global transform" of x.
- Mutual coherence parameter, $\mu := \max_{i \neq j} |A_i' A_j| / \|A_i\|_2 \|A_j\|_2$
- spark(A) $=$ smallest number of columns of A that are linearly dependent.
- Or, any set of $(spark(A) - 1)$ columns of A are always linearly independent.
- RIP, ROP
- many newer approaches...

Quantifying "incoherence": RIP

- A $K \times N$ matrix, A satisfies the S-Restricted Isometry Property if constant δ_S defined below is positive.
- Let A_T, $T \subset \{1, 2, \ldots N\}$ be the sub-matrix obtained by extracting the columns of A corresponding to the indices in T. Then δ_S is the smallest real number s.t.

$$(1 - \delta_S)\|c\|^2 \leq \|A_T c\|^2 \leq (1 + \delta_S)\|c\|^2$$

for all subsets $T \subset \{1, 2, \ldots N\}$ of size $|T| \leq S$ and for all $c \in \Re^{|T|}$.

  - In other words, every set of S or less columns of A has singular values b/w $\sqrt{1 \pm \delta_S}$
  - $\Leftrightarrow$ every set of S or less columns of A approximately orthogonal
  - $\Leftrightarrow$ A is approximately orthogonal for any S-sparse vector, c.

Examples of RIP

- If A is a random Gaussian, random Bernoulli, or Partial Fourier matrix with about $O(S \log N)$ rows, it will satisfy RIP(S) w.h.p.
- Partial Fourier * Wavelet: somewhat "incoherent"

Use for spectral estimation and comparison with MUSIC

- Given a periodic signal with period N that is a sparse sum of S sinusoids, i.e.

$$x[n] = \sum_k X[k] e^{j2\pi kn/N}$$

  where the DFT vector, X, is a $2S$-sparse vector.
- In other words, x[n] does not contain sinusoids at arbitrary frequencies (as allowed by MUSIC), but only contains harmonics of $2\pi/N$ and the fundamental period N is known.
- In matrix form, $x = F^* X$ where F is the DFT matrix and $F^{-1} = F^*$.

- Suppose we only receive samples of $x[n]$ at random times, i.e. we receive $y = Hx$ where H is an "undersampling matrix" (exactly one 1 in each row and at most one 1 in each column)
- With random time samples it is not possible to compute covariance of $\underline{x}[n] := [x[n], x[n-1], \ldots x[n-M]]'$, so cannot use MUSIC or the other standard spectral estimation methods.
- But can use CS. We are given $y = HF^*X$ and we know X is sparse. Also, $A := HF^*$ is the conjugate of the partial Fourier matrix and thus satisfies RIP w.h.p.
- If have $O(S \log N)$ random samples, we can find X exactly by solving

$$\min_X \|X\|_1 \; s.t. \; y = HF^*X$$

Quantifying "incoherence": ROP

- θ_{S_1,S_2}: measures the angle b/w subspaces spanned by A_{T_1}, A_{T_2} for disjoint sets, T_1, T_2 of sizes less than/equal to S_1, S_2 respectively
- $\theta_{S1,S2}$ is the smallest real number such that

$$|c1' A'_{T1} A_{T2} c2| < \theta_{S1,S2} \ \|c1\| \ \|c2\|$$

  for all $c1, c2$ and all sets $T1$ with $|T1| \le S1$ and all sets $T2$ with $|T2| \le S2$
- In other words

$$\theta_{S1,S2} = \min_{T1,T2:|T1|\le S1,|T2|\le S2} \min_{c1,c2} \frac{|c1' A'_{T1} A_{T2} c2|}{\|c1\| \ \|c2\|}$$

- Can show that δ_S is non-decreasing in S, θ is non-decreasing in $S1, S2$
- Also $\theta_{S1,S2} \le \delta_{S1+S2}$
- Also, $\|{A_{T_1}}' A_{T_2}\| \le \theta_{|T_1|,|T_2|}$

Theoretical Results

- If x is S-sparse, $y = Ax$, and if $\delta_S + \theta_{S,2S} < 1$, then basis pursuit exactly recovers x
- If x is S-sparse, $y = Ax + w$ with $\|w\|_2 \leq \epsilon$, and $\delta_{2S} < (\sqrt{2} - 1)$, then solution of basis-pursuit-noisy, $\hat{x}$ satisfies
$$\|x - \hat{x}\| \leq C_1(\delta_{2S})\epsilon$$
- basis-pursuit-noisy:
$$\min_x \|x\|_1 \ s.t. \ \|y - Ax\|_2 \leq \epsilon$$

MP and OMP

Applications

DSP applications

- Fourier sparse signals
  - Random sample in time
  - Random demodulator + integrator + uniform sample with low rate A-to-D
- N length signal that is sparse in any given basis Φ
  - Circularly convolve with an N-tap all-pass filter with random phase
  - Random sample in time or use random demodulator architecture

Compressibility: one definition

Papers to Read

- Decoding by Linear Programming (CS without noise, sparse signals)
- Dantzig Selector (CS with noise)
- Near Optimal Signal Recovery (CS for compressible signals)
- Applications of interest for DSP
  - Beyond Nyquist:... Tropp et al
  - Sparse MRI: ... Lustig et al
  - Single pixel camera: Rice, Baranuik's group
  - Compressive sampling by random convolution : Romberg

Sparse Recon. with Partial Support Knowledge

- Modified-CS (our group's work)
- Weighted ℓ_1
- von-Borries et al

Treating Outliers as Sparse Vectors

- Dense Error Correction via ell-1 minimization
- “Robust” PCA
- Recursive “Robust” PCA (our group’s work)

Phase Retrieval

- Recover an n-length signal $\mathbf{x}^*$ from its phaseless (magnitude-only) linear projections

$$\mathbf{y}_i := |\langle \mathbf{a}_i, \mathbf{x}^* \rangle|, \quad i = 1, 2, \ldots, m$$

  when the design vectors $\mathbf{a}_i$ are known.
  - applications: Fourier imaging problems where the phase is impossible or hard to obtain, e.g., sub-diffraction imaging, Fourier ptychography
- Provable solutions: PhaseLift [Candes et al'13], Non-convex iterative algorithms: AltMinPhase [Netrapalli et al,NIPS'13], Wirtinger Flow [Candes et al, T-IT'15], Truncated WF [Chen et al,NIPS'15], Reshaped WF [Zhang et al,NIPS'16])
  - all assume i.i.d. Gaussian $\mathbf{a}_i$'s
- TWF and later work on RWF
  - achieve optimal sample complexity: $m \geq Cn$
  - best time complexity: $mn \log(1/\epsilon)$

Notation

- MATLAB notation: $'$ denotes transpose, $\|.\|$ denotes the l_2 norm
- Phase-invariant distance: $\mathrm{dist}(\mathbf{x}^*, \hat{\mathbf{x}}) := \min_\theta \|\mathbf{x}^* e^{j\theta} - \hat{\mathbf{x}}\|$.
For real-valued signals, this simplifies to
$$\mathrm{dist}(\mathbf{x}^*, \hat{\mathbf{x}}) = \min(\|\mathbf{x}^* - \hat{\mathbf{x}}\|, \|\mathbf{x}^* + \hat{\mathbf{x}}\|).$$
- Subspace Distance: *2-norm of sines of principal angles b/w the subspaces*
$$\mathrm{SubsDist}(\mathbf{U}_1, \mathbf{U}_2) := \|(\mathbf{I} - \mathbf{U}_1 \mathbf{U}_1{}^\top)\mathbf{U}_2\|_F.$$
where $\mathbf{U}_1, \mathbf{U}_2$ are *"basis" matrices (tall matrices with orthonormal columns that span the corresponding subspace)*
- **Any $n \times q$ rank-r matrix $\mathbf{X}$ can be written as $\mathbf{X} = \mathbf{U}\mathbf{B}$**
    - where $\mathbf{U}$ is a tall $n \times r$ basis matrix; $\mathbf{B}$ is $r \times q$
- Guarantees assume $\mathbf{a}_{ik}$ i.i.d. standard Gaussian vectors

Low Rank PR (LRPR): Recover an $n \times q$ rank-r matrix $\mathbf{X}^* = [\mathbf{x}_1^*, \mathbf{x}_2^*, \ldots, \mathbf{x}_q^*]$ from

$$\mathbf{y}_{ik} := |\langle \mathbf{a}_{ik}, \mathbf{x}_k^* \rangle|, \;\; i = 1, \ldots, m, \;\; k = 1, \ldots, q.$$

Key application: fast dynamic Fourier ptychography

Linear LRPR (Compressive PCA): Recover $\mathbf{X}^*$ from

$$\mathbf{z}_{ik} := \langle \mathbf{a}_{ik}, \mathbf{x}_k^* \rangle, \;\; i = 1, \ldots, m, \;\; k = 1, \ldots, q.$$

Key application: fast dynamic MRI [Zhi Pei Lian et a], [Mathews Jacob et al]

Question: When does $m \ll n$ suffice?

- Even linear LRPR has received little attention in theoretical literature so far
  - Our ICML 2019 paper (Phaseless PCA): first useful guarantee for it
  - A convex relaxation approach in NeurIPS 2019 [Srinivasa et al,Neurips,2019]

Related Problems

Our problem (linear version): Recover an $n \times q$ rank-r matrix $\mathbf{X}^* = [\mathbf{x}_1^*, \mathbf{x}_2^*, \ldots, \mathbf{x}_q^*]$ from

$$\mathbf{z}_{ik} := \langle \mathbf{a}_{ik}, \mathbf{x}_k^* \rangle, \;\; i \in [m], k \in [q]$$

- Different from all 3 well-studied low-rank (LR) recovery problems
  - Multivariate Regression: above prob with $\mathbf{a}_{ik} = \mathbf{a}_i$ for all columns k [Neghaban-Wainwright,2011]
    - ★ no independence over $k \Rightarrow$ impossible to recover $\mathbf{X}^*$ with $m < n$.
  - LR matrix sensing: recover $\mathbf{X}^*$ from $\mathbf{y}_i = \langle \mathbf{A}_i, \mathbf{X}^* \rangle$ with $\mathbf{A}_i$'s dense [Netrapalli et al,2013]
    - ★ global measurements (y_i depends on entire $\mathbf{X}^*$): easier problem
  - LR matrix completion: recover $\mathbf{X}^*$ from a subset of its entries [Keshavan et al,2010],[Netrapalli et al,2013]
    - ★ completely local measurements
    - ★ need rows & cols to be dense' to allow for correct "interpolation" : ensured by assuming incoherence of left and right singular vectors
- Our problem
  - non-global measurements of $\mathbf{X}^*$, but global for each column
  - only need denseness of rows (incoherence of right singular vectors)

Right incoherence: incoherence of right singular vectors

Recall $\mathbf{X}^*$ is $n \times q$. Let

$$\mathbf{X}^* \overset{\mathrm{SVD}}{=} \mathbf{U}^* \underbrace{\mathbf{\Sigma}^* \mathbf{B}^*}_{\tilde{\mathbf{B}}^*}$$

be the r-SVD. Here $\mathbf{B}^* = \mathbf{V}^{*\prime}$ from std. SVD notation. Thus rows of $\mathbf{B}^*$ are unit 2-norm

Assume that

$$\max_k \|\mathbf{b}_k^*\|^2 \leq \mu^2 \frac{r}{q}$$

with $\mu \geq 1$ but not too large; assume constant w.r.t. n, q, r.

Right incoherence: incoherence of right singular vectors

Recall $\mathbf{X}^*$ is $n \times q$. Let

$$\mathbf{X}^* \overset{\mathrm{SVD}}{=} \mathbf{U}^* \underbrace{\mathbf{\Sigma}^* \mathbf{B}^*}_{\tilde{\mathbf{B}}^*}$$

be the r-SVD. Here $\mathbf{B}^* = \mathbf{V}^{*\prime}$ from std. SVD notation. Thus rows of $\mathbf{B}^*$ are unit 2-norm

Assume that

$$\max_k \|\mathbf{b}_k^*\|^2 \leq \mu^2 \frac{r}{q}$$

with $\mu \geq 1$ but not too large; assume constant w.r.t. n, q, r.

Implications:

- Unit 2-norm rows and above $\Rightarrow$ rows of $\mathbf{B}^*$ (and hence of $\mathbf{X}^*$) are dense, i.e., no entry too large, most are nonzero
- Also, above implies

$$\max_k \|\mathbf{x}_k^*\|^2 \leq (\kappa^2 \mu^2) \frac{\|\mathbf{X}^*\|_F^2}{q}$$

(assuming small κ, μ, this means that no signal's energy is too much larger than the average energy over all q signals)

Table 1: *The 3 well-studied low-rank matrix recovery problems and ours. All need to recover an $n \times q$ rank-r matrix $\mathbf{X}^* \overset{\text{SVD}}{=} \mathbf{U}^* \boldsymbol{\Sigma}^* \mathbf{V}^{*\top}$ from measurements as specified above.*

Problem	Measurement Model	Assumptions	Global Meas.?	Indep. Meas.?	Ident. Distr. Meas.?	Symmetric (rows, cols)
LRMS	$\mathbf{y}_i = \langle \mathbf{A}_i, \mathbf{X}^* \rangle$, $i \in [m]$ $\mathbf{A}_i$ i.i.d Gauss.		Yes	Yes	Yes	Yes
LRMC	$\mathbf{y}_{jk} = \langle \delta_{jk} \mathbf{e}_j \mathbf{e}_k{}^\top, \mathbf{X}^* \rangle$ $\delta_{jk} \overset{\text{iid}}{\sim} Bern(p)$, $j \in [n]$, $k \in [q]$	$\max_j \|(\mathbf{U}^{*\prime})_j\|^2 \le \mu^2 \frac{r}{n}$, $\max_k \|(\mathbf{V}^{*\prime})_k\|^2 \le \mu^2 \frac{r}{q}$	No	Yes	No	Yes
Multiv. Regr. (MVR)	$\mathbf{y}_{ik} = \langle \mathbf{a}_i \mathbf{e}_k{}^\top, \mathbf{X}^* \rangle = \langle \mathbf{a}_i, \mathbf{x}_k^* \rangle$ $\mathbf{a}_i \overset{\text{iid}}{\sim} \mathcal{N}(0, \mathbf{I}_n)$, $i \in [m]$	-	No	No	No	No
Linear LRPR (our problem)	$\mathbf{y}_{ik} = \langle \mathbf{a}_{ik} \mathbf{e}_k{}^\top, \mathbf{X}^* \rangle = \langle \mathbf{a}_{ik}, \mathbf{x}_k^* \rangle$ $\mathbf{a}_{ik} \overset{\text{iid}}{\sim} \mathcal{N}(0, \mathbf{I}_n)$, $i \in [m], k \in [q]$	$\max_k \|(\mathbf{V}^{*\prime})_k\|^2 \le \mu^2 \frac{r}{q}$	No	Yes	No	No

- LRMS: can prove RIP $\Rightarrow$ simplifies analysis
- MVR: cannot use law of large numbers over k $\Rightarrow$ *m has to be* more than n. Negahban-Wainwright,2011 prove $m \ge nr$ necessary and sufficient
- LRMC and our problem:
  - Most similar: both have non-global, indep. but not ident. dist. meas. $\Rightarrow$ both need incoherence assumps.
  - Difference: LRMC measurements symmetric (across rows, columns) & bounded.

Consider Linear LRPR: recover $\mathbf{X}^* = \mathbf{U}^*\tilde{\mathbf{B}}^*$ from

$$\mathbf{z}_{ik} := \langle \mathbf{a}_{ik}, \mathbf{x}_k^* \rangle = \langle (\mathbf{U}^{*\prime}\mathbf{a}_{ik}), \tilde{\mathbf{b}}_k^* \rangle = \langle (\mathbf{a}_{ik} \otimes \tilde{\mathbf{b}}_k^*), \mathbf{U}_{vec}^* \rangle, \;\; i \in [m], k \in [q]$$

($\mathbf{U}_{vec}^*$ vectorizes the $n \times r$ matrix $\mathbf{U}^*$ into an nr length vector; $\otimes$: Kronecker product)

- Given an estimate $\mathbf{U}$, can update each $\tilde{\mathbf{b}}_k^*$ by r-dimensional LS:

$$\hat{\mathbf{b}}_k = \arg\min_{\hat{\mathbf{b}}} \|(\mathbf{A}_k{}'\mathbf{U})\hat{\mathbf{b}} - \mathbf{z}_k\|_2^2, \;\; k = 1, 2, \dots, q$$

  - $\mathbf{U}'\mathbf{A}_k$ is $m \times r$: this step only needs $m \gtrsim r$; and fast: needs time $mqnr$
- Given estimates $\mathbf{b}_k$, can update $\mathbf{U}_{vec}^*$ by nr-dimensional LS:

$$\hat{\mathbf{U}}_{vec} = \arg\min_{\hat{\mathbf{U}}_{vec}} \left\| \begin{bmatrix} (\mathbf{A}_1 \otimes \mathbf{b}_1) \\ (\mathbf{A}_2 \otimes \mathbf{b}_2) \\ \vdots \\ (\mathbf{A}_q \otimes \mathbf{b}_q) \end{bmatrix}' \hat{\mathbf{U}}_{vec} - \begin{bmatrix} \mathbf{z}_1 \\ \mathbf{z}_2 \\ \vdots \\ \mathbf{z}_q \end{bmatrix} \right\|_2^2$$

  followed by QR decomposition on $\hat{\mathbf{U}}$ to get a basis matrix $\mathbf{U}$

  - slower step: needs time $mqnr \log(1/\epsilon)$

Proposed algorithm: LRPR idea [Nayer et al,Provable LRPR,T-IT,2020]

Consider LRPR: recover $\mathbf{X}^*$ from

$$\mathbf{y}_{ik} := |\mathbf{z}_{ik}|, \;\; \mathbf{z}_{ik} = \langle \mathbf{a}_{ik}, \mathbf{x}_k^* \rangle$$

Modifications needed for LRPR:

- Recovery of $\hat{\mathbf{b}}_k$'s now becomes an r-dimensional standard PR problem
  - outputs $\hat{\mathbf{b}}_k$, $\hat{\mathbf{x}}_k = \mathbf{U}\hat{\mathbf{b}}_k$, measurements' phase $\hat{c}_{ik} = \mathrm{phase}(\langle \mathbf{a}_{ik}, \mathbf{U}\hat{\mathbf{b}}_k \rangle)$
- Before the LS step for updating $\mathbf{U}$, estimate $\mathbf{z}_{ik}$ as
  - $\hat{\mathbf{z}}_{ik} = \hat{c}_{ik}\mathbf{y}_{ik}$

Proposed algorithm: LRPR idea [Nayer et al,Provable LRPR,T-IT,2020]

Consider LRPR: recover $\mathbf{X}^*$ from

$$\mathbf{y}_{ik} := |\mathbf{z}_{ik}|, \ \mathbf{z}_{ik} = \langle \mathbf{a}_{ik}, \mathbf{x}_k^* \rangle$$

Modifications needed for LRPR:

- Recovery of $\hat{\mathbf{b}}_k$'s now becomes an r-dimensional standard PR problem
  - outputs $\hat{\mathbf{b}}_k$, $\hat{\mathbf{x}}_k = \mathbf{U}\hat{\mathbf{b}}_k$, measurements' phase $\hat{c}_{ik} = \mathrm{phase}(\langle \mathbf{a}_{ik}, \mathbf{U}\hat{\mathbf{b}}_k \rangle)$
- Before the LS step for updating $\mathbf{U}$, estimate $\mathbf{z}_{ik}$ as
  - $\hat{\mathbf{z}}_{ik} = \hat{c}_{ik}\mathbf{y}_{ik}$

Key insight to get a sample complexity gain over standard PR:

- conditioned on $\mathbf{X}^*$, we have mq mutually independent measurements
- these are not identically distributed, but right incoherence $\Rightarrow$ similar enough so that concentration bounds are applicable over all the mq summands in the error terms

Proposed algorithm: Initialization for linear case [Nayer-Vaswani, arXiv 2021]

Consider

$$\mathbf{X}_{init,full} = \frac{1}{mq} \sum_{k=1}^{q} \sum_{i=1}^{m} \mathbf{a}_{ik} \mathbf{z}_{ik} \mathbf{e}_k{}'$$

Not too hard to see that

$$\mathbb{E}[\mathbf{X}_{init,full}] = \mathbf{X}^*$$

Thus, if mq large enough, the span of top r left singular vectors of $\mathbf{X}_{init,full}$ should also be a good estimate of $\mathrm{span}(\mathbf{U}^*)$ w.h.p. But

- $\mathbf{a}_{ik}\mathbf{z}_{ik}$'s are sub-exponential with max sub-expo norm $\sigma^*_{\max}\sqrt{r/q}$: not small enough
  - heavy tailed; a few "bad" measurements can bias the average a lot
- Fix: truncate (sum over only the "good" $\mathbf{z}_{ik}$s) [Candes, Chen, NIPS'15 (TWF)] – converts the r.v.s to sub-Gaussian
- Compute $\mathbf{U}_{\mathrm{init}}$ as top r left singular vectors of

$$\mathbf{X}_{init} := \frac{1}{mq} \sum_{i,k:\mathbf{z}_{ik}^2 \le \frac{9}{mq}\sum_{ik}\mathbf{z}_{ik}^2} \mathbf{a}_{ik} \mathbf{z}_{ik} \mathbf{e}_k{}'$$

Previous slide idea does not work because of phaseless measurements.
Consider

$$\mathbf{Y}_{U,\text{full}} = \frac{1}{mq} \sum_{k=1}^{q} \sum_{i=1}^{m} \mathbf{y}_{ik}^2 \mathbf{a}_{ik} \mathbf{a}_{ik}{}'$$

Not too hard to see that its expected value equals

$$\frac{1}{q} [\mathbf{U}^* ({\boldsymbol{\Sigma}^*}^2) \mathbf{U}^{*\prime} + 2\text{trace}({\boldsymbol{\Sigma}^*}^2) \mathbf{I}]$$

Thus, if mq large enough, the span of top r left singular vectors of $\mathbf{Y}_{U,\text{full}}$ should also be a good estimate of $\text{span}(\mathbf{U}^*)$ w.h.p.

- Since $\mathbf{y}_{ik}^2 \mathbf{a}_{ik} \mathbf{a}_{ik}{}'$ is heavy-tailed, a few "bad" (large) measurements can bias the average – need large mq to be robust to this.
- Fix: truncate (throw away "bad" $\mathbf{y}_{ik}$s) [Candes, Chen, NIPS'15 (TWF)]
- Compute $\mathbf{U}_{\text{init}}$ as top r left singular vectors of

$$\mathbf{Y}_U := \frac{1}{mq} \sum_{i,k: \mathbf{y}_{ik}^2 \leq \frac{9}{mq} \sum_{ik} \mathbf{y}_{ik}^2} \mathbf{y}_{ik}^2 \mathbf{a}_{ik} \mathbf{a}_{ik}{}'$$

AltMin for Linear LRPR

$\mathbf{U} \leftarrow$ top r singular vectors of $\mathbf{X}_{init} := \dfrac{1}{mq} \displaystyle\sum_{i,k: \mathbf{z}_{ik}^2 \leq \frac{9}{mq}\sum_{ik}\mathbf{z}_{ik}^2} \mathbf{y}_{ik}\mathbf{a}_{ik}\mathbf{e}_k{}'$

for $t = 0 : T$ **do**

$\quad \hat{\mathbf{b}}_k \leftarrow \mathrm{LS}(\{\mathbf{y}_k, \mathbf{U}'\mathbf{A}_k\})$ for each $k = 1, 2, \cdots, q$

$\quad \hat{\mathbf{X}}^t \leftarrow \mathbf{U}\hat{\mathbf{B}}$ where $\hat{\mathbf{B}} = [\hat{\mathbf{b}}_1, \hat{\mathbf{b}}_2, \ldots \hat{\mathbf{b}}_q]$

$\quad$ QR decomposition: $\hat{\mathbf{B}} \overset{\mathrm{QR}}{=} \mathbf{R}_B\mathbf{B}$

$\quad \hat{\mathbf{U}} \leftarrow \arg\min_{\tilde{\mathbf{U}}} \sum_{k=1}^{q}\sum_{i=1}^{m}(\mathbf{z}_{ik} - \mathbf{a}_{ik}{}'\tilde{\mathbf{U}}\mathbf{b}_k)^2$

$\quad$ QR decomp: $\hat{\mathbf{U}} \overset{\mathrm{QR}}{=} \mathbf{U}\mathbf{R}_U$

end for

AltMin-LowRaP: Alt-Min for Low Rank PR [Nayer et al, Provable LRPR, T-IT,2020]

$\mathbf{U} \leftarrow$ top r singular vectors of $\mathbf{Y}_U := \frac{1}{mq} \sum_{i,k:\mathbf{y}_{ik}^2 \leq \frac{9}{mq}\sum_{ik}\mathbf{y}_{ik}^2} \mathbf{y}_{ik}^2 \mathbf{a}_{ik}\mathbf{a}_{ik}{}'$

for $t = 0 : T$ **do**

  $\hat{\mathbf{b}}_k \leftarrow \mathrm{PR}(\{\mathbf{y}_k, \mathbf{U}'\mathbf{A}_k\})$ for each $k = 1, 2, \cdots, q$

  $\hat{\mathbf{X}}^t \leftarrow \mathbf{U}\hat{\mathbf{B}}$ where $\hat{\mathbf{B}} = [\hat{\mathbf{b}}_1, \hat{\mathbf{b}}_2, \ldots \hat{\mathbf{b}}_q]$

  QR decomposition: $\hat{\mathbf{B}} \overset{\mathrm{QR}}{=} \mathbf{R}_B \mathbf{B}$

  $\hat{\mathbf{c}}_{ik} \leftarrow phase(\langle \mathbf{a}_{ik}, \hat{\mathbf{x}}_{ik} \rangle)$, $i = 1, 2, \ldots, m$, $k = 1, 2, \cdots, q$

  $\hat{\mathbf{U}} \leftarrow \arg\min_{\tilde{\mathbf{U}}} \sum_{k=1}^q \sum_{i=1}^m (\hat{\mathbf{c}}_{ik}\mathbf{y}_{ik} - \mathbf{a}_{ik}{}'\tilde{\mathbf{U}}\mathbf{b}_k)^2$

  QR decomp: $\hat{\mathbf{U}} \overset{\mathrm{QR}}{=} \mathbf{U}\mathbf{R}_U$

end for

PR: use any of the standard PR methods, e.g. Truncated or Reshaped Wirtinger Flow. Estimate r: can be done using $\mathbf{Y}_U$ as well., Recover $\mathbf{X}^*$ ($n \times q$ matrix with rank r) from $\mathbf{z}_{ik} = \langle \mathbf{a}_{ik}, \mathbf{x}_k^* \rangle$, $i \in [1, m]$, $k \in [1, q]$.

Theorem (Guarantee for AltMin for linear LRPR)

Assume μ-incoherence of right singular vectors of $\mathbf{X}^*$. Set $T := C \log(1/\epsilon)$. Assume that, for each new update step, we use a new (independent) set of mq measurements with m satisfying

$$mq \geq C_{\kappa,\mu} \; nr^2$$

and $m \geq C \max(r, \log q, \log n)$. Then, w.p. at least $1 - 10n^{-10}$,

$$\mathrm{SubsDist}(\mathbf{U}^T, \mathbf{U}^*) \leq \epsilon \;\; \textit{and} \;\; \|\hat{\mathbf{X}} - \mathbf{X}^*\|_F \leq \epsilon \|\mathbf{X}^*\|_F,$$

Also, the errors decay exponentially with iteration t.

Time complexity: $mqnr \log^2(1/\epsilon)$. Sample complexity: $C_{\kappa,\mu} \; nr^2 \log(1/\epsilon)$.

Recover $\mathbf{X}^*$ ($n \times q$ matrix with rank r) from $\mathbf{y}_{ik} = |\langle \mathbf{a}_{ik}, \mathbf{x}^*_k \rangle|$, $i \in [1, m]$, $k \in [1, q]$.

Theorem (Guarantee for AltMinLowRaP fpr LRPR)

Assume μ-incoherence of right singular vectors of $\mathbf{X}^*$. Set $T := C \log(1/\epsilon)$. Assume new (independent) measurements as before with m satisfying

$$mq \geq C_{\kappa,\mu} \ nr^3 \text{ (for initialization)}$$
$$mq \geq C_{\kappa,\mu} \ nr^2 \text{ (for AltMin iterations)}$$

and $m \geq C \max(r, \log q, \log n)$. Then, w.p. at least $1 - 10n^{-10}$,

$$\mathrm{SubsDist}(\mathbf{U}^T, \mathbf{U}^*) \leq \epsilon, \ \ dist(\hat{\mathbf{x}}_k^T, \mathbf{x}_k^*) \leq \epsilon \|\mathbf{x}_k^*\|, \ \sum_{k=1}^{q} dist^2(\hat{\mathbf{x}}_k^T, \mathbf{x}_k^*) \leq \epsilon \|\mathbf{X}^*\|_F^2$$

Time complexity: $Cmqnr \log^2(1/\epsilon)$. Sample complexity: $C\kappa^6\mu^2 \ nr^2(r + \log(1/\epsilon))$.

Linear LRPR:

	Sample Comp. $mq \gtrsim C \cdot$	Time Comp.
SL-ECM [Krishnamurthy et al,Asilomar'14]	$n\sqrt{r}\sqrt{q}\frac{1}{\epsilon^2}$ **Too large**	$nq(r+m)$
Convex [Srinivasa et al, NeurIPS'19]	$(n+q)r\frac{1}{\epsilon^2}$ **Large**	$mqnr\frac{1}{\sqrt{\epsilon}}$ **Too slow**
AltMin (our old w.) [Nayer et al, Phaseless PCA] [Nayer et al, Provable LRPR	$nr^4\log(1/\epsilon)$	$mqnr\log^2(1/\epsilon)$
AltMin (this talk) [Nayer-Vaswani, Sample-Eff LRPR]	$nr^2\log(1/\epsilon)$	$mqnr\log^2(1/\epsilon)$

LRPR: no other work

Discussion: Linear LRPR

Problem: Recover a rank-r $n \times q$ matrix $\mathbf{X}^*$ from $\mathbf{z}_{ik} = \langle \mathbf{a}_{ik}, \mathbf{x}^*_k \rangle$, $i \in [m]$, $k \in [q]$.
Solution: AltMin with Spectral Init (non-convex approach)

- Treating κ, μ as constants, sample complexity is

$$m_{\text{tot}} q \geq C \; n r^2 \cdot \log(1/\epsilon)$$

Number of unknowns in $\mathbf{X}^*$ is $(q+n)r \approx 2nr$

- Sub-optimal by a factor of r
- reason: non-global meas. $\Rightarrow$ need to use incoherence to show that the mq scalar meas. are similar enough so that concentration bounds can be applied jointly over all mq summands in the error term
- Non-convex LRMC solutions' complexity is also sub-optimal by a factor of r for the same reason

Discussion: Linear LRPR

Problem: Recover a rank-r $n \times q$ matrix $\mathbf{X}^*$ from $\mathbf{z}_{ik} = \langle \mathbf{a}_{ik}, \mathbf{x}_k^* \rangle$, $i \in [m]$, $k \in [q]$.
Solution: AltMin with Spectral Init (non-convex approach)

- Treating κ, μ as constants, sample complexity is

$$m_{\text{tot}} q \geq C \; nr^2 \cdot \log(1/\epsilon)$$

  Number of unknowns in $\mathbf{X}^*$ is $(q+n)r \approx 2nr$

  - Sub-optimal by a factor of r
  - reason: non-global meas. $\Rightarrow$ need to use incoherence to show that the mq scalar meas. are similar enough so that concentration bounds can be applied jointly over all mq summands in the error term
  - Non-convex LRMC solutions' complexity is also sub-optimal by a factor of r for the same reason

- Related Work
  - Our older work on Provable LRPR (ICML,2019; T'IT,2020):
    - ★ above result improves it by a factor of r^2
  - Neurips 2019 paper (convex relaxation solution):
    - ★ above result improves it by a factor of $\frac{1}{r\epsilon^2}$; also much slower approach

Discussion: LRPR

Recover a rank-r $n \times q$ matrix $\mathbf{X}^*$ from $\mathbf{y}_{ik} = |\langle \mathbf{a}_{ik}, \mathbf{x}_k^* \rangle|$, $i \in [m]$, $k \in [q]$.

- Treating κ, μ as constants, sample complexity is

$$m_{\text{tot}} q \geq C \ nr^2(r + \log(1/\epsilon))$$

  Number of unknowns in $\mathbf{X}^*$ is $(q + n)r \approx 2nr$

  - initialization comp. sub-optimal by a factor of r^2; reason:
    - non-global meas. (same as before) and phaseless meas $\Rightarrow$ cannot define a matrix $\mathbf{X}_{init}$ whose expected value is equal or close to $\mathbf{X}^*$
    - need to define a "squared" matrix $\mathbf{Y}_U$ whose expected value is close to $\mathbf{X}^* \mathbf{X}^{*\prime} + c\mathbf{I}$
    - the same problem for sparse PR as well
  - iterations comp. sub-optimal by a factor of r;
    - well-known fact from PR and sparse PR: once carefully initialized PR problems similar to linear ones.

Discussion: LRPR

Recover a rank-r $n \times q$ matrix $\mathbf{X}^*$ from $\mathbf{y}_{ik} = |\langle \mathbf{a}_{ik}, \mathbf{x}^*_k \rangle|$, $i \in [m]$, $k \in [q]$.

- Treating κ, μ as constants, sample complexity is
$$m_{\mathrm{tot}} q \geq C \; nr^2 (r + \log(1/\epsilon))$$
  Number of unknowns in $\mathbf{X}^*$ is $(q+n)r \approx 2nr$
  - initialization comp. sub-optimal by a factor of r^2; reason:
    - ★ non-global meas. (same as before) and phaseless meas $\Rightarrow$ cannot define a matrix $\mathbf{X}_{init}$ whose expected value is equal or close to $\mathbf{X}^*$
    - ★ need to define a "squared" matrix $\mathbf{Y}_U$ whose expected value is close to $\mathbf{X}^* \mathbf{X}^{*\prime} + c\mathbf{I}$
    - ★ the same problem for sparse PR as well
  - iterations comp. sub-optimal by a factor of r;
    - ★ well-known fact from PR and sparse PR: once carefully initialized PR problems similar to linear ones.
- No existing useful guarantees for our problem. Closest well studied problems:
  - LRMC: also has non-global meas., but LRMC is linear
    - ★ best non-convex LRMC guarantees sub-optimal by a factor of r
  - sparse PR: also has phaseless meas., but meas. are global
    - ★ best sparse PR guarantees sub-optimal by a factor of s (sparsity size).

Discussion: LRPR

Recover a rank-r $n \times q$ matrix $\mathbf{X}^*$ from $\mathbf{y}_{ik} = |\langle \mathbf{a}_{ik}, \mathbf{x}^*_k \rangle|$, $i \in [m]$, $k \in [q]$.

- Treating κ, μ as constants, sample complexity is

$$m_{\mathrm{tot}} q \geq C \; nr^2(r + \log(1/\epsilon))$$

  Number of unknowns in $\mathbf{X}^*$ is $(q+n)r \approx 2nr$
  - initialization comp. sub-optimal by a factor of r^2; reason:
    - ★ non-global meas. (same as before) and phaseless meas $\Rightarrow$ cannot define a matrix $\mathbf{X}_{init}$ whose expected value is equal or close to $\mathbf{X}^*$
    - ★ need to define a "squared" matrix $\mathbf{Y}_U$ whose expected value is close to $\mathbf{X}^*\mathbf{X}^{*\prime} + c\mathbf{I}$
    - ★ the same problem for sparse PR as well
  - iterations comp. sub-optimal by a factor of r;
    - ★ well-known fact from PR and sparse PR: once carefully initialized PR problems similar to linear ones.
- No existing useful guarantees for our problem. Closest well studied problems:
  - LRMC: also has non-global meas., but LRMC is linear
    - ★ best non-convex LRMC guarantees sub-optimal by a factor of r
  - sparse PR: also has phaseless meas., but meas. are global
    - ★ best sparse PR guarantees sub-optimal by a factor of s (sparsity size).
- Comparison with standard (unstructured) PR
  - Standard PR sample complexity is nq: much larger when $r^2 \ll q$

The theorem is an immediate consequence of the following two claims.

Lemma (Initialization)

Pick a $\delta_0 < 0.25$. *Then, w.p. at least* $1 - 2\exp\left(n - c\frac{\delta_0^2 mq}{\kappa^4 r^3}\right) - 2\exp\left(-c\frac{\delta_0^3 mq}{\kappa^4 \mu^2 r^3}\right)$,

$$\mathrm{SubsDist}(\mathbf{U}^*, \mathbf{U}^0) \leq \delta_0$$

Proof: See Provable Low Rank Phase Retrieval, T-IT 20

Lemma (AltMin Iterations – exponential decay of error)

Assume that $\mathrm{SubsDist}(\mathbf{U}^*, \mathbf{U}^0) \leq \delta_0 = c/\kappa^2$ *and* $\mathrm{SubsDist}(\mathbf{U}^*, \mathbf{U}^t) \leq \delta_t := 0.2^t \delta_0$ *Then, w.p. at least* $1 - \left(\exp\left(nr - c\frac{mq}{\kappa^2\mu^4 r}\right) - \exp(\log q + r - cm)\right)$,

$$\mathrm{SubsDist}(\mathbf{U}^*, \mathbf{U}^{t+1}) \leq \delta_{t+1} := 0.2\delta_t \tag{1}$$

Proof: See next few slides

Proof of the second AltMin iterations lemma I

Let $\mathbf{U} \equiv \mathbf{U}^t$, $\hat{\mathbf{B}} \equiv \hat{\mathbf{B}}^t$, $B \equiv \mathbf{B}^t$.

Lemma (Algebra lemma)

$$\mathrm{SubsDist}(\mathbf{U}^*, \mathbf{U}^{t+1}) \le \frac{\mathrm{MainTerm}}{\sigma_{\min}(\mathbf{X}^* \mathbf{B}^\top) - \mathrm{MainTerm}} \quad \textit{where}$$

$$\mathrm{MainTerm} := \frac{\max_{\mathbf{W} \in \mathcal{S}_W} |\mathrm{Term1}(\mathbf{W})| + \max_{\mathbf{W} \in \mathcal{S}_W} |\mathrm{Term2}(\mathbf{W})|}{\min_{\mathbf{W} \in \mathcal{S}_W} \mathrm{Term3}(\mathbf{W})},$$

$$\mathrm{Term1}(\mathbf{W}) := \sum_{ik} \mathbf{b}_k{}^\top \mathbf{W}^\top \mathbf{a}_{ik} \mathbf{a}_{ik}{}^\top (\mathbf{X}^* \mathbf{B}^\top \mathbf{b}_k - \mathbf{x}_k^*)$$

$$\mathrm{Term2}(\mathbf{W}) := \sum_{ik} (\bar{\mathbf{c}}_{ik}^* \hat{\mathbf{c}}_{ik} - 1)(\mathbf{x}_k^*{}^\top \mathbf{a}_{ik})(\mathbf{a}_{ik}{}^\top \mathbf{W} \mathbf{b}_k),$$

$$\mathrm{Term3}(\mathbf{W}) := \sum_{ik} |\mathbf{a}_{ik}{}^\top \mathbf{W} \mathbf{b}_k|^2,$$

$$\mathcal{S}_W := \{\mathbf{W} \in \Re^{n \times r} : \ \|\mathbf{W}\|_F = 1\},$$

and $\mathbf{c}_{ik}^*, \hat{\mathbf{c}}_{ik}$ *are the phases of* $\mathbf{a}_{ik}{}^\top \mathbf{x}_k^*$ *and* $\mathbf{a}_{ik}{}^\top \hat{\mathbf{x}}_k$*.*

Lemma (High probability bounds on the $\mathrm{SubsDist}$ bound terms)

Assume that $\mathrm{SubsDist}(\mathbf{U}^*, \mathbf{U}^t) \leq \delta_t$ *with* $\delta_t < \delta_0 = c/\kappa^2$. *Then, it is possible to show that*

1. *w.p. at least* $1 - 2\exp\left(nr(\log 17) - c\frac{mq\epsilon_1^2}{\kappa^2\mu^2 r}\right) - \exp(\log q + r - cm)$,

$$\max_{\mathbf{W} \in \mathcal{S}_W} |\mathrm{Term1}(\mathbf{W})| \leq Cm\epsilon_1 \delta_t \sigma^*_{\max},$$

2. *w.p. at least* $1 - 2\exp\left(nr\log(17) - c\frac{mq\epsilon_2^2}{\mu^2\kappa r}\right) - \exp(\log q + r - cm)$,

$$\max_{\mathbf{W} \in \mathcal{S}_W} |\mathrm{Term2}(\mathbf{W})| \leq Cm(\epsilon_2 + \sqrt{\delta_t})\delta_t \sigma^*_{\max},$$

3. *w.p. at least* $1 - 2\exp\left(nr(\log 17) - c\frac{\epsilon_3^2 mq}{\mu^2\kappa^2 r}\right) - \exp\left(\log q + r - cm\right)$,

$$\min_{\mathbf{W} \in \mathcal{S}_W} \mathrm{Term3}(\mathbf{W}) \geq 0.5(1-\epsilon_3)m,$$

4. $\sigma_{\min}(\mathbf{X}^*\mathbf{B}^\top) \geq \sigma^*_{\min}$.

The proof uses

1. "$\hat{\mathbf{B}}$ lemma" given next (bounds $\|\mathbf{g}_k - \hat{\mathbf{b}}_k\|$ and $\|\mathbf{G} - \hat{\mathbf{B}}\|_F$ with $\mathbf{g}_k = \mathbf{U}'\mathbf{x}_k^*$),
2. right singular vectors' incoherence: $\max_k \|\mathbf{b}_k^*\| \leq \mu\sqrt{r/q}$, and
3. sub-exponential Bernstein inequality

Assume $\mathrm{SubsDist}(\mathbf{U}^*, \mathbf{U}) \le \delta_t$. Define a rotated version of $\tilde{\mathbf{b}}_k^*$,

$$\mathbf{g}_k = \mathbf{U}'\mathbf{x}_k^* = \mathbf{U}'\mathbf{U}^*\tilde{\mathbf{b}}_k^*$$

- Our estimate $\hat{\mathbf{b}}_k$ can be shown to be an approx of $\mathbf{g}_k$: $\|\mathbf{g}_k - \hat{\mathbf{b}}_k\| \lesssim \|(\mathbf{I} - \mathbf{U}\mathbf{U}^\top)\mathbf{U}^*\tilde{\mathbf{b}}_k^*\|$
  - proof idea: Given an estimate $\mathbf{U}$, recovery of each $\tilde{\mathbf{b}}_k^*$ is an r-dimensional noisy PR problem with noise proportional to $\|(\mathbf{I} - \mathbf{U}\mathbf{U}^\top)\mathbf{U}^*\tilde{\mathbf{b}}_k^*\|$; using a noisy PR result, can show
    $$\|\mathbf{g}_k - \hat{\mathbf{b}}_k\|^2 \le C\|(\mathbf{I} - \mathbf{U}\mathbf{U}^\top)\mathbf{U}^*\tilde{\mathbf{b}}_k^*\|^2$$

Implications

- $\|\mathbf{g}_k - \hat{\mathbf{b}}_k\| \le C\delta_t\|\mathbf{x}_k^*\|$
- $\|\hat{\mathbf{B}} - \mathbf{G}\|_F \lesssim \delta_t \sigma_{\max}^*$
  - proof idea: use $\sum_k \|\mathbf{M}\tilde{\mathbf{b}}_k^*\|^2 = \|\mathbf{M}\tilde{\mathbf{B}}^*\|_F^2 \le \|\mathbf{M}\|_F^2\|\tilde{\mathbf{B}}^*\|^2 = \|\mathbf{M}\|_F^2 {\sigma_{\max}^*}^2$
- Also implies $\hat{\mathbf{x}}_k = \mathbf{U}\hat{\mathbf{b}}_k$ is an approx of $\mathbf{x}_k^* = \mathbf{U}\mathbf{g}_k + (\mathbf{I} - \mathbf{U}\mathbf{U}^\top)\mathbf{U}^*\tilde{\mathbf{b}}_k^*$, can show
- This implies $\|\hat{\mathbf{X}} - \mathbf{X}^*\|_F \lesssim \delta_t \sigma_{\max}^*$
- Above claims and right incoherence imply incoherence of $\mathbf{B}$: $\|\mathbf{b}_k\| \le \kappa\mu\sqrt{r/q}$.

mat-dist$(\mathbf{G}, \hat{\mathbf{B}}) := \sqrt{\sum_{k=1}^{q} \operatorname{dist}^2(\mathbf{g}_k, \hat{\mathbf{b}}_k)}$

Lemma ($\hat{\mathbf{B}}$ Lemma)

Let $\mathbf{g}_k := \mathbf{U}'\mathbf{U}^* \tilde{\mathbf{b}}_k^*$. *Assume that* $\operatorname{SubsDist}(\mathbf{U}^t, \mathbf{U}^*) \le \delta_t$ *with* $\delta_t < \delta_0 = c/\kappa^2$. *Then, w.p. at least* $1 - \exp(\log q + r - cm)$,

$$\operatorname{dist}\left(\mathbf{g}_k, \hat{\mathbf{b}}_k\right) \le \operatorname{dist}\left(\mathbf{x}_k^*, \hat{\mathbf{x}}_k\right) \le C\delta_t \|\tilde{\mathbf{b}}_k^*\|$$

$$\textit{mat-dist}(\mathbf{G}, \hat{\mathbf{B}}) \le \textit{mat-dist}(\mathbf{X}^*, \hat{\mathbf{X}}) \le C\delta_t \sigma_{\max}^*$$

$$\|\mathbf{b}_k\| \le \frac{\operatorname{dist}(\hat{\mathbf{b}}_k, \mathbf{g}_k) + \|\mathbf{g}_k\|}{0.95\sigma_{\min}^* - \textit{mat-dist}(\mathbf{G}, \hat{\mathbf{B}})} \le 2\kappa\mu\sqrt{r/q}.$$

Bounding Term1

Let $\mathbf{p}_k := (\mathbf{X}^* \mathbf{B}^\top \mathbf{b}_k - \mathbf{x}_k^*)$, then

$$Term1(\mathbf{W}) = \sum_{ik} \mathbf{b}_k{}^\top \mathbf{W}^\top \mathbf{a}_{ik} \mathbf{a}_{ik}{}^\top \mathbf{p}_k = trace(\sum_{ik} \mathbf{p}_k \mathbf{b}_k{}^\top \mathbf{W}^\top \mathbf{a}_{ik} \mathbf{a}_{ik}{}^\top)$$

Need to bound $\max_{\mathbf{W} \in \mathcal{S}_W} |Term1(\mathbf{W})|$ where $\mathcal{S}_W = \{\mathbf{W} : \|\mathbf{W}\|_F = 1\}$. Do this as follows

1. $\mathbf{a}_{ik} \sim \mathcal{N}(0, \mathbf{I})$ and $\mathbf{B}\mathbf{B}^\top = \mathbf{I} \Rightarrow \sum_k \mathbf{p}_k \mathbf{b}_k{}^\top = 0 \Rightarrow \mathbb{E}[\mathrm{Term1}] = 0$.
2. Use sub-exponential Bernstein and following ideas to bound $|Term1(\mathbf{W})|$ for a fixed $\mathbf{W}$
   - $\hat{\mathbf{X}} \mathbf{B}^\top \mathbf{b}_k = \hat{\mathbf{x}}_k$ and the previous lemma ($\hat{\mathbf{B}}$ bound) $\Rightarrow$

     $$\|\mathbf{p}_k\| \le \|\mathbf{X}^* - \hat{\mathbf{X}}\|_F \|\mathbf{b}_k\| + \|\hat{\mathbf{x}}_k - \mathbf{x}_k^*\| \lesssim \delta_t \sigma_{\max}^* \kappa \mu \sqrt{r/q}$$

   - $\hat{\mathbf{X}}(\mathbf{B}^\top \mathbf{B} - \mathbf{I}) = \mathbf{0} \Rightarrow$

     $$\|[\mathbf{p}_1, \mathbf{p}_2, \dots \mathbf{p}_q]\|_F = \|\mathbf{X}^*(\mathbf{B}^\top \mathbf{B} - \mathbf{I})\|_F \lesssim \delta_t \sigma_{\max}^*.$$

3. Define an epsilon-net, $\bar{\mathcal{S}}_W$ on $\mathcal{S}_W$ with say $\epsilon = 1/8$. Can show (see Vershynin) that the number of discrete points on $\bar{\mathcal{S}}_W$, is at most $(1 + 2/\epsilon)^{nr} = 17^{nr}$
4. Use union bound to extend the $Term1(\mathbf{W})$ bound from *a* fixed $\mathbf{W} \in \bar{\mathcal{S}}_W$ to *all* $\mathbf{W} \in \bar{\mathcal{S}}_W$.
5. Develop an "epsilon-net" argument to extend from $\bar{\mathcal{S}}_W$ to $\mathcal{S}_W$.
   - Can show that if $\max_{\mathbf{W} \in \bar{\mathcal{S}}_W} Term1(\mathbf{W}) \le d$, then $\max_{\mathbf{W} \in \mathcal{S}_W} Term1(\mathbf{W}) \le 1.5d$

Term3 bound:

1. $\mathbf{B}\mathbf{B}^\top = \mathbf{I}$ and $\|\mathbf{W}\|_F = 1 \Rightarrow \mathbb{E}[Term3] = m\sum_k \|\mathbf{W}\mathbf{b}_k\|^2 = m$
2. right incoherence and sub-exponential Bernstein inequality
3. epsilon-net argument as above: a little more complicated for $\min_{\mathbf{W}\in\mathcal{S}_W}$.

Term2 bound:

1. $\sum_{ik} \mathbb{E}[\mathbb{1}_{\mathbf{c}^*_{ik}\neq\hat{\mathbf{c}}_{ik}} (\mathbf{a}_{ik}{}^\top \mathbf{x}^*_k)^2] \lesssim \sum_{ik} \frac{\mathrm{dist}(\mathbf{x}^*_k, \hat{\mathbf{x}}_k)^3}{\|\mathbf{x}^*_k\|} \lesssim \delta_t^3 {\sigma^*_{\max}}^2$