# Scalable Algorithms for Tractable Schatten Quasi-Norm Minimization 

Fanhua Shang, Yuanyuan Liu, James Cheng<br>Department of Computer Science and Engineering, The Chinese University of Hong Kong<br>\{fhshang, yyliu, jcheng\} @cse.cuhk.edu.hk


#### Abstract

The Schatten- $p$ quasi-norm $(0<p<1)$ is usually used to replace the standard nuclear norm in order to approximate the rank function more accurately. However, existing Schatten$p$ quasi-norm minimization algorithms involve singular value decomposition (SVD) or eigenvalue decomposition (EVD) in each iteration, and thus may become very slow and impractical for large-scale problems. In this paper, we first define two tractable Schatten quasi-norms, i.e., the Frobenius/nuclear hybrid and bi-nuclear quasi-norms, and then prove that they are in essence the Schatten- $2 / 3$ and $1 / 2$ quasi-norms, respectively, which lead to the design of very efficient algorithms that only need to update two much smaller factor matrices. We also design two efficient proximal alternating linearized minimization algorithms for solving representative matrix completion problems. Finally, we provide the global convergence and performance guarantees for our algorithms, which have better convergence properties than existing algorithms. Experimental results on synthetic and real-world data show that our algorithms are more accurate than the state-of-the-art methods, and are orders of magnitude faster.


## Introduction

In recent years, the matrix rank minimization problem arises in a wide range of applications such as matrix completion, robust principal component analysis, low-rank representation, multivariate regression and multi-task learning. To solve such problems, Fazel, Hindi, and Boyd; Candès and Tao; Recht, Fazel, and Parrilo (2001; 2010; 2010) have suggested to relax the rank function by its convex envelope, i.e., the nuclear norm. In fact, the nuclear norm is equivalent to the $\ell_{1}$-norm on singular values of a matrix, and thus it promotes a low-rank solution. However, it has been shown in (Fan and Li 2001) that the $\ell_{1}$-norm regularization over-penalizes large entries of vectors, and results in a biased solution. By realizing the intimate relationship between them, the nuclear norm penalty also over-penalizes large singular values, that is, it may make the solution deviate from the original solution as the $\ell_{1}$-norm does (Nie, Huang, and Ding 2012; Lu et al. 2015). Compared with the nuclear norm, the Schatten- $p$ quasi-norm for $0<p<1$ makes a closer approximation to the rank function. Consequently, the

[^0]Schatten- $p$ quasi-norm minimization has attracted a great deal of attention in images recovery (Lu and Zhang 2014; Lu et al. 2014), collaborative filtering (Nie et al. 2012; Lu et al. 2015; Mohan and Fazel 2012) and MRI analysis (Majumdar and Ward 2011). In addition, many non-convex surrogate functions of the $\ell_{0}$-norm listed in (Lu et al. 2014; 2015) have been extended to approximate the rank function, such as SCAD (Fan and Li 2001) and MCP (Zhang 2010).

All non-convex surrogate functions mentioned above for low-rank minimization lead to some non-convex, nonsmooth, even non-Lipschitz optimization problems. Therefore, it is crucial to develop fast and scalable algorithms which are specialized to solve some alternative formulations. So far, Lai, Xu, and Yin (2013) proposed an iterative reweighted lease squares (IRucLq) algorithm to approximate the Schatten- $p$ quasi-norm minimization problem, and proved that the limit point of any convergent subsequence generated by their algorithm is a critical point. Moreover, Lu et al. (2014) proposed an iteratively reweighted nuclear norm (IRNN) algorithm to solve many non-convex surrogate minimization problems. For matrix completion problems, the Schatten- $p$ quasi-norm has been shown to be empirically superior to the nuclear norm (Marjanovic and Solo 2012). In addition, Zhang, Huang, and Zhang (2013) theoretically proved that the Schatten- $p$ quasi-norm minimization with small $p$ requires significantly fewer measurements than the convex nuclear norm minimization. However, all existing algorithms have to be solved iteratively and involve SVD or EVD in each iteration, which incurs high computational cost and is too expensive for solving large-scale problems (Cai and Osher 2013; Liu et al. 2014).

In contrast, as an alternative non-convex formulation of the nuclear norm, the bilinear spectral regularization as in (Srebro, Rennie, and Jaakkola 2004; Recht, Fazel, and Parrilo 2010) has been successfully applied in many largescale applications, e.g., collaborative filtering (Mitra, Sheorey, and Chellappa 2010). As the Schatten- $p$ quasi-norm is equivalent to the $\ell_{p}$ quasi-norm on singular values of a matrix, it is natural to ask the following question: can we design equivalent matrix factorization forms for the cases of the Schatten quasi-norm, e.g., $p=2 / 3$ or $1 / 2$ ?

In order to answer the above question, in this paper we first define two tractable Schatten quasi-norms, i.e., the Frobenius/nuclear hybrid and bi-nuclear quasi-norms. We
then prove that they are in essence the Schatten $-2 / 3$ and $1 / 2$ quasi-norms, respectively, for solving whose minimization we only need to perform SVDs on two much smaller factor matrices as contrary to the larger ones used in existing algorithms, e.g., IRNN. Therefore, our method is particularly useful for many "big data" applications that need to deal with large, high dimensional data with missing values. To the best of our knowledge, this is the first paper to scale Schatten quasi-norm solvers to the Netflix dataset. Moreover, we provide the global convergence and recovery performance guarantees for our algorithms. In other words, this is the best guaranteed convergence for algorithms that solve such challenging problems.

## Notations and Background

The Schatten- $p$ norm $(0<p<\infty)$ of a matrix $X \in \mathbb{R}^{m \times n}$ ( $m \geq n$ ) is defined as

$$
\|X\|_{S_{p}} \triangleq\left(\sum_{i=1}^{n} \sigma_{i}^{p}(X)\right)^{1 / p}
$$

where $\sigma_{i}(X)$ denotes the $i$-th singular value of $X$. When $p=1$, the Schatten-1 norm is the well-known nuclear norm, $\|X\|_{*}$. In addition, as the non-convex surrogate for the rank function, the Schatten- $p$ quasi-norm with $0<p<1$ is a better approximation than the nuclear norm (Zhang, Huang, and Zhang 2013) (analogous to the superiority of the $\ell_{p}$ quasinorm to the $\ell_{1}$-norm (Daubechies et al. 2010)).

To recover a low-rank matrix from some linear observations $b \in \mathbb{R}^{s}$, we consider the following general Schatten- $p$ quasi-norm minimization problem,

$$
\begin{equation*}
\min _{X} \lambda\|X\|_{S_{p}}^{p}+f(\mathcal{A}(X)-b) \tag{1}
\end{equation*}
$$

where $\mathcal{A}: \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{s}$ denotes the linear measurement operator, $\lambda>0$ is a regularization parameter, and the loss function $f(\cdot): \mathbb{R}^{s} \rightarrow \mathbb{R}$ generally denotes certain measurement for characterizing $\mathcal{A}(X)-b$. The above formulation can address a wide range of problems, such as matrix completion (Marjanovic and Solo 2012; Rohde and Tsybakov 2011) $\left(\mathcal{A}\right.$ is the sampling operator and $\left.f(\cdot)=\|\cdot\|_{2}^{2}\right)$, robust principal component analysis (Candès et al. 2011; Wang, Liu, and Zhang 2013; Shang et al. 2014) ( $\mathcal{A}$ is the identity operator and $f(\cdot)=\|\cdot\|_{1}$ ), and multivariate regression (Hsieh and Olsen 2014) $(\mathcal{A}(X)=A X$ with $A$ being a given matrix, and $f(\cdot)=\|\cdot\|_{F}^{2}$ ). Furthermore, $f(\cdot)$ may be also chosen as the Hinge loss in (Srebro, Rennie, and Jaakkola 2004) or the $\ell_{p}$ quasi-norm in (Nie et al. 2012).

Analogous to the $\ell_{p}$ quasi-norm, the Schatten- $p$ quasinorm is also non-convex for $p<1$, and its minimization is generally NP-hard (Lai, Xu, and Yin 2013). Therefore, it is crucial to develop efficient algorithms to solve some alternative formulations of Schatten- $p$ quasi-norm minimization (1). So far, only few algorithms (Lai, Xu, and Yin 2013; Mohan and Fazel 2012; Nie et al. 2012; Lu et al. 2014) have been developed to solve such problems. Furthermore, since all existing Schatten- $p$ quasi-norm minimization algorithms involve SVD or EVD in each iteration, they suffer from a high computational cost of $O\left(n^{2} m\right)$, which severely limits
their applicability to large-scale problems. Although there have been many efforts towards fast SVD or EVD computation such as partial SVD (Larsen 2005), the performance of those methods is still unsatisfactory for real-life applications (Cai and Osher 2013).

## Tractable Schatten Quasi-Norms

As in (Srebro, Rennie, and Jaakkola 2004), the nuclear norm has the following alternative non-convex formulations.
Lemma 1. Given a matrix $X \in \mathbb{R}^{m \times n}$ with $\operatorname{rank}(X)=$ $r \leq d$, the following holds:

$$
\begin{aligned}
\|X\|_{*} & =\min _{U \in \mathbb{R}^{m \times d}, V \in \mathbb{R}^{n \times d}: X=U V^{T}}\|U\|_{F}\|V\|_{F} \\
& =\min _{U, V: X=U V^{T}} \frac{\|U\|_{F}^{2}+\|V\|_{F}^{2}}{2} .
\end{aligned}
$$

## Frobenius/Nuclear Hybrid Quasi-Norm

Motivated by the equivalence relation between the nuclear norm and the bilinear spectral regularization (please refer to (Srebro, Rennie, and Jaakkola 2004; Recht, Fazel, and Parrilo 2010)), we define a Frobenius/nuclear hybrid (F/N) norm as fellows
Definition 1. For any matrix $X \in \mathbb{R}^{m \times n}$ with $\operatorname{rank}(X)=$ $r \leq d$, we can factorize it into two much smaller matrices $U \in \mathbb{R}^{m \times d}$ and $V \in \mathbb{R}^{n \times d}$ such that $X=U V^{T}$. Then the Frobenius/nuclear hybrid norm of $X$ is defined as

$$
\|X\|_{\mathrm{F} / \mathrm{N}}:=\min _{X=U V^{T}}\|U\|_{*}\|V\|_{F}
$$

In fact, the Frobenius/nuclear hybrid norm is not a real norm, because it is non-convex and does not satisfy the triangle inequality of a norm. Similar to the well-known Schatten- $p$ quasi-norm $(0<p<1)$, the Frobenius/nuclear hybrid norm is also a quasi-norm, and their relationship is stated in the following theorem.
Theorem 1. The Frobenius/nuclear hybrid norm $\|\cdot\|_{\mathrm{F} / \mathrm{N}}$ is a quasi-norm. Surprisingly, it is also the Schatten-2/3 quasinorm, i.e.,

$$
\|X\|_{\mathrm{F} / \mathbf{N}}=\|X\|_{S_{2 / 3}},
$$

where $\|X\|_{S_{2 / 3}}$ denotes the Schatten-2/3 quasi-norm of $X$.
Property 1. For any matrix $X \in \mathbb{R}^{m \times n}$ with $\operatorname{rank}(X)=$ $r \leq d$, the following holds:

$$
\begin{aligned}
\|X\|_{\mathrm{F} / \mathrm{N}} & =\min _{U \in \mathbb{R}^{m \times d}, V \in \mathbb{R}^{n \times d}: X=U V^{T}}\|U\|_{*}\|V\|_{F} \\
& =\min _{X=U V^{T}}\left(\frac{2\|U\|_{*}+\|V\|_{F}^{2}}{3}\right)^{3 / 2}
\end{aligned}
$$

The proofs of Property 1 and Theorem 1 can be found in the Supplementary Materials.

## Bi-Nuclear Quasi-Norm

Similar to the definition of the above Frobenius/nuclear hybrid norm, our bi-nuclear (BiN) norm is naturally defined as follows.

Definition 2. For any matrix $X \in \mathbb{R}^{m \times n}$ with $\operatorname{rank}(X)=$ $r \leq d$, we can factorize it into two much smaller matrices $U \in \mathbb{R}^{m \times d}$ and $V \in \mathbb{R}^{n \times d}$ such that $X=U V^{T}$. Then the bi-nuclear norm of $X$ is defined as

$$
\|X\|_{\mathrm{BiN}}:=\min _{X=U V^{T}}\|U\|_{*}\|V\|_{*} .
$$

Similar to the Frobenius/nuclear hybrid norm, the binuclear norm is also a quasi-norm, as stated in the following theorem.

Theorem 2. The bi-nuclear norm $\|\cdot\|_{\text {BiN }}$ is a quasi-norm. In addition, it is also the Schatten-1/2 quasi-norm, i.e.,

$$
\|X\|_{\mathrm{BiN}}=\|X\|_{S_{1 / 2}}
$$

The proof of Theorem 2 can be found in the Supplementary Materials. Due to the relationship between the binuclear quasi-norm and the Schatten- $1 / 2$ quasi-norm, it is easy to verify that the bi-nuclear quasi-norm possesses the following properties.
Property 2. For any matrix $X \in \mathbb{R}^{m \times n}$ with $\operatorname{rank}(X)=$ $r \leq d$, the following holds:

$$
\begin{aligned}
\|X\|_{\mathrm{BiN}} & =\min _{X=U V^{T}}\|U\|_{*}\|V\|_{*}=\min _{X=U V^{T}} \frac{\|U\|_{*}^{2}+\|V\|_{*}^{2}}{2} \\
& =\min _{X=U V^{T}}\left(\frac{\|U\|_{*}+\|V\|_{*}}{2}\right)^{2} .
\end{aligned}
$$

The following relationship between the nuclear norm and the Frobenius norm is well known: $\|X\|_{F} \leq\|X\|_{*} \leq$ $\sqrt{r}\|X\|_{F}$. Similarly, the analogous bounds hold for the Frobenius/nuclear hybrid and bi-nuclear quasi-norms, as stated in the following property.
Property 3. For any matrix $X \in \mathbb{R}^{m \times n}$ with $\operatorname{rank}(X)=r$, the following inequalities hold:

$$
\begin{gathered}
\|X\|_{*} \leq\|X\|_{\mathrm{F} / \mathrm{N}} \leq \sqrt{r}\|X\|_{*}, \\
\|X\|_{*} \leq\|X\|_{\mathrm{F} / \mathrm{N}} \leq\|X\|_{\mathrm{BiN}} \leq r\|X\|_{*} .
\end{gathered}
$$

The proof of Property 3 can be found in the Supplementary Materials. It is easy to see that Property 3 in turn implies that any low Frobenius/nuclear hybrid or bi-nuclear norm approximation is also a low nuclear norm approximation.

## Optimization Algorithms

## Problem Formulations

To bound the Schatten- $2 / 3$ or $-1 / 2$ quasi-norm of $X$ by $\frac{1}{3}\left(2\|U\|_{*}+\|V\|_{F}^{2}\right)$ or $\frac{1}{2}\left(\|U\|_{*}+\|V\|_{*}\right)$, we mainly consider the following general structured matrix factorization formulation as in (Haeffele, Young, and Vidal 2014),

$$
\begin{equation*}
\min _{U, V} \lambda \varphi(U, V)+f\left(\mathcal{A}\left(U V^{T}\right)-b\right) \tag{2}
\end{equation*}
$$

where the regularization term $\varphi(U, V)$ denotes $\frac{1}{3}\left(2\|U\|_{*}+\right.$ $\left.\|V\|_{F}^{2}\right)$ or $\frac{1}{2}\left(\|U\|_{*}+\|V\|_{*}\right)$.

As mentioned above, there are many Schatten- $p$ quasinorm minimization problems for various real-world applications. Therefore, we propose two efficient algorithms to solve the following low-rank matrix completion problems:

$$
\begin{equation*}
\min _{U, V} \frac{\lambda\left(2\|U\|_{*}+\|V\|_{F}^{2}\right)}{3}+\frac{1}{2}\left\|\mathcal{P}_{\Omega}\left(U V^{T}\right)-\mathcal{P}_{\Omega}(D)\right\|_{F}^{2}, \tag{3}
\end{equation*}
$$

$$
\begin{equation*}
\min _{U, V} \frac{\lambda\left(\|U\|_{*}+\|V\|_{*}\right)}{2}+\frac{1}{2}\left\|\mathcal{P}_{\Omega}\left(U V^{T}\right)-\mathcal{P}_{\Omega}(D)\right\|_{F}^{2} \tag{4}
\end{equation*}
$$

where $\mathcal{P}_{\Omega}$ denotes the linear projection operator, i.e., $\mathcal{P}_{\Omega}(D)_{i j}=D_{i j}$ if $(i, j) \in \Omega$, and $\mathcal{P}_{\Omega}(D)_{i j}=0$ otherwise. Due to the operator $\mathcal{P}_{\Omega}$ in (3) and (4), we usually need to introduce some auxiliary variables for solving them. To avoid introducing auxiliary variables, motivated by the proximal alternating linearized minimization (PALM) method proposed in (Bolte, Sabach, and Teboulle 2014), we propose two fast PALM algorithms to efficiently solve (3) and (4). The space limitation refrains us from fully describing each algorithm, but we try to give enough details of a representative algorithm for solving (3) and discussing their differences.

## Updating $U_{k+1}$ and $V_{k+1}$ with Linearization Techniques

Let $g_{k}(U):=\left\|\mathcal{P}_{\Omega}\left(U V_{k}^{T}\right)-\mathcal{P}_{\Omega}(D)\right\|_{F}^{2} / 2$, and then its gradient is Lipschitz continuous with constant $l_{k+1}^{g}$, meaning that $\left\|\nabla g_{k}\left(U_{1}\right)-\nabla g_{k}\left(U_{2}\right)\right\|_{F} \leq l_{k+1}^{g}\left\|U_{1}-U_{2}\right\|_{F}$ for any $U_{1}, U_{2} \in \mathbb{R}^{m \times d}$. By linearizing $g_{k}(U)$ at $U_{k}$ and adding a proximal term, then we have the following approximation:

$$
\begin{equation*}
\widehat{g_{k}}\left(U, U_{k}\right)=g_{k}\left(U_{k}\right)+\left\langle\nabla g_{k}\left(U_{k}\right), U-U_{k}\right\rangle+\frac{l_{k+1}^{g}}{2}\left\|U-U_{k}\right\|_{F}^{2} \tag{5}
\end{equation*}
$$

Thus, we have

$$
\begin{align*}
& U_{k+1}=\underset{U}{\arg \min } \frac{2 \lambda}{3}\|U\|_{*}+\widehat{g_{k}}\left(U, U_{k}\right) \\
= & \underset{U}{\arg \min } \frac{2 \lambda}{3}\|U\|_{*}+\frac{l_{k+1}^{g}}{2}\left\|U-U_{k}+\frac{\nabla g_{k}\left(U_{k}\right)}{l_{k+1}^{g}}\right\|_{F}^{2} . \tag{6}
\end{align*}
$$

Similarly, we have

$$
\begin{equation*}
V_{k+1}=\underset{V}{\arg \min } \frac{\lambda}{3}\|V\|_{F}^{2}+\frac{l_{k+1}^{h}}{2}\left\|V-V_{k}+\nabla h_{k}\left(V_{k}\right) / l_{k+1}^{h}\right\|_{F}^{2}, \tag{7}
\end{equation*}
$$

where $h_{k}(V):=\left\|\mathcal{P}_{\Omega}\left(U_{k+1} V^{T}\right)-\mathcal{P}_{\Omega}(D)\right\|_{F}^{2} / 2$ with the Lipschitz constant $l_{k+1}^{h}$. The problems (6) and (7) are known to have closed-form solutions, which of the former is given by the so-called matrix shrinkage operator (Cai, Candès, and Shen 2010). In contrast, for solving (4), $U_{k+1}$ is computed in the same way as (6), and $V_{k+1}$ is given by

$$
\begin{equation*}
V_{k+1}=\underset{V}{\arg \min } \frac{\lambda}{2}\|V\|_{*}+\frac{l_{k+1}^{h}}{2}\left\|V-V_{k}+\nabla h_{k}\left(V_{k}\right) / l_{k+1}^{h}\right\|_{F}^{2} \tag{8}
\end{equation*}
$$

## Updating Lipschitz Constants

Next we compute the Lipschitz constants $l_{k+1}^{g}$ and $l_{k+1}^{h}$ at the $(k+1)$-iteration.

$$
\begin{aligned}
& \quad\left\|\nabla g_{k}\left(U_{1}\right)-\nabla g_{k}\left(U_{2}\right)\right\|_{F}=\left\|\left[\mathcal{P}_{\Omega}\left(U_{1} V_{k}^{T}-U_{2} V_{k}^{T}\right)\right] V_{k}\right\|_{F} \\
& \leq\left\|V_{k}\right\|_{2}^{2}\left\|U_{1}-U_{2}\right\|_{F}, \\
& \left\|\nabla h_{k}\left(V_{1}\right)-\nabla h_{k}\left(V_{2}\right)\right\|_{F}=\left\|U_{k+1}^{T}\left[\mathcal{P}_{\Omega}\left(U_{k+1}\left(V_{1}^{T}-V_{2}^{T}\right)\right)\right]\right\|_{F} \\
& \leq\left\|U_{k+1}\right\|_{2}^{2}\left\|V_{1}-V_{2}\right\|_{F} .
\end{aligned}
$$

Hence, both Lipschitz constants are updated by

$$
\begin{equation*}
l_{k+1}^{g}=\left\|V_{k}\right\|_{2}^{2} \text { and } l_{k+1}^{h}=\left\|U_{k+1}\right\|_{2}^{2} \tag{9}
\end{equation*}
$$

```
Algorithm 1 Solving (3) via PALM
Input: \(\mathcal{P}_{\Omega}(D)\), the given rank \(d\) and \(\lambda\).
Initialize: \(U_{0}, V_{0}, \varepsilon\) and \(k=0\).
    while not converged do
        Update \(l_{k+1}^{g}\) and \(U_{k+1}\) by (9) and (6), respectively.
        Update \(l_{k+1}^{h}\) and \(V_{k+1}\) by (9) and (7), respectively.
        Check the convergence condition,
            \(\max \left\{\left\|U_{k+1}-U_{k}\right\|_{F},\left\|V_{k+1}-V_{k}\right\|_{F}\right\}<\varepsilon\).
    end while
Output: \(U_{k+1}, V_{k+1}\).
```


## PALM Algorithms

Based on the above development, our algorithm for solving (3) is given in Algorithm 1. Similarly, we also design an efficient PALM algorithm for solving (4). The running time of Algorithm 1 is dominated by performing matrix multiplications. The total time complexity of Algorithm 1, as well as the algorithm for solving (4), is $O(n m d)$, where $d \ll m, n$.

## Algorithm Analysis

We now provide the global convergence and low-rank matrix recovery guarantees for Algorithm 1, and the similar results can be obtained for the algorithm for solving (4).

## Global Convergence

Before analyzing the global convergence of Algorithm 1, we first introduce the definition of the critical points of a non-convex function given in (Bolte, Sabach, and Teboulle 2014).

Definition 3. Let a non-convex function $f: \mathbb{R}^{n} \rightarrow$ $(-\infty,+\infty]$ be a proper and lower semi-continuous function, and $\operatorname{dom} f=\left\{x \in \mathbb{R}^{n}: f(x)<+\infty\right\}$.

- For any $x \in \operatorname{dom} f$, the Frèchet sub-differential of $f$ at $x$ is defined as

$$
\widehat{\partial} f(x)=\left\{u \in \mathbb{R}^{n}: \lim _{y \neq x} \inf _{y \rightarrow x} \frac{f(y)-f(x)-\langle u, y-x\rangle}{\|y-x\|_{2}} \geq 0\right\}
$$

$$
\text { and } \widehat{\partial} f(x)=\emptyset \text { if } x \notin \operatorname{dom} f \text {. }
$$

- The limiting sub-differential of $f$ at $x$ is defined as

$$
\begin{aligned}
\partial f(x)= & \left\{u \in \mathbb{R}^{n}: \exists x^{k} \rightarrow x, f\left(x^{k}\right) \rightarrow f(x)\right. \\
& \text { and } \left.u^{k} \in \widehat{\partial} f\left(x^{k}\right) \rightarrow u \text { as } k \rightarrow \infty\right\} .
\end{aligned}
$$

- The points whose sub-differential contains 0 are called critical points. For instance, the point $x$ is a critical point of $f$ if $0 \in \partial f(x)$.
Theorem 3 (Global Convergence). Let $\left\{\left(U_{k}, V_{k}\right)\right\}$ be a sequence generated by Algorithm 1, then it is a Cauchy sequence and converges to a critical point of (3).

The proof of the theorem can be found in the Supplementary Materials. Theorem 3 shows the global convergence of Algorithm 1. We emphasize that, different from the general subsequence convergence property, the global convergence property is given by $\left(U_{k}, V_{k}\right) \rightarrow(\widehat{U}, \widehat{V})$ as the number of iteration $k \rightarrow+\infty$, where $(\widehat{U}, \widehat{V})$ is a critical point
of (3). As we have stated, existing algorithms for solving the non-convex and non-smooth problem, such as IRucLq and IRNN, have only subsequence convergence ( Xu and Yin 2014). According to (Attouch and Bolte 2009), we know that the convergence rate of Algorithm 1 is at least sub-linear, as stated in the following theorem.
Theorem 4 (Convergence Rate). The sequence $\left\{\left(U_{k}, V_{k}\right)\right\}$ generated by Algorithm 1 converges to a critical point $(\widehat{U}, \widehat{V})$ of (3) at least in the sub-linear convergence rate, that is, there exists $C>0$ and $\theta \in(1 / 2,1)$ such that

$$
\left\|\left[U_{k}^{T}, V_{k}^{T}\right]-\left[\widehat{U}^{T}, \widehat{V}^{T}\right]\right\|_{F} \leq C k^{-\frac{1-\theta}{2 \theta-1}}
$$

## Recovery Guarantee

In the following, we show that when sufficiently many entries are observed, the critical point generated by our algorithms recovers a low-rank matrix "close to" the groundtruth one. Without loss of generality, assume that $D=$ $Z+E \in \mathbb{R}^{m \times n}$, where $Z$ is a true matrix, and $E$ denotes a random gaussian noise.
Theorem 5. Let $(\widehat{U}, \widehat{V})$ be a critical point of the problem (3) with given rank $d$, and $m \geq n$. Then there exists an absolute constant $C_{1}$, such that with probability at least $1-$ $2 \exp (-m)$,
$\frac{\left\|Z-\widehat{U} \widehat{V}^{T}\right\|_{F}}{\sqrt{m n}} \leq \frac{\|E\|_{F}}{\sqrt{m n}}+C_{1} \beta\left(\frac{m d \log (m)}{|\Omega|}\right)^{1 / 4}+\frac{2 \sqrt{d} \lambda}{3 C_{2} \sqrt{|\Omega|}}$,
where $\beta=\max _{i, j}\left|D_{i, j}\right|$ and $C_{2}=\frac{\left\|\mathcal{P}_{\Omega}\left(D-\hat{U} \hat{V}^{T}\right) \hat{V}\right\|_{F}}{\left\|\mathcal{P}_{\Omega}\left(D-\hat{U} \hat{V}^{T}\right)\right\|_{F}}$.
The proof of the theorem and the analysis of lowerboundedness of $C_{2}$ can be found in the Supplementary Materials. When the samples size $|\Omega| \gg m d \log (m)$, the second and third terms diminish, and the recovery error is essentially bounded by the "average" magnitude of entries of the noise matrix $E$. In other words, only $O(m d \log (m))$ observed entries are needed, which is significantly lower than $O\left(m r \log ^{2}(m)\right)$ in standard matrix completion theories (Candès and Recht 2009; Keshavan, Montanari, and Oh 2010; Recht 2011). We will confirm this result by our experiments in the following section.

## Experimental Results

We now evaluate both the effectiveness and efficiency of our algorithms for solving matrix completion problems, such as collaborative filtering and image recovery. All experiments were conducted on an Intel Xeon E7-4830V2 2.20GHz CPU with 64G RAM.

Algorithms for Comparison We compared our algorithms, BiN and $\mathrm{F} / \mathrm{N}$, with the following state-of-the-art methods: $\mathrm{IRucLq}^{1}$ (Lai, Xu, and Yin 2013): In IRucLq, $p$ varies from 0.1 to 1 with increment 0.1 , and the parameters $\lambda$ and $\alpha$ are set to $10^{-6}$ and 0.9 , respectively. In addition, the rank parameter of the algorithm is updated dynamically as in (Lai, Xu, and Yin 2013), that is, it only needs to compute the partial EVD. IRNN ${ }^{2}$ (Lu et al. 2014): We choose

[^1]

Figure 1: The recovery accuracy of IRucLq, IRNN and our algorithms on noisy random matrices of size $100 \times 100$.


Figure 2: The running time (seconds) and RSE results vs. sizes of matrices (left) and noise factors (right).
the $\ell_{p}$-norm, SCAD and MCP penalties as the regularization term among eight non-convex penalty functions, where $p$ is chosen from the range of $\{0.1,0.2, \ldots, 1\}$. At each iteration, the parameter $\lambda$ is dynamically decreased by $\lambda_{k}=0.7 \lambda_{k-1}$, where $\lambda_{0}=10\left\|\mathcal{P}_{\Omega}(D)\right\|_{\infty}$.

For our algorithms, we set the regularization parameter $\lambda=5$ or $\lambda=100$ for noisy synthetic and real-world data, respectively. Note that the rank parameter $d$ is estimated by the strategy in (Wen, Yin, and Zhang 2012). In addition, we evaluate the performance of matrix recovery by the relative squared error ( RSE ) and the root mean square error (RMSE), i.e., RSE $:=\|X-Z\|_{F} /\|Z\|_{F}$ and RMSE $:=$ $\frac{1}{|T|} \sqrt{\Sigma_{(i, j) \in T}\left(X_{i j}-D_{i j}\right)^{2}}$, where $T$ is the test set.

## Synthetic Matrix Completion

The synthetic matrices $Z \in \mathbb{R}^{m \times n}$ with rank $r$ are generated by the following procedure: the entries of both $U \in \mathbb{R}^{m \times r}$ and $V \in \mathbb{R}^{n \times r}$ are first generated as independent and identically distributed (i.i.d.) numbers, and then $Z=U V^{T}$ is assembled. Since all these algorithms have very similar recovery performance on noiseless matrices, we only conducted experiments on noisy matrices with different noise levels, i.e., $\mathcal{P}_{\Omega}(D)=\mathcal{P}_{\Omega}(Z+n f * E)$, where $n f$ denotes the noise factor. In other worlds, the observed subset is corrupted by i.i.d. standard Gaussian random noise as in (Lu et al. 2014). In addition, only $20 \%$ or $30 \%$ entries of $D$ are sampled uniformly at random as training data, i.e., sampling ratio $(\mathrm{SR})=20 \%$ or $30 \%$. The rank parameter $d$ of our algorithms is set to $\lfloor 1.25 r\rfloor$ as in (Wen, Yin, and Zhang 2012).

The average RSE results of 100 independent runs on noisy random matrices are shown in Figure 1, which shows that

Table 1: Testing RMSE on MovieLens1M, MovieLens10M and Netflix.

| Datasets | MovieLens1M | MovieLens10M | Netflix |
| :--- | :---: | :---: | :---: |
| $\%$ SR | $50 \% / 70 \% / 90 \%$ | $50 \% / 70 \% / 90 \%$ | $50 \% / 70 \% / 90 \%$ |
| APGL | $1.2564 / 1.1431 / 0.9897$ | $1.1138 / 0.9455 / 0.8769$ | $1.0806 / 0.9885 / 0.9370$ |
| LMaFit | $0.9138 / 0.9019 / 0.8845$ | $0.8705 / 0.8496 / 0.8244$ | $0.9062 / 0.8923 / 0.8668$ |
| IRucLq | $0.9099 / 0.8918 / 0.8786$ | $-/-/-$ | $-/-/-$ |
| IRNN | $0.9418 / 0.9275 / 0.9032$ | $-/-/-$ | $-/-/-$ |
| BiN | $\mathbf{0 . 8 7 4 1 / 0 . 8 5 9 3 / 0 . 8 4 8 5}$ | $0.8274 / 0.8115 / 0.7989$ | $0.8650 / 0.8487 / 0.8413$ |
| F/N | $0.8764 / \mathbf{0 . 8 5 6 2 / \mathbf { 0 . 8 4 4 1 }}$ | $\mathbf{0 . 8 1 5 8 / 0 . 8 0 2 1 / \mathbf { 0 . 7 9 2 1 }}$ | $\mathbf{0 . 8 6 1 8 / 0 . 8 4 5 9 / 0 . 8 4 0 4}$ |

if $p$ varies from 0.1 to 0.7 , IRucLq and IRNN-Lp achieve similar recovery performance as IRNN-SCAD, IRNN-MCP and our algorithms; otherwise, IRucLq and IRNN-Lp usually perform much worse than the other four methods, especially $p=1$. We also report the running time of all the methods with $20 \%$ SR as the size of noisy matrices increases, as shown in Figure 2. Moreover, we present the RSE results of those methods and $\mathrm{APGL}^{3}$ (Toh and Yun 2010) (which is one of the nuclear norm solvers) with different noise factors. Figure 2 shows that our algorithms are significantly faster than the other methods, while the running time of IRucLq and IRNN increases dramatically when the size of matrices increases, and they could not yield experimental results within 48 hours when the size of matrices is $50,000 \times 50,000$. This further justifies that both our algorithms have very good scalability and can address large-scale problems. In addition, with only $20 \%$ SR, all Schatten quasi-norm methods significantly outperform APGL in terms of RSE.

## Collaborative Filtering

We tested our algorithms on three real-world recommendation system data sets: the MovieLens1M, MovieLens10M ${ }^{4}$ and Netflix datasets (KDDCup 2007). We randomly chose $50 \%, 70 \%$ and $90 \%$ as the training set and the remaining as the testing set, and the experimental results are reported over 10 independent runs. In addition to the methods used above, we also compared our algorithms with one of the fastest existing methods, LMaFit ${ }^{5}$ (Wen, Yin, and Zhang 2012). The testing RMSE of all these methods on the three

[^2]

Figure 3: The testing RMSE of LMaFit and our algorithms with ranks varying from 5 to 20 and $70 \%$ SR.


Figure 4: The running time (seconds) on three data sets (left, best viewed in colors) and MovieLens1M (right).
data sets is reported in Table 1, which shows that all those methods with non-convex penalty functions perform significantly better than the convex nuclear norm solver, APGL. In addition, our algorithms consistently outperform the other methods in terms of prediction accuracy. This further confirms that our two Schatten quasi-norm regularized models can provide a good estimation of a low-rank matrix. Moreover, we report the average testing RMSE and running time of our algorithms on these three data sets in Figures 3 and 4, where the rank varies from 5 to 20 and SR is set to $70 \%$. Note that IRucLq and IRNN-Lp could not run on the two larger data sets due to runtime exceptions. It is clear that our algorithms are much faster than AGPL, IRucLq and IRNNLp on all these data sets. They perform much more robust with respect to ranks than LMaFit, and are comparable in speed with it. This shows that our algorithms have very good scalability and are suitable for real-world applications.

## Image Recovery

We also applied our algorithms to gray-scale image recovery on the Boat image of size $512 \times 512$, where $50 \%$ of pixels in the input image were replaced by random Gaussian noise, as shown in Figure 5(b). In addition, we employed the well known peak signal-to-noise ratio (PSNR) to measure the recovery performance. The rank parameter of our algorithms and IRucLq was set to 100. Due to limited space, we only report the best results (PSNR and CPU time) of APGL, LMaFit, IRucLq and IRNN-Lp in Figure 5, which shows that our two algorithms achieve much better recovery performance than the other methods in terms of PSNR. And impressively, both our algorithms are significantly faster than the other methods except LMaFit and at least 70 times faster


Figure 5: Comparison of image recovery on the Boat image of size $512 \times 512$ : (a) Original image; (b) Image with Gaussian noise; (c) APGL (PSNR: 24.93, Time: 15.47 sec ); (d) LMaFit (PSNR: 25.89, Time: 6.95 sec ); (e) IRucLq (PSNR: 26.36, Time: 805.81 sec ); (f) IRNN-Lp (PSNR: 26.21, Time: 943.28 sec ); (g) BiN (PSNR: 26.94, Time: 8.93 sec ); (h) F/N (PSNR: 27.62, Time: 10.80 sec ).
than IRucLq and IRNN-Lp.

## Conclusions

In this paper we defined two tractable Schatten quasi-norms, i.e., the Frobenius/nuclear hybrid and bi-nuclear quasinorms, and proved that they are in essence the Schatten$2 / 3$ and $1 / 2$ quasi-norms, respectively. Then we designed two efficient proximal alternating linearized minimization algorithms to solve our Schatten quasi-norm minimization for matrix completion problems, and also proved that each bounded sequence generated by our algorithms globally converges to a critical point. In other words, our algorithms not only have better convergence properties than existing algorithms, e.g., IRucLq and IRNN, but also reduce the computational complexity from $O\left(m n^{2}\right)$ to $O(m n d)$, with $d$ being the estimated rank $(d \ll m, n)$. We also provided the recovery guarantee for our algorithms, which implies that they need only $O(m d \log (m))$ observed entries to recover a lowrank matrix with high probability. Our experiments showed that our algorithms outperform the state-of-the-art methods in terms of both efficiency and effectiveness.

## Acknowledgements

We thank the reviewers for their constructive comments. The authors are partially supported by the SHIAE fund 8115048 and the Hong Kong GRF 2150851.

## References

Attouch, H., and Bolte, J. 2009. On the convergence of the proximal algorithm for nonsmooth functions involving analytic features. Math. Program. 116(1-2):5-16.
Bolte, J.; Sabach, S.; and Teboulle, M. 2014. Proximal alternating linearized minimization for nonconvex and nonsmooth problems. Math. Program. 146:459-494.
Cai, J.-F., and Osher, S. 2013. Fast singular value thresholding without singular value decomposition. Methods Anal. appl. 20(4):335-352.
Cai, J.-F. C.; Candès, E.; and Shen, Z. 2010. A singular value thresholding algorithm for matrix completion. SIAM J. Optim. 20(4):1956-1982.
Candès, E., and Recht, B. 2009. Exact matrix completion via convex optimization. Found. Comput. Math. 9(6):717-772.
Candès, E., and Tao, T. 2010. The power of convex relaxation: Near-optimal matrix completion. IEEE Trans. Inf. Theory 56(5):2053-2080.
Candès, E.; Li, X.; Ma, Y.; and Wright, J. 2011. Robust principal component analysis? J. ACM 58(3):1-37.
Daubechies, I.; DeVore, R.; Fornasier, M.; and Guntuk, C. 2010. Iteratively reweighted least squares minimization for sparse recovery. Commun. Pure Appl. Math. 63:1-38.
Fan, J., and Li, R. 2001. Variable selection via nonconcave penalized likelihood and its Oracle properties. J. Am. Statist. Assoc. 96:1348-1361.
Fazel, M.; Hindi, H.; and Boyd, S. 2001. A rank minimization heuristic with application to minimum order system approximation. In Proc. IEEE Amer. Control Conf., 4734-4739.
Haeffele, B. D.; Young, E. D.; and Vidal, R. 2014. Structured low-rank matrix factorization: Optimality, algorithm, and applications to image processing. In Proc. 31st Int. Conf. Mach. Learn. (ICML), 2007-2015.
Hsieh, C.-J., and Olsen, P. A. 2014. Nuclear norm minimization via active subspace selection. In Proc. 31st Int. Conf. Mach. Learn. (ICML), 575-583.
KDDCup. 2007. ACM SIGKDD and Netflix. In Proc. KDD Cup and Workshop.
Keshavan, R.; Montanari, A.; and Oh, S. 2010. Matrix completion from a few entries. IEEE Trans. Inf. Theory 56(6):29802998.

Lai, M.; Xu, Y.; and Yin, W. 2013. Improved iteratively rewighted least squares for unconstrained smoothed $\ell_{p}$ minimization. SIAM J. Numer. Anal. 51(2):927-957.
Larsen, R. 2005. PROPACK-software for large and sparse SVD calculations.
Liu, Y.; Shang, F.; Cheng, H.; and Cheng, J. 2014. A Grassmannian manifold algorithm for nuclear norm regularized least squares problems. In Proc. 30th Conf. Uncert. in Art. Intel. (UAI), 515-524.
Lu, Z., and Zhang, Y. 2014. Iterative reweighted singular value minimization methods for $\ell_{p}$ regularized unconstrained matrix minimization. arXiv:1401.0869.

Lu, C.; Tang, J.; Yan, S.; and Lin, Z. 2014. Generalized nonconvex nonsmooth low-rank minimization. In Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), 4130-4137.
Lu, C.; Zhu, C.; Xu, C.; Yan, S.; and Lin, Z. 2015. Generalized singular value thresholding. In Proc. AAAI Conf. Artif. Intell. (AAAI), 1805-1811.
Majumdar, A., and Ward, R. K. 2011. An algorithm for sparse MRI reconstruction by Schatten $p$-norm minimization. Magn. Reson. Imaging 29:408-417.
Marjanovic, G., and Solo, V. 2012. On $\ell_{p}$ optimization and matrix completion. IEEE Trans. Signal Process. 60(11):57145724.

Mitra, K.; Sheorey, S.; and Chellappa, R. 2010. Large-scale matrix factorization with missing data under additional constraints. In Proc. Adv. Neural Inf. Process. Syst. (NIPS), 1642-1650.
Mohan, K., and Fazel, M. 2012. Iterative reweighted algorithms for matrix rank minimization. J. Mach. Learn. Res. 13:34413473.

Nie, F.; Wang, H.; Cai, X.; Huang, H.; and Ding, C. 2012. Robust matrix completion via joint Schatten $p$-norm and $L_{p}$-norm minimization. In Proc. 12th IEEE Int. Conf. Data Min. (ICDM), 566-574.
Nie, F.; Huang, H.; and Ding, C. 2012. Low-rank matrix recovery via efficient Schatten p-norm minimization. In Proc. AAAI Conf. Artif. Intell. (AAAI), 655-661.
Recht, B.; Fazel, M.; and Parrilo, P. A. 2010. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. SIAM Rev. 52(3):471-501.
Recht, B. 2011. A simpler approach to matrix completion. J. Mach. Learn. Res. 12:3413-3430.
Rohde, A., and Tsybakov, A. B. 2011. Estimation of highdimensional low-rank matrices. Ann. Statist. 39(2):887-930.
Shang, F.; Liu, Y.; Cheng, J.; and Cheng, H. 2014. Robust principal component analysis with missing data. In Proc. 23rd ACM Int. Conf. Inf. Knowl. Manag. (CIKM), 1149-1158.
Srebro, N.; Rennie, J.; and Jaakkola, T. 2004. Maximum-margin matrix factorization. In Proc. Adv. Neural Inf. Process. Syst. (NIPS), 1329-1336.
Toh, K.-C., and Yun, S. 2010. An accelerated proximal gradient algorithm for nuclear norm regularized least squares problems. Pac. J. Optim. 6:615-640.
Wang, S.; Liu, D.; and Zhang, Z. 2013. Nonconvex relaxation approaches to robust matrix recovery. In Proc. 23rd Int. Joint Conf. Artif. Intell. (IJCAI), 1764-1770.
Wen, Z.; Yin, W.; and Zhang, Y. 2012. Solving a low-rank factorization model for matrix completion by a nonlinear successive over-relaxation algorithm. Math. Prog. Comp. 4(4):333361.

Xu, Y., and Yin, W. 2014. A globally convergent algorithm for nonconvex optimization based on block coordinate update. arXiv:1410.1386.
Zhang, M.; Huang, Z.; and Zhang, Y. 2013. Restricted pisometry properties of nonconvex matrix recovery. IEEE Trans. Inf. Theory 59(7):4316-4323.
Zhang, C. H. 2010. Nearly unbiased variable selection under minimax concave penalty. Ann. Statist. 38(2):894-942.


[^0]:    Copyright © $\mathfrak{C}$ 2016, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

[^1]:    ${ }^{1}$ http://www.math.ucla.edu/~wotaoyin/
    ${ }^{2}$ https://sites.google.com/site/canyilu/

[^2]:    ${ }^{3}$ http://www.math.nus.edu.sg/~mattohkc/
    ${ }^{4} \mathrm{http}: / /$ www.grouplens.org/node/73
    ${ }^{5}$ http://lmafit.blogs.rice.edu/.

