Basis and Spectral Representations: Matrix Completion

Gil Shabat

School of Computer Science
Tel-Aviv University

November 3, 2012

Introduction

- Matrix completion recently became a very popular research topic
- Has many applications in collaborative filtering, computer vision and machine learning
- The “Holy Grail” of matrix completion problems is the Low-Rank completion, which is an NP-hard problem

Agenda

In this lesson we will discuss:

- Motivation for the matrix completion problem
- Short review on required mathematical background (SVD, Norms...)
- Nuclear norm derivative
- Projected gradient method
- Approximation techniques and geometry
- Approximation for completion
- The SVT algorithm

Matrix Approximation and Completion

- Matrix completion is a popular research field with many applications to biology, machine learning, image processing and other scientific fields.
- In particular - Low Rank completions
- Example: The Netflix problem for movie recommendations:
  - Rows correspond to viewers, columns to movies
  - Entry $X_{ij} \in \{1, ..., 5\}$ is the rating
  - $480,000$ viewers $\times 18,000$ movies $\Rightarrow 8.9 \times 10^9$ entries
- Each viewer recommends 200 movies on average, so only 1.2% of the entries contain data
- The task is to predict the ratings that viewers will give to movies they have not yet rated

Matrix Approximation and Completion

Mathematically, the problem is:

$$\begin{aligned} &\text{minimize } \operatorname{rank}(\mathbf{X}) \\ &\text{subject to } X_{ij} = M_{ij}, \quad (i,j) \in \Omega \end{aligned} \tag{1}$$

We are looking for a matrix with minimal rank, such that the given entries do not change. The problem in Eq. 7 is NP-hard Complete by minizing the rank:

$$\mathbf{M} = \begin{bmatrix} M_{1,1} & M_{1,2} \\ M_{2,1} & ? \end{bmatrix} = \begin{bmatrix} 1 & 2 \\ 3 & ? \end{bmatrix}.$$

The missing entry should be 6 and the rank will be 1. In this case, the solution is unique.

Matrix Approximation and Completion

In the course we will discuss some matrix completion related topics:

- Least squares approximation of matrices under different spectral regularizations
- Methods for approximating just some of the entries
- Discuss convergence of different methods: local and global convergence
- Discuss today's state of the art methods for matrix completion and their drawbacks

Matrix Approximation and Completion

An example for matrix approximation method:

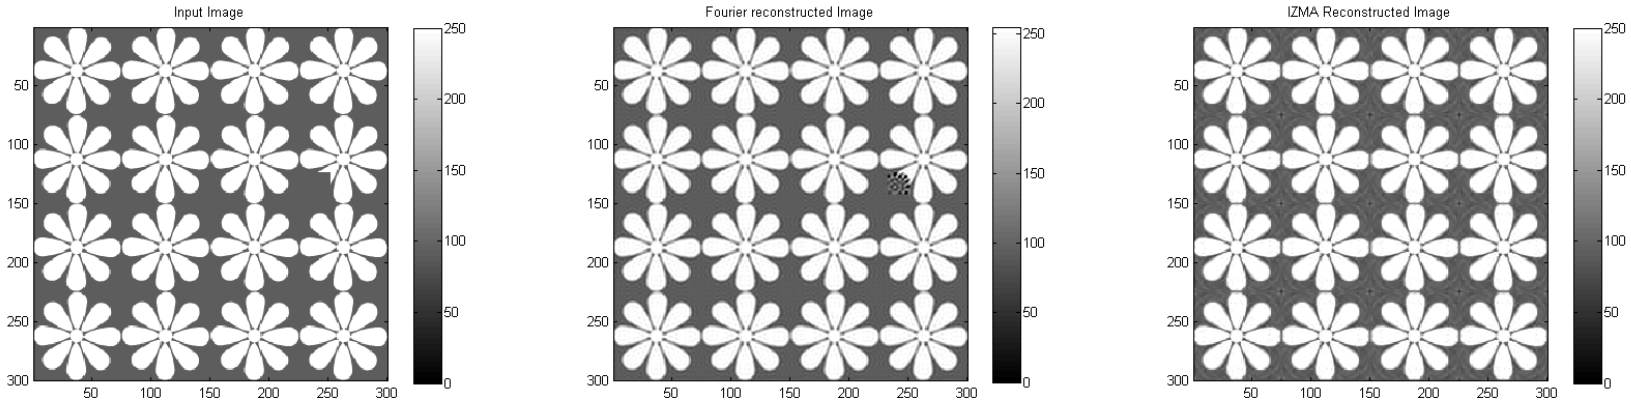

Original, FFT and Rank. The approximation by rank successfully constructed the flower.

Matrix Norms

- Frobenius norm:

$$\|\mathbf{A}\|_F^2 = trace(\mathbf{A}^*\mathbf{A}) = \sum_{i,j} |a_{ij}|^2 = \sum_i \sigma_i^2$$

- Spectral norm (induced norm for p=2):

$$\|\mathbf{A}\|_2 = \frac{sup\|\mathbf{A}\mathbf{x}\|_2}{\|\mathbf{x}\|_2} = \sigma_1$$

$\|\mathbf{A}\|_2$ is equal to the largest singular values of $\mathbf{A}$

Matrix Norms - More General

- Ky-Fan norm norm:

$$\|\mathbf{A}\|_k = \sum_{i=1}^{k} \sigma_i$$

- Schatten norm:

$$\|\mathbf{A}\|_p = (\sum_{i=1}^{n} \sigma_i^p)^{\frac{1}{p}}$$

  Uses all the singular values.
- A special case for $p = 1$, Nuclear norm:

$$\|\mathbf{A}\|_* = \sum_{i=1}^{n} \sigma_i$$

- In general, it should be a gauge function of the singular values.

Why Low Rank?

It is very often occurs that data depends on a small number of factors. This is typical for physical data or real life images.

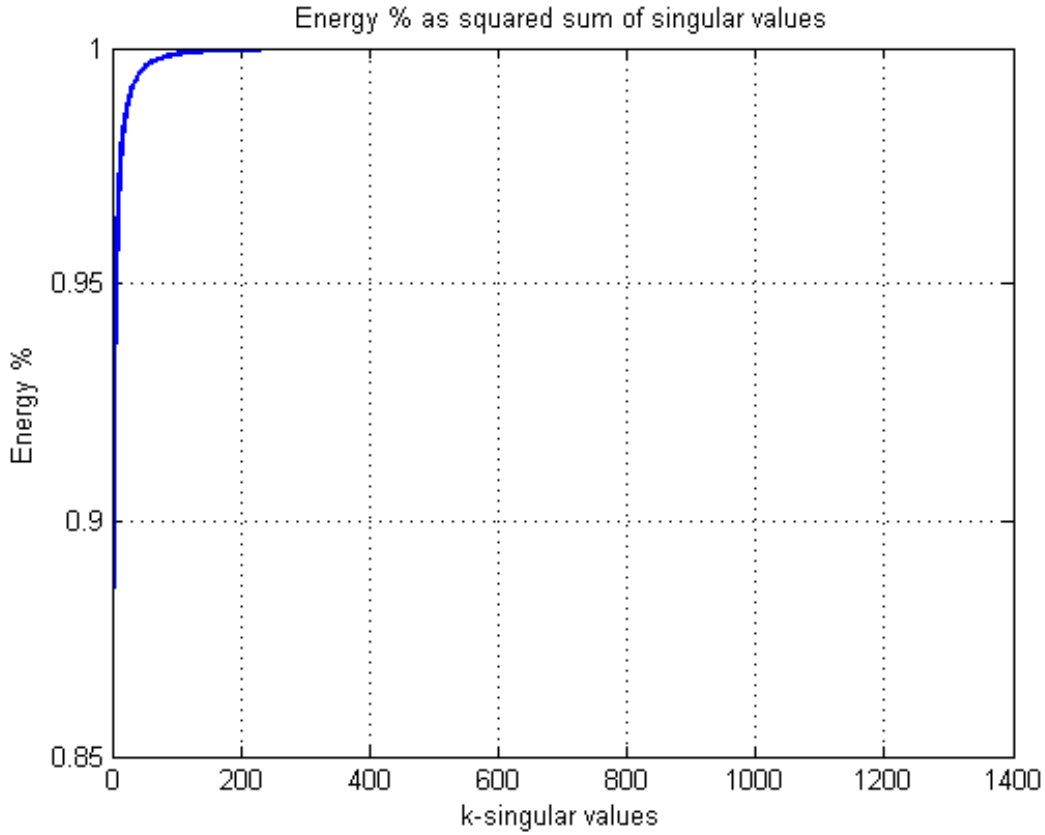

Approximation Theorems

The most famous approximation theorem is the Eckart-Young theorem (1936). The theorem gives the best low rank approximation of a given matrix of minimizing $\|\tilde{\mathbf{X}} - \mathbf{X}\|_F$ (or $\|\tilde{\mathbf{X}} - \mathbf{X}\|_2$) s.t. $rank(\tilde{\mathbf{X}}) = k$

Theorem (Eckart-Young)

The best (Frobenius norm) approximation to a matrix $\mathbf{X}$ *by a rank* k *matrix, is given by* $\tilde{\mathbf{X}} = \mathbf{U}\tilde{\mathbf{S}}\mathbf{V}^*$ *where* $diag(\mathbf{S}) = (s_1, .., s_k, 0, .., 0)$ *and* $\mathbf{USV}^*$ *is the SVD of* $\mathbf{X}$

The same applies to the spectral norm as well:

Theorem (Spectral norm minimization)

The best spectral norm approximation to a matrix $\mathbf{X}$ *by a rank* k *matrix, is given by* $\tilde{\mathbf{X}} = \mathbf{U}\tilde{\mathbf{S}}\mathbf{V}^*$ *where* $diag(\mathbf{S}) = (s_1, .., s_k, 0, .., 0)$ *and* $\mathbf{USV}^*$ *is the SVD of* $\mathbf{X}$

Approximation Theorems

Approximation under Frobenius norm:

$$\begin{aligned} &\text{minimize } \|\mathbf{X}-\mathbf{M}\|_F \\ &\text{subject to } \|\mathbf{X}\|_F \leq \lambda. \end{aligned}$$

The solution is given by: $\mathbf{X} = \frac{\mathbf{M}}{\|\mathbf{M}\|_\mathbf{F}} \min(\|\mathbf{M}\|_\mathbf{F}, \lambda)$. The proof is done by observing on $\mathbf{M}$ as if it were a point in an $m \times n$ dimensional space. We are looking for the closest point on the ball $\|\mathbf{X}\|_F \leq \lambda$ Suppose we want to take only some of the entries into account. That is:

$$\begin{aligned} &\text{minimize } \|\mathcal{P}\mathbf{X}-\mathcal{P}\mathbf{M}\|_F \\ &\text{subject to } \|\mathbf{X}\|_F \leq \lambda. \end{aligned}$$

The proof is similar to the previous one but here we are looking for a point $\mathbf{X}$ on the sphere that is the closest to a line whose points $\mathbf{X}' \in \mathcal{H}$ satisfy $\mathcal{P}\mathbf{X}' = \mathcal{P}\mathbf{M}$. By geometrical considerations, this point is given by $\mathbf{X} = \frac{\mathcal{P}\mathbf{M}}{\|\mathcal{P}\mathbf{M}\|_F}\lambda$.

More approximation theorems

The Procustrates theorem:

$$\begin{aligned} &\text{minimize } \|\mathbf{X} - \mathbf{M}\|_F \\ &\text{subject to } \mathbf{X}^*\mathbf{X} = \mathbf{I}. \end{aligned}$$

Can be extended to:

$$\begin{aligned} &\text{minimize } \|\mathbf{X} - \mathbf{M}\|_F \\ &\text{subject to } \mathbf{X}^*\mathbf{X} = \mathbf{D^2}. \end{aligned}$$

($\mathbf{D}$ diagonal). Can be converted by using: $\mathbf{X} = \mathbf{YD}$, $\mathbf{Y}^*\mathbf{Y} = \mathbf{I}$. When $\mathbf{D}$ is unknown, an iterative solution exists - the constraints is to have $\mathbf{X}^*\mathbf{X}$ diagonal.

Pinching Theorem

The “Pinching” theorem states that:

$$\|diag(\mathbf{A})\| \leq \|\mathbf{A}\|$$

For every norm, satisfying $\|\mathbf{UAV}\| = \|\mathbf{A}\|$ for $\mathbf{U}, \mathbf{V}$ orthogonal. Hence, the theorem is applicable to all the matrix norm we discussed so far (Frobenius, Schatten, Ky-Fan). We will see a simpler prove, but more specific to the nuclear and spectral norms using Jacobi rotations.

Jacobi Rotations

Jacobi rotations are used to reduce a symmetric matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ to a diagonal form using rotation matrices. The idea is to reduce the norm of the off-diagonal entries of $\mathbf{A}$ by using the rotation matrix $\mathbf{Q}$. $\mathbf{Q}$ is an $n \times n$ matrix that is equal to the identity matrix except for four entries, given by:

$$\begin{aligned} q_{kk} &= q_{ll} = \cos\theta \\ q_{kl} &= \sin\theta \\ q_{lk} &= -\sin\theta \end{aligned} \tag{2}$$

$$\mathbf{B} = \mathbf{Q}^{\mathsf{T}}\mathbf{A}\mathbf{Q} \tag{3}$$

where θ is chosen to minimize the off-diagonal part of $\mathbf{B}$ that it is given by:

$$\tau \triangleq \cot\theta = \frac{a_{ll} - a_{kk}}{2a_{kl}}, \quad t \triangleq \tan\theta = \frac{\operatorname{sign}(\tau)}{|\tau| + \sqrt{1+\tau^2}}. \tag{4}$$

Jacobi Rotations

Theorem (The main theorem)

Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ *be a symmetric matrix and let* $\mathbf{B} = \mathbf{J}^{\mathbf{T}}\mathbf{A}\mathbf{J}$ *be its Jacobi rotation for the entries* (k, l)*. Assume that* $a_{kk} \geq a_{ll}$*. Then,* $b_{kk} \geq a_{kk}$ *and* $b_{ll} \leq a_{ll}$*. More precisely,* $b_{kk} = a_{kk} + \delta$ *and* $b_{ll} = a_{ll} - \delta$*,* $(\delta \geq 0)$*.*

Jacobi Rotations

Proof.

The proof uses the Jacobi roation. In each application of Jacobi rotation matrix to **A**, the norm of the off-diagonal part is getting smaller and the diagonal part changes as well. By simple calculations, it is possible to find the following updated equations for the new diagonal:

$$\begin{aligned} b_{kk} &= a_{kk} - ta_{kl} \\ b_{ll} &= a_{ll} + ta_{kl}. \end{aligned} \tag{5}$$

The sign of t is equal to the sign of τ. The sign of τ depends on the three entries a_{kk}, a_{kl}, a_{ll}, as shown from the expression $\tau = \cot\theta = \frac{a_{ll}-a_{kk}}{2a_{kl}}$. We divide it into four cases: □

Jacobi Rotations

Proof.

1. $a_{kl} > 0$, $a_{ll} > a_{kk}$: In this case, τ is positive and therefore t is positive. $b_{kk} < a_{kk}$ and $b_{ll} > a_{ll}$. The smallest entry a_{kk} becomes even smaller and the largest entry a_{ll} becomes even larger;
2. $a_{kl} > 0$, $a_{ll} < a_{kk}$: Here t is negative and according to the update equations, a_{kk} is getting larger and a_{ll} is getting smaller ($b_{kk} > a_{kk}$ and $b_{ll} < a_{ll}$);
3. $a_{kl} < 0$, $a_{ll} > a_{kk}$: t is negative, a_{kk} is getting smaller and a_{ll} is getting larger ($b_{ll} > a_{ll}$ and $b_{kk} < a_{kk}$);
4. $a_{kl} < 0$, $a_{ll} < a_{kk}$: t is positive, a_{ll} is getting smaller and a_{kk} is getting larger ($b_{ll} < a_{ll}$ and $b_{kk} > a_{kk}$).

The conclusion from the application of the Jacobi rotation is that the largest diagonal entry becomes even larger and the smallest

Proving the Pinching theorem

Theorem (Pinching for the Schatten norm (Symmetric))

Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ *be symmetric matrix. Then,* $\|diag(\mathbf{A})\|_p \leq \|\mathbf{A}\|_p$.

Proof.

We apply the Jacobi rotation to $\mathbf{A}$ such that $\mathbf{B} = \mathbf{J^T A J}$ while operating on entry (k, l). Suppose $a_{kk} \geq a_{ll}$ and $\delta \geq 0$. We examine the expression
$|b_{kk}|^p + |b_{ll}|^p = |a_{kk} + \delta|^p + |a_{ll} - \delta|^p \geq |a_{kk}|^p + |a_{ll}|^p$. Each iteration increases the l_p norm of the diagonal until it reaches the Schatten norm of $\mathbf{A}$. □

Pinching theorem

Lemma (Extension to real matrices)

Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ *be a square matrix. Then,* $\|diag(\mathbf{A})\|_p \leq \|\mathbf{A}\|_p$.

Proof.

From the triangle inequality $\|\mathbf{A} + \mathbf{A}^\mathbf{T}\| \leq \|\mathbf{A}\| + \|\mathbf{A}^\mathbf{T}\| = 2\|\mathbf{A}\|$.
Hence $\|\mathbf{A}\| \geq \frac{1}{2}\|\mathbf{A} + \mathbf{A}^\mathbf{T}\|$. Since $\mathbf{A} + \mathbf{A}^\mathbf{T}$ is symmetric, we use Theorem 4 that yields
$\|\mathbf{A}\| \geq \frac{1}{2}\|\mathbf{A} + \mathbf{A}^\mathbf{T}\| \geq \frac{1}{2}\|\text{diag}(\mathbf{A} + \mathbf{A}^\mathbf{T})\| = \|\text{diag}(\mathbf{A})\|$. □

Pinching theorem

Lemma (Extension to complex matrices with real diagonal)

Let $\mathbf{A}$ *be a square matrix with real diagonal. Then,* $\|diag(\mathbf{A})\|_p \leq \|\mathbf{A}\|_p$.

Proof.

The proof is similar to the proof of Lemma 5. From the triangle inequality we get $\|\mathbf{A}\| \geq \frac{1}{2}\|\mathbf{A} + \text{conj}(\mathbf{A})\|$. By using Lemma 5 we get $\frac{1}{2}\|\mathbf{A} + \text{conj}(\mathbf{A})\| \geq \frac{1}{2}\|\text{diag}(\mathbf{A} + \text{conj}(\mathbf{A})\| = \|\text{diag}(\mathbf{A})\|$. Here we used the fact that diag($\mathbf{A}$) is real. □

Pinching theorem

Theorem (Extension to complex matrices)

Let $\mathbf{A} \in \mathbb{C}^{n \times n}$ *be a square matrix,. Then,* $\|diag(\mathbf{A})\|_p \leq \|\mathbf{A}\|_p$.

Proof.

Let $\mathbf{U}$ be a diagonal unitary (square) matrix whose elements are $u_j = e^{-i\theta_j}$ where θ_j is the phase of a_{jj}. Because of the structure of $\mathbf{U}$, diag($\mathbf{UA}$) is real. Since $|u_j| = 1$ we get $\|\text{diag}(\mathbf{A})\| = \|\text{diag}(\mathbf{UA})\|$. From Lemma 6 we get $\|\text{diag}(\mathbf{A})\| = \|\text{diag}(\mathbf{UA})\| \leq \|\mathbf{UA}\| = \|\mathbf{A}\|$. □

Now that we have the Pinching theorem, we can use it to prove matrix approximation theorems to nuclear and spectral norms.

Approximation under spectral norm

Lemma (Minimization of the Frobenius norm under the spectral norm constraint)

Assume the SVD of $\mathbf{M}$ *is given by* $\mathbf{M} = \mathbf{USV}^*$ *where* $\mathbf{S} = diag(\sigma_1, .., \sigma_n)$. *Then, the matrix* $\mathbf{X}$, *which minimizes* $\|\mathbf{X} - \mathbf{M}\|_F$ *such that* $\|\mathbf{X}\|_2 \leq \lambda$, *is given by* $\mathbf{X} = \mathbf{U\tilde{S}V}^*$ *where* $\tilde{\sigma}_i$ *are the singular values of* $\tilde{S}$ *and* $\tilde{\sigma}_i = \min(\sigma_i, \lambda), i = 1, \ldots k,\ k \leq n.$

Proof.

$\|\mathbf{X} - \mathbf{M}\|_F = \|\mathbf{X} - \mathbf{USV}^*\|_F = \|\mathbf{U}^*\mathbf{XV} - \mathbf{S}\|_F$. Since $\mathbf{S}$ is diagonal, $\|\text{diag}(\mathbf{U}^*\mathbf{XV}) - \mathbf{S}\|_F \leq \|\mathbf{U}^*\mathbf{XV} - \mathbf{S}\|_F$. From P.T. $\|\text{diag}(\mathbf{U}^*\mathbf{XV})\|_2 \leq \|\mathbf{U}^*\mathbf{XV}\|_2$. Therefore, $\mathbf{U}^*\mathbf{XV}$ has to be diagonal and the best minimizer under the spectral norm constraint is achieved by minimizing each element separately yielding $\mathbf{U}^*\mathbf{XV} = \text{diag}(\min(\sigma_i, \lambda)),\ i = 1, \ldots k, k \leq n$ and $\mathbf{X} = \mathbf{U\tilde{S}V}^*$. □

Approximation under nuclear norm

The same argument that states that $\mathbf{U}^*\mathbf{XV}$ has to be diagonal can also be applied when the constraint is given by the nuclear norm. Define $\tilde{\mathbf{S}} = \mathbf{U}^*\mathbf{XV}$. We wish to minimize $\|\tilde{\mathbf{S}} - \mathbf{S}\|_F = \sum_i (\tilde{\sigma}_i - \sigma_i)^2$ s.t. $\|\mathbf{X}\|_* = \|\tilde{\mathbf{S}}\|_* = \sum_i |\tilde{\sigma}_i| \leq \lambda, i = 1, \ldots k, k \leq n$. Note that $\tilde{\sigma}_i$ has to be nonnegative otherwise it will increase the Frobenius norm but will not change the nuclear norm. Hence, the problem can now be formulated as:

$$\begin{aligned} &\text{minimize } \textstyle\sum_i (\tilde{\sigma}_i - \sigma_i)^2 \\ &\text{subject to } \textstyle\sum_i \tilde{\sigma}_i \leq \lambda \\ &\text{subject to } \tilde{\sigma}_i \geq 0. \end{aligned} \tag{6}$$

This is a standard convex optimization problem that can be solved by methods such as semidefinite programming

Approximation of certain entries

We have seen how to approximate full matrices, can we use them to approximate only in certain entries?

$$\begin{aligned} &\text{minimize } \|\mathcal{P}\mathbf{X} - \mathcal{P}\mathbf{M}\|_F \\ &\text{subject to } f(\mathbf{X}) \leq 0 \end{aligned}$$

Where $\mathcal{P}$ is a projection operator that indicates the entries we wish to approximate and zeros the others.

Solving Iteratively

Suppose we try to reduce the error between an initial guess and the matrix we try to approximate at the specific entries, by applying the projected gradient method:

$$\mathbf{X}_{n+1} = \mathcal{D}(\mathbf{X}_n - \mu_n \mathcal{P}(\mathbf{X}_n - \mathbf{M}))$$

Where $\mathcal{D}$ is a new operator that returns the best approximation for the full matrix. For example, if we are interested in the rank approximation, then $\mathcal{D}$ will be applying the EY theorem.
It rises the following questions:

- How to determine the step size μ_n ?
- Does the algorithm converge?
- Does it converges to the correct solution?

Solving Iteratively

For simplicity, assume $\mu_n = 1$ for every n. The error at the nth iteration is defined by:

$$\epsilon_n = \|\mathcal{P}\mathbf{X}_n - \mathcal{P}\mathbf{M}\|_F$$

Does ϵ_n converge? Since the error is non-negative, it is sufficient to show it is monotonically decreasing.

Theorem (Algorithm convergence)

The error ϵ_n is monotonically decreasing in each iteration, for $\mu_n = 1$. Hence, the algorithm converges

We will now proof the theorem.

Solving Iteratively

Definitions:

- $\mathcal{P}$ - Projection operator: $\mathcal{P}\mathbf{X} \triangleq \mathbf{B} \otimes \mathbf{X}$, $B_{i,j} \in \{0, 1\}$
- $\mathcal{D}$ - Best approximation operator: $\mathcal{D}\mathbf{X}$ is the best approximation of $\mathbf{X}$
- $\mathcal{W}$ - Entries correction operator: $\mathcal{W}\mathbf{X} \triangleq (\mathcal{I} - \mathcal{P})\mathbf{X} + \mathcal{P}\mathbf{M}$, has the following properties:
  - $\mathcal{P}\mathcal{W}\mathbf{X} = \mathcal{P}\mathbf{M}$;
  - $(\mathcal{I} - \mathcal{P})\mathcal{W}\mathbf{X} = (\mathcal{I} - \mathcal{P})\mathbf{X}$;
  - $\mathbf{X} - \mathcal{W}\mathbf{X} = \mathcal{P}\mathbf{X} - \mathcal{P}\mathbf{M}$.
- $\mathcal{T}$ - For convenience: $\mathcal{T}\mathbf{X} \triangleq \mathcal{D}\mathcal{W}\mathbf{X}$

Solving Iteratively

Proof.

Let $\mathbf{Q}$ be the line of all matrices with zero error, and let $\mathcal{T}^n\mathbf{X}$ be the point at the nth iteration. $\mathcal{T}^n\mathbf{X}$ is mapped to $\mathcal{W}\mathcal{T}^n\mathbf{X} \in \mathbf{Q}$ and has zero error and then to $\mathcal{T}^{n+1}\mathbf{X}$. $\mathcal{T}^{n+1}\mathbf{X}$ must be inside a ball centered at $\mathcal{W}\mathcal{T}^n\mathbf{X}$ and whose radius is $\|\mathcal{W}\mathcal{T}^n\mathbf{X} - \mathcal{T}^n\mathbf{X}\|$ otherwise $\mathcal{T}^n\mathbf{X}$ is a better low rank approximation to $\mathcal{W}\mathcal{T}^n\mathbf{X}$ and this contradicts the Eckart-Young Theorem. □

Matrix Approximations I

Mathematically, it is written as:

$$\begin{aligned} \|\mathcal{T}^{n+1}\mathbf{X} - \mathcal{W}\mathcal{T}^{n}\mathbf{X}\|^2 &= \|(\mathcal{I}-\mathcal{P})\mathcal{T}^{n+1}\mathbf{X} - (\mathcal{I}-\mathcal{P})\mathcal{W}\mathcal{T}^{n}\mathbf{X}\|^2 + \\ \|\mathcal{P}\mathcal{T}^{n+1}\mathbf{X} - \mathcal{P}\mathcal{W}\mathcal{T}^{n}\mathbf{X}\|^2 &\leq \|\mathcal{T}^{n}\mathbf{X} - \mathcal{W}\mathcal{T}^{n}\mathbf{X}\|^2 \\ &= \|\mathcal{P}\mathcal{T}^{n}\mathbf{X} - \mathcal{P}\mathbf{M}\|^2 \end{aligned} \tag{7}$$

where in Eq. 7 we used the third property of $\mathcal{W}$
and since (third property):
$\|\mathcal{P}\mathcal{T}^{n+1}\mathbf{X} - \mathcal{P}\mathcal{W}\mathcal{T}^{n}\mathbf{X}\| = \|\mathcal{P}\mathcal{T}^{n+1}\mathbf{X} - \mathcal{P}\mathbf{M}\|$
we finally get:

$$\|\mathcal{P}\mathcal{T}^{n+1}\mathbf{X} - \mathcal{P}\mathbf{M}\| \leq \|\mathcal{P}\mathcal{T}^{n}\mathbf{X} - \mathcal{P}\mathbf{M}\| \tag{8}$$

Solving Iteratively

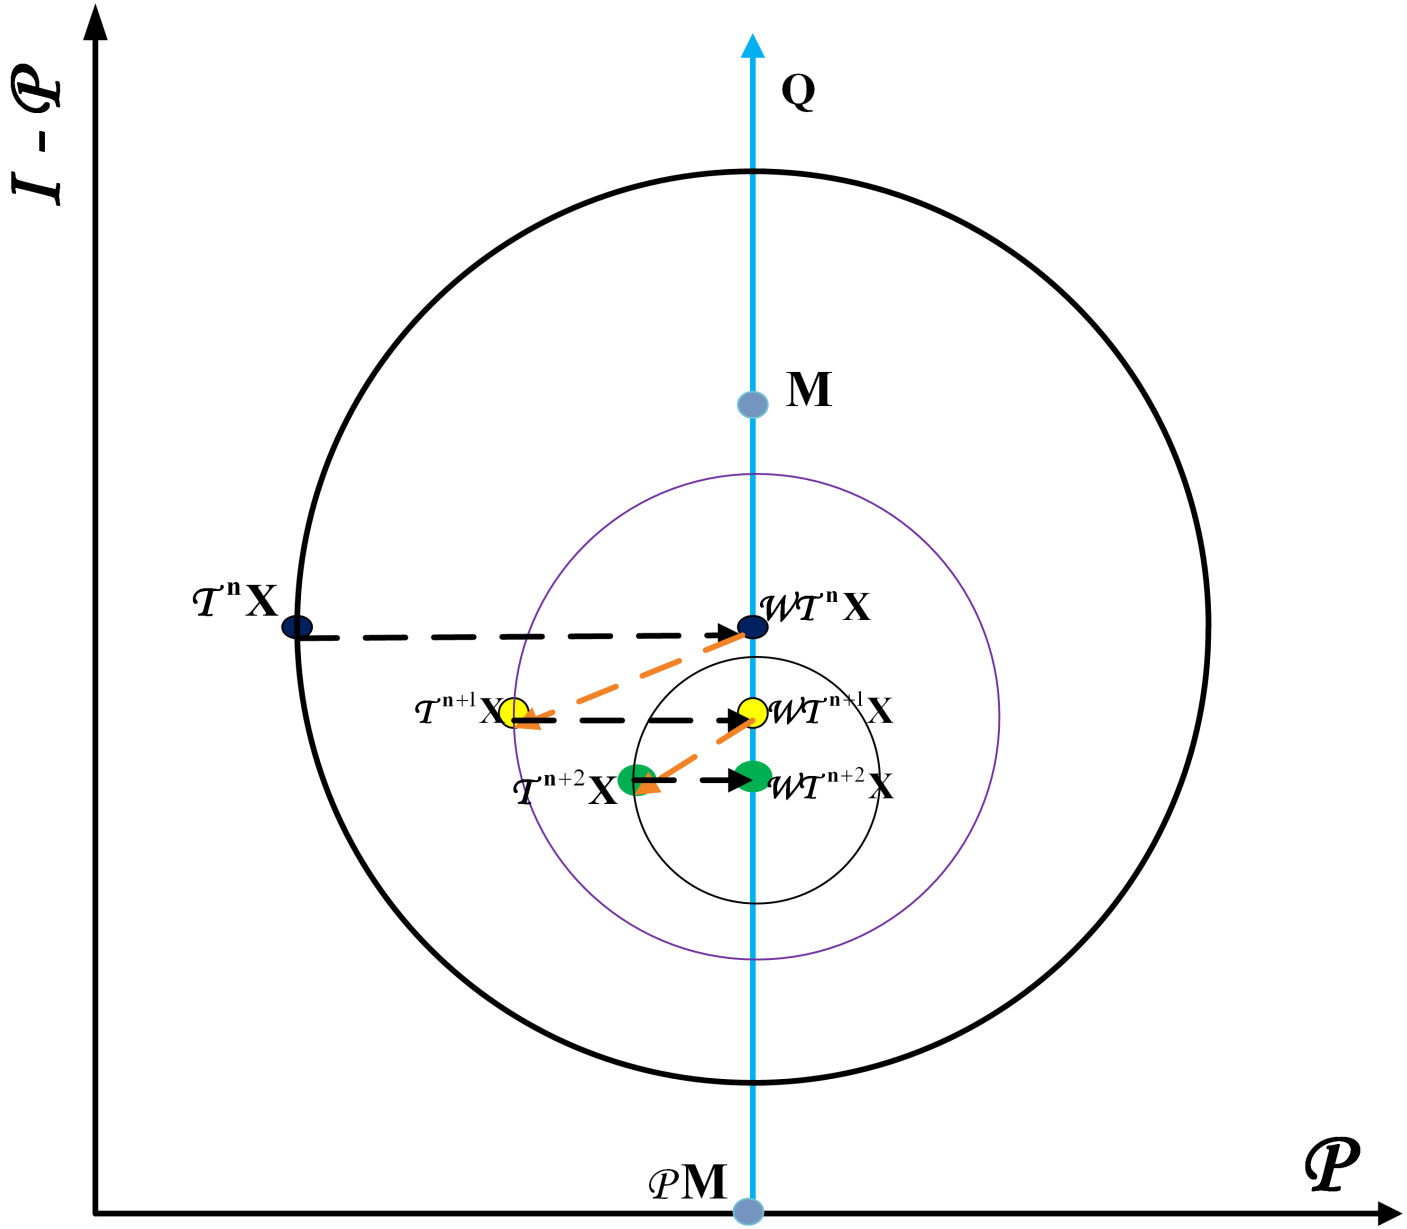

Solving Iteratively

We showed that the algorithm converges for $\mu_n = 1$ but is it optimal? Optimal step size is done by applying the Armijo rule in each iteration. In particular, it looks like that:

$$l[n] = \text{argmin}_{j \in \mathcal{Z}_{\geq 0}} : f(\mathbf{X}_{n,j}) \leq f(\mathbf{X}_n) - \sigma \text{trace}(\nabla f(\mathbf{X}_n)^T(\mathbf{X}_n - \mathbf{Z}_{n,j}))$$
$$\mathbf{Z}_{n,j} = \mathcal{D}(\mathbf{X}_n - \tilde{\mu} 2^{-j} \nabla f(\mathbf{X}^n))$$
$$\mu_n = \tilde{\mu} 2^{-l[n]}$$

Where $f(X) = \frac{1}{2}\|\mathcal{P}\mathbf{X} - \mathcal{P}\mathbf{M}\|_F^2$, $\tilde{\mu} > 0$ and $\sigma \in (0, 1)$ Therefore, finding the optimal step size is has significant computational cost, as it requires several applications of the $\mathcal{D}$.

Solving Iteratively

We now show that in this case, the optimal step size is given by $\mu_n = 1$.

Theorem (Optimal step size)

For the interest zone matrix approximation problem, the optimal step size is universal and given by $\mu = 1$

proof. Let $\mathbf{X}_n$ be a current point in the iterative process that satisfies the constraint (i.e. $n \geq 1$) and let $\mathbf{Q}$ be the geometric region of all the matrices $\mathbf{X}$ which satisfies $\|\mathcal{P}\mathbf{X} - \mathcal{P}\mathbf{M}\| = 0$. The geometric interpretation of an error for a given point $\mathbf{X}$ is the horizontal distance between $\mathbf{X}$ and $\mathbf{Q}$. Let $\mathbf{Y} = \mathbf{X}_n - \mu\mathcal{P}(\mathbf{X}_n - \mathbf{M})$ with $0 < \mu < 1$ and let $\tilde{\mathbf{Y}} = \mathbf{X}_n - \mathcal{P}(\mathbf{X}_n - \mathbf{M})$. Note that the difference between $\mathbf{Y}$ and $\tilde{\mathbf{Y}}$ is strictly on the $\mathcal{P}$ axis and that $\mathbf{Y}$ is between

Solving Iteratively

$\mathbf{X}_n$ and $\mathbf{Q}$. $\mathcal{D}$ maps $\mathbf{Y}$ to $\mathcal{D}\mathbf{Y}$ which is the best approximation to $\mathbf{Y}$ which satisfies the constraint. This point must be inside ball $\mathcal{B}_1$ centered at $\mathbf{Y}$ with radius $\|\mathbf{Y} - \mathbf{X_n}\|_F$. On the other hand, $\mathcal{D}\tilde{\mathbf{Y}}$ is in ball $\mathcal{B}_2$, centered in $\tilde{\mathbf{Y}}$ whose radius is $\|\tilde{\mathbf{Y}} - \mathbf{X}_n\|_F$ and is the best approximation to $\tilde{\mathbf{Y}}$. Because $\mathcal{D}\mathbf{Y}$ satisfies the constraint then $\mathbf{D}\tilde{\mathbf{Y}}$ must be inside a smaller ball, whose radius is $\|\tilde{\mathbf{Y}} - \mathcal{D}\mathbf{Y}\|_F$. Note that in ball $\mathcal{B}_3$ whose center is $\mathbf{Y}$ and its radius $\|\mathbf{Y} - \mathcal{D}\mathbf{Y}\|_F$, there are no points to satisfy the constraint, hence $\mathcal{D}\tilde{\mathbf{Y}} \notin \mathcal{B}_3$. Along with the fact that the line connecting $\mathbf{Y}$ and $\tilde{\mathbf{Y}}$ is parallel to the $\mathcal{P}$ axis, we get that $\|\mathcal{P}\mathcal{D}\tilde{\mathbf{Y}} - \mathcal{P}\mathbf{M}\|_F \leq \|\mathcal{P}\mathcal{D}\mathbf{Y} - \mathcal{P}\mathbf{M}\|_F$, which means that in every iteration, choosing $\mu < 1$ will lead to an error greater (or equal) to the error achieved for choosing $\mu = 1$. This completes the proof showing $\mu = 1$ is the best choice.

Solving Iteratively

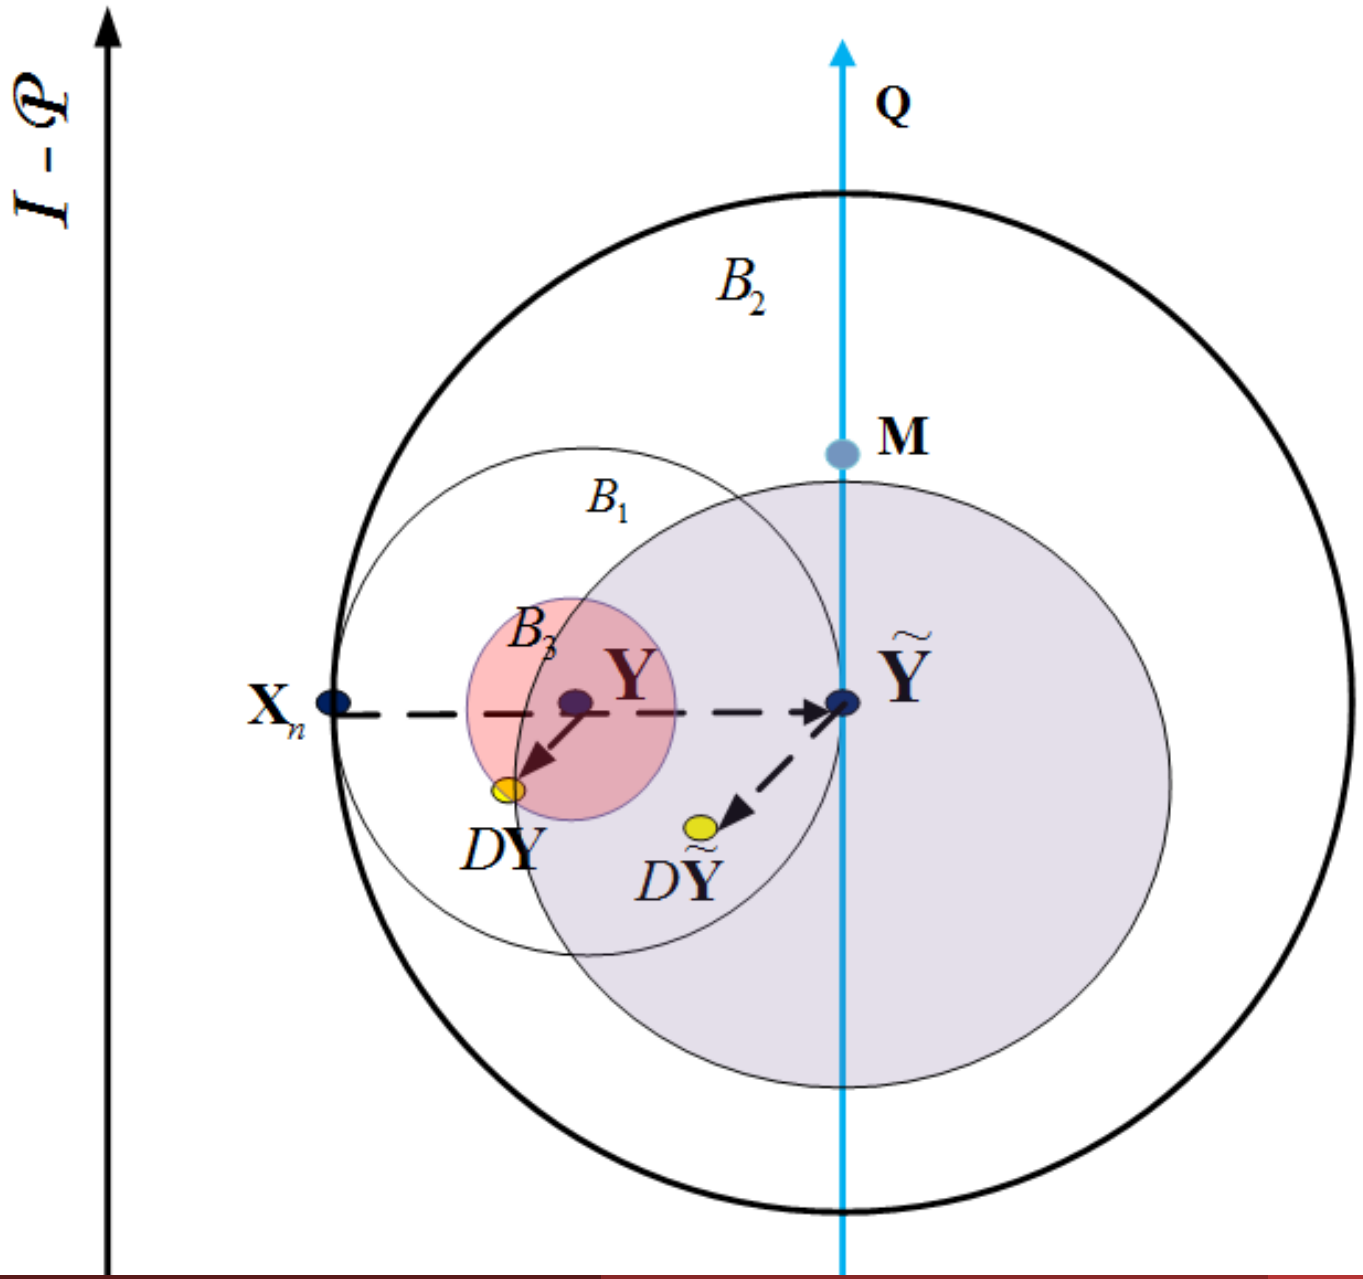

Solving Iteratively

Does the algorithm converge to the global solution? A new theorem for the convergence of the projected gradient states:

Theorem (Convergence to global solution)

If the projection is convex and orthogonal, then by applying the optimal step size in each iteration, the algorithm will converge to the optimal solution

In our case: $\mathcal{D}$ is a projection, it is convex (at least for the norms), we use the optimal step size, but something is missing... $\mathcal{D}$ is not orthogonal: $\mathcal{D} \neq \mathcal{D}^*$ However, this is overkill. For some cases convergence to global solution is guarantee and even if not, it finds the global solution with high probability. The power of this method lies in its versatility.

Using This Approach for Completion

Assuming we converge to the global solution, a simple and robust algorithm can be created using binary search. The minimum λ is 0 and the maximum is the norm of any matrix satisfying the constraint. Then, binary search is applied to find the matrix with the minimal value of λ that satisfies the constraints.

Example

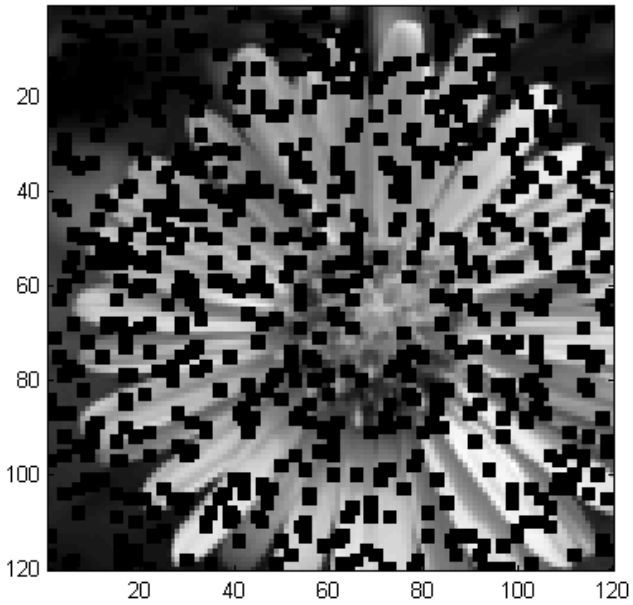

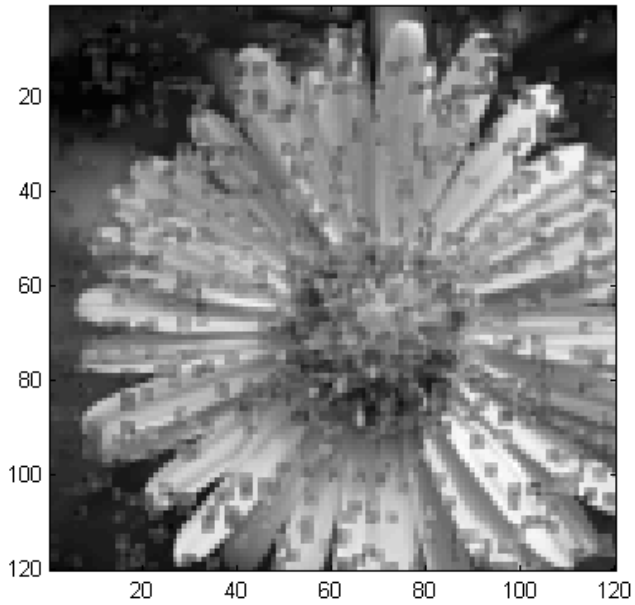

Co

Image vs Reconstructed Image using nuclear norm minimization

Singular Value Thresholding

A recently popular algorithm is the Singular Value Thresholding (SVT).

- We are interested in the solving the low rank matrix completion problem $\rightarrow$ NP-hard
- Instead, solve the easier nuclear norm minimization problem
- Given a matrix of rank r and size $n \times n$ to recover, then if the number of samples m taken uniformly satisfies:

$$m \geq Cn^{1.2} r \log n$$

Then the low rank matrix can be reconstructed with high probability.

Singular Value Thresholding

- Of course, the samples must be in a way that reconstruction is possible. For example, if no row or column is sampled, it is impossible to reconstruct the true matrix, even if it is rank 1, no matter what algorithm is used.
- However, the theorem is still important, since for low rank matrices, $\mathcal{O}(n^{1.2})$ entries is needed to achieve reconstruction with high probability, must less than $\mathcal{O}(n^2)$.
- Note that an $n \times n$ matrix of rank r depends on $r(2n - r)$ degrees of freedom

Shrinkage Operator

Consider the SVD of $\mathbf{X}$:

$$\mathbf{X} = \mathbf{U}\mathbf{S}\mathbf{V}^*$$

The shrinkage operator, $\mathcal{D}_\tau$ is defined by:

$$\mathcal{D}_\tau(\mathbf{X}) = \mathbf{U}\mathcal{D}_\tau(\mathbf{S})\mathbf{V}^*$$

Where $\mathcal{D}\mathbf{S} = \text{diag}(\sigma_i - \tau)_+$. In words, zeroing the singular values bigger than τ.

Singular Value Thresholding

Theorem

For each $\tau > 0$ *and* $\mathbf{Y} \in \mathcal{R}^{n_1 \times n_2}$ *the shrinkage operator obeys:*

$$\mathcal{D}_\tau(\mathbf{Y}) = argmin_{\mathbf{X}} \frac{1}{2}\|\mathbf{X} - \mathbf{Y}\|_F^2 + \tau\|\mathbf{X}\|_*$$

Proof (Sketch) The proof is based on the fact that $\frac{1}{2}\|\mathbf{X} - \mathbf{Y}\|_F^2 + \tau\|\mathbf{X}\|_*$ is strictly convex. Hence, it has a unique minimum. Showing the minimum is given by $\mathcal{D}_\tau$ is done using the subdifferential $\partial\|\mathbf{X}\|_*$

Singular Value Thresholding

The idea is to minimize

$$\frac{1}{2}\|\mathbf{X}\|_F^2 + \tau\|\mathbf{X}\|_*$$

such that $\mathcal{P}\mathbf{X} = \mathcal{P}\mathbf{M}$
As $\tau \to \infty$ we will get the solution we desire, of minimizing the nuclear norm.
The iterations are:

- $\mathbf{X}^k = \mathcal{D}_\tau(\mathbf{Y}^{k-1})$
- $\mathbf{Y}^k = \mathbf{Y}^{k-1} + \delta_k \mathcal{P}(\mathbf{M} - \mathbf{X}^k)$

Singular Value Thresholding

A few notes:

- The solution found by SVT is tend to have low rank. This is an empirical result but the reason is simple: zeroing the small singular values pushes the solution to have low rank
- Sparsity - If $\mathcal{P}$ is mostly zeros, the iterations can use this property to save storage (on the $\mathbf{Y}$ iterations)
- It is possible not to compute the entire SVD, since only the singular values bigger than τ are interesting. A good strategy is to use Lanczos method. Because $\mathbf{Y}$ is sparse, it can be applied to vectors rapidly.
- Can be extended to minimize the nuclear norm under different convex constraints.