

# Scalable Convex Methods for Low-Rank Matrix Problems

Alp Yurtsever<sup>†</sup>, Quoc Tran-Dinh<sup>‡</sup> and Volkan Cevher<sup>†</sup>

<sup>†</sup>Laboratory for Information and Inference Systems, EPFL, Switzerland

<sup>‡</sup>Department of Statistics and Operations Research, UNC, USA

**Abstract**—We describe our new scalable primal-dual convex optimization framework to solve low-rank matrix recovery problems. The main characteristics of our framework is the cheap per-iteration complexity and the low-memory footprint. We demonstrate the flexibility and scalability of our framework by solving matrix completion, quantum tomography and phase retrieval problems.

## I. INTRODUCTION

Low-rank matrix recovery problems have seen many applications recently, which include but are not limited to the matrix completion [1], clustering [2], robust subspace learning [3], quantum tomography [4] and the phase retrieval via convex semidefinite relaxations [5].

In convex optimization, low-rankness is induced by the nuclear norm (a.k.a., Schatten 1-norm,  $\|\cdot\|_{S_1}$ ), which can appear in the objective function or in the constraints. However, computation and storage bottlenecks restrict the scalability of the state of the art convex solvers for these problems. In this work, we describe a scalable convex primal-dual algorithmic framework to solve low-rank matrix problems. More details about this work can be found in [6] and [7].

## II. ALGORITHM & CONTRIBUTIONS

Our algorithmic framework (universal primal-dual gradient method, abbreviated as UniPDGrad) extends Nesterov’s universal gradient methods [8] for the primal-dual setting in a non-trivial fashion. We develop a new accelerated variant (AccUniPDGrad) adopting FISTA scheme, which requires less computation compared to the fast scheme in [8]. See [6] for detailed description of our algorithmic scheme and the convergence guarantees. We summarize the characteristics of our framework below.

### A. Computational efficiency

Conditional gradient methods (CGM, a.k.a. Frank-Wolfe-type methods [9]) attracted increasing interest in recent years (see [10], [11] and the references therein). Its popularity comes from its cheap iteration cost, especially when the feasible set is a simple polytope.

For low-rank matrix problems, at each iteration, the main competitors of CGM (projected or proximal gradient methods) often require a full singular value decomposition. On the other side, CGM leverages the so called linear minimization oracles, which corresponds to the computing a leading eigenvector which is much cheaper [10].

Linear minimization oracle of CGM is indeed a special case of a the more general Fenchel-type operators:

$$\text{sharp}_f(x) = \arg \max_s \{ \langle s, x \rangle - f(s) \}.$$

Our algorithmic framework also leverages Fenchel-type operators, hence it has similar iteration cost as CGM.

### B. Flexibility

CGM relies on some assumptions about the smoothness of the objective function [10], [11]. Hence, when using CGM, we typically consider the following formulation:

$$\underset{X \in \mathcal{X}}{\text{minimize}} \quad \frac{1}{2} \|A(X) - b\|_2^2 \quad \text{subj. to} \quad \|X\|_{S_1} \leq \kappa, \quad (1)$$

where  $A : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^d$  is a linear map,  $\kappa$  is a tuning parameter, and  $\mathcal{X}$  is a convex set.

Our algorithmic framework also applies to the following flipped formulation which is beyond the scope of CGM:

$$\underset{X \in \mathcal{X}}{\text{minimize}} \quad \frac{1}{2} \|X\|_{S_1}^2 \quad \text{subj. to} \quad A(X) - b \in \mathcal{K}, \quad (2)$$

where  $\mathcal{K}$  is a simple convex set. Note that this formulation can be easier to tune with an accurate noise model.

### C. Storage efficiency

While the solution has  $\mathcal{O}(r(m+n))$  degrees of freedom ( $r$  denotes the rank of solution), a random dense iterate requires  $\mathcal{O}(mn)$  storage. Our framework has a dual description that requires less storage. Note however that we need to keep track of the decision variable which comes as a stream of rank-1 updates. Hence, we can use the thin SVD updating method described in [12] to maintain the SVD of the decision variable rather than computing it in the ambient dimensions.

## III. NUMERICAL EXPERIMENTS

We assess the empirical performance of our algorithm solving matrix completion, quantum tomography and phase retrieval problems.

### A. Matrix Completion with MovieLens Dataset

We estimate a low-rank matrix  $X \in \mathbb{R}^{m \times n}$  from its subsampled entries  $b \in \mathbb{R}^d$ , where  $A(\cdot)$  is the sampling operator. We apply our algorithms to (1) and (2) using the MovieLens 100K dataset using the default `ub` test and training data partitions. CGM cannot handle (2) and only solve (1). Our metrics are the normalized objective residual and the root mean squared error (RMSE) calculated for the test data. Results are shown in Figure 1. See [6, Section 5.2] for details.

### B. Quantum Tomography with Pauli Operators

A  $q$ -qubit quantum system is mathematically characterized by its density matrix, which is a complex  $p \times p$  positive semidefinite Hermitian matrix  $X^\natural$ , where  $p = 2^q$ . We can deduce the state from performing compressive linear measurements  $b = A(X) \in \mathbb{C}^d$  based on Pauli operators [4]. We generate a random pure quantum state (e.g., rank-1  $X^\natural$ ), and we take  $n = 2p \log p$  random Pauli measurements. For  $q = 14$  qubits system, this corresponds to a  $268'435'456$  dimensional problem with  $n = 317'983$  measurements. We compare our algorithm against CGM. Results are shown in Figure 2. See [6, Section 5.1] for details.

### C. Phase Retrieval with Coded Random Diffraction Pattern

We consider coded diffraction pattern measurements with octonary modulation, using the random design of the modulating waveforms from [13] with  $L = 20$  random waveforms. We compare the computational time to reach  $10^{-2}$  reconstruction error with different approaches. We also test our framework on HD and full HD images. Results are shown in Figures 3 to 5. See [7] for more details.

## ACKNOWLEDGMENT

This work was supported in part by ERC Future Proof, SNF 200021-146750, SNF CRSII2-147633, and NCCR-Marvel.

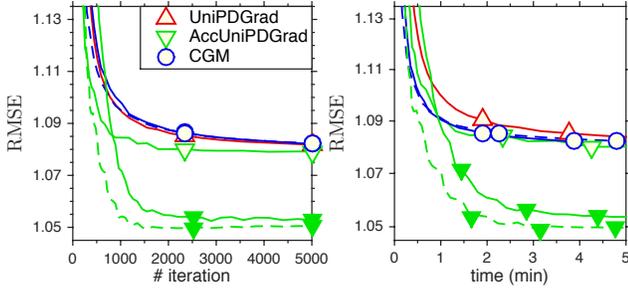


Fig. 1. The performance of the algorithms for the MC problems. The dashed lines correspond to the line-search (in the weighting step) variants. The empty and filled markers correspond to the formulations (1) and (2) respectively, with  $\mathcal{X} = \mathbb{R}^{m \times n}$ . For (1), we choose the tuning parameter  $\kappa = 9975/2$  as in [14]. For (2), we consider equality constraint in linear inclusion, i.e.,  $\mathcal{K} = \{0\}$ . We set accuracy input  $\epsilon = 10^{-3}$  for (Acc)UniPDGrad.

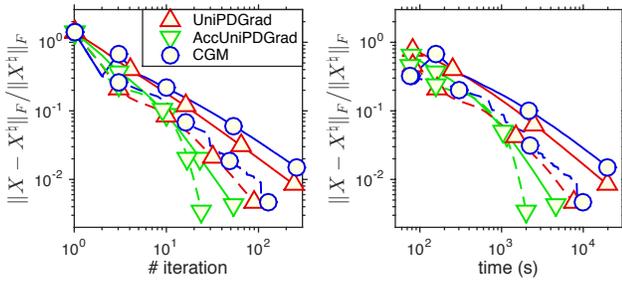


Fig. 2. The convergence behavior of algorithms for the  $q = 14$  qubits quantum tomography problem. We solve a variant of (1) where  $\mathcal{X}$  is the positive semidefinite cone in  $\mathbb{C}^{n \times n}$ , and the nuclear norm constraint is replaced with  $\text{trace}(X) = 1$ . The solid lines correspond to the theoretical weighting scheme, and the dashed lines correspond to a variant which uses greedy weights for the primal averaging step. We set accuracy input  $\epsilon = 10^{-2}$  for (Acc)UniPDGrad.

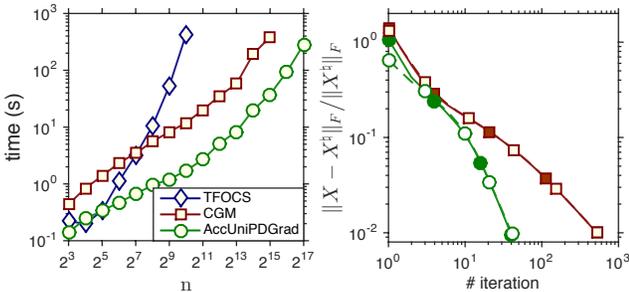


Fig. 3. The empirical comparison of the algorithms for solving phase retrieval problem. We solve a (1),  $\mathcal{X}$  is the positive semidefinite cone in  $\mathbb{C}^{n \times n}$ , and we set  $\kappa = \text{mean}(b)$  as suggested in [7]. (Left) presents the time required by each algorithm to reach  $10^{-2}$  reconstruction error with respect to the data size. We compare against the CGM and the solver provided in [13] which is based on TFOCS software package [15]. (Right) illustrates the convergence behavior of the algorithms in the tests with real images. Empty and filled markers correspond to EPFL and Milky Way images in Figures 4 and 5 respectively.

## REFERENCES

[1] E. J. Candès, and B. Recht, *Exact Matrix Completion via Convex Optimization*, Foundations of Computational Mathematics, vol. 9, no. 6, pp. 717, 2009.

[2] P. Awasthi, A. S. Bandeira, M. Charikar, R. Krishnaswamy, S. Villar, and R. Ward, *Relax, No Need to Round: Integrality of Clustering Formulations* in Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science, pp. 191–200, New York, USA, 2015.



Fig. 4. EPFL image of size  $1280 \times 720$ , reconstructed by solving a convex problem in  $\mathbb{C}^{n \times n}$  with  $n = 921'600$ , by 41 iterations of the AccUniPDGrad that takes 20 minutes: PSNR = 45.54 dB.



Fig. 5. Milky Way image of size  $1920 \times 1080$ , reconstructed by solving a convex problem in  $\mathbb{C}^{n \times n}$  with  $n = 2'073'600$ , by 40 iterations of the AccUniPDGrad that takes 42 minutes: PSNR = 54.44 dB.

[3] Y.-M. Cheung and J. Lou, *Scalable Spectral  $k$ -Support Norm Regularization for Robust Low Rank Subspace Learning*, in Proceedings of the 25th ACM International on Conference on Information and Knowledge Management, pp. 1151–1160, New York, USA, 2016.

[4] D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker, and J. Eisert, *Quantum State Tomography via Compressed Sensing*, Physical Review Letters, vol. 105, pp. 150401, Oct. 2010.

[5] E. J. Candès, T. Strohmer, and V. Voroninski, *PhaseLift: Exact and Stable Signal Recovery from Magnitude Measurements via Convex Programming*, IEEE Transactions on Signal Processing, vol. 60, pp. 2422–2432, 2012.

[6] A. Yurtsever, Q. Tran-Dinh, and V. Cevher, *A Universal Primal-Dual Convex Optimization Framework*, In Advances in Neural Information Processing Systems 28, pp. 3132–3140, 2015.

[7] 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 (CAMSAP), 2015.

[8] Y. Nesterov, *Universal Gradient Methods for Convex Optimization Problems*, Mathematical Programming, vol. 152, pp. 381–404, Aug. 2015.

[9] M. Frank and P. Wolfe, *An Algorithm for Quadratic Programming*, Naval Research Logistics Quarterly, vol. 3, pp. 95–110, 1956.

[10] M. Jaggi, *Revisiting Frank-Wolfe: Projection-Free Sparse Convex Optimization*, Journal of Machine Learning Research Workshop and Conference Proceedings, vol. 28, pp. 427–435, 2013.

[11] Y. Nesterov, *Complexity bounds for primal-dual methods minimizing the model of objective function*, CORE, University Catholique Louvain, Belgium, Tech. Rep., 2015.

[12] M. Brand, *Fast low rank modifications of thin singular value decomposition*, Linear Algebra and its Applications, vol. 415, pp. 20–30, 2006.

[13] E. J. Candès, X. Li, and M. Soltanolkotabi, *Phase retrieval from coded diffraction patterns*, Appl. Comput. Harmon. Anal., vol. 39, pp. pp. 277–299, 2015.

[14] M. Jaggi and M. Sulovský, *A Simple Algorithm for Nuclear Norm Regularized Problems*, in Proceedings of the 27th International Conference on Machine Learning (ICML2010), pp. 471–478, 2010.

[15] S. R. Becker, E. J. Candès, and M. C. Grant, *Templates for convex cone problems with applications to sparse signal recovery*, Mathematical Programming Computation, vol. 3, pp. 165–218, 2011.