

Accelerated Outlier Detection in Low-Rank and Structured Data: Robust PCA and Extensions

HanQin Cai

Department of Mathematics
University of California at Los Angeles

Seminar on Applied Math and Data Science
March 25th, 2020

Outline

- 1 Introduction
- 2 Existing Approaches with Theoretical Guarantee
- 3 Proposed Approach
- 4 Numerical Experiments
- 5 Extension to Low-Rank Hankel Recovery
- 6 Conclusion and Future Work

Principal Component Analysis (PCA)

- Principal Component Analysis (PCA), a.k.a. the best low rank approximation:

$$\begin{aligned} & \underset{\mathbf{L}'}{\text{minimize}} \|\mathbf{D} - \mathbf{L}'\|_F \\ & \text{subject to } \text{rank}(\mathbf{L}') \leq r \end{aligned}$$

– Fundamental tool for dimension reduction



- Can be efficiently solved via truncated singular-value decomposition (SVD) \mathcal{H}_r :

$$\mathcal{H}_r(\mathbf{D}) = \mathcal{H}_r(\mathbf{U}\mathbf{\Sigma}\mathbf{V}^T) = \mathbf{U}\mathbf{\Sigma}_r\mathbf{V}^T$$

where $\mathbf{\Sigma}_r$ keeps only the first r singular values.


Principal Component Analysis (PCA)

- PCA, a.k.a. the best low rank approximation.
- Solved efficiently via truncated SVD \mathcal{H}_r .
- Tolerates small (zero-mean) noise, but over-sensitive to outliers.



Robust PCA

- Given $D = L + S \in \mathbb{R}^{m \times n}$, where L is low rank and S is sparse.






- Given $\mathbf{D} = \mathbf{L} + \mathbf{S} \in \mathbb{R}^{m \times n}$, where \mathbf{L} is low rank and \mathbf{S} is sparse.
- Split/recover \mathbf{L} , \mathbf{S} from \mathbf{D} :

$$\begin{aligned} & \underset{\mathbf{L}', \mathbf{S}'}{\text{minimize}} \|\mathbf{D} - \mathbf{L}' - \mathbf{S}'\|_F \\ & \text{subject to } \text{rank}(\mathbf{L}') \leq r, \|\mathbf{S}'\|_0 \leq \alpha mn \end{aligned} \tag{1}$$

where α is the sparsity level of \mathbf{S} , compared to its size.


- Non-convex problem
- Tolerates outliers better.

- Video background subtraction with static background



Applications

- Video background subtraction with static background
- Face recognition



Applications

- Video background subtraction with static background
- Face recognition
- System identification
- Fault isolation
- Netflix problem
- and more ...

RPCA is Ill-Posed without Assumptions

Cannot have a unique decomposition if \mathbf{D} is both low rank and sparse.

$$\begin{aligned} \begin{bmatrix} 0 & 0 & 0 \\ 0 & 5 & 0 \\ 0 & 0 & 0 \end{bmatrix} &= \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 \\ 0 & 4 & 0 \\ 0 & 0 & 0 \end{bmatrix} \\ &= \begin{bmatrix} 0 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 0 \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 0 \end{bmatrix} \\ &= \dots \end{aligned}$$

Assumption A1 for Uniqueness of Solution

- $\mathbf{L} \in \mathbb{R}^{m \times n}$ is rank- r with μ -incoherence, i.e.,

$$\max_i \|\mathbf{e}_i^T \mathbf{U}\|_2 \leq \sqrt{\frac{\mu r}{m}}$$

$$\max_j \|\mathbf{e}_j^T \mathbf{V}\|_2 \leq \sqrt{\frac{\mu r}{n}}$$

where $\mathbf{L} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$ is the SVD of \mathbf{L} .

- $\mu \in [1, \max\{m, n\}/r]$
- Smaller μ is better, which implies the entries of \mathbf{L} are not too spiky/sparse. - Energy is evenly distributed among the entries of \mathbf{L} .

Assumption A2 for Uniqueness of Solution

- $\mathbf{S} \in \mathbb{R}^{m \times n}$ is α -sparse, i.e., \mathbf{S} has at most αn non-zero entries in each row, and at most αm non-zero entries in each column.
- Non-zero entries are not locally dense.
- Satisfied with high probability if the support of \mathbf{S} is drawn by some stochastic methods.
 - Bernoulli process
 - Uniform sampling

Outline

- 1 Introduction
- 2 Existing Approaches with Theoretical Guarantee**
 - Convex Methods
 - Non-Convex Methods
- 3 Proposed Approach
- 4 Numerical Experiments
- 5 Extension to Low-Rank Hankel Recovery
- 6 Conclusion and Future Work

Outline

- 1 Introduction
- 2 Existing Approaches with Theoretical Guarantee
 - Convex Methods
 - Non-Convex Methods
- 3 Proposed Approach
- 4 Numerical Experiments
- 5 Extension to Low-Rank Hankel Recovery
- 6 Conclusion and Future Work

- Consider convex Principal Component Pursuit (PCP):

$$\begin{aligned} & \underset{\mathbf{L}', \mathbf{S}'}{\text{minimize}} \|\mathbf{L}'\|_* + \lambda \|\mathbf{S}'\|_1 \\ & \text{subject to } \mathbf{L}' + \mathbf{S}' = \mathbf{D} \end{aligned} \tag{2}$$

- Nuclear norm term for low rank
 - $\|\mathbf{L}\|_* = \sum_{i=1}^r \sigma_i(\mathbf{L})$
- Vector ℓ_1 -norm term for sparsity, where \mathbf{S} is viewed as a long vector
 - $\|\mathbf{S}\|_1 = \sum_{ij} |S_{ij}|$

Recovery Guarantee of PCP

If $\lambda = \sqrt{1/\max\{m, n\}}$ is chosen, then, with high probability, the solution of convex PCP

$$\begin{aligned} & \underset{\mathbf{L}', \mathbf{S}'}{\text{minimize}} \|\mathbf{L}'\|_* + \lambda \|\mathbf{S}'\|_1 \\ & \text{subject to } \mathbf{L}' + \mathbf{S}' = \mathbf{D} \end{aligned}$$

is exact the solution of original non-convex Robust PCA problem

$$\begin{aligned} & \underset{\mathbf{L}', \mathbf{S}'}{\text{minimize}} \|\mathbf{D} - \mathbf{L}' - \mathbf{S}'\|_F \\ & \text{subject to } \text{rank}(\mathbf{L}') \leq r, \|\mathbf{S}'\|_0 \leq \alpha mn \end{aligned}$$

under some natural conditions. [Candès, Li, Ma, Wright, 2009]

- Off-the-shelf solver
 - CVX package for Matlab
 - Based on semidefinite programming
 - $O(n^6)$ - expensive
 - Cannot handle problem larger than 200×200 dimension

- Off-the-shelf solver
 - CVX package for Matlab
 - Based on semidefinite programming
 - $O(n^6)$ - expensive
 - Cannot handle problem larger than 200×200 dimension
- Alternating direction method of multipliers (ADMM)
 - Consider the augmented Lagrangian

$$\ell(\mathbf{L}, \mathbf{S}, \mathbf{Y}) = \|\mathbf{L}\|_* + \lambda \|\mathbf{S}\|_1 + \langle \mathbf{Y}, \mathbf{D} - \mathbf{L} - \mathbf{S} \rangle + \frac{\eta}{2} \|\mathbf{D} - \mathbf{L} - \mathbf{S}\|_F^2$$

- Updating \mathbf{L} , \mathbf{S} and dual variable \mathbf{Y} alternatively

- Off-the-shelf solver
 - CVX package for Matlab
 - Based on semidefinite programming
 - $O(n^6)$ - expensive
 - Cannot handle problem larger than 200×200 dimension
- Alternating direction method of multipliers (ADMM)
 - Consider the augmented Lagrangian

$$\ell(\mathbf{L}, \mathbf{S}, \mathbf{Y}) = \|\mathbf{L}\|_* + \lambda \|\mathbf{S}\|_1 + \langle \mathbf{Y}, \mathbf{D} - \mathbf{L} - \mathbf{S} \rangle + \frac{\eta}{2} \|\mathbf{D} - \mathbf{L} - \mathbf{S}\|_F^2$$

- Updating \mathbf{L} , \mathbf{S} and dual variable \mathbf{Y} alternatively
- Convergence of ADMM has been well studied
- Good per-iteration complexity, but sublinear convergence rate

Outline

- 1 Introduction
- 2 Existing Approaches with Theoretical Guarantee
 - Convex Methods
 - Non-Convex Methods
 - Gradient Descent
 - Alternating Projections
- 3 Proposed Approach
- 4 Numerical Experiments
- 5 Extension to Low-Rank Hankel Recovery
- 6 Conclusion and Future Work

Gradient Descent (GD)

- Updating L :

- Consider $L = U\Sigma V^T = (U\Sigma^{1/2})(\Sigma^{1/2}V^T) := PQ^T$
- Gradient descent on P , Q separately, based on objective function:

$$\underbrace{\frac{1}{2}\|PQ^T + S - D\|_F^2}_{\text{loss function}} + \underbrace{\frac{1}{8}\|P^T P - Q^T Q\|_F^2}_{\text{keep scale being close}} \quad (3)$$

- Enforce incoherence on P , Q after gradient descent

- Updating S :

- Sparsification operator \mathcal{F}_α
 - Keeps only the largest α -fraction elements per row and column
 - Partial sorting on each row and column - expensive with larger α

[Yi, Park, Chen, Caramanis, 2016]

Recovery Guarantee of GD

- The output $(\mathbf{P}_k, \mathbf{Q}_k)$ of GD satisfies

$$\|\mathbf{P}_k \mathbf{Q}_k^T - \mathbf{L}\|_F \leq \varepsilon$$

in $O(\kappa \log(1/\varepsilon))$ iterations if followings are satisfied:

- \mathbf{L} is μ -incoherent
 - \mathbf{S} is α -sparse with $\alpha \leq O\left(\min\left\{\frac{1}{\mu r^{1.5} \kappa^{1.5}}, \frac{1}{\mu r \kappa^2}\right\}\right)$
 - $\gamma = 2$: parameter for sparsification
 - $\eta \leq 1/36\sigma_1^L$: step size for gradient descent
- Linear convergence with rate of $1 - O(1/\kappa)$.

κ : condition number of \mathbf{L}

Alternating Projections (AltProj)

Consider two sets:

- $\mathcal{M}_r = \{\mathbf{L} \in \mathbb{R}^{m \times n} \mid \text{rank}(\mathbf{L}) \leq r\}$
- $\mathcal{S}_\alpha = \{\mathbf{S} \in \mathbb{R}^{m \times n} \mid \mathbf{S} \text{ is } \alpha\text{-sparse}\}$

Alternating Projections (AltProj)

Consider two sets:

- $\mathcal{M}_r = \{\mathbf{L} \in \mathbb{R}^{m \times n} \mid \text{rank}(\mathbf{L}) \leq r\}$
- $\mathcal{S}_\alpha = \{\mathbf{S} \in \mathbb{R}^{m \times n} \mid \mathbf{S} \text{ is } \alpha\text{-sparse}\}$
- Find $\{(\mathbf{L}, \mathbf{S}) \mid \mathbf{L} \in \mathcal{M}_r, \mathbf{S} \in \mathcal{S}_\alpha \text{ and } \mathbf{L} + \mathbf{S} = \mathbf{D}\}$

Alternating Projections (AltProj)

Consider two sets:


- $\mathcal{M}_r = \{\mathbf{L} \in \mathbb{R}^{m \times n} \mid \text{rank}(\mathbf{L}) \leq r\}$
- $\mathcal{S}_\alpha = \{\mathbf{S} \in \mathbb{R}^{m \times n} \mid \mathbf{S} \text{ is } \alpha\text{-sparse}\}$
- Find $\{(\mathbf{L}, \mathbf{S}) \mid \mathbf{L} \in \mathcal{M}_r, \mathbf{S} \in \mathcal{S}_\alpha \text{ and } \mathbf{L} + \mathbf{S} = \mathbf{D}\}$

Updating:

- $\mathbf{L}_{new} = \mathcal{H}_r(\mathbf{D} - \mathbf{S}_{old})$.
 - \mathcal{H}_r : truncated SVD - expensive
- $\mathbf{S}_{new} = \mathcal{T}_\zeta(\mathbf{D} - \mathbf{L}_{new})$.
 - $\mathcal{T}_\zeta(x) = x \cdot \mathbb{1}_{\{|x| > \zeta\}}$: hard-thresholding
 - Choose ζ such that $\text{supp}(\mathbf{S}_{new}) \subset \text{supp}(\mathbf{S})$

[Netrapalli, Niranjan, Sanghavi, Anandkumar, Jain, 2014]

Illustration of Alternating Projections



Alternating Projections (AltProj)

Consider two sets:

- $\mathcal{M}_r = \{\mathbf{L} \in \mathbb{R}^{m \times n} \mid \text{rank}(\mathbf{L}) \leq r\}$
- $\mathcal{S}_\alpha = \{\mathbf{S} \in \mathbb{R}^{m \times n} \mid \mathbf{S} \text{ is } \alpha\text{-sparse}\}$
- Find $\{(\mathbf{L}, \mathbf{S}) \mid \mathbf{L} \in \mathcal{M}_r, \mathbf{S} \in \mathcal{S}_\alpha \text{ and } \mathbf{L} + \mathbf{S} = \mathbf{D}\}$

Updating:

- $\mathbf{L}_{new} = \mathcal{H}_r(\mathbf{D} - \mathbf{S}_{old}); \quad \mathbf{S}_{new} = \mathcal{T}_\zeta(\mathbf{D} - \mathbf{L}_{new}).$

Break algorithm into r stages:

- At the t^{th} stage, use \mathcal{H}_t instead of \mathcal{H}_r
- Overcome the case of bad condition number of \mathbf{L}

[Netrapalli, Niranjan, Sanghavi, Anandkumar, Jain, 2014]

Recovery Guarantee of AltProj

- The output $(\mathbf{L}_{t,k}, \mathbf{S}_{t,k})$ of AltProj satisfies

$$\|\mathbf{L}_{t,k} - \mathbf{L}\|_F \leq \varepsilon, \quad \|\mathbf{S}_{t,k} - \mathbf{S}\|_\infty \leq \varepsilon/\sqrt{mn},$$

and $\text{supp}(\mathbf{S}_{t,k}) \subset \text{supp}(\mathbf{S})$

in $O(t \log_{\frac{1}{2}} \varepsilon)$ iterations if followings are satisfied:

- \mathbf{L} is μ -incoherent
 - \mathbf{S} is α -sparse with $\alpha \leq \frac{1}{512\mu r}$
 - $\beta = 4\mu r/\sqrt{mn}$: parameter for thresholding
- Linear convergence with rate of $1/2$, at t^{th} stage.

Outline

- 1 Introduction
- 2 Existing Approaches with Theoretical Guarantee
- 3 Proposed Approach**
 - Algorithm Design
 - Theoretical Guarantee
- 4 Numerical Experiments
- 5 Extension to Low-Rank Hankel Recovery
- 6 Conclusion and Future Work

Outline

- 1 Introduction
- 2 Existing Approaches with Theoretical Guarantee
- 3 Proposed Approach**
 - **Algorithm Design**
 - Theoretical Guarantee
- 4 Numerical Experiments
- 5 Extension to Low-Rank Hankel Recovery
- 6 Conclusion and Future Work

- AltProj is fast enough when updating \mathbf{S} with hard thresholding, but using truncated SVD for \mathbf{L} updating can still be very expensive when the problem size is larger.
- Accelerating by first projecting $\mathbf{D} - \mathbf{S}_{old}$ onto some local low dimensional subspace before obtaining a new estimate of \mathbf{L} via truncated SVD.

The Low Dimensional Subspace

- $\mathcal{M}_r = \{\mathbf{L} \in \mathbb{R}^{m \times n} \mid \text{rank}(\mathbf{L}) \leq r\}$ is a Riemannian manifold.
- Tangent space of \mathcal{M}_r at \mathbf{L}


$$T = \{\mathbf{U}\mathbf{A}^T + \mathbf{B}\mathbf{V}^T \mid \mathbf{A} \in \mathbb{R}^{n \times r}, \mathbf{B} \in \mathbb{R}^{m \times r}\}$$

where $\mathbf{L} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$ is the SVD of \mathbf{L} .

- The low dimensional subspace we want, but don't have in practice.
- Instead, use tangent space of \mathcal{M}_r at \mathbf{L}_{old} .
- Project an arbitrary matrix onto the tangent space T :

$$\mathcal{P}_T \mathbf{Z} = \mathbf{U}\mathbf{U}^T \mathbf{Z} + \mathbf{Z}\mathbf{V}\mathbf{V}^T - \mathbf{U}\mathbf{U}^T \mathbf{Z}\mathbf{V}\mathbf{V}^T \quad (4)$$

Illustration of Our Approach



RPCA by Accelerated Alternating Projections (AccAltProj)

Initialization

$k = 0$

while $\|D - L_k - S_k\|_F / \|D\|_F \geq \varepsilon$ **do**

$\tilde{L}_k = \text{Trim}(L_k, \mu)$ // Enforce Incoherence of Tangent Space

$L_{k+1} = \mathcal{H}_r(\mathcal{P}_{\tilde{T}_k}(D - S_k))$

$\zeta_{k+1} = \beta(\sigma_{r+1}(\mathcal{P}_{\tilde{T}_k}(D - S_k)) + \gamma^{k+1}\sigma_1(\mathcal{P}_{\tilde{T}_k}(D - S_k)))$

$S_{k+1} = \mathcal{T}_{\zeta_{k+1}}(D - L_{k+1})$

$k = k + 1$

end while

Accelerating Truncated SVD

Denote $\mathbf{Z} := \mathbf{D} - \mathbf{S}_k$. Computing $\mathcal{H}_r(\mathcal{P}_{\tilde{T}_k} \mathbf{Z})$ efficiently:

- $$\mathcal{P}_{\tilde{T}_k} \mathbf{Z} = \underbrace{\begin{bmatrix} \tilde{\mathbf{U}}_k & \mathbf{Q}_1 \end{bmatrix}}_{\text{orthogonal}} \underbrace{\begin{bmatrix} \tilde{\mathbf{U}}_k^T \mathbf{Z} \tilde{\mathbf{V}}_k & \mathbf{R}_2^T \\ \mathbf{R}_1 & 0 \end{bmatrix}}_{2r \times 2r} \underbrace{\begin{bmatrix} \tilde{\mathbf{V}}_k & \mathbf{Q}_2 \end{bmatrix}^T}_{\text{orthogonal}}$$
 - $\mathbf{Q}_1 \mathbf{R}_1 = (\mathbf{I} - \tilde{\mathbf{U}}_k \tilde{\mathbf{U}}_k^T) \mathbf{Z} \tilde{\mathbf{V}}_k \in \mathbb{R}^{m \times r}$
 - $\mathbf{Q}_2 \mathbf{R}_2 = (\mathbf{I} - \tilde{\mathbf{V}}_k \tilde{\mathbf{V}}_k^T) \mathbf{Z} \tilde{\mathbf{U}}_k \in \mathbb{R}^{n \times r}$

Accelerating Truncated SVD

Denote $\mathbf{Z} := \mathbf{D} - \mathbf{S}_k$. Computing $\mathcal{H}_r(\mathcal{P}_{\tilde{T}_k} \mathbf{Z})$ efficiently:

- $\mathcal{P}_{\tilde{T}_k} \mathbf{Z} = \underbrace{\begin{bmatrix} \tilde{\mathbf{U}}_k & \mathbf{Q}_1 \end{bmatrix}}_{\text{orthogonal}} \underbrace{\begin{bmatrix} \tilde{\mathbf{U}}_k^T \mathbf{Z} \tilde{\mathbf{V}}_k & \mathbf{R}_2^T \\ \mathbf{R}_1 & 0 \end{bmatrix}}_{2r \times 2r} \underbrace{\begin{bmatrix} \tilde{\mathbf{V}}_k & \mathbf{Q}_2 \end{bmatrix}^T}_{\text{orthogonal}}$
 - $\mathbf{Q}_1 \mathbf{R}_1 = (\mathbf{I} - \tilde{\mathbf{U}}_k \tilde{\mathbf{U}}_k^T) \mathbf{Z} \tilde{\mathbf{V}}_k \in \mathbb{R}^{m \times r}$
 - $\mathbf{Q}_2 \mathbf{R}_2 = (\mathbf{I} - \tilde{\mathbf{V}}_k \tilde{\mathbf{V}}_k^T) \mathbf{Z} \tilde{\mathbf{U}}_k \in \mathbb{R}^{n \times r}$
- $\mathcal{H}_r(\mathcal{P}_{\tilde{T}_k} \mathbf{Z}) = \begin{bmatrix} \tilde{\mathbf{U}}_k & \mathbf{Q}_1 \end{bmatrix} \mathcal{H}_r \left(\begin{bmatrix} \tilde{\mathbf{U}}_k^T \mathbf{Z} \tilde{\mathbf{V}}_k & \mathbf{R}_2^T \\ \mathbf{R}_1 & 0 \end{bmatrix} \right) \begin{bmatrix} \tilde{\mathbf{V}}_k & \mathbf{Q}_2 \end{bmatrix}^T$
- Only need a SVD of $2r \times 2r$ matrix + two QR-decompositions

Accelerating Truncated SVD

Denote $\mathbf{Z} := \mathbf{D} - \mathbf{S}_k$. Computing $\mathcal{H}_r(\mathcal{P}_{\tilde{T}_k} \mathbf{Z})$ efficiently:

$$\bullet \mathcal{P}_{\tilde{T}_k} \mathbf{Z} = \underbrace{\begin{bmatrix} \tilde{\mathbf{U}}_k & \mathbf{Q}_1 \end{bmatrix}}_{\text{orthogonal}} \underbrace{\begin{bmatrix} \tilde{\mathbf{U}}_k^T \mathbf{Z} \tilde{\mathbf{V}}_k & \mathbf{R}_2^T \\ \mathbf{R}_1 & 0 \end{bmatrix}}_{2r \times 2r} \underbrace{\begin{bmatrix} \tilde{\mathbf{V}}_k & \mathbf{Q}_2 \end{bmatrix}^T}_{\text{orthogonal}}$$

$$- \mathbf{Q}_1 \mathbf{R}_1 = (\mathbf{I} - \tilde{\mathbf{U}}_k \tilde{\mathbf{U}}_k^T) \mathbf{Z} \tilde{\mathbf{V}}_k \in \mathbb{R}^{m \times r}$$

$$- \mathbf{Q}_2 \mathbf{R}_2 = (\mathbf{I} - \tilde{\mathbf{V}}_k \tilde{\mathbf{V}}_k^T) \mathbf{Z} \tilde{\mathbf{U}}_k \in \mathbb{R}^{n \times r}$$

$$\bullet \mathcal{H}_r(\mathcal{P}_{\tilde{T}_k} \mathbf{Z}) = \begin{bmatrix} \tilde{\mathbf{U}}_k & \mathbf{Q}_1 \end{bmatrix} \mathcal{H}_r \left(\begin{bmatrix} \tilde{\mathbf{U}}_k^T \mathbf{Z} \tilde{\mathbf{V}}_k & \mathbf{R}_2^T \\ \mathbf{R}_1 & 0 \end{bmatrix} \right) \begin{bmatrix} \tilde{\mathbf{V}}_k & \mathbf{Q}_2 \end{bmatrix}^T$$

• Only need a SVD of $2r \times 2r$ matrix + two QR-decompositions

• Complexities: $4n^2r + n^2 + O(nr^2 + r^3)$ $[\mathcal{H}_r \mathcal{P}_{\tilde{T}_k}]$

$O(n^2r)$ with large hidden constant $[\mathcal{H}_r]$

Efficient \mathbf{S} Updating with Hard Thresholding

- Updating \mathbf{S}_{k+1} with $\mathcal{T}_{\zeta_{k+1}}(\mathbf{D} - \mathbf{L}_{k+1})$, where

$$\zeta_{k+1} = \beta \left(\sigma_{r+1} \left(\mathcal{P}_{\tilde{T}_k}(\mathbf{D} - \mathbf{S}_k) \right) + \gamma^{k+1} \sigma_1 \left(\mathcal{P}_{\tilde{T}_k}(\mathbf{D} - \mathbf{S}_k) \right) \right),$$

so that we have $\text{supp}(\mathbf{S}_{k+1}) \subset \text{supp}(\mathbf{S})$ in theory.

Efficient \mathbf{S} Updating with Hard Thresholding

- Updating \mathbf{S}_{k+1} with $\mathcal{T}_{\zeta_{k+1}}(\mathbf{D} - \mathbf{L}_{k+1})$, where

$$\zeta_{k+1} = \beta \left(\sigma_{r+1} \left(\mathcal{P}_{\tilde{T}_k}(\mathbf{D} - \mathbf{S}_k) \right) + \gamma^{k+1} \sigma_1 \left(\mathcal{P}_{\tilde{T}_k}(\mathbf{D} - \mathbf{S}_k) \right) \right),$$

so that we have $\text{supp}(\mathbf{S}_{k+1}) \subset \text{supp}(\mathbf{S})$ in theory.

- The singular values of $\mathcal{P}_{\tilde{T}_k}(\mathbf{D} - \mathbf{S}_k)$ were already computed in \mathbf{L} updating.
 - Cost $O(1)$ for computing ζ_{k+1}

Efficient \mathbf{S} Updating with Hard Thresholding

- Updating \mathbf{S}_{k+1} with $\mathcal{T}_{\zeta_{k+1}}(\mathbf{D} - \mathbf{L}_{k+1})$, where

$$\zeta_{k+1} = \beta \left(\sigma_{r+1} \left(\mathcal{P}_{\tilde{T}_k}(\mathbf{D} - \mathbf{S}_k) \right) + \gamma^{k+1} \sigma_1 \left(\mathcal{P}_{\tilde{T}_k}(\mathbf{D} - \mathbf{S}_k) \right) \right),$$

so that we have $\text{supp}(\mathbf{S}_{k+1}) \subset \text{supp}(\mathbf{S})$ in theory.

- The singular values of $\mathcal{P}_{\tilde{T}_k}(\mathbf{D} - \mathbf{S}_k)$ were already computed in \mathbf{L} updating.
 - Cost $O(1)$ for computing ζ_{k+1}
- Complexity: $2n^2 + O(1)$
 - No expensive partial sorting needed.

Initialization

- If choose $\mathbf{L}_0 = 0$, then $\mathbf{U}_0 = \mathbf{V}_0 = 0$.
 - We don't want this happen.

- If choose $\mathbf{L}_0 = 0$, then $\mathbf{U}_0 = \mathbf{V}_0 = 0$.
 - We don't want this happen.

Initialization by Two Steps of AltProj

$$\mathbf{L}_{-1} = 0$$

$$\zeta_{-1} = \beta_{init} \cdot \sigma_1(\mathbf{D})$$

$$\mathbf{S}_{-1} = \mathcal{T}_{\zeta_{-1}}(\mathbf{D} - \mathbf{L}_{-1})$$


$$\mathbf{L}_0 = \mathcal{H}_r(\mathbf{D} - \mathbf{S}_{-1})$$

$$\zeta_0 = \beta \cdot \sigma_1(\mathbf{D} - \mathbf{S}_{-1})$$

$$\mathbf{S}_0 = \mathcal{T}_{\zeta_0}(\mathbf{D} - \mathbf{L}_0)$$

- Two steps of AltProj give us a close enough tangent space to start with.

Illustration of Our Approach



Outline

- 1 Introduction
- 2 Existing Approaches with Theoretical Guarantee
- 3 Proposed Approach**
 - Algorithm Design
 - Theoretical Guarantee**
- 4 Numerical Experiments
- 5 Extension to Low-Rank Hankel Recovery
- 6 Conclusion and Future Work

Main Theoretical Result

Recovery Guarantee of AccAltProj

Let \mathbf{L} and \mathbf{S} be two matrices satisfying Assumptions A1 and A2 with

$$\alpha \leq O \left(\min \left\{ \frac{1}{\mu r^2 \kappa^3}, \frac{1}{\mu^{1.5} r^2 \kappa}, \frac{1}{\mu^2 r^2} \right\} \right).$$

If the thresholding parameters obey $\frac{\mu r \sigma_1^L}{\sqrt{mn} \sigma_1^D} \leq \beta_{init} \leq \frac{3\mu r \sigma_1^L}{\sqrt{mn} \sigma_1^D}$ and $\beta = \frac{\mu r}{2\sqrt{mn}}$, along with the convergence rate parameter $\gamma \in (\frac{1}{\sqrt{12}}, 1)$, then the outputs of AccAltProj satisfy

$$\|\mathbf{L} - \mathbf{L}_k\|_F \leq \varepsilon \sigma_1^L, \|\mathbf{S} - \mathbf{S}_k\|_\infty \leq \frac{\varepsilon}{\sqrt{mn}} \sigma_1^L, \text{ and } \text{supp}(\mathbf{S}_k) \subset \text{supp}(\mathbf{S})$$

in $O(\log_\gamma \varepsilon)$ iterations.

- Linear convergence with rate of γ .

Proof Sketch

- Ⓐ Show any $m \times n$ RPCA problem can be reduced to a symmetric RPCA problem.

Proof Sketch

- A Show any $m \times n$ RPCA problem can be reduced to a symmetric RPCA problem.
- B Prove the symmetric case by mathematical induction:

Proof Sketch

- Ⓐ Show any $m \times n$ RPCA problem can be reduced to a symmetric RPCA problem.
- Ⓑ Prove the symmetric case by mathematical induction:
 - Base case: Prove initialization by 2-step AltProj satisfies

$$\|\mathbf{L}_0 - \mathbf{L}\|_2 \leq 8\alpha\mu r\sigma_1^L, \|\mathbf{S}_0 - \mathbf{S}\|_\infty \leq \frac{\mu r}{n}\sigma_1^L, \text{ and } \text{supp}(\mathbf{S}_0) \subset \text{supp}(\mathbf{S}).$$

Proof Sketch

- Ⓐ Show any $m \times n$ RPCA problem can be reduced to a symmetric RPCA problem.
- Ⓑ Prove the symmetric case by mathematical induction:
 - Base case: Prove initialization by 2-step AltProj satisfies

$$\|\mathbf{L}_0 - \mathbf{L}\|_2 \leq 8\alpha\mu r\sigma_1^L, \|\mathbf{S}_0 - \mathbf{S}\|_\infty \leq \frac{\mu r}{n}\sigma_1^L, \text{ and } \text{supp}(\mathbf{S}_0) \subset \text{supp}(\mathbf{S}).$$

- Induction Step: Prove following two lemmas

① If $\|\mathbf{L}_k - \mathbf{L}\|_2 \leq 8\alpha\mu r\gamma^k\sigma_1^L$, $\|\mathbf{S}_k - \mathbf{S}\|_\infty \leq \frac{\mu r}{n}\gamma^k\sigma_1^L$, and $\text{supp}(\mathbf{S}_k) \subset \text{supp}(\mathbf{S})$, then

$$\|\mathbf{L}_{k+1} - \mathbf{L}\|_2 \leq 8\alpha\mu r\gamma^{k+1}\sigma_1^L \text{ and}$$

$$\|\mathbf{L}_{k+1} - \mathbf{L}\|_\infty \leq \frac{(1-8\alpha\mu r\kappa)\mu r}{2n}\gamma^{k+1}\sigma_1^L.$$

Proof Sketch

- Ⓐ Show any $m \times n$ RPCA problem can be reduced to a symmetric RPCA problem.
- Ⓑ Prove the symmetric case by mathematical induction:
 - Base case: Prove initialization by 2-step AltProj satisfies

$$\|\mathbf{L}_0 - \mathbf{L}\|_2 \leq 8\alpha\mu r\sigma_1^L, \|\mathbf{S}_0 - \mathbf{S}\|_\infty \leq \frac{\mu r}{n}\sigma_1^L, \text{ and } \text{supp}(\mathbf{S}_0) \subset \text{supp}(\mathbf{S}).$$

- Induction Step: Prove following two lemmas

① If $\|\mathbf{L}_k - \mathbf{L}\|_2 \leq 8\alpha\mu r\gamma^k\sigma_1^L$, $\|\mathbf{S}_k - \mathbf{S}\|_\infty \leq \frac{\mu r}{n}\gamma^k\sigma_1^L$, and $\text{supp}(\mathbf{S}_k) \subset \text{supp}(\mathbf{S})$, then

$$\|\mathbf{L}_{k+1} - \mathbf{L}\|_2 \leq 8\alpha\mu r\gamma^{k+1}\sigma_1^L \text{ and}$$

$$\|\mathbf{L}_{k+1} - \mathbf{L}\|_\infty \leq \frac{(1-8\alpha\mu r\kappa)\mu r}{2n}\gamma^{k+1}\sigma_1^L.$$

② If $\|\mathbf{L}_{k+1} - \mathbf{L}\|_\infty \leq \frac{(1-8\alpha\mu r\kappa)\mu r}{2n}\gamma^{k+1}\sigma_1^L$, then


$$\|\mathbf{S}_{k+1} - \mathbf{S}\|_\infty \leq \frac{\mu r}{n}\gamma^{k+1}\sigma_1^L \text{ and } \text{supp}(\mathbf{S}_{k+1}) \subset \text{supp}(\mathbf{S}).$$

Key Properties of Tangent Space Projection (1)

- If L , L_k are two rank- r matrices, then

$$\|(\mathcal{I} - \mathcal{P}_{T_k})(L - L_k)\|_2 \leq \frac{\|L - L_k\|_2^2}{\sigma_r^L}.$$

- Make the local convergence analysis possible.



Key Properties of Tangent Space Projection (2)

- If \mathbf{L}_k is μ -incoherent and \mathbf{S} is α -sparse, then

$$\|\mathcal{P}_{T_k} \mathbf{S}\|_\infty \leq 3\alpha\mu r \|\mathbf{S}\|_\infty.$$

- $\mathcal{P}_{T_k} \mathbf{S}$ is not sparse.
- Ensures sparsity can still be used for error bound after tangent space projection.

Algorithm	Updating S	Updating L	Tolerance of p	Iterations needed
AccAltProj	$O(n^2)$	$O(rn^2)$	$O\left(\frac{1}{\max\{\mu r^2 \kappa^3, \mu^{1.5} r^2 \kappa, \mu^2 r^2\}}\right)$	$O\left(\log\left(\frac{1}{\epsilon}\right)\right)$
AltProj	$O(rn^2)$	$O(r^2 n^2)$	$O\left(\frac{1}{\mu r}\right)$	$O\left(\log\left(\frac{1}{\epsilon}\right)\right)$
GD	$O(n^2 + pn^2 \log(pn))$	$O(rn^2)$	$O\left(\frac{1}{\max\{\mu r^{1.5} \kappa^{1.5}, \mu r \kappa^2\}}\right)$	$O\left(\kappa \log\left(\frac{1}{\epsilon}\right)\right)$

*complexities in table are based on $\mathbf{D} \in \mathbb{R}^{n \times n}$

We have the best per-iteration complexity, with the best order of convergence, in the class of provable non-convex Robust PCA algorithms.

Outline

- 1 Introduction
- 2 Existing Approaches with Theoretical Guarantee
- 3 Proposed Approach
- 4 Numerical Experiments**
 - Synthetic Datasets
 - Video Background Subtraction
- 5 Extension to Low-Rank Hankel Recovery
- 6 Conclusion and Future Work

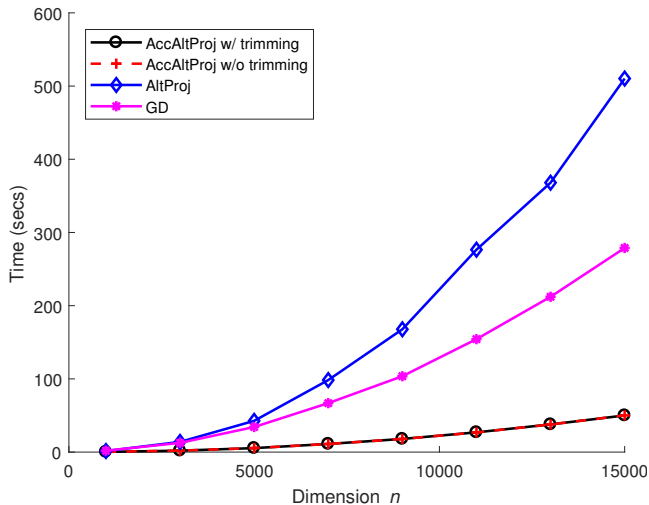
Outline

- 1 Introduction
- 2 Existing Approaches with Theoretical Guarantee
- 3 Proposed Approach
- 4 Numerical Experiments**
 - Synthetic Datasets
 - Video Background Subtraction
- 5 Extension to Low-Rank Hankel Recovery
- 6 Conclusion and Future Work

- $L = PQ^T$, where $P, Q \in \mathbb{R}^{n \times r}$ are random matrices where their entries are drawn i.i.d. from standard normal distribution.
- $\text{supp}(\mathbf{S})$ is sampled uniformly without replacement.
- The values of non-zero entries of \mathbf{S} are drawn i.i.d from uniform distribution over the domain $[-c \cdot \mathbb{E}(|L_{ij}|), c \cdot \mathbb{E}(|L_{ij}|)]$.
 - Smaller c gives a harder splitting task (unless it becomes too small to be outliers, like 10^{-8})

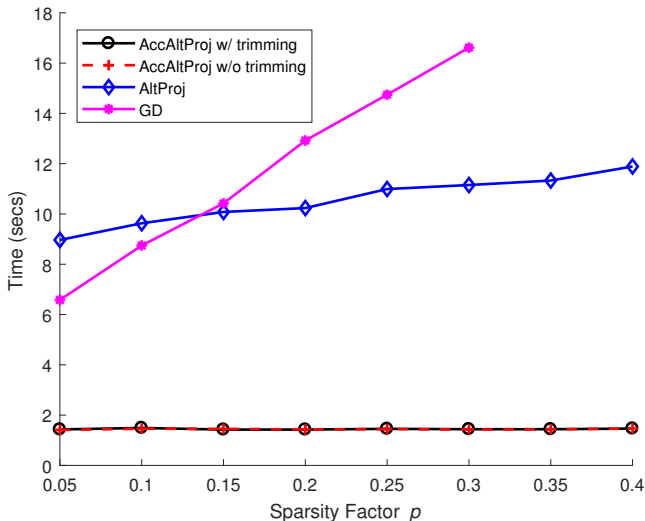
- Compare among three non-convex algorithms:
 - GD
 - Use 1.1α for sparsification \mathcal{F}
 - AltProj
 - Only time the r^{th} stage
 - AccAltProj with Trimming
 - Use 1.1μ for trimming
 - AccAltProj without Trimming
 - Skip the step $\tilde{\mathbf{L}}_k = \text{Trim}(\mathbf{L}_k, \mu)$
- PROPACK is used for large size truncated SVD
 - Good for AltProj which uses it every iteration

Speed Comparison: Dimension vs Runtime



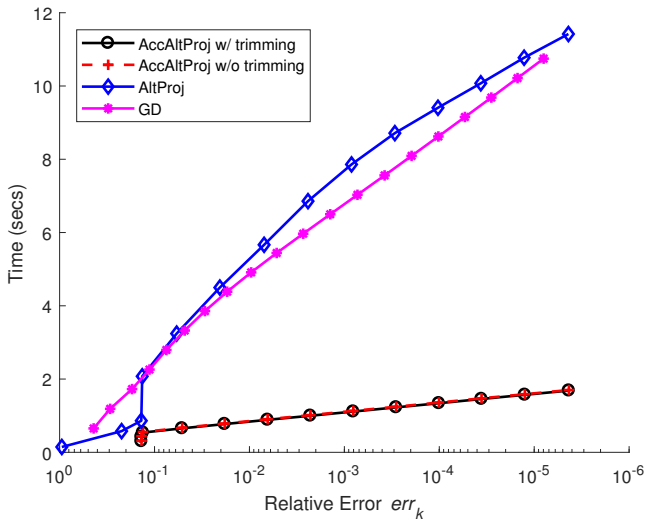
$r = 5, \alpha = 0.1, c = 1$, stop at $err_k < 10^{-4}$.

Speed Comparison: Sparsity vs Runtime



$r = 5, c = 1, n = 2500$, stop at $err_k < 10^{-4}$.

Speed Comparison: Relative Error vs Runtime



$r = 5, \alpha = 0.1, c = 1, n = 2500.$

Rate of Success for Varying c and α


$c = 0.2$	0.3	0.35	0.4	0.45	0.5	0.55	0.6
AccAltProj w/ trimming	10	10	10	10	10	10	10
AccAltProj w/o trimming	10	10	10	10	10	10	10
AltProj	10	10	10	10	10	10	10
GD	10	10	10	0	0	0	0
$c = 1$	0.3	0.35	0.4	0.45	0.5	0.55	0.6
AccAltProj w/ trimming	10	10	10	10	10	10	10
AccAltProj w/o trimming	10	10	10	10	10	10	10
AltProj	10	10	10	10	10	10	10
GD	10	10	10	10	9	0	0
$c = 5$	0.3	0.35	0.4	0.45	0.5	0.55	0.6
AccAltProj w/ trimming	10	10	10	10	10	10	10
AccAltProj w/o trimming	10	10	10	10	10	10	10
AltProj	10	10	10	10	10	10	10
GD	10	10	10	10	10	10	7

$r = 5, n = 2500.$

Outline

- 1 Introduction
- 2 Existing Approaches with Theoretical Guarantee
- 3 Proposed Approach
- 4 Numerical Experiments**
 - Synthetic Datasets
 - **Video Background Subtraction**
- 5 Extension to Low-Rank Hankel Recovery
- 6 Conclusion and Future Work

- Each frame in the video is vectorized, and becomes a row (or a column) in a data matrix.
- The static background is the low rank component of the matrix.
- The moving objects, i.e., foreground, forms the sparse components of the matrix.



<https://youtu.be/k6uaeQky2sc>

Comparison of Video Outputs

- Shoppingmall(S): 256×320 frame size; 1000 frames
- Restaurant(R): 120×160 frame size; 3055 frames

	AccAltProj w/ trim		AccAltProj w/o trim		AltProj		GD	
	runtime	μ	runtime	μ	runtime	μ	runtime	μ
S	38.98s	2.12	38.79s	2.26	82.97s	2.13	161.1s	2.85
R	28.09s	5.16	27.94s	5.25	69.12s	5.28	107.3s	6.07

Trimming helps the consistency a litter bit among the frames of the background, while only uses slightly more time.

AccAltProj has been published as:

Cai, HanQin, Jian-Feng Cai, and Ke Wei. "Accelerated alternating projections for robust principal component analysis." *The Journal of Machine Learning Research* 20.1 (2019): 685-717.

Code can be found at:

https://github.com/caesarcai/AccAltProj_for_RPCA

Outline

- 1 Introduction
- 2 Existing Approaches with Theoretical Guarantee
- 3 Proposed Approach
- 4 Numerical Experiments
- 5 Extension to Low-Rank Hankel Recovery**
- 6 Conclusion and Future Work

Form a Hankel Matrix

Given a complex vector $\mathbf{x} \in \mathbb{C}^n$, we can form a corresponding Hankel matrix $\mathcal{H}(\mathbf{x}) \in \mathbb{C}^{n_1 \times n_2}$ by

$$\mathcal{H}(\mathbf{x}) = \begin{bmatrix} x_0 & x_1 & x_2 & \cdots & x_{n_2-1} \\ x_1 & x_2 & x_3 & \cdots & x_{n_2} \\ x_2 & x_3 & x_4 & \cdots & x_{n_2+1} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ x_{n_1-1} & x_{n_1} & x_{n_1+1} & \cdots & x_{n-1} \end{bmatrix}$$

where $n_1 + n_2 = n + 1$.

- We try to keep the matrix (almost) square: when n is odd, we choose $n_1 = n_2 = \frac{n+1}{2}$; when n is even, we choose $n_1 = \frac{n}{2} + 1$ and $n_2 = \frac{n}{2}$.

*We try to using consistent notation with the paper, which is unfortunately inconsistent with the notation in previous sections.

Low-Rank Hankel Matrix

Not all Hankel matrices are low-rank, but low-rank Hankel matrix appears in many applications:

- Spectrally sparse signal
 - Magnetic resonance image (MRI)
 - Nuclear magnetic resonance (NMR) spectroscopy
 - Fluorescence microscopy
 - Analog-to-digital conversion
 - Earthquake-induced vibration data
- Dynamic sampling
- Auto-regression

Hankel Outlier Detection Problem

- To recover the vector \mathbf{x} from the corrupted observation $\mathbf{y} = \mathbf{x} + \mathbf{s} \in \mathbb{C}^n$ (or equivalently $\mathcal{H}(\mathbf{y}) = \mathcal{H}(\mathbf{x}) + \mathcal{H}(\mathbf{s}) \in \mathbb{C}^{n_1 \times n_2}$), we want to solve the minimization problem:

$$\begin{aligned} & \underset{\mathbf{x}, \mathbf{s}}{\text{minimize}} \quad \|\mathbf{y} - \mathbf{x} - \mathbf{s}\|_2 \\ & \text{subject to} \quad \text{rank}(\mathcal{H}(\mathbf{x})) \leq r, \|\mathbf{s}\|_0 \leq \alpha n \end{aligned} \tag{5}$$

- The sparse vector \mathbf{s} can be:
 - Impulse noise in MRI/NMR signals.
 - Corrupted time snap in dynamic sampling
 - Black swan events on financial marketing

Accelerated Alternating Projections for Hankel Recovery

Initialization, set $k = 0$

while $\|\mathbf{y} - \mathbf{x}_k - \mathbf{s}_k\| / \|\mathbf{z}\| \geq \epsilon$ **do**

$$\tilde{\mathbf{L}}_k = \text{Trim}(\mathbf{L}_k, \mu)$$

$$\mathbf{L}_{k+1} = \mathcal{D}_r \mathcal{P}_{\tilde{\mathcal{T}}_k} \mathcal{H}(\mathbf{y} - \mathbf{s}_k)$$

$$\mathbf{x}_{k+1} = \mathcal{H}^\dagger(\mathbf{L}_{k+1})$$

$$\zeta_{k+1} = \beta \left(\sigma_{r+1} \left(\mathcal{P}_{\tilde{\mathcal{T}}_k} \mathcal{H}(\mathbf{z} - \mathbf{s}_k) \right) \right) + \gamma^{k+1} \sigma_1 \left(\mathcal{P}_{\tilde{\mathcal{T}}_k} \mathcal{H}(\mathbf{z} - \mathbf{s}_k) \right)$$

$$\mathbf{s}_{k+1} = \mathcal{T}_{\zeta_{k+1}}(\mathbf{z} - \mathbf{x}_{k+1})$$

$$k = k + 1$$

end while

- Trim: Enforce incoherence on a low rank matrix
- \mathcal{D}_r : the best rank r approximation
- $\mathcal{P}_{\tilde{\mathcal{T}}_k}$: Projection onto subspace $\tilde{\mathcal{T}}_k = \{\tilde{\mathbf{U}}_k \mathbf{A}^T + \mathbf{B} \tilde{\mathbf{V}}_k^* \mid \mathbf{A} \in \mathbb{R}^{n_2 \times r}, \mathbf{B} \in \mathbb{R}^{n_1 \times r}\}$
- \mathcal{H}^\dagger , normalized dual of \mathcal{H} such that $\mathcal{H}^\dagger \mathcal{H} = \mathcal{I}$
- \mathcal{T}_ζ : Hard thresholding

Computational Complexity of Updating x

- Due to the construction of Hankel matrix, a Hankel matrix and vector multiplication can be re-formula to a vector-vector convolution, which can be computed via FFT. Hence, a Hankel matrix and matrix multiplication only needs $O(rn \log(n))$ flops.
- With a SVD of $2r \times 2r$ matrix + two QR-decompositions, we need $O(nr^2 + n \log(n)r + r^3)$ flops for updating \mathbf{L} .

Computational Complexity of Updating \mathbf{x}

- Due to the construction of Hankel matrix, a Hankel matrix and vector multiplication can be re-formula to a vector-vector convolution, which can be computed via FFT. Hence, a Hankel matrix and matrix multiplication only needs $O(rn \log(n))$ flops.
- With a SVD of $2r \times 2r$ matrix + two QR-decompositions, we need $O(nr^2 + n \log(n)r + r^3)$ flops for updating \mathbf{L} .
- $\mathbf{x}_{k+1} = \mathcal{H}^\dagger(\mathbf{L}_{k+1}) = \sum_{i=1}^r \boldsymbol{\Sigma}_{k+1}^{(i,i)} \omega_j \sum_{a+b=j} \mathbf{U}_{k+1}^{(a,i)} \overline{\mathbf{V}_{k+1}^{(b,i)}}$, $\forall j$
 - ω_j is the weight for averaging j^{th} anti-diagonal
 - Convolutions can be computed by FFT. r of the convolutions result $O(n \log(n)r)$ complexity.
- Overall, need $O(nr^2 + n \log(n)r + r^3)$ flops for updating \mathbf{x} .

Assumptions

A1 $\mathcal{H}(\mathbf{x}) \in \mathbb{C}^{n_1 \times n_2}$ is μ -incoherence: there exists a numerical constant $\mu_0 > 0$ such that

$$\sigma_r(\mathbf{E}_L^* \mathbf{E}_L) \geq \frac{n_1}{\mu_0}, \quad \sigma_r(\mathbf{E}_R^* \mathbf{E}_R) \geq \frac{n_2}{\mu_0},$$

where $\mathcal{H}(\mathbf{x}) = \mathbf{E}_L \mathbf{D} \mathbf{E}_R^T$ is its Vandermonde decomposition.

- In the undamped spectrally sparse signals, this assumption is guaranteed with at least $\frac{2}{n}$ separations between the non-zero frequencies [LiaoFannjiang'2016].

A2 $\mathbf{s} \in \mathbb{C}^n$ is α -sparse: no more than αn non-zero entries in the vector.

- this implies that $\mathcal{H}(\mathbf{s})$ is no more than 2α -sparse matrix.

Recovery Guarantee of AAP-Hankel

Let \mathbf{x} and \mathbf{s} be two matrices satisfying Assumptions A1 and A2 with

$$\alpha \lesssim O\left(\min\left\{\frac{1}{\mu r^2 \kappa^3}, \frac{1}{\mu^{1.5} r^2 \kappa}, \frac{1}{\mu^2 r^2}\right\}\right).$$

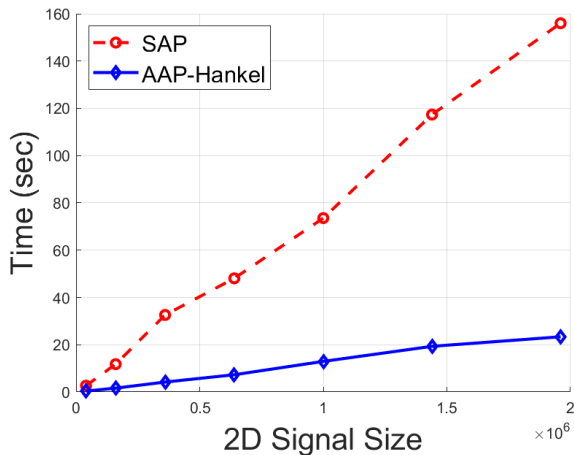
If the thresholding parameters obey $\frac{\mu r \sigma_1^x}{\sqrt{n_1 n_2 \sigma_1^z}} \leq \beta_{init} \leq \frac{3\mu r \sigma_1^x}{\sqrt{n_1 n_2 \sigma_1^z}}$ and $\beta = \frac{\mu r}{2\sqrt{n_1 n_2}}$, along with the convergence rate parameter $\gamma \in (\frac{1}{\sqrt{12}}, 1)$, then the outputs of Algorithm 1 satisfy

$$\|\mathcal{H}(\mathbf{x} - \mathbf{x}_k)\|_F \leq \varepsilon \sigma_1^x, \|\mathbf{s} - \mathbf{s}_k\|_\infty \leq \frac{\varepsilon}{\sqrt{mn}} \sigma_1^x, \text{ and } \text{supp}(\mathbf{s}_k) \subset \text{supp}(\mathbf{s})$$

in $O(\log_\gamma \varepsilon)$ iterations.

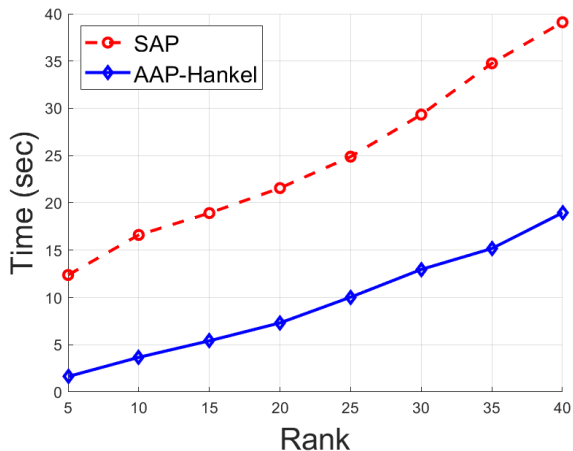
- σ_1^x denotes the largest singular value of $\mathcal{H}(\mathbf{x})$.
- Linear convergence with rate of γ .

Speed Comparison: Dimension vs Runtime



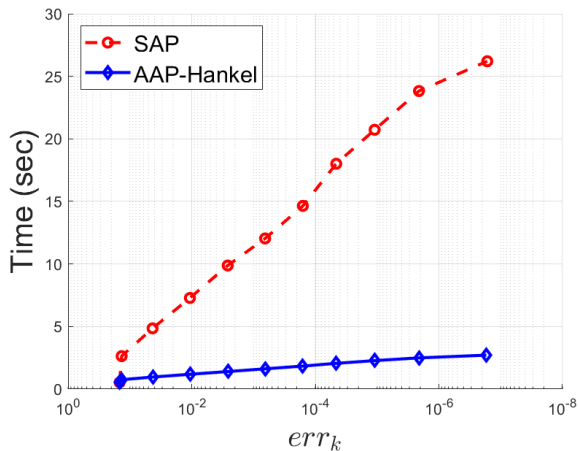
$r = 5, \alpha = 0.1, c = 1$, stop at $err_k < 10^{-4}$.

Speed Comparison: Sparsity vs Runtime






$n = 400^2$, $\alpha = 0.1$, $c = 1$, stop at $err_k < 10^{-4}$.

Speed Comparison: Relative Error vs Runtime

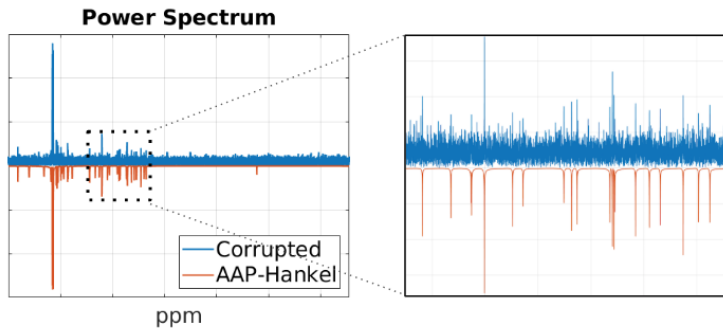


$n = 400^2, r = 5, \alpha = 0.1, c = 1.$

Phase Transition



Impulse Corrupted NMR Singal



$n = 32768$, 50% corruption.

AAP-Hankel is under review, you can find it on arXiv:

Cai, HanQin, Jian-Feng Cai, Tianming Wang, and Guojian Yin. "Fast and Robust Spectrally Sparse Signal Recovery: A Provable Non-Convex Approach via Robust Low-Rank Hankel Matrix Reconstruction." *arXiv preprint arXiv:1910.05859* (2019).

Code can be found at:

<https://github.com/caesarcai/AAP-Hankel>

Outline

- 1 Introduction
- 2 Existing Approaches with Theoretical Guarantee
- 3 Proposed Approach
- 4 Numerical Experiments
- 5 Extension to Low-Rank Hankel Recovery
- 6 Conclusion and Future Work**

We present a new algorithm for Robust PCA and extend it to low-rank Hankel recovery problem.

- Non-convex with theoretical guarantee
- Best per-iteration complexity
- Guaranteed linear convergence
- Significant speed advantage over the other state-of-the-art algorithms
- High robustness in practice

- Recovery guarantee proof for “without Trimming”
- Relax tolerance of α to match our experimental results
- Study partial observed Robust PCA
 - Acceleration with guaranteed recovery
- Study recovery stability with additive noise
 - Theoretically and practically

Thank you!

Any Question?

