
Sketchy Decisions: Convex Low-Rank Matrix Optimization with Optimal Storage

Including supplementary appendix

Alp Yurtsever
EPFL

Madeleine Udell
Cornell

Joel A. Tropp
Caltech

Volkan Cevher
EPFL

Abstract

This paper concerns a fundamental class of convex matrix optimization problems. It presents the first algorithm that uses optimal storage and provably computes a low-rank approximation of a solution. In particular, when all solutions have low rank, the algorithm converges to a solution. This algorithm, SketchyCGM, modifies a standard convex optimization scheme, the conditional gradient method, to store only a small randomized sketch of the matrix variable. After the optimization terminates, the algorithm extracts a low-rank approximation of the solution from the sketch. In contrast to non-convex heuristics, the guarantees for SketchyCGM do not rely on statistical models for the problem data. Numerical work demonstrates the benefits of SketchyCGM over heuristics.

1 MOTIVATION

This paper discusses a fundamental class of convex matrix optimization problems with low-rank solutions. We argue that the main obstacle that prevents us from solving these problems at scale is not arithmetic, but storage. We exhibit the first provably correct algorithm for these problems with optimal storage.

1.1 Vignette: Matrix Completion

To explain the challenge, we consider the problem of low-rank matrix completion.

Let $\mathbf{X}_\natural \in \mathbb{R}^{m \times n}$ be an unknown matrix, but assume that a bound r on the rank of \mathbf{X}_\natural is available, where

Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS) 2017, Fort Lauderdale, Florida, USA. JMLR: W&CP volume 54. Copyright 2017 by the author(s).

$r \ll \min\{m, n\}$. Suppose that we record noisy observations of a subset E of entries from the matrix:

$$b_{ij} = (\mathbf{X}_\natural)_{ij} + \xi_{ij} \quad \text{for } (i, j) \in E.$$

The variables $\xi_{ij} \in \mathbb{R}$ model (unknown) noise. The goal is to approximate the full matrix \mathbf{X}_\natural .

Matrix completion arises in machine learning applications, such as recommendation systems [34].

We can frame the matrix completion problem as a rank-constrained optimization:

$$\underset{\mathbf{X} \in \mathbb{R}^{m \times n}}{\text{minimize}} \quad \sum_{(i,j) \in E} (x_{ij} - b_{ij})^2 \quad \text{s.t.} \quad \text{rank } \mathbf{X} \leq r. \quad (1)$$

In general, the formulation (1) is intractable. Instead, we retrench to a tractable convex problem [7, 34]:

$$\underset{\mathbf{X} \in \mathbb{R}^{m \times n}}{\text{minimize}} \quad \sum_{(i,j) \in E} (x_{ij} - b_{ij})^2 \quad \text{s.t.} \quad \|\mathbf{X}\|_{S_1} \leq \alpha. \quad (2)$$

The Schatten 1-norm $\|\cdot\|_{S_1}$ returns the sum of the singular values of its argument; it is an effective proxy for the rank [13]. Adjusting the value of the parameter α modulates the rank of a solution \mathbf{X}_\star of (2). If we have enough data and choose α well, we expect that each solution \mathbf{X}_\star approximates the target matrix \mathbf{X}_\natural .

The convex problem (2) is often a good model for matrix completion when the number of observations $|E| = \tilde{O}(r(m+n))$, where \tilde{O} suppresses log-like factors; see [7, 34]. We can write a rank- r approximation to a solution \mathbf{X}_\star using $\Theta(r(m+n))$ parameters. Thus, we can express the problem and an approximate solution with $\tilde{O}(r(m+n))$ storage.

Nevertheless, we need fully mn numbers to express the decision variable \mathbf{X} for the optimization problem (2). The cost of storing the decision variable prevents us from solving large-scale instances of (2), even without worrying about arithmetic.

This discrepancy raises a question: **Is there an algorithm that computes an approximate solution to (2) using the optimal storage $\tilde{O}(r(m+n))$?**

1.2 Vignette: Phase Retrieval

Here is another instance of the same predicament.

Fix a vector $\mathbf{x}_\natural \in \mathbb{C}^n$. Suppose that we acquire d noisy quadratic measurements of \mathbf{x}_\natural with the form

$$b_i = |\langle \mathbf{a}_i, \mathbf{x}_\natural \rangle|^2 + \xi_i \quad \text{for } i = 1, 2, \dots, d. \quad (3)$$

The $\mathbf{a}_i \in \mathbb{C}^n$ are known measurement vectors, and the $\xi_i \in \mathbb{R}$ model measurement noise. Given the data \mathbf{b} and the vectors \mathbf{a}_i , the phase retrieval problem asks us to reconstruct \mathbf{x}_\natural up to a global phase shift.

Phase retrieval problems are prevalent in imaging science because it is easier to measure the intensity of light than its phase. In practice, the vectors \mathbf{a}_i are structured because they reflect the physics of the imaging system. For more details, see Appendix B.3 and the examples in [3, 8, 10, 22].

Let us outline a convex approach [3, 8, 10, 22] to the phase retrieval problem. The data (3) satisfies

$$b_i = \mathbf{a}_i^* \mathbf{X}_\natural \mathbf{a}_i + \xi_i \quad \text{where } \mathbf{X}_\natural = \mathbf{x}_\natural \mathbf{x}_\natural^*.$$

Thus, we can formulate phase retrieval as

$$\begin{aligned} & \underset{\mathbf{X} \in \mathbb{C}^{n \times n}}{\text{minimize}} && \sum_{i=1}^d (\mathbf{a}_i^* \mathbf{X} \mathbf{a}_i - b_i)^2 \\ & \text{s.t.} && \text{rank } \mathbf{X} = 1, \quad \mathbf{X} \succeq \mathbf{0}. \end{aligned} \quad (4)$$

Now, pass to the convex problem

$$\begin{aligned} & \underset{\mathbf{X} \in \mathbb{C}^{n \times n}}{\text{minimize}} && \sum_{i=1}^d (\mathbf{a}_i^* \mathbf{X} \mathbf{a}_i - b_i)^2 \\ & \text{s.t.} && \text{tr } \mathbf{X} \leq \alpha, \quad \mathbf{X} \succeq \mathbf{0}. \end{aligned} \quad (5)$$

We can estimate the parameter $\alpha \in \mathbb{R}_+$ from \mathbf{a}_i and \mathbf{b} ; see [39, Sec. II]. To approximate the true vector \mathbf{x}_\natural , we compute a top eigenvector \mathbf{x}_\star of a solution to (5).

This procedure is often an effective approach for phase retrieval when the number of measurements $d = \Theta(n)$; see [36, Sec. 2.8]. Once again, we recognize a discrepancy. The problem data $\mathbf{b} \in \mathbb{R}^d$ and the approximate solution $\mathbf{x}_\star \in \mathbb{C}^n$ use storage $\Theta(n)$, but the matrix variable in (5) requires $\Theta(n^2)$ storage.

We may ask: **Is there an algorithm that computes an approximate solution to (5) using the optimal storage $\Theta(n)$?**

1.3 Low-Rank Matrix Optimization Methods

Matrix completion and phase retrieval are examples of *convex low-rank matrix optimization* (CLRO) problems. Informally, this class contains convex optimization problems whose decision variable is a matrix and whose solutions are (close to) low rank. These

problems often arise as convex relaxations of rank-constrained problems; however, the convex formulations are important in their own right.

There has been extensive empirical and theoretical work to validate the use of CLROs in a spectrum of applications. For example, see [7, 8, 13, 16, 22].

Over the last 20 years, optimization researchers have developed a diverse collection of algorithms for CLRO problems. Surprisingly, every extant method lacks guarantees on storage or convergence (or both).

Convex optimization algorithms dominated the early literature on algorithms for CLRO. The initial efforts, such as [21], focused on interior-point methods, whose storage and arithmetic costs are forbidding. To resolve this issue, researchers turned to first-order convex algorithms, including bundle methods [20], (accelerated) proximal gradient methods [2, 32, 35], and the conditional gradient method (CGM) [11, 15, 19, 23, 25].

Convex algorithms are guaranteed to solve a CLRO. They come with a complete theory, including rigorous stopping criteria and bounds on convergence rates. They enjoy robust performance in practice. On the other hand, convex algorithms from the literature do not scale well enough to solve large CLRO problems because they operate on and store full-size matrices.

The CGM iteration is sometimes touted as a low-storage method for CLRO [23]. Indeed, CGM is guaranteed to increase the rank of an iterate by at most one per iteration. Nevertheless, the algorithm converges slowly, so intermediate iterates can have very high rank. CGM variants, such as [31, 39], that control the rank of iterates lack storage guarantees or may not converge to a global optimum.

Recently, many investigators have sought recourse in nonconvex heuristics for solving CLROs. This line of work depends on the factorization idea of Burer & Monteiro [6], which rewrites the matrix variable as a product of two low-rank factors. There are many heuristic procedures, e.g., [4–6, 24], that use clever initialization and nonlinear programming schemes in an attempt to optimize the factors directly. The resulting algorithms can have optimal storage costs, and they may achieve a fast rate of local convergence.

There has been an intensive effort to justify the application of nonconvex heuristics for CLRO. To do so, researchers often frame unverifiable statistical assumptions on the problem data. For example, in the matrix completion problem (2), it is common to assume that the entries of the matrix are revealed according to some ideal probability distribution [7, 24]. When these assumptions fail, nonconvex heuristics can converge to the wrong point, or they may even diverge.

Contributions. This paper explains how to extend the convex optimization algorithm CGM to obtain an approximate solution to a class of CLRO problems using optimal storage. Our algorithm operates much like CGM, but it never forms the matrix variable explicitly. Instead, we maintain a small randomized sketch of the matrix variable over the course of the iteration by means of a bespoke sketching method [37]. After the optimization method converges, we extract an approximate solution from the sketch. This technique achieves optimal storage, yet it converges under the same conditions and with the same guarantees as CGM.

In summary, this paper presents a solution to the problems posed above: the first algorithm for convex low-rank matrix optimization problems that provably uses optimal storage to compute an approximate solution.

1.4 Notation

We write $\|\cdot\|$ for the Euclidean norm, $\|\cdot\|_F$ for the Frobenius norm, and $\|\cdot\|_{S_1}$ for the Schatten 1-norm (aka the *trace norm* or the *nuclear norm*). Depending on context, $\langle \cdot, \cdot \rangle$ refers to the Euclidean or Frobenius inner product. The symbol $*$ denotes the conjugate transpose of a vector or matrix, as well as the adjoint of a linear map. The dagger \dagger refers to the pseudoinverse. The symbol $[M]_r$ stands for a best rank- r Frobenius-norm approximation of the matrix M . The function $\text{dist}_F(M; S)$ returns the minimum Frobenius-norm distance from M to a set S . The symbol \succcurlyeq denotes the semidefinite order. We use the computer science interpretation of the order notation $\mathcal{O}, \tilde{\mathcal{O}}, \Omega, \Theta$.

2 A LOW-RANK MATRIX OPTIMIZATION PROBLEM

Let us begin with a generalization of the convex matrix completion formulation (2). In §5.5, we return to the psd setting of the phase retrieval problem (5).

We consider a convex program with a matrix variable:

$$\underset{\mathbf{X} \in \mathbb{R}^{m \times n}}{\text{minimize}} \ f(\mathcal{A}\mathbf{X}) \quad \text{s.t.} \quad \|\mathbf{X}\|_{S_1} \leq \alpha. \quad (6)$$

The linear operator $\mathcal{A} : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^d$ and its adjoint $\mathcal{A}^* : \mathbb{R}^d \rightarrow \mathbb{R}^{m \times n}$ take the form

$$\begin{aligned} \mathcal{A}\mathbf{X} &= [\langle \mathbf{A}_1, \mathbf{X} \rangle \quad \dots \quad \langle \mathbf{A}_d, \mathbf{X} \rangle]; \\ \mathcal{A}^*\mathbf{z} &= \sum_{i=1}^d z_i \mathbf{A}_i. \end{aligned} \quad (7)$$

Each coefficient matrix $\mathbf{A}_i \in \mathbb{R}^{m \times n}$.

We interpret $\mathcal{A}\mathbf{X}$ as a set of linear measurements of the matrix \mathbf{X} . For example, in the matrix completion problem (2), the image $\mathcal{A}\mathbf{X}$ lists the entries of \mathbf{X} indexed by the set E .

The function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is convex and continuously differentiable. In many situations, it is natural to regard the objective function as a loss: $f(\mathcal{A}\mathbf{X}) = \text{loss}(\mathcal{A}\mathbf{X}; \mathbf{b})$ for a vector $\mathbf{b} \in \mathbb{R}^d$ of measured data.

By choosing the parameter $\alpha \in \mathbb{R}_+$ to be sufficiently small, we can often ensure that each minimizer of (6) is low-rank or close to low-rank.

Our goal is to develop a practical algorithm that provably computes a low-rank approximation of a solution to the problem (6).

To validate (6) as a model for a given application, one must undertake a separate empirical or theoretical study. We do not engage this question in our work.

2.1 Storage Issues

Suppose that we want to produce a low-rank approximation to a solution of a generic instance of the problem (6). What kind of storage can we hope to achieve?

It is clear that we need $\Theta(r(m+n))$ numbers to express a rank- r approximate solution to (6). We must also understand how much extra storage is incurred because of the specific problem instance (\mathcal{A}, f) .

It is natural to instate a *black-box model* for the linear map \mathcal{A} , its adjoint \mathcal{A}^* , and the objective function f . For arbitrary vectors $\mathbf{u} \in \mathbb{R}^m$ and $\mathbf{v} \in \mathbb{R}^n$ and $\mathbf{z} \in \mathbb{R}^d$, assume we have routines that can compute

$$\mathcal{A}(\mathbf{u}\mathbf{v}^*) \quad \text{and} \quad \mathbf{u}^*(\mathcal{A}^*\mathbf{z}) \quad \text{and} \quad (\mathcal{A}^*\mathbf{z})\mathbf{v}. \quad (8)$$

We also assume routines for evaluating the function f and its gradient ∇f for any argument $\mathbf{z} \in \mathbb{R}^d$. We may neglect the storage used to compute these primitives. Every algorithm based on these primitives must use storage $\Omega(m+n+d)$ just to represent their outputs.

Thus, under the black-box model, any algorithm that produces a rank- r solution to a generic instance of (6) must use storage $\Omega(d+r(m+n))$. We say that an algorithm is *storage optimal* if it achieves this bound.

The parameter d often reflects the amount of data that we have acquired, and it is usually far smaller than the dimension mn of the matrix variable in (6).

The problems that concern us are data-limited; that is, $d \ll mn$. This is the situation where a strong structural prior (e.g., low rank or small Schatten 1-norm) is essential for fitting the data. This challenge is common in machine learning problems (e.g., matrix completion for recommendation systems), as well as in scientific applications (e.g., phase retrieval).

To the best of our knowledge, no extant algorithm for (6) is guaranteed to produce an approximation of an optimal point and also enjoys optimal storage cost.

3 CONDITIONAL GRADIENT

To develop our algorithm for the model problem (6), we must first describe a standard algorithm called the *conditional gradient method* (CGM). Classic and contemporary references include [11, 15, 19, 23, 25].

3.1 The CGM Iteration

Here is the CGM algorithm for (6). Start with a feasible point, such as

$$\mathbf{X}_0 = \mathbf{0} \in \mathbb{R}^{m \times n}. \quad (9)$$

At each iteration $t = 0, 1, 2, \dots$, compute an update direction \mathbf{H}_t using the formulas

$$\begin{aligned} (\mathbf{u}_t, \mathbf{v}_t) &= \text{MaxSingVec}(\mathcal{A}^*(\nabla f(\mathcal{A}\mathbf{X}_t))); \\ \mathbf{H}_t &= -\alpha \mathbf{u}_t \mathbf{v}_t^*. \end{aligned} \quad (10)$$

MaxSingVec returns a left/right pair of maximum singular vectors. Update the decision variable:

$$\mathbf{X}_{t+1} = (1 - \eta_t)\mathbf{X}_t + \eta_t \mathbf{H}_t \quad (11)$$

where $\eta_t = 2/(t+2)$. The convex combination (11) remains feasible for (6) because \mathbf{X}_t and \mathbf{H}_t are feasible.

CGM is a valuable algorithm for (6) because we can efficiently find the rank-one update direction \mathbf{H}_t by means of the singular vector computation (10). The weak point of CGM is that the rank of \mathbf{X}_t typically increases with t , and the peak rank of an iterate \mathbf{X}_t is often much larger than the rank of the solution of (6). See Figures 6 and 7 for an illustration.

3.2 The CGM Stopping Rule

The CGM algorithm admits a simple stopping criterion. Given a suboptimality parameter $\varepsilon > 0$, we halt the CGM iteration when the duality gap $\delta_t \leq \varepsilon$:

$$\delta_t = \langle \mathcal{A}\mathbf{X}_t - \mathcal{A}\mathbf{H}_t, \nabla f(\mathcal{A}\mathbf{X}_t) \rangle \leq \varepsilon. \quad (12)$$

Let \mathbf{X}_\star be an optimal point for (6). It is not hard to show [23, Sec. 2] that

$$f(\mathcal{A}\mathbf{X}_t) - f(\mathcal{A}\mathbf{X}_\star) \leq \delta_t. \quad (13)$$

Thus, the condition (12) ensures that the objective value $f(\mathcal{A}\mathbf{X}_t)$ is ε -suboptimal. The CGM iterates satisfy (12) within $O(\varepsilon^{-1})$ iterations [23, Thm. 1].

3.3 The Opportunity

The CGM iteration (9)–(11) requires $\Theta(mn)$ storage because it maintains the $m \times n$ matrix decision variable \mathbf{X}_t . We develop a remarkable extension of CGM that provably computes a rank- r approximate solution to (6) with working storage $\Theta(d + r(m+n))$. Our approach depends on two efficiencies:

- We use the low-dimensional “dual” variable $\mathbf{z}_t = \mathcal{A}\mathbf{X}_t \in \mathbb{R}^d$ to drive the iteration.
- Instead of storing \mathbf{X}_t , we maintain a small randomized sketch with size $\Theta(r(m+n))$.

It is easy to express the CGM iteration in terms of the “dual” variable $\mathbf{z}_t = \mathcal{A}\mathbf{X}_t$. We can obviously rewrite the formula (10) for computing the rank-one update direction \mathbf{H}_t in terms of \mathbf{z}_t . We obtain an update rule for \mathbf{z}_t by applying the linear map \mathcal{A} to (11). Likewise, the stopping criterion (12) can be evaluated using \mathbf{z}_t and \mathbf{H}_t . Under the black-box model (8), the dual formulation of CGM has storage cost $\Theta(m+n+d)$.

Yet the dual formulation has a flaw: it “solves” the problem (6), but we do not know the solution!

Indeed, we must also track the evolution (11) of the primal decision variable \mathbf{X}_t . In the next subsection, we summarize a randomized sketching method [37] that allows us to compute an accurate rank- r approximation of \mathbf{X}_t but operates with storage $\Theta(r(m+n))$.

4 MATRIX SKETCHING

Suppose that $\mathbf{X} \in \mathbb{R}^{m \times n}$ is a matrix that is presented to us as a stream of linear updates, as in (11). For a parameter $r \ll \min\{m, n\}$, we wish to maintain a small sketch that allows us to compute a rank- r approximation of the final matrix \mathbf{X} . Let us summarize an approach developed in our paper [37].

4.1 The Randomized Sketch

Draw and fix two independent standard normal matrices $\mathbf{\Omega}$ and $\mathbf{\Psi}$ where

$$\begin{aligned} \mathbf{\Omega} &\in \mathbb{R}^{n \times k} && \text{with } k = 2r + 1; \\ \mathbf{\Psi} &\in \mathbb{R}^{\ell \times m} && \text{with } \ell = 4r + 3. \end{aligned} \quad (14)$$

The sketch consists of two matrices \mathbf{Y} and \mathbf{W} that capture the range and co-range of \mathbf{X} :

$$\mathbf{Y} = \mathbf{X}\mathbf{\Omega} \in \mathbb{R}^{m \times k} \quad \text{and} \quad \mathbf{W} = \mathbf{\Psi}\mathbf{X} \in \mathbb{R}^{\ell \times n}. \quad (15)$$

We can efficiently update the sketch (\mathbf{Y}, \mathbf{W}) to reflect a rank-one linear update to \mathbf{X} of the form

$$\mathbf{X} \leftarrow \beta_1 \mathbf{X} + \beta_2 \mathbf{u}\mathbf{v}^*. \quad (16)$$

Both the storage cost for the sketch and the arithmetic cost of an update are $\Theta(r(m+n))$.

4.2 The Reconstruction Algorithm

The following procedure yields a rank- r approximation $\hat{\mathbf{X}}$ of the matrix \mathbf{X} stored in the sketch (15).

$$\mathbf{Q} = \text{orth}(\mathbf{Y}); \quad \mathbf{B} = (\mathbf{\Psi}\mathbf{Q})^\dagger \mathbf{W}; \quad \hat{\mathbf{X}} = \mathbf{Q}[\mathbf{B}]_r. \quad (17)$$

The matrix \mathbf{Q} has orthonormal columns that span the range of \mathbf{Y} . The extra storage costs of the reconstruction are negligible; its arithmetic cost is $\Theta(r^2(m+n))$. See [37, §4.2] for the intuition behind this method. It achieves the following error bound.

Theorem 1 (Reconstruction error [37, Thm. 5.1]). *Fix a target rank r . Let \mathbf{X} be a matrix, and let (\mathbf{Y}, \mathbf{W}) be a sketch of \mathbf{X} of the form (14)–(15). The procedure (17) yields a rank- r matrix $\hat{\mathbf{X}}$ with*

$$\mathbb{E} \|\mathbf{X} - \hat{\mathbf{X}}\|_{\text{F}} \leq 3\sqrt{2} \|\mathbf{X} - [\mathbf{X}]_r\|_{\text{F}}.$$

Similar bounds hold with high probability.

Remarks. The sketch size parameters (k, ℓ) appearing in (14) are recommended to balance storage against reconstruction quality. See [37] and our follow-up work for more details and for other sketching methods.

5 SKETCHING + CGM

We are now prepared to present SketchyCGM, a storage-optimal extension of the CGM algorithm for the convex problem (6). This method delivers a provably accurate low-rank approximation to a solution of (6). See Algorithm 1 for complete pseudocode.

5.1 The SketchyCGM Iteration

Fix the suboptimality ε and the rank r . Draw and fix standard normal matrices $\mathbf{\Omega} \in \mathbb{R}^{n \times k}$ and $\mathbf{\Psi} \in \mathbb{R}^{\ell \times m}$ as in (14). Initialize the iterate and the sketches:

$$\mathbf{z}_0 = \mathbf{0}_d; \quad \mathbf{Y}_0 = \mathbf{0}_{m \times k}; \quad \text{and} \quad \mathbf{W}_0 = \mathbf{0}_{\ell \times n}. \quad (18)$$

At each iteration $t = 0, 1, 2, \dots$, compute an update direction via Lanczos or via randomized SVD [17]:

$$\begin{aligned} (\mathbf{u}_t, \mathbf{v}_t) &= \text{MaxSingVec}(\mathcal{A}^*(\nabla f(\mathbf{z}_t))); \\ \mathbf{h}_t &= \mathcal{A}(-\alpha \mathbf{u}_t \mathbf{v}_t^*). \end{aligned} \quad (19)$$

Set the learning rate $\eta_t = 2/(t+2)$. Update the iterate and the two sketches:

$$\begin{aligned} \mathbf{z}_{t+1} &= (1 - \eta_t)\mathbf{z}_t + \eta_t \mathbf{h}_t; \\ \mathbf{Y}_{t+1} &= (1 - \eta_t)\mathbf{Y}_t + \eta_t(-\alpha \mathbf{u}_t \mathbf{v}_t^*)\mathbf{\Omega}; \\ \mathbf{W}_{t+1} &= (1 - \eta_t)\mathbf{W}_t + \eta_t \mathbf{\Psi}(-\alpha \mathbf{u}_t \mathbf{v}_t^*). \end{aligned} \quad (20)$$

The iteration continues until it triggers the stopping criterion:

$$\langle \mathbf{z}_t - \mathbf{h}_t, \nabla f(\mathbf{z}_t) \rangle \leq \varepsilon. \quad (21)$$

At any iteration t , we can form a rank- r approximate solution $\hat{\mathbf{X}}_t$ of the model problem (6) by invoking the procedure (17) with $\mathbf{Y} = \mathbf{Y}_t$ and $\mathbf{W} = \mathbf{W}_t$.

5.2 Guarantees

Suppose that the CGM iteration (9)–(11) generates the sequence $(\mathbf{X}_t : t = 0, 1, 2, \dots)$ of decision variables and the sequence $(\mathbf{H}_t : t = 0, 1, 2, \dots)$ of update directions. It is easy to verify that the SketchyCGM iteration (18)–(20) maintains the loop invariants

$$\begin{aligned} \mathbf{z}_t &= \mathcal{A}\mathbf{X}_t \quad \text{and} \quad \mathbf{h}_t = \mathcal{A}\mathbf{H}_t; \\ \mathbf{Y}_t &= \mathbf{X}_t\mathbf{\Omega} \quad \text{and} \quad \mathbf{W}_t = \mathbf{\Psi}\mathbf{X}_t. \end{aligned} \quad (22)$$

In view of the inequality (13) and the invariant (22), the stopping rule (21) ensures that \mathbf{X}_t is an ε -suboptimal solution to (6) when the iteration halts. Furthermore, Theorem 1 ensures that the computed solution $\hat{\mathbf{X}}_t$ is a near-optimal rank- r approximation of \mathbf{X}_t at each time t .

5.3 Storage Costs

The total storage cost is $\Theta(d + r(m+n))$ for the dual variable \mathbf{z}_t , the random matrices $(\mathbf{\Omega}, \mathbf{\Psi})$, and the sketch (\mathbf{Y}, \mathbf{W}) . Owing to the black-box assumption (8), the algorithm completes the singular vector computations in (19) with $\Theta(d+m+n)$ working storage. At no point during the iteration do we instantiate an $m \times n$ matrix! Arithmetic costs are on the same order as the standard version of CGM.

5.4 Convergence Results for SketchyCGM

SketchyCGM is a provably correct method for computing a low-rank approximation of a solution to (6).

Theorem 2. *Suppose that the iterates \mathbf{X}_t from the CGM iteration (9)–(11) converge to a matrix \mathbf{X}_{cgm} . Let $\hat{\mathbf{X}}_t$ be the rank- r reconstruction of \mathbf{X}_t produced by SketchyCGM. Then*

$$\lim_{t \rightarrow \infty} \mathbb{E} \|\hat{\mathbf{X}}_t - \mathbf{X}_{\text{cgm}}\|_{\text{F}} \leq 3\sqrt{2} \|\mathbf{X}_{\text{cgm}} - [\mathbf{X}_{\text{cgm}}]_r\|_{\text{F}}.$$

In particular, if $\text{rank}(\mathbf{X}_{\text{cgm}}) \leq r$, then

$$\mathbb{E} \|\hat{\mathbf{X}}_t - \mathbf{X}_{\text{cgm}}\|_{\text{F}} \rightarrow 0.$$

SketchyCGM always works in the fundamental case where each solution of (6) has low rank.

Theorem 3. *Suppose that the solution set S_* of the problem (6) contains only matrices with rank r or less. Then SketchyCGM attains $\mathbb{E} \text{dist}_{\text{F}}(\hat{\mathbf{X}}_t, S_*) \rightarrow 0$.*

Suppose that the optimal point of (6) is stable in the sense that the value of the objective function increases as we depart from optimality. Then the SketchyCGM estimates converge at a controlled rate.

Theorem 4. *Fix $\kappa > 0$ and $\nu > 0$. Suppose the (unique) solution \mathbf{X}_* of (6) has $\text{rank}(\mathbf{X}_*) \leq r$ and*

$$f(\mathcal{A}\mathbf{X}) - f(\mathcal{A}\mathbf{X}_*) \geq \kappa \|\mathbf{X} - \mathbf{X}_*\|_{\text{F}}^{\nu} \quad (23)$$

Algorithm 1 SketchyCGM for model problem (6)

Input: Data for (6); suboptimality ε ; target rank r
Output: Rank- r approximate solution $\hat{\mathbf{X}} = \mathbf{U}\Sigma\mathbf{V}^*$ of (6) in factored form

```

1  function SKETCHYCGM
2      SKETCH.INIT( $m, n, r$ )
3       $\mathbf{z} \leftarrow \mathbf{0}$ 
4      for  $t \leftarrow 0, 1, 2, 3, \dots$  do
5           $(\mathbf{u}, \mathbf{v}) \leftarrow \text{MaxSingVec}(\mathcal{A}^*(\nabla f(\mathbf{z})))$ 
6           $\mathbf{h} \leftarrow \mathcal{A}(-\alpha\mathbf{u}\mathbf{v}^*)$ 
7          if  $\langle \mathbf{z} - \mathbf{h}, \nabla f(\mathbf{z}) \rangle \leq \varepsilon$  then break for
8               $\eta \leftarrow 2/(t+2)$ 
9               $\mathbf{z} \leftarrow (1-\eta)\mathbf{z} + \eta\mathbf{h}$ 
10             SKETCH.CGMUPDATE( $-\alpha\mathbf{u}, \mathbf{v}, \eta$ )
11          $(\mathbf{U}, \Sigma, \mathbf{V}) \leftarrow \text{SKETCH.RECONSTRUCT}()$ 
12         return  $(\mathbf{U}, \Sigma, \mathbf{V})$ 

```

— Methods for SKETCH object —

```

13 function SKETCH.INIT( $m, n, r$ )
14      $k \leftarrow 2r+1$  and  $\ell \leftarrow 4r+3$ 
15      $\Omega \leftarrow \text{randn}(n, k)$  and  $\Psi \leftarrow \text{randn}(\ell, m)$ 
16      $\mathbf{Y} \leftarrow \text{zeros}(m, k)$  and  $\mathbf{W} \leftarrow \text{zeros}(\ell, n)$ 
17 function SKETCH.CGMUPDATE( $\mathbf{u}, \mathbf{v}, \eta$ )
18      $\mathbf{Y} \leftarrow (1-\eta)\mathbf{Y} + \eta\mathbf{u}(\mathbf{v}^*\Omega)$ 
19      $\mathbf{W} \leftarrow (1-\eta)\mathbf{W} + \eta(\Psi\mathbf{u})\mathbf{v}^*$ 
20 function SKETCH.RECONSTRUCT()
21      $\mathbf{Q} \leftarrow \text{orth}(\mathbf{Y})$ 
22      $\mathbf{B} \leftarrow (\Psi\mathbf{Q}) \setminus \mathbf{W}$ 
23      $(\mathbf{U}, \Sigma, \mathbf{V}) \leftarrow \text{svds}(\mathbf{B}, r)$ 
24     return  $(\mathbf{Q}\mathbf{U}, \Sigma, \mathbf{V})$ 

```

for all feasible \mathbf{X} . Then we have the error bound

$$\mathbb{E} \|\hat{\mathbf{X}}_t - \mathbf{X}_\star\|_{\text{F}} \leq 6 \left(\frac{2C\kappa^{-1}}{t+2} \right)^{1/\nu} \quad \text{for } t = 0, 1, 2, \dots$$

where C is the curvature constant [23, Eqn. (3)] of the problem (6).

See Appendix A for the proofs of these results.

5.5 SketchyCGM for PSD Optimization

Next, we present a generalization of the convex phase retrieval problem (5). Consider the convex template

$$\underset{\mathbf{X} \in \mathbb{C}^{n \times n}}{\text{minimize}} f(\mathcal{A}\mathbf{X}) \quad \text{s.t.} \quad \text{tr } \mathbf{X} \leq \alpha, \quad \mathbf{X} \succeq \mathbf{0}. \quad (24)$$

As before, $\mathcal{A} : \mathbb{C}^{n \times n} \rightarrow \mathbb{C}^d$ is a linear map, and $f : \mathbb{C}^d \rightarrow \mathbb{R}$ is a differentiable convex function.

It is easy to adapt SketchyCGM to handle (24) instead of (6). To sketch the complex psd matrix variable,

we follow the approach described in [37, Sec. 7.3]. We also make small changes to the SketchyCGM iteration. Replace the computation (19) with

$$(\lambda_t, \mathbf{u}_t) = \text{MinEig}(\mathcal{A}^*(\nabla f(\mathbf{z}_t)));$$

$$\mathbf{h}_t = \begin{cases} \mathcal{A}(\alpha\mathbf{u}_t\mathbf{u}_t^*), & \lambda_t \leq 0 \\ \mathbf{0}, & \text{otherwise.} \end{cases}$$

MinEig returns the minimum eigenvalue λ_t and an associated eigenvector \mathbf{u}_t of a conjugate symmetric matrix. This variant behaves the same as SketchyCGM.

6 COMPUTATIONAL EVIDENCE

In this section, we demonstrate that SketchyCGM is a practical algorithm for convex low-rank matrix optimization. We focus on phase retrieval problems because they provide a dramatic illustration of the power of storage-optimal convex optimization. Appendix B adduces additional examples, including some matrix completion problems.

6.1 Synthetic Phase Retrieval Problems

To begin, we show that our approach to solving the convex phase retrieval problem (5) has better memory scaling than other convex optimization methods.

We compare five convex optimization algorithms: the classic proximal gradient method (PGM) [32]; the Auslender–Teboulle (AT) accelerated method [2]; the classic CGM algorithm [23]; a storage-efficient CGM variant (ThinCGM) [39] based on low-rank SVD updating; and the psd variant of SketchyCGM from §5.5 with rank parameter $r = 1$.

All five methods solve (24) reliably. The proximal methods (PGM and AT) perform a full eigenvalue decomposition of the iterate at each step, but they can be accelerated by adaptively choosing the number of eigenvectors to compute. The methods based on CGM only need the top eigenvector, so they perform less arithmetic per iteration.

To compare the storage costs of the five algorithms, let us consider a synthetic phase retrieval problem. We draw a vector $\mathbf{x}_\dagger \in \mathbb{C}^n$ from the complex standard normal distribution. Then we acquire $d = 10n$ phaseless measurements (3), corrupted with independent Gaussian noise so that the SNR is 20 dB. The measurement vectors \mathbf{a}_i derive from a coded diffraction pattern; see §B.3.2 for details. We solve the convex problem (5) with α equal to the average of the measurements b_i ; see [39, Sec. II]. Then we compute a top eigenvector \mathbf{x}_\star of the solution.

Figure 1(A) displays storage costs for each algorithm as the signal length n increases. See §B.3.2 for the

numerical data. We approximate memory usage by reporting the total workspace allocated by MATLAB for the algorithm. PGM, AT, and CGM have static allocations, but they use a matrix variable of size n^2 . ThinCGM attempts to maintain a low-rank approximation of the decision variable, but the rank increases steadily, so the algorithm fails after $n = 10^5$. In contrast, SketchyCGM has a static memory allocation of $\Theta(n)$. It already offers the best memory footprint for $n = 10^2$, and it still works when $n = 10^6$.

In fact, SketchyCGM can tackle even larger problems. We were able to reconstruct an image with $n = 3,264 \times 2,448 \approx 7.99 \cdot 10^6$ pixels, treated as a vector $\mathbf{x}_\dagger \in \mathbb{C}^n$, given $d = 20n$ noiseless coded diffraction measurements, as above. Figure 1(B) plots the convergence of the relative error: $\min_{\phi \in \mathbb{R}} \|e^{i\phi} \hat{\mathbf{x}}_t - \mathbf{x}_\dagger\| / \|\mathbf{x}_\dagger\|$, where $\hat{\mathbf{x}}_t$ is a top eigenvector of the SketchyCGM iterate $\hat{\mathbf{X}}_t$. After 150 iterations, the algorithm produced an image with relative error 0.0290 and with PSNR 36.19 dB. Thus, we can solve (24) when the psd matrix variable \mathbf{X} has $6.38 \cdot 10^{13}$ complex entries!

As compared to other convex optimization algorithms, the main weakness of SketchyCGM is the $\mathcal{O}(\varepsilon^{-1})$ iteration count. Some algorithms, such as AT, can achieve $\mathcal{O}(\varepsilon^{-1/2})$ iteration count, but they are limited to smaller problems. Closing this gap is an important question for future work.

6.2 Fourier Ptychography

Up to now, it has not been possible to attack phase retrieval problems of a realistic size by solving the convex formulation (5). As we have shown, current convex optimization algorithms cannot achieve scale. Instead, many researchers apply nonconvex heuristics to solve phase retrieval problems [9, 14, 22, 27]. These heuristics can produce reasonable solutions, but they require extensive tuning and have limited effectiveness. In this section, we demonstrate that, without any modification, SketchyCGM can solve a phase retrieval problem from an application in microscopy. Furthermore, it produces an image that is superior to the results obtained using major nonconvex heuristics.

We study a phase retrieval problem that arises from an imaging modality called Fourier ptychography (FP) [22]. The authors of [22] provided measurements of a slide containing human blood cells from a working FP system. We treat the sample as an image with $n = 25,600$ pixels, which we represent as a vector $\mathbf{x}_\dagger \in \mathbb{C}^n$. The goal is to reconstruct the *phase* of the image, which roughly corresponds with the thickness of the sample at a given location.

The data consists of 29 illuminations, each contain-

ing 6,400 pixels. The number of measurements $d = 185,600$. The measurement vectors \mathbf{a}_i are obtained from windowed discrete Fourier transforms. We can formulate the problem of reconstructing the sample \mathbf{x}_\dagger using the convex phase retrieval template (5) with the parameter $\alpha = 1,400$. In this instance, the psd matrix variable $\mathbf{X} \in \mathbb{C}^{n \times n}$ has $6.55 \cdot 10^8$ complex entries.

To solve (5), we run the SketchyCGM variant from §5.5 with the rank parameter $r = 1$ for 10,000 iterations. We factor the rank-one matrix output to obtain an approximation \mathbf{x}_\star of the sample. Figure 2(A) displays the phase of the reconstruction \mathbf{x}_\star .

Figure 2 also includes comparisons with two nonconvex heuristics. The authors of [22] provided a reconstruction obtained via the Burer–Monteiro method [6]. We also applied Wirtinger Flow [9] with the recommended parameters. SketchyCGM yields a smooth and detailed phase reconstruction. Burer–Monteiro produces severe artifacts, which suggest an unphysical oscillation in the thickness of the sample. Wirtinger Flow fails completely. These results are consistent with [22, Fig. 4], which indicates 5–10 dB improvement of convex optimization over heuristics.

The quality of phase reconstruction can be essential for scientific purposes. In this particular example, some of the blood cells are infected with malaria parasites (Figure 2(A), red boxes). Diagnosis is easier when the visual acuity of the reconstruction is high.

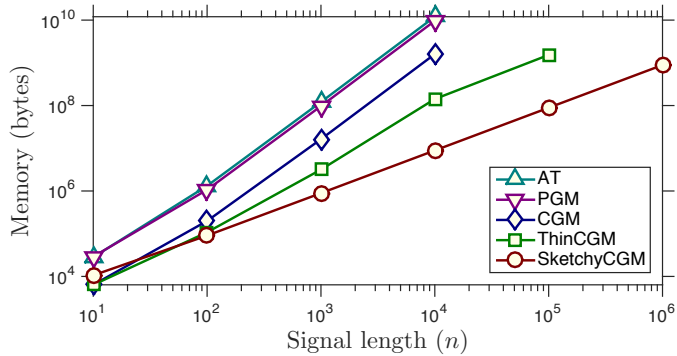
Appendix B.3 contains additional details and results on Fourier ptychographic imaging.

7 DISCUSSION

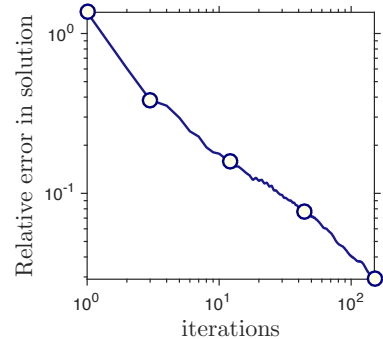
We have shown that it is possible to construct a low-rank approximate solution to a large-scale matrix optimization problem by sketching the decision variable. Let us contrast our approach against other low-storage techniques for large-scale optimization.

Sketchy Decisions. To the best of our knowledge, there are no direct precedents for the idea and realization of an optimization algorithm that sketches the decision variable. This work does partake in a broader vision that randomization can be used to design numerical algorithms [17, 26, 38].

Researchers have considered **sketching the problem data** as a way to reduce the size of a problem specification in exchange for additional error. This idea dates to the paper of Sarlós [33]; see also [26, 30, 38]. There are also several papers [1, 12, 29] in which researchers try to improve the computational or storage footprint of convex optimization methods by **sketching internal variables**, such as Hessians.

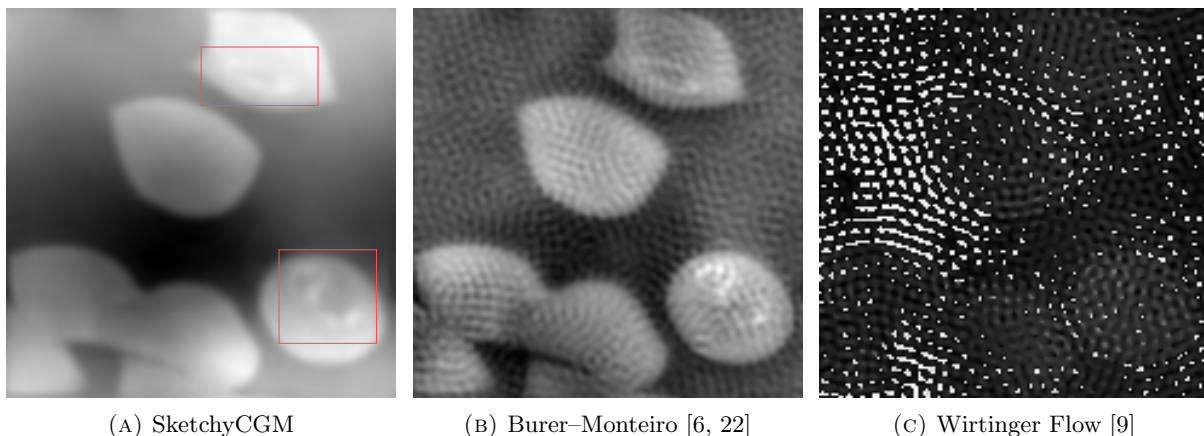


(A) Memory usage for five algorithms



(B) Convergence on largest problem

FIGURE 1: *Memory Usage and Convergence.* (A) Memory scaling for five convex optimization algorithms applied to a synthetic instance of the convex phase retrieval problem (5). (B) Relative ℓ_2 error achieved by SketchyCGM for convex phase retrieval of a synthetic signal of length $n = 8 \cdot 10^6$. See §6.2 for further details.



(A) SketchyCGM

(B) Burer–Monteiro [6, 22]

(C) Wirtinger Flow [9]

FIGURE 2: *Imaging by Fourier Ptychography.* Three algorithms for Fourier ptychographic imaging via phase retrieval. Brightness indicates the complex phase of a pixel, which roughly corresponds with the thickness of the sample. Only relative differences in brightness are meaningful. Red boxes mark malaria parasites in blood cells.

None of these approaches address the core issue that concerns us: the decision variable may require much more storage than the solution.

Dropping Nonconvexity. We have already discussed a major trend in which researchers develop algorithms that attack **nonconvex reformulations** of a problem. For example, see [4–6, 24]. The main advantage is to reduce the size of the decision variable; some methods also have the ancillary benefit of rapid local convergence. On the other hand, these algorithms are provably correct only under strong statistical assumptions on the problem data.

Prospects. Our work shows that convex optimization need not have high storage complexity for problems with a compact specification and a simple solution. In particular, for low-rank matrix optimization, storage is no reason to drop convexity.

It has not escaped our notice that the specific pairing of sketching and CGM that we have postulated immediately suggests a possible mechanism for solving other structured convex programs using optimal storage.

Acknowledgments

JAT and MU were supported in part by ONR Awards N00014-11-1-0025 and N00014-17-1-2146 and the Gordon & Betty Moore Foundation. VC and AY were supported in part by the European Commission under Grant ERC Future Proof, SNF 200021-146750, and SNF CRSII2-147633.

A Convergence of SketchyCGM

In this appendix, we develop a basic convergence theory for the SketchyCGM algorithm. We focus on situations where the low-rank estimates produced by SketchyCGM converge to a low-rank solution of the model problem (6).

Preliminaries. We rely on the following standard convergence result for CGM.

Fact 5 (CGM: Convergence rate for objective). *Let \mathbf{X}_\star be an arbitrary solution to (6). For each $t \geq 0$, the matrix \mathbf{X}_t given by the CGM iteration (9)–(11) satisfies*

$$f(\mathcal{A}\mathbf{X}_t) - f(\mathcal{A}\mathbf{X}_\star) \leq \frac{2C}{2+t}.$$

The number C is called the curvature constant [23, Eqn. (3)] of the problem (6).

Theorem 2: A basic convergence result. First, we study the case where the standard CGM iteration (9)–(11) converges. In this setting, we can show that SketchyCGM produces iterates that tend toward a matrix close to the limiting value of CGM.

Proof of Theorem 2. According to the triangle inequality,

$$\mathbb{E} \|\hat{\mathbf{X}}_t - \mathbf{X}_{\text{cgm}}\|_{\text{F}} \leq \mathbb{E} \|\hat{\mathbf{X}}_t - \mathbf{X}_t\|_{\text{F}} + \|\mathbf{X}_t - \mathbf{X}_{\text{cgm}}\|_{\text{F}}.$$

The error bound for reconstruction from the sketch, Theorem 1, shows that

$$\begin{aligned} \mathbb{E} \|\hat{\mathbf{X}}_t - \mathbf{X}_t\|_{\text{F}} &\leq 3\sqrt{2} \|\mathbf{X}_t - [\mathbf{X}_t]_r\|_{\text{F}} \\ &\leq 3\sqrt{2} \|\mathbf{X}_t - [\mathbf{X}_{\text{cgm}}]_r\|_{\text{F}}. \end{aligned}$$

The second inequality holds because $[\mathbf{X}_t]$ is a best rank- r approximation of \mathbf{X}_t with respect to Frobenius norm, whereas $[\mathbf{X}_{\text{cgm}}]_r$ is an undistinguished rank- r matrix.

Combine the last two displays to obtain

$$\begin{aligned} \lim_{t \rightarrow \infty} \mathbb{E} \|\hat{\mathbf{X}}_t - \mathbf{X}_{\text{cgm}}\|_{\text{F}} &\leq \lim_{t \rightarrow \infty} (\mathbb{E} \|\hat{\mathbf{X}}_t - \mathbf{X}_t\|_{\text{F}} + \|\mathbf{X}_t - \mathbf{X}_{\text{cgm}}\|_{\text{F}}) \\ &\leq \lim_{t \rightarrow \infty} 3\sqrt{2} \|\mathbf{X}_t - [\mathbf{X}_{\text{cgm}}]_r\|_{\text{F}} \\ &= 3\sqrt{2} \|\mathbf{X}_{\text{cgm}} - [\mathbf{X}_{\text{cgm}}]_r\|_{\text{F}}. \end{aligned}$$

The second inequality and the last line follow from the limit $\mathbf{X}_t \rightarrow \mathbf{X}_{\text{cgm}}$.

If $\text{rank}(\mathbf{X}_{\text{cgm}}) \leq r$, then $\mathbf{X}_{\text{cgm}} = [\mathbf{X}_{\text{cgm}}]_r$. Therefore, we may conclude that $\mathbb{E} \|\hat{\mathbf{X}}_t - \mathbf{X}_{\text{cgm}}\|_{\text{F}} \rightarrow 0$. \square

Theorem 3: When all solutions are low rank.

Next, we examine the situation where all of the solutions to the model problem (6) have low rank. In this case, we can show that SketchyCGM produces a sequence of approximations that approaches the solution set of the problem. This point does not follow formally from Theorem 2 because CGM need not produce a convergent sequence of iterates.

Proof of Theorem 3. As a consequence of the triangle inequality,

$$\mathbb{E} \text{dist}_{\text{F}}(\hat{\mathbf{X}}_t, S_\star) \leq \mathbb{E} \|\hat{\mathbf{X}}_t - \mathbf{X}_t\|_{\text{F}} + \text{dist}_{\text{F}}(\mathbf{X}_t, S_\star).$$

We claim that the second term $\text{dist}_{\text{F}}(\mathbf{X}_t, S_\star) \rightarrow 0$. If so, then the first term converges to zero:

$$\begin{aligned} \mathbb{E} \|\hat{\mathbf{X}}_t - \mathbf{X}_t\|_{\text{F}} &\leq 3\sqrt{2} \|\mathbf{X}_t - [\mathbf{X}_t]_r\|_{\text{F}} \\ &\leq 3\sqrt{2} \text{dist}_{\text{F}}(\mathbf{X}_t, S_\star) \rightarrow 0. \end{aligned}$$

The first inequality is Theorem 1. The second bound holds because $[\mathbf{X}_t]_r$ is a best rank- r approximation of \mathbf{X}_t , while S_\star is an unremarkable set of rank- r matrices. These observations complete the proof.

Let us turn to the claim. Abbreviate the objective function of (6) as $g = f \circ \mathcal{A}$. The continuous function g attains its minimal value g_\star on the compact feasible set of (6). The standard convergence result for CGM, Fact 5, shows that $g(\mathbf{X}_t) \rightarrow g_\star$. Now, fix a number $\delta > 0$. Define

$$\begin{aligned} E &= \{\mathbf{X} \in \mathbb{R}^{m \times n} : \|\mathbf{X}\|_{S_1} \leq \alpha; \text{dist}(\mathbf{X}, S_\star) \geq \delta\}; \\ v &= \inf\{g(\mathbf{X}) : \mathbf{X} \in E\}. \end{aligned}$$

If E is empty, then $v = +\infty$. Otherwise, the continuous function g attains the value v on the compact set E . In either case, $v > g_\star$ because E contains no optimal point of (6). Since $g(\mathbf{X}_t) \rightarrow g_\star$, we must have $g(\mathbf{X}_t) < v$ for all sufficiently large t . Therefore, $\mathbf{X}_t \notin E$ for large t . We conclude that $\text{dist}_{\text{F}}(\mathbf{X}_t, S_\star) < \delta$ for large t , as required. \square

Theorem 4: Rate of convergence.

Finally, we identify a setting where we can bound the rate at which the estimates produced by SketchyCGM converge to an optimal point of (6). To do so, we assume that the optimal point is stable in the sense that feasible perturbations away from optimality are reflected in increases in the value of the objective function.

Proof of Theorem 4. Since the CGM iterate \mathbf{X}_t is feasible for (6), we can use the stability hypothesis (23) to calculate that

$$\begin{aligned} f(\mathcal{A}\mathbf{X}_t) - f(\mathcal{A}\mathbf{X}_\star) &\geq \kappa \|\mathbf{X}_t - \mathbf{X}_\star\|_{\text{F}}^\nu \\ &\geq \kappa \|\mathbf{X}_t - [\mathbf{X}_t]_r\|_{\text{F}}^\nu \\ &\geq \kappa [(3\sqrt{2})^{-1} \mathbb{E} \|\mathbf{X}_t - \hat{\mathbf{X}}_t\|_{\text{F}}]^\nu. \end{aligned}$$

The second inequality holds because \mathbf{X}_\star has rank r , while $[\mathbf{X}_t]_r$ is a better rank- r approximation of \mathbf{X}_t . The last inequality follows from Theorem 1.

The latter display implies that

$$\begin{aligned} \mathbb{E} \|\hat{\mathbf{X}}_t - \mathbf{X}_\star\| &\leq \mathbb{E} \|\hat{\mathbf{X}}_t - \mathbf{X}_t\| + \|\mathbf{X}_t - \mathbf{X}_\star\| \\ &\leq (3\sqrt{2} + 1) [\kappa^{-1} (f(\mathcal{A}\mathbf{X}_t) - f(\mathcal{A}\mathbf{X}_\star))]^{1/\nu}. \end{aligned}$$

To complete the proof, invoke the standard convergence result for CGM, Fact 5, to bound the quantity $f(\mathcal{A}\mathbf{X}_t) - f(\mathcal{A}\mathbf{X}_\star)$. Last, simplify the constant. \square

B More computational evidence

This appendix elaborates on the phase retrieval experiments described in §6. It also presents additional experiments on matrix completion and phase retrieval.

B.1 Loss functions

Our experiments involve a number of different models for data, so we require several elementary loss functions. Each of these maps is an extended convex function $\psi : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$. Define

$$\begin{aligned} \psi_{\text{gauss}}(z; b) &= \frac{1}{2}(z - b)^2; \\ \psi_{\text{huber}}(z; b) &= \text{huber}(z, b); \\ \psi_{\text{logistic}}(z; b) &= \log(1 + \exp(-bz)); \\ \psi_{\text{poisson}}(z; b) &= z - b \log z. \end{aligned}$$

The objectives correspond, respectively, to the negative log-likelihood of observing the data b under a Gaussian model, a Gauss–Laplace model, a Bernoulli model, and a Poisson model.

B.2 Matrix completion

One principal advantage of SketchyCGM is its flexibility. Without any modification, the algorithm can provably solve any convex optimization problem with a smooth objective and a Schatten 1-norm constraint. To demonstrate this point numerically, we present the results of fitting the benchmark MovieLens 100K and 10M datasets [18] with three different loss functions.

The MovieLens N dataset contains (about) N ratings that users of a website assigned to a collection of movies. Suppose that there are m users and n movies. The data consists of triples (i, j, b) where $i \in \{1, \dots, m\}$ is a user, $j \in \{1, \dots, n\}$ is a movie, and $b \in \{1, 2, 3, 4, 5\}$ is the rating of movie j by user i .

We preprocess the data minimally. We remove the movies that are not rated by any user, as well as the users that have not provided any ratings. To fit a

logistic model, we also binarize the ratings by replacing all values above 3.5 with 1 and the rest with -1 . Thus, the logistic model seeks a classifier that separates high ratings (4 and 5) from low ratings (1, 2, and 3).

We can use low-rank matrix completion to fit a model to the MovieLens data. To see why, introduce a target matrix $\mathbf{X}_\dagger \in \mathbb{R}^{m \times n}$ that tabulates all ratings (known and unknown) of movies by users. One may imagine that \mathbf{X}_\dagger has low rank because a lot of the variation in the ratings is explained by the quality of the movie, its genre, and a user’s preference for that genre. We only observe a subset of the entries of \mathbf{X}_\dagger , and our goal is to predict the remaining entries.

We model this matrix completion problem using the formulation (6). Let E be a training set of user–movie pairs, and let $\mathbf{b} \in \mathbb{R}^E$ list the associated ratings. Introduce the linear map $\mathcal{A} : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^E$ where

$$\mathcal{A}\mathbf{X} = \{x_{ij} : (i, j) \in E\}.$$

The objective function $f : \mathbb{R}^E \rightarrow \mathbb{R}$ has the form

$$f_\ell(\mathbf{z}) = \frac{1}{|E|} \sum_{(i,j) \in E} \psi_\ell(z_{ij}, b_{ij}),$$

where $\ell \in \{\text{gauss}, \text{huber}, \text{logistic}\}$.

Suppose that $\mathbf{X}_\star \in \mathbb{R}^{m \times n}$ is an estimate for the target matrix \mathbf{X}_\dagger . Let E' be the test set of user–movie pairs, with ratings listed in $\mathbf{b}' \in \mathbb{R}^{E'}$. Define the test error

$$\text{test}_\ell(\mathbf{X}_\star) = \frac{1}{|E'|} \sum_{(i,j) \in E'} \psi_\ell((\mathbf{X}_\star)_{ij}, b'_{ij}).$$

Once again, $\ell \in \{\text{gauss}, \text{huber}, \text{logistic}\}$.

For the 100K dataset, we use the default `ub` partition of the data into training and test sets. For each loss function, we sweep α from 3,000 to 10,000 in steps of 500. As expected, the rank of a solution \mathbf{X}_\star of the problem (6) increases with α . We select the value of α that provides the best test error after 10,000 iterations of CGM.

A similar procedure applies to the 10M dataset, with the default `rb` partition. This time, we sweep α from 50,000 to 250,000 in steps of 25,000.

For each dataset and each loss function, we fit the model (6) with the designated value of α by applying SketchyCGM. Figures 3 and 4 show how the test error for the SketchyCGM reconstruction varies as we change the rank parameter r in SketchyCGM. For the 100K dataset, rank $r = 50$ yields test error similar with the CGM solution. For the 10M dataset, rank $r = 200$ yields equivalent performance.

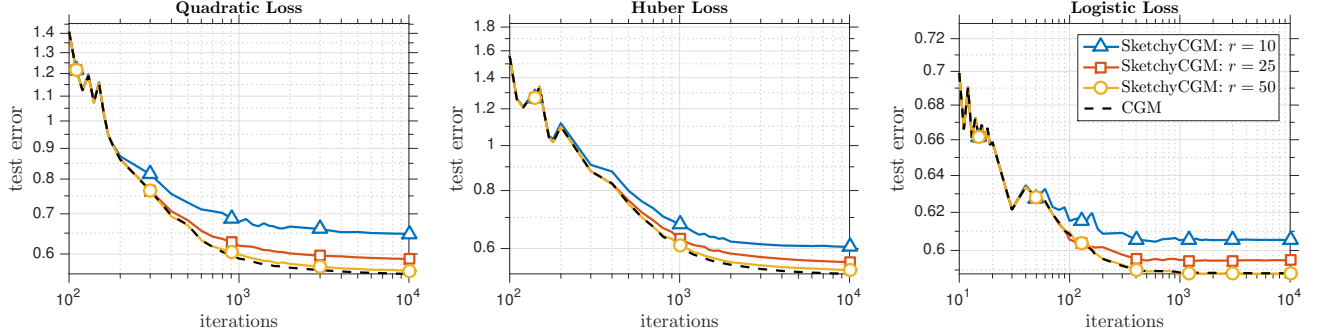


FIGURE 3: *MovieLens 100K Dataset*. Convergence of test error for CGM and SketchyCGM solutions to (6) with three loss functions. The parameter α is chosen by cross-validation on the final CGM solution: $\alpha = 7,000$ for quadratic loss, 7,500 for Huber loss, and 4,500 for logistic loss. See §B.2 for details.

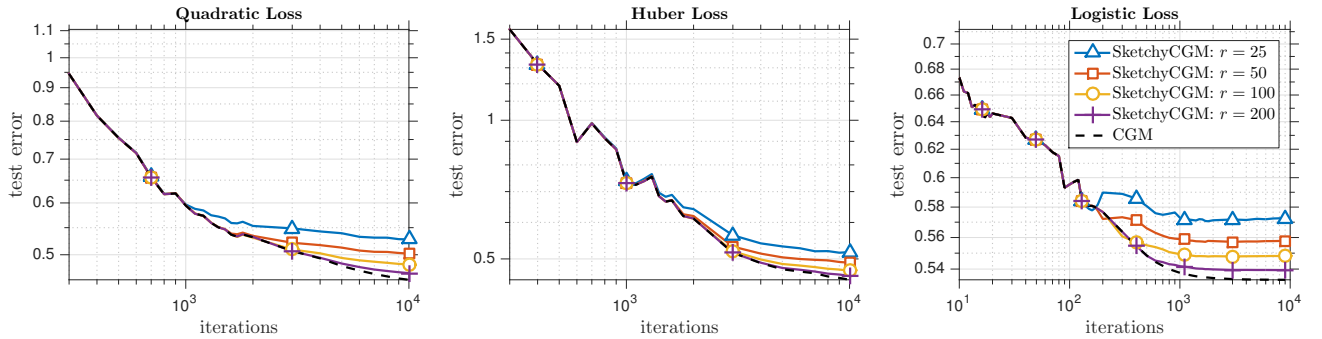


FIGURE 4: *MovieLens 10M Dataset*. Convergence of test error for CGM and SketchyCGM solutions to (6) with three loss functions. The parameter α is chosen by cross-validation on the final CGM solution: $\alpha = 150,000$ for quadratic and logistic losses, and 175,000 for Huber loss. See §B.2 for details.

B.3 Phase retrieval problems

Section 6 showcases the empirical performance of SketchyCGM on a family of phase retrieval problems. This section provides details of these experiments, as well as some related examples.

B.3.1 Overview

The setup is the same as in §1.2. Let $\mathbf{x}_\dagger \in \mathbb{C}^n$ be a vector, and suppose we acquire measurements

$$b_i = |\langle \mathbf{a}_i, \mathbf{x}_\dagger \rangle|^2 + \xi_i \quad \text{for } i = 1, \dots, d. \quad (25)$$

We can modify the measurement vectors $\mathbf{a}_i \in \mathbb{C}^n$ to obtain a range of problems. We can also adjust the distribution of the noise ξ_i .

To model the measurement process (25), it is convenient to form the matrix $\mathbf{A} \in \mathbb{C}^{d \times n}$ whose rows are the measurement vectors \mathbf{a}_i^* . Then define a linear map $\mathcal{A} : \mathbb{C}^{n \times n} \rightarrow \mathbb{C}^d$ and its adjoint $\mathcal{A}^* : \mathbb{C}^d \rightarrow \mathbb{C}^{n \times n}$ via

$$\begin{aligned} \mathcal{A}\mathbf{X} &= \text{diag}(\mathbf{A}\mathbf{X}\mathbf{A}^*); \\ \mathcal{A}^*\mathbf{z} &= \mathbf{A}^* \text{diag}^*(\mathbf{z})\mathbf{A}. \end{aligned} \quad (26)$$

The map $\text{diag} : \mathbb{C}^{d \times d} \rightarrow \mathbb{C}^d$ extracts the diagonal of

a matrix; $\text{diag}^* : \mathbb{C}^d \rightarrow \mathbb{C}^{d \times d}$ maps a vector into a diagonal matrix. When \mathbf{X} is psd, note that $\mathcal{A}\mathbf{X} \in \mathbb{R}_+^d$.

We instantiate the convex optimization template (24) with the linear map (26) and the objective function

$$f_\ell(\mathbf{z}) = \sum_{i=1}^d \psi_\ell(z_i, b_i).$$

Here, the parameter $\ell \in \{\text{gauss}, \text{poisson}\}$. Following [39, Sec. II], we usually set $\alpha = d^{-1} \sum_{i=1}^d b_i$. We approximate the true vector \mathbf{x}_\dagger by means of a maximum eigenvector \mathbf{x}_* of a solution \mathbf{X}_* to (24).

B.3.2 Synthetic phase retrieval

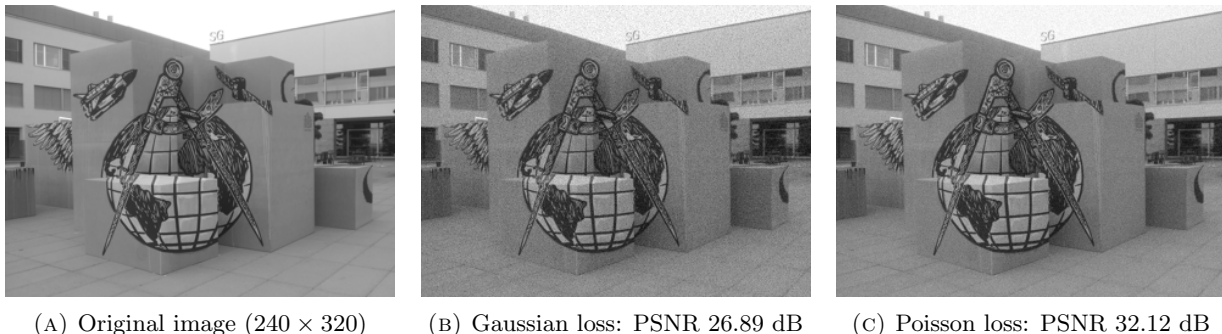
In §6.1, we considered a measurement model based on random coded diffraction patterns. This is a synthetic setup inspired by an imaging system where one can modulate the image before diffraction occurs [9].

For this example, the matrix $\mathbf{A} \in \mathbb{C}^{d \times n}$ appearing in (26) takes the form

$$\mathbf{A} = \begin{bmatrix} \mathbf{F}_n & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{F}_n & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & \mathbf{F}_n \end{bmatrix} \begin{bmatrix} \mathbf{D}_1 \\ \mathbf{D}_2 \\ \vdots \\ \mathbf{D}_s \end{bmatrix}. \quad (27)$$

TABLE 1: *Memory usage.* Approximate memory usage (in bytes) for five convex solvers applied to the convex phase retrieval problem (5), as shown in Figure 1. The dash — indicates that an algorithm ran out of memory.

Signal length (n)	10	10^2	10^3	10^4	10^5	10^6
AT	$2.82 \cdot 10^4$	$1.32 \cdot 10^6$	$1.21 \cdot 10^8$	$1.20 \cdot 10^{10}$	—	—
PGM	$2.90 \cdot 10^4$	$1.07 \cdot 10^6$	$9.70 \cdot 10^7$	$9.61 \cdot 10^9$	—	—
CGM	$6.36 \cdot 10^3$	$1.99 \cdot 10^5$	$1.63 \cdot 10^7$	$1.60 \cdot 10^9$	—	—
ThinCGM	$6.36 \cdot 10^3$	$1.08 \cdot 10^5$	$3.17 \cdot 10^6$	$1.42 \cdot 10^8$	$1.53 \cdot 10^9$	—
SketchyCGM	$1.06 \cdot 10^4$	$9.06 \cdot 10^4$	$8.90 \cdot 10^5$	$8.88 \cdot 10^6$	$8.88 \cdot 10^7$	$8.88 \cdot 10^8$


 FIGURE 5: *Gaussian and Poisson phase retrieval under Poisson noise.* See §B.3.3 for details.

In this expression, $\mathbf{F}_n \in \mathbb{C}^{n \times n}$ is the discrete Fourier transform (DFT) matrix, and $\mathbf{D}_i \in \mathbb{C}^{n \times n}$ are diagonal matrices that describe modulating waveforms. The parameter s represents the number of n -dimensional views of the target vector $\mathbf{x}_i \in \mathbb{C}^n$ that we acquire. The total number of measurements $d = sn$.

We generate each diagonal entry of each matrix \mathbf{D}_i independently at random. Each one is the product of two independent random variables U_1 and U_2 , where U_1 is chosen uniformly from $\{1, i, -1, -i\}$ and U_2 is drawn from $\{\sqrt{2}/2, \sqrt{3}\}$ with probabilities 0.8 and 0.2.

In this setting, we can represent the linear map \mathcal{A} in (26) using $3d$ bits, and we can apply it efficiently using the FFT algorithm.

For the scaling experiments in §6.1, the measurements take the form (25) where the \mathbf{a}_i^* are the rows of (27). We solve (24) with the loss f_{gauss} , the linear map (26)–(27), and α set to the average of the data b_i . The remaining details appear in §6.1. Table 1 summarizes the storage costs for solving this type of synthetic phase retrieval problem with five different convex optimization algorithms.

B.3.3 Synthetic Poisson phase retrieval

In many imaging systems, a Poisson noise model is more appropriate than a Gaussian noise model. Let us demonstrate that the SketchyCGM algorithm can

solve synthetic phase retrieval problems with the loss f_{poisson} . This work highlights the importance of adapting the loss function to the noise distribution.

The setup is similar to §B.3.2. Fix a small image $\mathbf{x}_i \in \mathbb{C}^n$ with $n = 240 \times 320 = 76,800$ pixels. Acquire $d = 20n$ measurements of the form (25) using the coded diffraction model (26)–(27). Each realization $\xi_i \in \mathbb{R}^d$ of the noise is drawn iid from a POISSON distribution, whose mean is chosen so that the SNR of the measurements is 20 dB.

We formulate the phase retrieval problem using the template (24) with the loss f_{poisson} , the linear map (26)–(27), and with α set to the average of the measurements b_i .

To solve this problem via SketchyCGM, it is necessary to make some small modifications [28]. We initialize the algorithm with the dual vector $\mathbf{z}_0 = d^{-1/2}\mathbf{1}$ and set the learning rate $\eta_t = 2/(3+t)$. The rank parameter $r = 1$, and the algorithm runs for 100 iterations.

Figure 5 displays the results of this computation. We also compare the output with a reconstruction obtained by applying the unmodified version of SketchyCGM to solve (24) with the (mismatched) loss f_{gauss} , the same linear map, and the same value of α . As expected, the Poisson formulation performs better.

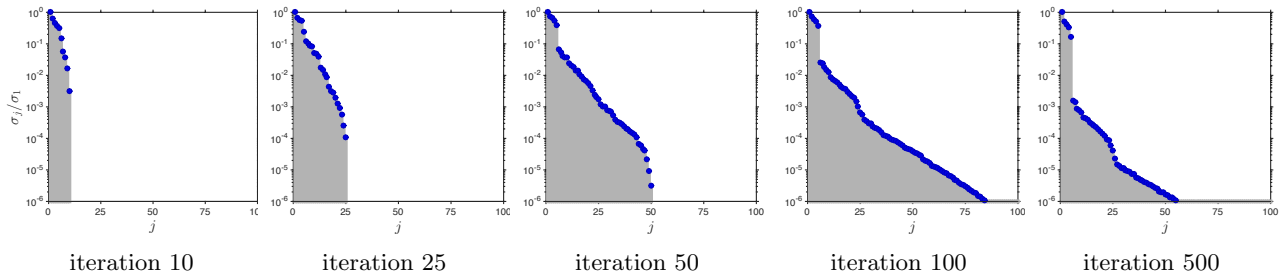


FIGURE 6: *Spectrum of the CGM Iterates.* CGM is applied to the phase retrieval problem (5) with the blood cell image data described in §6.2. These plots show the relative singular value spectrum σ_j/σ_1 for five specific iterates. The solution to (5) appears to have rank five. See §B.3.4 for discussion.

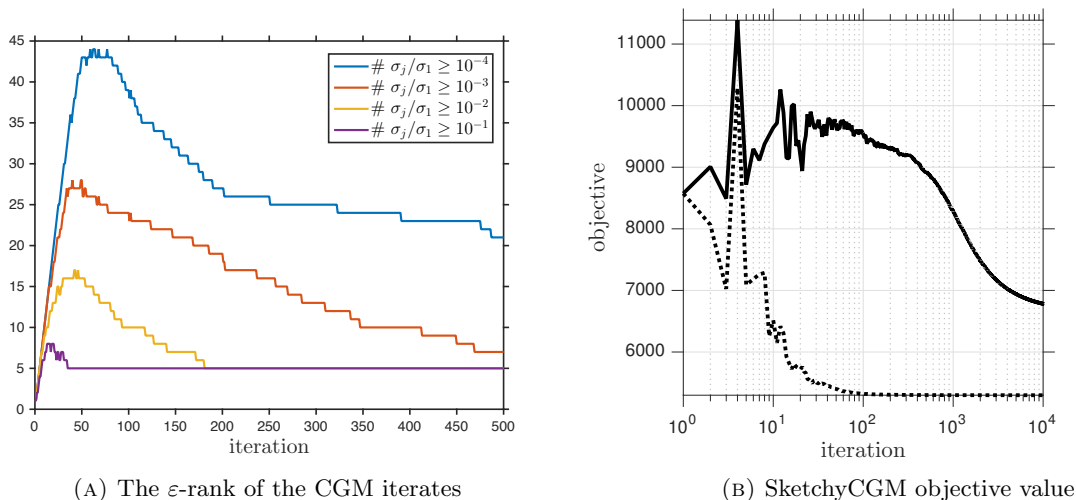


FIGURE 7: *Convergence of CGM and SketchyCGM.* CGM and SketchyCGM are applied to the phase retrieval problem (5) with the blood cell image data described in §6.2. (A) This panel displays the evolution of the ε -rank of the CGM iterates for several choices of ε . (B) This panel shows the objective values achieved by the CGM iterates \mathbf{X}_t (dashed) and the rank-one SketchyCGM estimates $\hat{\mathbf{X}}_t$ (solid). There is a gap because the solution to (5) appears to have rank five; see (A).

B.3.4 Fourier ptychography

In §6.2, we discussed a real-world phase retrieval problem arising from Fourier ptychography [22]. Let us explain the mathematical model for this problem and present some additional numerical work.

For Fourier ptychography, the matrix $\mathbf{A} \in \mathbb{C}^{d \times n}$ appearing in (26) has the following structure.

$$\mathbf{A} = \begin{bmatrix} \mathbf{F}_q^* & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{F}_q^* & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & \mathbf{F}_q^* \end{bmatrix} \begin{bmatrix} \mathbf{D}_1 \\ \mathbf{D}_2 \\ \vdots \\ \mathbf{D}_s \end{bmatrix} \mathbf{F}_n. \quad (28)$$

In this expression, $\mathbf{F}_n \in \mathbb{C}^{n \times n}$ is the 2D discrete Fourier transform, and $\mathbf{F}_q^* \in \mathbb{C}^{q \times q}$ is a low-dimensional 2D discrete Fourier transform. The sparse matrices $\mathbf{D}_i \in \mathbb{C}^{q \times n}$ describe bandpass filters; each column of \mathbf{D}_i has at most one nonzero entry. The

number of measurements $d = sq$. See [22] for details about the physical setup and the mathematical model.

Section 6.2 describes a specific instance of Fourier ptychography imaging, applied to a slide containing red blood cells. To perform phase retrieval, we use the optimization problem (24) with loss f_{gauss} and with $\alpha = 1,400$. The measurement vectors \mathbf{a}_i^* are the rows of (28). We apply several algorithms, including a moderate number of iterations of CGM, SketchyCGM with rank parameter $r = 1$, the Burer–Monteiro method [6, 22], and Wirtinger flow [9].

Figure 6 and 7(A) provide information about the spectrum of the CGM iterates. For several values of ε , we observe that the ε -rank¹ of the iterates becomes large before declining to the value five. This type of behavior is typical for CGM, and it scuttles CGM variants

¹The ε -rank of a matrix is the number of singular values that exceed $\varepsilon\sigma_1$, where σ_j is the j th largest singular value.

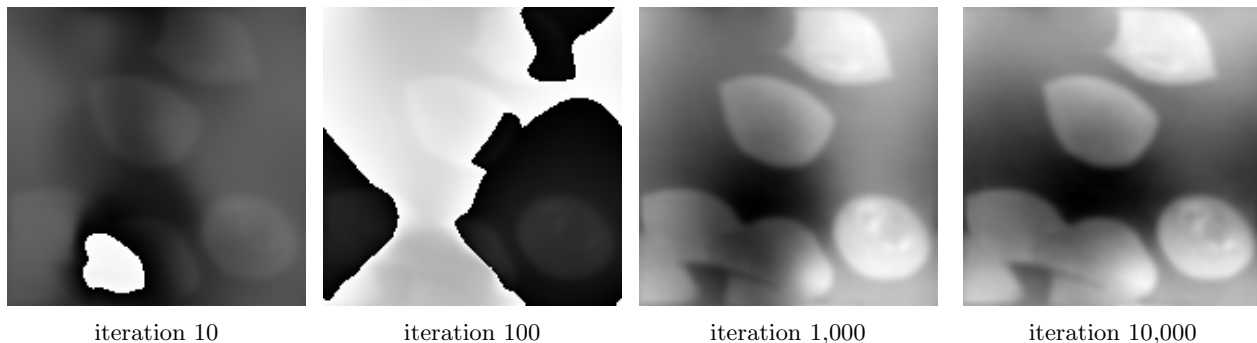


FIGURE 8: *Phase Reconstruction from Fourier Ptychographic Data.* SketchyCGM is applied to the phase retrieval problem (5) with the blood cell image data described in §6.2. These panels show the evolution of the phase reconstruction as a function of the iteration. Brightness indicates the complex phase; only relative differences are meaningful. The image has diagnostic quality after 1,000 iterations, but it continues to sharpen.

that try to control the rank of the iterates directly.

Ideally, the solution \mathbf{X}_* to the convex formulation (5) of a phase retrieval problem has rank one. But these computations suggest that, for the blood cell data, the solution actually has rank five. The increase in rank is due to nonidealities in the imaging system, such as the spatial incoherence of the light source. In essence, the measurements capture a superposition of several slightly different images. Regardless, the convex model (5) is still effective, and a top eigenvector of the solution \mathbf{X}_* still provides a good approximation to the image [22].

Figure 7(B) charts the objective value $f(\mathcal{A}\hat{\mathbf{X}}_t)$ attained by the SketchyCGM iterates. We can implicitly compute the objective value $f(\mathcal{A}\mathbf{X}_t)$ for the CGM iterate \mathbf{X}_t using the loop invariant (22). Note that \mathbf{X}_t achieves a much smaller objective value than the rank-one approximation $\hat{\mathbf{X}}_t$ produced by SketchyCGM. The discrepancy is due to the fact that the solution to the optimization problem has approximate rank five.

Figure 8 displays snapshots of the SketchyCGM iterates as the algorithm proceeds. We see that SketchyCGM already achieve a diagnostic quality image after 1,000 iterations, but the algorithm continues to resolve the image as it runs.

Last, Figure 9 shows the phase gradient of the solution to (5) obtained with three different algorithms; these plots provide an alternative view of Figure 2. Roughly, the phase gradient indicates the change in the thickness of the sample at a given location. Therefore, absolute changes in the value of the phase gradient are meaningful. Note the unphysical oscillations in the reconstruction via Burer–Monteiro [6, 22]. The reconstruction via Wirtinger Flow [9] contains no information at all.

References

- [1] N. Agarwal, B. Bullins, and E. Hazan. Second-order stochastic optimization in linear time. Available at <http://arXiv.org/abs/1602.03943>, Feb. 2016.
- [2] A. Auslender and M. Teboulle. Interior gradient and proximal methods for convex and conic optimization. *SIAM J. Optim.*, 16(3):697–725 (electronic), 2006.
- [3] R. Balan, B. G. Bodmann, P. G. Casazza, and D. Edidin. Painless reconstruction from magnitudes of frame coefficients. *J. Fourier Anal. Appl.*, 15(4): 488–501, 2009.
- [4] S. Bhojanapalli, A. Kyrillidis, and S. Sanghavi. Dropping convexity for faster semi-definite optimization. *J. Mach. Learn. Res.*, 49:1–53, 2016.
- [5] N. Boumal, V. Voroninski, and A.S. Bandeira. The non-convex Burer-Monteiro approach works on smooth semidefinite programs. In *Adv. Neural Information Processing Systems 29*, pages 2757–2765. Curran Associates, Inc., 2016.
- [6] S. Burer and R. D. C. Monteiro. A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization. *Math. Program.*, 95(2, Ser. B): 329–357, 2003.
- [7] E. J. Candès and Y. Plan. Matrix completion with noise. *Proc. IEEE*, 98(6):925–936, 2010.
- [8] E. J. Candès, Y. C. Eldar, T. Strohmer, and V. Voroninski. Phase retrieval via matrix completion. *SIAM J. Imaging Sci.*, 6(1):199–225, 2013.
- [9] E. J. Candès, X. Li, and M. Soltanolkotabi. Phase retrieval via Wirtinger Flow: Theory and algorithms. *IEEE Trans. Inform. Theory*, 61(4):1985–2007, 2015.
- [10] A. Chai, M. Moscoso, and G. Papanicolaou. Array imaging using intensity-only measurements. *Inverse Problems*, 27(1):015005, 2011.
- [11] K. L. Clarkson. Coresets, sparse greedy approximation, and the Frank-Wolfe algorithm. *ACM Trans. Algorithms*, 6(4):Art. 63, 30, 2010.
- [12] M. A. Erdoğan and A. Montanari. Convergence rates of sub-sampled Newton methods. In *Adv. Neural Information Processing Systems 28*, pages 3052–3060. Curran Associates, Inc., 2015.

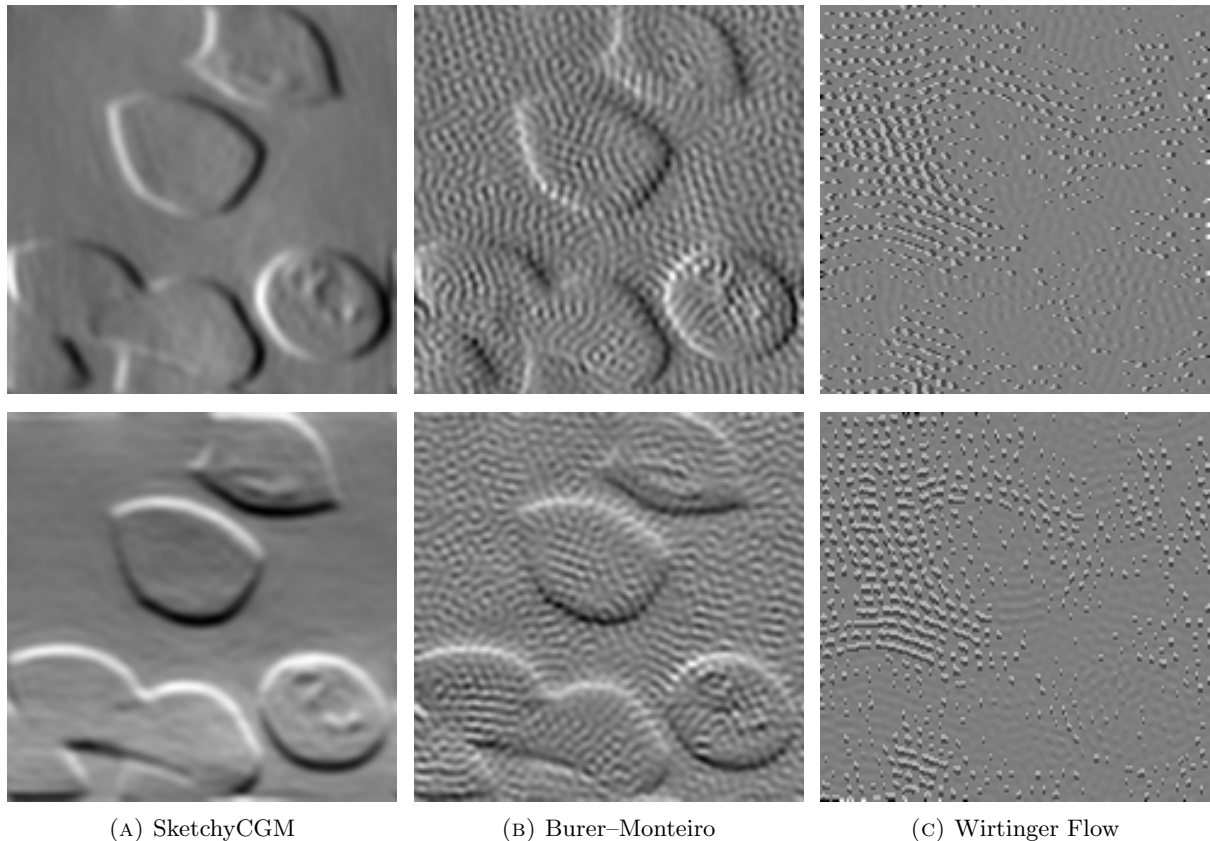


FIGURE 9: *Phase Gradients for Fourier Ptychographic Imaging.* Three algorithms are applied to the phase retrieval problem (5) with the blood cell image data described in §6.2. [top] The horizontal differences of the phase maps presented in Figure 2. [bottom] The vertical differences of the phase maps. Brightness roughly corresponds to a change in thickness of the sample.

- [13] M. Fazel. *Matrix rank minimization with applications.* PhD Dissertation, Stanford Univ., Palo Alto, 2002.
- [14] J. R. Fienup. Phase retrieval algorithms: a comparison. *Appl. Optics*, 21(15):2758–2769, 1982.
- [15] M. Frank and P. Wolfe. An algorithm for quadratic programming. *Naval Res. Logist. Quart.*, 3:95–110, 1956.
- [16] M. X. Goemans and D. P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. Assoc. Comput. Mach.*, 42(6):1115–1145, 1995.
- [17] N. Halko, P. G. Martinsson, and J. A. Tropp. Finding structure with randomness: probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Rev.*, 53(2):217–288, 2011.
- [18] F. M. Harper and J. A. Konstantin. The MovieLens datasets: history and context. *ACM. Trans. Interactive Intell. Sys.*, 5(4):Article 19, Dec. 2015. doi: <http://dx.doi.org/10.1145/2827872>.
- [19] E. Hazan. Sparse approximate solutions to semidefinite programs. In *Proc. 8th Latin American Theoretical Informatics Symposium*, Apr. 2008.
- [20] C. Helmberg and F. Rendl. A spectral bundle method for semidefinite programming. *SIAM J. Optim.*, 10(3):673–696, 2000.
- [21] C. Helmberg, F. Rendl, R. J. Vanderbei, and H. Wolkowicz. An interior-point method for semidefinite programming. *SIAM J. Optim.*, 6(2):342–361, 1996.
- [22] R. Horstmeyer, R. Y. Chen, X. Ou, B. Ames, J. A. Tropp, and C. Yang. Solving ptychography with a convex relaxation. *New J. Physics*, 17(5):053044, 2015.
- [23] M. Jaggi. Revisiting Frank–Wolfe: Projection-free sparse convex optimization. In *Proc. 30th Intl. Conf. Machine Learning*, Atlanta, 2013.
- [24] P. Jain, P. Netrapalli, and S. Sanghavi. Low-rank matrix completion using alternating minimization. In *STOC '13: Proc. 45th Ann. ACM Symp. Theory of Computing*, pages 665–674, Palo Alto, Dec. 2013.
- [25] E. S. Levitin and B. T. Poljak. Minimization methods in the presence of constraints. *Ž. Vychisl. Mat. i Mat. Fiz.*, 6:787–823, 1966. ISSN 0044-4669.
- [26] M. W. Mahoney. Randomized algorithms for matrices and data. *Foundations and Trends® in Machine Learning*, 3(2):123–224, 2011.
- [27] Netrapalli, P. Jain, and S. Sanghavi. Phase retrieval using alternating minimization. *IEEE Trans. Signal Processing*, 63(18):4814–4826, 2015.
- [28] G. Odor, Y.-H. Li, A. Yurtsever, Y.-P. Hsieh, Q. Tran-Dinh, M. El Halabi, and V. Cevher. Frank-Wolfe

- works for non-Lipschitz continuous gradient objectives: scalable Poisson phase retrieval. In *41st IEEE Intl. Conf. Acoustics, Speech and Signal Processing (ICASSP)*, 2016.
- [29] M. Pilanci and M. Wainwright. Newton sketch: A linear-time optimization algorithm with linear-quadratic convergence. Available at <http://arXiv.org/abs/1505.02250>, May 2015.
- [30] M. Pilanci and M. J. Wainwright. Randomized sketches of convex programs with sharp guarantees. *IEEE Trans. Inform. Theory*, 61(9):5096–5115, 2015.
- [31] N. Rao, P. Shah, and S. Wright. Conditional gradient with enhancement and truncation for atomic-norm regularization. Available at http://people.inf.ethz.ch/jaggim/NIPS-workshop-FW-greedy/papers/rao_shah_wright_final.pdf, 2015.
- [32] R. T. Rockafellar. Monotone operators and the proximal point algorithm. *SIAM J. Control Optim.*, 14(5):877–898, 1976.
- [33] T. Sarlós. Improved approximation algorithms for large matrices via random projections. In *Proc. 47th Ann. IEEE Symp. Foundations of Computer Science*, Berkeley, 2006.
- [34] N. Srebro, J. Rennie, and T. Jaakkola. Maximum-margin matrix factorizations. In *Adv. Neural Information Processing Systems 17*, Vancouver, Dec. 2004.
- [35] K.-C. Toh and S. Yun. An accelerated proximal gradient algorithm for nuclear norm regularized linear least squares problems. *Pacific J. Optim.*, 6(615-640):15, 2010.
- [36] J. A. Tropp. Convex recovery of a structured signal from independent random linear measurements. In G. Pfander, editor, *Sampling Theory, a Renaissance*, chapter 2, pages 67–101. Birkhäuser, 2015.
- [37] J. A. Tropp, A. Yurtsever, M. Udell, and V. Cevher. Randomized single-view algorithms for low-rank matrix reconstruction. ACM Report 2017-01, Caltech, Jan. 2017. Available at <http://arXiv.org/abs/1609.00048>.
- [38] D. P. Woodruff. Sketching as a tool for numerical linear algebra. *Found. Trends Theor. Comput. Sci.*, 10(1-2):iv+157, 2014.
- [39] A. Yurtsever, Y.-P. Hsieh, and V. Cevher. Scalable convex methods for phase retrieval. In *6th IEEE Intl. Workshop on Computational Advances in Multi-Sensor Adaptive Processing*, 2015.