
Robust Recovery via Implicit Bias of Discrepant Learning Rates for Double Over-parameterization

Chong You^{†*} Zhihui Zhu^{‡*} Qing Qu[#] Yi Ma[†]

[†]Department of EECS, University of California, Berkeley

[‡]Department of Electrical and Computer Engineering, University of Denver

[#]Center for Data Science, New York University

Abstract

Recent advances have shown that implicit bias of gradient descent on over-parameterized models enables the recovery of low-rank matrices from linear measurements, even with no prior knowledge on the intrinsic rank. In contrast, for *robust* low-rank matrix recovery from *grossly corrupted* measurements, over-parameterization leads to overfitting without prior knowledge on both the intrinsic rank and sparsity of corruption. This paper shows that with a *double over-parameterization* for both the low-rank matrix and sparse corruption, gradient descent with *discrepant learning rates* provably recovers the underlying matrix even without prior knowledge on neither rank of the matrix nor sparsity of the corruption. We further extend our approach for the robust recovery of natural images by over-parameterizing images with deep convolutional networks. Experiments show that our method handles different test images and varying corruption levels with a single learning pipeline where the network width and termination conditions do not need to be adjusted on a case-by-case basis. Underlying the success is again the implicit bias with discrepant learning rates on different over-parameterized parameters, which may bear on broader applications. Our code is available at <https://github.com/ChongYou/robust-image-recovery>.

1 Introduction

Learning *over-parameterized models*, which have more parameters than the problem’s intrinsic dimension, is becoming a crucial topic in machine learning [1–11]. While classical learning theories suggest that over-parameterized models tend to *overfit* [12], recent advances showed that an optimization algorithm may produce an *implicit bias* that regularizes the solution with desired properties. This type of results has led to new insights and better understandings on gradient descent for solving several fundamental problems, including logistic regression on linearly separated data [1], compressive sensing [2, 3], sparse phase retrieval [4], nonlinear least-squares [5], low-rank (deep) matrix factorization [6–9, 13], and deep linear neural networks [10, 11], etc.

Inspired by these recent advances [1–11], in this work we present a new type of practical methods for *robust* recovery of *structured* signals via model over-parameterization. In particular, we aim to learn an unknown signal $\mathbf{X}_* \in \mathbb{R}^{n \times n}$ from its *grossly corrupted* linear measurements

$$\mathbf{y} = \mathcal{A}(\mathbf{X}_*) + \mathbf{s}_*, \quad (1)$$

*Correspondence to Chong You (cyou@berkeley.edu) and Zhihui Zhu (zhihui.zhu@du.edu). The first two authors contributed equally to this work.

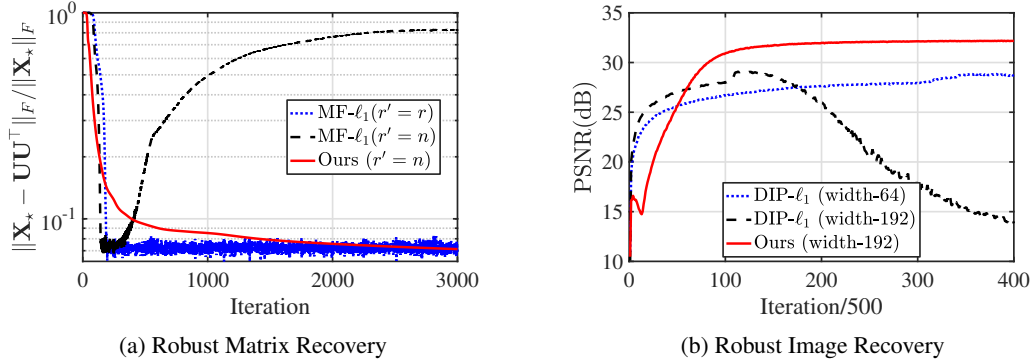


Figure 1: **Learning curves for robust recovery of low-rank matrices (a) and natural images (b).** (a) Classical matrix factorization (MF) method with ℓ_1 penalty requires exact parameterization (left blue), otherwise over-parameterization leads to overfitting without early termination (left black). (b) Previous deep image prior (DIP) method with ℓ_1 penalty requires tuning network width (with width = 64, right blue) or early termination (with width = 192, right black). For both problems, our double over-parameterization (DOP) method achieves superior performance and requires neither early termination nor precise parameterization (red curves).

where the linear operator $\mathcal{A}(\cdot) : \mathbb{R}^{n \times n} \mapsto \mathbb{R}^m$, and $\mathbf{s}_* \in \mathbb{R}^m$ is a (sparse) corruption vector. Variants of the problem appear ubiquitously in signal processing and machine learning [14–19].

Robust recovery of low-rank matrices. The recovery of a low-rank matrix \mathbf{X}_* has broad applications in face recognition (where self-shadowing, specularity, or saturation in brightness can be modeled as outliers), video surveillance (where the foreground objects are usually modeled as outliers) and beyond [14]. A classical method for low-rank matrix recovery is via *nuclear norm* minimization¹, which is provably correct under certain incoherent conditions [14, 20]. However, minimizing nuclear norm involves expensive computations of singular value decomposition (SVD) of n -by- n matrices [21], which prohibits its application to problem size of practical interest.

The computational challenge has been addressed by recent development of matrix factorization (MF) methods [18, 22]. Such methods are based on parameterizing the signal $\mathbf{X} \in \mathbb{R}^{n \times n}$ via the factorization $\mathbf{X} = \mathbf{U}\mathbf{U}^\top$ [23], and solving the associated nonconvex optimization problem

$$\min_{\mathbf{U} \in \mathbb{R}^{n \times r}} \|\mathcal{A}(\mathbf{U}\mathbf{U}^\top) - \mathbf{y}\|_1. \quad (2)$$

In above, the ℓ_1 -norm penalty is adopted to account for sparse noise² \mathbf{s}_* , and r denotes the rank of \mathbf{X}_* . While theoretical guarantees for the correctness of (2) can be established [18, 29–31] (see Figure 1a, dotted blue curve), the result relies on the exact information of the rank of \mathbf{X}_* , which is usually *not* available *a priori* in practice. On the other hand, if $\mathbf{U} \in \mathbb{R}^{n \times r}$ in (2) is replaced with an *over-parameterized* factor $\mathbf{U} \in \mathbb{R}^{n \times n}$, for which r is no longer required, the method easily *overfits* the corruptions and does not produce the desired solution \mathbf{X}_* , unless we do some *early stopping* [32, 33] (see Figure 1a, dashed black curve).

Robust recovery of natural images. Robust recovery of natural images \mathbf{X}_* is often considered as a challenging task due to the lack of *universal* mathematical modeling for natural images. While sparse and low-rank based methods have been demonstrated to be effective for years [34–39], the state-of-the-art performance is obtained by learned priors with deep convolutional neural networks. Such methods operate by end-to-end training of neural networks from pairs of corrupted/clean images [40–42], and often *cannot* effectively handle test cases with corruption type and noise level that are different from those of the training data.

Recently, this challenge has been partially addressed by the so-called *deep image prior* (DIP) [43], which has shown impressive results on many image reconstruction tasks. The method is based on parameterizing an image \mathbf{X} by a deep network $\mathbf{X} = \phi(\boldsymbol{\theta})$, which turns out to be a flexible prior for

¹Nuclear norm (the tightest convex envelope to matrix rank) is defined as the sum of singular values.

²This is a commonly used approach for robust optimization problems, such as robust regression [24], robust subspace learning [25–27], robust phase retrieval [17, 28], robust matrix recovery [16, 19], and many more.

modeling the underlying distributions of a variety of natural images. The network $\phi(\theta)$ has a U-shaped architecture and may be viewed as a multi-layer, nonlinear extension of the low-rank matrix factorization $\mathbf{U}\mathbf{U}^\top$. Hence, it may not come as a surprise that DIP inherits the drawbacks of the MF approaches for low-rank matrix recovery. Namely, it requires either a meticulous choice of network width [44] (see Figure 1b, dotted blue curve) or early termination of the learning process [33] (see Figure 1b, dashed black curve).

Overview of our methods and contributions. We introduce a *double over-parameterization (DOP)* framework for robust recovery that effectively addresses the issues associated with MF and DIP. For the low-rank matrix recovery problem, our method is based on over-parameterizing both the low-rank matrix \mathbf{X}_* (with matrix factorization) and the sparse corruption \mathbf{s}_* (with Hadamard product). This gives rise to a highly over-parameterized formulation that has infinitely many global solutions. We further present a gradient descent algorithm with *discrepant learning rates* on different optimization variables, and show that it introduces an implicit algorithmic bias that enables correct recovery of $(\mathbf{X}_*, \mathbf{s}_*)$ with provable guarantees. The benefit of such an approach *w.r.t.* the state of the art is summarized as follows (see also Table 1):

- *More scalable.* Our method is based on gradient descent only and does not require computing an SVD per iteration as in convex approaches [14]. Hence it is significantly more scalable.
- *Prior knowledge free.* Our method requires *no* prior knowledge on neither the rank of \mathbf{X}_* nor the sparsity of \mathbf{s}_* , and is free of parameter tuning. This is *in contrast* to existing nonconvex approaches [16, 19, 30, 45] for which the true rank and/or sparsity are required *a priori*.
- *Provably correct.* Under similar conditions of the convex approach, our method converges to the ground truth $(\mathbf{X}_*, \mathbf{s}_*)$ asymptotically.

Underlying the success of our method is the notion of *implicit bias of discrepant learning rates*. The idea is that the algorithmic low-rank and sparse regularizations need to be balanced for the purpose of identifying the underlying rank and sparsity. With a lack of means of tuning a regularization parameter, we show in Section 2.3 that

the desired balance can be obtained by using *different* learning rates for different optimization parameters. Such a property may be of separate interest and bear on a broader range of problems.

Finally, we demonstrate the broad applicability of the DOP framework for the task of the robust recovery of natural images. We only need to replace the over-parameterization $\mathbf{U}\mathbf{U}^\top$ for low-rank matrices by a U-shaped network $\phi(\theta)$ for natural images (as adopted in DIP [43]) and solve the optimization problem by gradient descent with discrepant learning rates. This produces a powerful and easy-to-use method with favorable properties when compared to the original DIP (see Figure 1b). To be precise, our method handles different test images and varying corruption levels with a single learning pipeline, in which network width and termination conditions do not need to be adjusted on a case-by-case basis.

2 Main Results and Algorithms

2.1 A Double Over-Parameterization (DOP) Formulation

As precluded in (1), we first consider the problem of recovering a rank- r ($r \ll n$) positive semidefinite matrix³ $\mathbf{X}_* \in \mathbb{R}^{n \times n}$ from its corrupted linear measurements $\mathbf{y} = \mathcal{A}(\mathbf{X}_*) + \mathbf{s}_*$, where $\mathbf{s}_* \in \mathbb{R}^m$ is a vector of sparse corruptions. Recent advances on algorithmic implicit bias for optimizing over-parameterized models [2, 3, 6–9] motivate us to consider a nonlinear least squares for robust matrix recovery, with over-parameterization $\mathbf{X} = \mathbf{U}\mathbf{U}^\top$ and $\mathbf{s} = \mathbf{g} \circ \mathbf{g} - \mathbf{h} \circ \mathbf{h}$:

$$\min_{\mathbf{U} \in \mathbb{R}^{n \times r}, \{\mathbf{g}, \mathbf{h}\} \subseteq \mathbb{R}^m} f(\mathbf{U}, \mathbf{g}, \mathbf{h}) := \frac{1}{4} \left\| \mathcal{A}(\mathbf{U}\mathbf{U}^\top) + (\mathbf{g} \circ \mathbf{g} - \mathbf{h} \circ \mathbf{h}) - \mathbf{y} \right\|_2^2, \quad (3)$$

³By a lifting trick such as [46, 47], our method can be extended to handling arbitrary rectangular matrices.

where the dimensional parameter $r' \geq r$ and “ \circ ” denotes the Hadamard (i.e., entrywise) product. In practice, the choice of r' depends on how much prior information we have for \mathbf{X}_* : r' can be either taken as an estimated upper bound for r , or taken as $r' = n$ with no prior knowledge. In the following, we provide more backgrounds for the choice of our formulation (3).

- *Implicit low-rank prior via matrix factorization.* For the *vanilla* low rank matrix recovery problem with no outlier (i.e., $\mathbf{s}_* = \mathbf{0}$), the low-rank matrix \mathbf{X}_* can be recovered via over-parameterization $\mathbf{X} = \mathbf{U}\mathbf{U}^\top$ [6–9, 48]. In particular, the work [6] showed that with small initialization and infinitesimal learning rate, gradient descent on

$$\min_{\mathbf{U} \in \mathbb{R}^{n \times n}} \frac{1}{2} \|\mathcal{A}(\mathbf{U}\mathbf{U}^\top) - \mathbf{y}\|_2^2 \quad (4)$$

converges to the minimum nuclear norm solution under certain commute conditions for $\mathcal{A}(\cdot)$.

- *Implicit sparse prior via Hadamard multiplication.* For the classical sparse recovery problem [49, 50] which aims to recover a sparse $\mathbf{s}_* \in \mathbb{R}^m$ from its linear measurement $\mathbf{b} = \mathbf{A}\mathbf{s}_*$, recent work [2, 3] showed that it can also be dealt with via over-parameterization $\mathbf{s} = \mathbf{g} \circ \mathbf{g} - \mathbf{h} \circ \mathbf{h}$. In particular, the work [2] showed that with small initialization and infinitesimal learning rate, gradient descent on

$$\min_{\{\mathbf{g}, \mathbf{h}\} \subseteq \mathbb{R}^m} \|\mathbf{A}(\mathbf{g} \circ \mathbf{g} - \mathbf{h} \circ \mathbf{h}) - \mathbf{b}\|_2^2 \quad (5)$$

correctly recovers the sparse \mathbf{s}_* when \mathbf{A} satisfies certain restricted isometric properties [51].

The benefit of DOP in (3), as we shall see, is that it allows specific algorithms to *automatically* identify both the intrinsic rank of \mathbf{X}_* and the sparsity level of \mathbf{s}_* without any prior knowledge.

2.2 Algorithmic Regularizations via Gradient Descent

Obviously, over-parameterization leads to *under-determined* problems which can have infinite number of solutions, so that not all solutions of (3) correspond to the desired $(\mathbf{X}_*, \mathbf{s}_*)$. For example, for any given \mathbf{U} , one can always construct a pair (\mathbf{g}, \mathbf{h}) for (3) such that $(\mathbf{U}, \mathbf{g}, \mathbf{h})$ achieves the global minimum value $f = 0$. Nonetheless, as we see in this work, the gradient descent iteration on (3),

$$\begin{aligned} \mathbf{U}_{k+1} &= \mathbf{U}_k - \tau \cdot \nabla_{\mathbf{U}} f(\mathbf{U}_k, \mathbf{g}_k, \mathbf{h}_k) = \mathbf{U}_k - \tau \cdot \mathcal{A}^*(\mathbf{r}_k) \mathbf{U}_k, \\ \begin{bmatrix} \mathbf{g}_{k+1} \\ \mathbf{h}_{k+1} \end{bmatrix} &= \begin{bmatrix} \mathbf{g}_k \\ \mathbf{h}_k \end{bmatrix} - \alpha \cdot \tau \cdot \begin{bmatrix} \nabla_{\mathbf{g}} f(\mathbf{U}_k, \mathbf{g}_k, \mathbf{h}_k) \\ \nabla_{\mathbf{h}} f(\mathbf{U}_k, \mathbf{g}_k, \mathbf{h}_k) \end{bmatrix} = \begin{bmatrix} \mathbf{g}_k \\ \mathbf{h}_k \end{bmatrix} - \alpha \cdot \tau \cdot \begin{bmatrix} \mathbf{r}_k \circ \mathbf{g}_k \\ -\mathbf{r}_k \circ \mathbf{h}_k \end{bmatrix}, \end{aligned} \quad (6)$$

with properly chosen learning rates $(\tau, \alpha \cdot \tau)$ enforces implicit bias on the solution path, that it automatically identifies the desired, regularized solution $(\mathbf{X}_*, \mathbf{s}_*)$. Here, in (6) we have $\mathcal{A}^*(\cdot) : \mathbb{R}^m \mapsto \mathbb{R}^{n \times n}$ being the conjugate operator of $\mathcal{A}(\cdot)$ and $\mathbf{r}_k = \mathcal{A}(\mathbf{U}_k \mathbf{U}_k^\top) + \mathbf{g}_k \circ \mathbf{g}_k - \mathbf{h}_k \circ \mathbf{h}_k - \mathbf{y}$.

It should be noted that the scalar $\alpha > 0$, which controls the ratio of the learning rates for \mathbf{U} and (\mathbf{g}, \mathbf{h}) , plays a crucial role on the quality of the solution that the iterate in (6) converges to (see Figure 2a). We will discuss this in more details in the next subsection (i.e., Section 2.3).

Convergence to low-rank & sparse solutions. Based on our discussion in Section 2.1, it is expected that the gradient descent (6) converges to a solution $(\mathbf{U}, \mathbf{g}, \mathbf{h})$ such that

$$\mathbf{X} = \mathbf{U}\mathbf{U}^\top \quad \text{and} \quad \mathbf{s} = \mathbf{g} \circ \mathbf{g} - \mathbf{h} \circ \mathbf{h} \quad (7)$$

have the minimum nuclear norm and ℓ_1 norm, respectively. More specifically, we expect the solution (\mathbf{X}, \mathbf{s}) in (7) of the nonlinear least squares (3) also serves as a *global* solution to a convex problem

$$\min_{\mathbf{X} \in \mathbb{R}^{n \times n}, \mathbf{s} \in \mathbb{R}^m} \|\mathbf{X}\|_* + \lambda \cdot \|\mathbf{s}\|_1, \quad \text{s.t.} \quad \mathcal{A}(\mathbf{X}) + \mathbf{s} = \mathbf{y}, \quad \mathbf{X} \succeq \mathbf{0}, \quad (8)$$

for which we state more rigorously in Section 3 under a particular setting. However, it should be noted that obtaining the global solution of (8) does not necessarily mean that we find the desired solution $(\mathbf{X}_*, \mathbf{s}_*)$ — the penalty $\lambda > 0$ in (8), which balances the low-rank and the sparse regularizations, plays a crucial role on the quality of the solution to (8). For instance, when $\mathcal{A}(\cdot)$ is the identity operator, under proper conditions of $(\mathbf{X}_*, \mathbf{s}_*)$, the work [14] showed that the global solution of (8) is exactly the target solution $(\mathbf{X}_*, \mathbf{s}_*)$ *only* when $\lambda = 1/\sqrt{n}$.

2.3 Implicit Bias with Discrepant Learning Rates

As noted above, a remaining challenge is to control the implicit regularization of the gradient descent (6) so that the algorithm converges to the solution to (8) with the desired value $\lambda = 1/\sqrt{n}$. Without explicit regularization terms in our new objective (3), at first glance this might seem impossible. Nonetheless, we show that this can simply be achieved by adapting the ratio of learning rates α between \mathbf{U} and (\mathbf{g}, \mathbf{h}) in our gradient step (6), which is one of our key contributions in this work that could also bear broader interest. More specifically, this phenomenon can be captured by the following observation.

Observation. *With small enough learning rate τ and initialization $(\mathbf{U}_0, \mathbf{g}_0, \mathbf{h}_0)$ close enough to the origin, gradient descent (6) converges to a solution of (8) with $\lambda = 1/\alpha$.*

In comparison to the classical optimization theory [52] where learning rates *only* affect algorithm convergence rate but not the quality of the solution, our observation here is surprisingly different — using discrepant learning rates on different over-parameterized variables actually affects the specific solution that the algorithm converges to⁴. In the following, let us provide some intuitions of why this happens and discuss its implications. We leave a more rigorous mathematical treatment to Section 3.

Intuitions. The relation $\lambda = 1/\alpha$ implies that a large learning rate for one particular optimization variable in (6) leads to a small penalty on its implicit regularization counterpart in (8). From a high-level perspective, this happens because a larger learning rate allows the optimization variable to move faster away from its initial point (which is close to the origin), resulting in a weaker regularization effect (which penalizes the distance of the variable to the origin) on its solution path.

Implications. The implicit bias of discrepant learning rates provides a new and powerful way for controlling the regularization effects in over-parameterized models. For robust matrix recovery in particular, it reveals an equivalency between our method in (3) and the convex method in (8) with one-to-one correspondence between learning rate ratio α and the regularization parameter λ . By picking $\alpha = \sqrt{n}$, we may directly quote results from [14] and conclude that our method correctly recovers \mathbf{X}_* with information-theoretically optimal sampling complexity and sparsity levels (see Figure 2). Note that this is achieved with no prior knowledge on the rank of \mathbf{X}_* and sparsity of \mathbf{s}_* . Next, we show that such benefits have implications beyond robust low-rank matrix recovery.

2.4 Extension to Robust Recovery of Natural Images

Finally, we address the problem of robust recovery of a natural image⁵ $\mathbf{X}_* \in \mathbb{R}^{C \times H \times W}$ from its corrupted observation $\mathbf{y} = \mathbf{X}_* + \mathbf{s}_*$, for which the structure of \mathbf{X}_* cannot be fully captured by a low-rank model. Inspired by the work [43] on showing the effectiveness of an image prior from a deep convolutional network $\mathbf{X} = \phi(\boldsymbol{\theta})$ of particular architectures, where $\boldsymbol{\theta} \in \mathbb{R}^c$ denotes network parameters, we use the following formulation for image recovery:

$$\theta \in \mathbb{R}^c, \{\mathbf{g}, \mathbf{h}\} \subseteq \mathbb{R}^{C \times H \times W} \quad f(\boldsymbol{\theta}, \mathbf{g}, \mathbf{h}) = \frac{1}{4} \|\phi(\boldsymbol{\theta}) + (\mathbf{g} \circ \mathbf{g} - \mathbf{h} \circ \mathbf{h}) - \mathbf{y}\|_F^2. \quad (9)$$

As we will empirically demonstrate in Section 4.2, gradient descent on (9)

$$\begin{aligned} \boldsymbol{\theta}_{k+1} &= \boldsymbol{\theta}_k - \tau \cdot \nabla_{\boldsymbol{\theta}} f(\boldsymbol{\theta}_k, \mathbf{g}_k, \mathbf{h}_k), \\ \begin{bmatrix} \mathbf{g}_{k+1} \\ \mathbf{h}_{k+1} \end{bmatrix} &= \begin{bmatrix} \mathbf{g}_k \\ \mathbf{h}_k \end{bmatrix} - \alpha \cdot \tau \cdot \begin{bmatrix} \nabla_{\mathbf{g}} f(\boldsymbol{\theta}_k, \mathbf{g}_k, \mathbf{h}_k) \\ \nabla_{\mathbf{h}} f(\boldsymbol{\theta}_k, \mathbf{g}_k, \mathbf{h}_k) \end{bmatrix}, \end{aligned} \quad (10)$$

with a balanced learning rate ratio α also enforces implicit regularizations on the solution path to the desired solution. It should be noted that this happens even that the over-parameterization $\mathbf{X} = \phi(\boldsymbol{\theta})$ is a highly *nonlinear* network (in comparison with shallow linear network $\mathbf{X} = \mathbf{U}\mathbf{U}^\top$ [6] or deep linear network [8]), which raises several intriguing theoretical questions to be further investigated. For example, we empirically observe that the effect of α is analogous to that in low-rank matrix recovery, meaning that the optimal α is independent of the image \mathbf{X}_* and corruption level of \mathbf{s}_* . This provides an attractive feature of our method as it does not require any parameter tuning.

⁴Our result also differs from [53–55] which analyze the effect of initial large learning rates.

⁵Here, H, W are height and width of the image, respectively. C denotes the number of channels of the image, where $C = 1$ for a grayscale image and $C = 3$ for a color image with RGB channels.

3 Convergence Analysis of Gradient Flow Dynamics

We provide a dynamical analysis certifying our observation in Section 2.3. Similar to [6, 8], we consider a special case where the measurement matrices $\{\mathbf{A}_i\}_{i=1}^m$ associated with \mathcal{A} commute⁶, and investigate the trajectories of the discrete gradient iterate of \mathbf{U}_k , \mathbf{g}_k , and \mathbf{h}_k in (6) with *infinitesimal* learning rate τ (i.e., $\tau \rightarrow 0$). In other words, we study their *continuous* dynamics counterparts $\mathbf{U}_t(\gamma)$, $\mathbf{g}_t(\gamma)$, and $\mathbf{h}_t(\gamma)$, which are parameterized by $t \in [0, +\infty)$ and initialized at $t = 0$ with

$$\mathbf{U}_0(\gamma) = \gamma \mathbf{I}, \quad \mathbf{g}_0(\gamma) = \gamma \mathbf{1}, \quad \mathbf{h}_0(\gamma) = \gamma \mathbf{1}, \quad (11)$$

where $\gamma > 0$. Thus, analogous to (6), the behaviors of the continuous gradient flows can be captured via the following differential equations

$$\dot{\mathbf{U}}_t(\gamma) = \lim_{\tau \rightarrow 0} (\mathbf{U}_{t+\tau}(\gamma) - \mathbf{U}_t(\gamma)) / \tau = -\mathcal{A}^*(\mathbf{r}_t(\gamma)) \mathbf{U}_t(\gamma), \quad (12)$$

$$\begin{bmatrix} \dot{\mathbf{g}}_t(\gamma) \\ \dot{\mathbf{h}}_t(\gamma) \end{bmatrix} = \lim_{\tau \rightarrow 0} \left(\begin{bmatrix} \mathbf{g}_{t+\tau}(\gamma) \\ \mathbf{h}_{t+\tau}(\gamma) \end{bmatrix} - \begin{bmatrix} \mathbf{g}_t(\gamma) \\ \mathbf{h}_t(\gamma) \end{bmatrix} \right) / \tau = -\alpha \cdot \begin{bmatrix} \mathbf{r}_t(\gamma) \circ \mathbf{g}_t(\gamma) \\ -\mathbf{r}_t(\gamma) \circ \mathbf{h}_t(\gamma) \end{bmatrix}, \quad (13)$$

with $\mathbf{r}_t(\gamma) = \mathcal{A}(\mathbf{U}_t(\gamma) \mathbf{U}_t^\top(\gamma)) + \mathbf{g}_t(\gamma) \circ \mathbf{g}_t(\gamma) - \mathbf{h}_t(\gamma) \circ \mathbf{h}_t(\gamma) - \mathbf{y}$. Let $\mathbf{X}_t(\gamma) = \mathbf{U}_t(\gamma) \mathbf{U}_t^\top(\gamma)$, then we can derive the gradient flow for $\mathbf{X}_t(\gamma)$ via chain rule

$$\dot{\mathbf{X}}_t(\gamma) = \dot{\mathbf{U}}_t(\gamma) \mathbf{U}_t^\top(\gamma) + \mathbf{U}_t(\gamma) \dot{\mathbf{U}}_t^\top(\gamma) = -\mathcal{A}^*(\mathbf{r}_t(\gamma)) \mathbf{X}_t(\gamma) - \mathbf{X}_t(\gamma) \mathcal{A}^*(\mathbf{r}_t(\gamma)). \quad (14)$$

For any $\gamma > 0$, suppose the limits of $\mathbf{X}_t(\gamma)$, $\mathbf{g}_t(\gamma)$, and $\mathbf{h}_t(\gamma)$ as $t \rightarrow +\infty$ exist and denote by

$$\mathbf{X}_\infty(\gamma) := \lim_{t \rightarrow +\infty} \mathbf{X}_t(\gamma), \quad \mathbf{g}_\infty(\gamma) := \lim_{t \rightarrow +\infty} \mathbf{g}_t(\gamma), \quad \mathbf{h}_\infty(\gamma) := \lim_{t \rightarrow +\infty} \mathbf{h}_t(\gamma). \quad (15)$$

Then when the initialization is infinitesimally small with $\gamma \rightarrow 0$, we show that the following holds.

Theorem 1. *Assume that the measurement matrices $\{\mathbf{A}_i\}_{i=1}^m$ are symmetric and commutable with $\mathbf{A}_i \mathbf{A}_j = \mathbf{A}_j \mathbf{A}_i$ for $\forall 1 \leq i \neq j \leq m$, and the gradient flows of $\mathbf{U}_t(\gamma)$, $\mathbf{g}_t(\gamma)$, and $\mathbf{h}_t(\gamma)$ satisfy (12) and (13) and are initialized by (11). Let $\mathbf{X}_\infty(\gamma)$, $\mathbf{g}_\infty(\gamma)$, and $\mathbf{h}_\infty(\gamma)$ be the limit points as defined in (15). Suppose that our initialization is infinitesimally small such that*

$$\widehat{\mathbf{X}} := \lim_{\gamma \rightarrow 0} \mathbf{X}_\infty(\gamma), \quad \widehat{\mathbf{g}} := \lim_{\gamma \rightarrow 0} \mathbf{g}_\infty(\gamma), \quad \widehat{\mathbf{h}} := \lim_{\gamma \rightarrow 0} \mathbf{h}_\infty(\gamma)$$

exist and $(\widehat{\mathbf{X}}, \widehat{\mathbf{g}}, \widehat{\mathbf{h}})$ is a global optimal solution to (3) with

$$\mathcal{A}(\widehat{\mathbf{X}}) + \widehat{\mathbf{s}} = \mathbf{y} \quad \text{and} \quad \widehat{\mathbf{s}} = \widehat{\mathbf{g}} \circ \widehat{\mathbf{g}} - \widehat{\mathbf{h}} \circ \widehat{\mathbf{h}}.$$

Then we have $\widehat{\mathbf{g}} \circ \widehat{\mathbf{h}} = \mathbf{0}$, and $(\widehat{\mathbf{X}}, \widehat{\mathbf{s}})$ is also a global optimal solution to (8), with $\lambda = \alpha^{-1}$.

A proof of Theorem 1 is given in Appendix B. Our proof follows a similar procedure as that in [6] and is based on constructing a dual certificate for the convex problem in (8) with the algorithm dynamic of gradient descent. The major difference from [6] is on handling the additional sparse term \mathbf{s}_* and characterizing the discrepant learning rates in the verification of the dual certificate.

Our result shows that among the infinite many global solutions to (3), gradient descent biases towards the one with the minimum nuclear norm (for \mathbf{X}) and ℓ_1 norm (for \mathbf{s}) with relative weight controlled by α . In the following, we discuss the rationality of the assumptions for Theorem 1.

- Since gradient descent almost surely converges to a local minimum [56] and any local minimum is likely to be a global minimum in low-rank matrix optimization with over-parameterization (i.e., $\mathbf{U} \in \mathbb{R}^{n \times n}$) [57], the assumption that $(\widehat{\mathbf{X}}, \widehat{\mathbf{g}}, \widehat{\mathbf{h}})$ is a global optimal solution is generally satisfied.
- The commutative assumption on \mathcal{A} is commonly adopted in the analysis of over-parameterized low-rank matrix recovery. While it may rule out many practical cases (e.g., when \mathcal{A} is identity), it is empirically observed to be non-essential [6, 8]. A recent work [7] provides a more sophisticated analysis of the discrete dynamic under the restricted isometric assumption where the commutative assumption is not needed. We believe such analysis can be extended to our DOP setting as well, and leave it as future work.

⁶Any $\mathcal{A} : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^m$ can be written as $\mathcal{A}(\mathbf{Z}) = [\langle \mathbf{A}_1, \mathbf{Z} \rangle, \dots, \langle \mathbf{A}_m, \mathbf{Z} \rangle]$ for some $\{\mathbf{A}_i \in \mathbb{R}^{n \times n}\}_{i=1}^m$.




Figure 2: **Numerical results for robust PCA** (with $n = 50$). (a) Relative reconstruction error for the output $(\hat{\mathbf{X}}, \hat{\mathbf{S}})$ of our method with varying step size ratio α . Here, we set $r = 5$ and $p = 30\%$. (b, c) Probability of successful recovery over 10 trials with varying r (y-axis) and p (x-axis). Here, we fixed $\alpha = \sqrt{n}$. White indicates always successful recovery, while black means total failure.

4 Experiments

In this section, we provide experimental evidence for the implicit bias of the learning rate discussed in Section 2.3. Through experiments for low-rank matrix recovery, Section 4.1 shows that the learning rate ratio α affects the solution that gradient descent converges to, and that an optimal α does not depend on the rank of matrix and sparsity of corruption. Furthermore, Section 4.2 shows the effectiveness of implicit bias of learning rate in alleviating overfitting for robust image recovery, and demonstrates that our method produces better recovery quality when compared to DIP for varying test images and corruption levels, all with a single model and set of learning parameters.

4.1 Robust Recovery of Low-rank Matrices

We conduct experiments for a particular case of low-rank matrix recovery problem, namely the robust PCA problem, in which the operator \mathcal{A} is the identity operator. More specifically, the goal is to recover a low-rank matrix \mathbf{X}_* and a sparse matrix \mathbf{S}_* from the mixture $\mathbf{Y} = \mathbf{X}_* + \mathbf{S}_*$, possibly without prior knowledge on the rank r of \mathbf{X}_* and the percentage of nonzero entries p of \mathbf{S}_* . For any given r , we generate $\mathbf{X}_* \in \mathbb{R}^{n \times n}$ by setting $\mathbf{X}_* = \mathbf{U}_* \mathbf{U}_*^T$, where \mathbf{U}_* is an $n \times r$ matrix with i.i.d. entries drawn from standard Gaussian distribution. For any given p , we generate $\mathbf{S}_* \in \mathbb{R}^{n \times n}$ by sampling uniformly at random $n^2 \cdot p$ locations from the matrix and setting those entries by sampling i.i.d. from a zero-mean Gaussian distribution with standard deviation 10. We use $n = 50$.

We apply our DOP method in (3) for the robust PCA problem. Specifically, we initialize \mathbf{U} and \mathbf{g} by drawing i.i.d. entries from zero mean Gaussian distribution with standard deviation 10^{-4} , and initialize \mathbf{h} to be the same as \mathbf{g} . The learning rate for \mathbf{U} as well as for $\{\mathbf{g}, \mathbf{h}\}$ are set to τ and $\alpha \cdot \tau$, respectively, where $\tau = 10^{-4}$ for all experiments.

Effect of learning rate ratio α . We set $r = 5$ and $p = 30\%$, perform 2×10^4 gradient descent steps and compute reconstruction errors for the output of our algorithm $(\hat{\mathbf{X}}, \hat{\mathbf{S}})$ relative to the ground truth $(\mathbf{X}_*, \mathbf{S}_*)$. Figure 2a illustrates the performance with varying values of α . We observe that α affects the solution that the algorithm converges to, which verifies that the learning rate ratio has an implicit regularization effect. Moreover, the best performance is given by $\alpha = \sqrt{n}$, which is in accordance with our theoretical result in Theorem 1.

Effect of rank r and outlier ratio p and phase transition. We now fix $\alpha = \sqrt{n}$ and study the ability of our method for recovering matrices of varying rank r with varying percentage of corruption p . For each pair (r, p) , we apply our method and declare the recovery to be successful if the relative reconstruction error $\frac{\|\hat{\mathbf{X}} - \mathbf{X}_*\|_F}{\|\mathbf{X}_*\|_F}$ is less than 0.1. Figure 2c displays the fraction of successful recovery in 10 Monte Carlo trials. It shows that a single value of the parameter α leads to correct recovery for a wide range of r and p . Figure 2b shows the fraction of successful recovery for the convex method in (8) (with the Accelerated Proximal Gradient [58] solver⁷).

⁷http://people.eecs.berkeley.edu/~yima/matrix-rank/sample_code.html




Figure 3: **Robust image recovery for salt-and-pepper corruption.** PSNR of the results is displayed below the images. For our method, all cases use the same network width, learning rate, and termination condition. For DIP and DIP- ℓ_1 , case-dependent early stopping is used which is essential for their good performance. Despite that, DOP (our method) achieves the highest PSNRs and best visual quality.

4.2 Robust Recovery of Natural Images

Following [43], we evaluate the performance of our method for robust image recovery using four images, namely “F16”, “Peppers”, “Baboon” and “Kodim03” (see the second column of Figure 3), from a standard dataset⁸. Corruption for the images is synthesized by adding salt-and-pepper noise, where ratio p of randomly chosen pixels are replaced with either 1 or 0 (each with 50% probability).

The network $\phi(\theta)$ for our method in (9) is the same as the denoising network in [43], which has a U-shaped architecture with skip connections. Each layer of the network contains a convolutional layer, a nonlinear LeakyReLU layer and a batch normalization layer. We also follow [43] on the initialization of network parameters θ , which is the Kaiming initialization. Meanwhile, we initialize g and h by drawing i.i.d. entries from zero-mean Gaussian distribution with a standard deviation of 10^{-5} . The learning rate τ is set to 1.0. The learning rate ratio $\alpha = 500$ is tuned on a specific image (i.e., “F16”) under a specific corruption level (i.e., $p = 50\%$) for achieving the highest PSNR, and is subsequently fixed for all experiments. Finally, our method is *untrained* (i.e., it does *not* require any training data), so for fair comparisons we do not compare with trained image denoising methods such as [40–42].

No need to tune model width or terminate early. We compare our method with a variant of DIP that we call DIP- ℓ_1 , which is based on solving $\min_{\theta} \ell(\phi(\theta) - \mathbf{y})$ with $\ell(\cdot)$ being $\|\cdot\|_1$. As shown in Figure 1b, DIP- ℓ_1 requires either choosing appropriate network width or early termination to

⁸http://www.cs.tut.fi/~foi/GCF-BM3D/index.html#ref_results

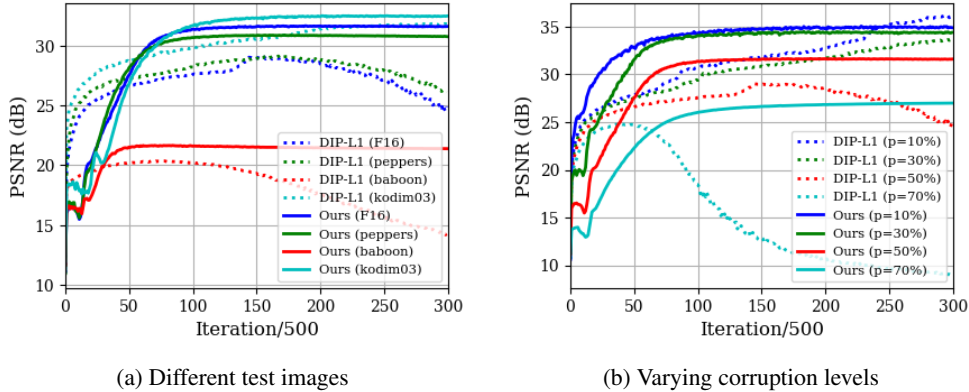


Figure 4: **Learning curves for robust image recovery with different test images (a) and varying corruption levels (b).** DIP- ℓ_1 requires case-dependent early stopping while DOP (our method) does not.

avoid overfitting. Note that neither of these can be carried out in practice without the true (clean) image. On the other hand, our method does not require tuning network width or early termination. Its performance continues to improve as training proceeds until stabilises.

Handling different images and varying corruption levels. The benefit mentioned above enables our method to handle different images and varying corruption levels *without* the need to tune network width, termination condition and any other learning parameters. In our experiments, we fix the number of channels in each convolutional layer to be 128, run 150,000 gradient descent iterations and display the output image. The results are shown in Figure 3 (for four different test images) and Figure 5 (for four corruption levels, see Appendix C). For DIP and DIP- ℓ_1 , we report the result with the highest PSNR in the learning process (i.e., we perform the best early termination for evaluation purposes – note that this cannot be carried out in practice). Despite that, our method obtains better recovery quality in terms of PSNR for all cases. We display the learning curves for these experiments in Figure 4a and Figure 4b, which show that our method does not overfit for all cases while DIP- ℓ_1 requires a case-dependent early termination.

5 Conclusion

In this work, we have shown both theoretically and empirically that the benefits of implicit bias of gradient descent can be extended to over-parameterization of two low-dimensional structures. The key to the success is the choice of discrepant learning rates that can properly regulate the optimization path so that it converges to the desired optimal solution. Such a framework frees us from the need of prior knowledge in the structures or from the lack of scalability of previous approaches. This has led to state of the art recovery results for both low-rank matrices and natural images. We hope this work may encourage people to investigate in the future if the same framework and idea generalize to a mixture of multiple and broader families of structures.

Broader Impact

Robust learning of structured signals from high-dimensional data has a wide range of applications, including imaging processing, computer vision, recommender systems, generative models and many more. In this work, we presented a new type of practical methods and provided improved understandings of solving these problems via over-parameterized models. In particular, our method exploits the implicit bias introduced by the learning algorithm, with the underlying driving force being the intrinsic structure of the data itself rather than human handcrafting. Such a design methodology helps to eliminate human bias in the design process, hence provides the basis for developing truly fair machine learning systems.

Acknowledgments and Disclosure of Funding

CY and YM acknowledge support from Tsinghua-Berkeley Shenzhen Institute (TBSI) Research Fund. YM acknowledges support from ONR grant N00014-20-1-2002 and the joint Simons

Foundation-NSF DMS grant #2031899, as well as support from Berkeley AI Research (BAIR), Berkeley FHL Vive Center for Enhanced Reality, and Berkeley Center for Augmented Cognition. ZZ acknowledges support from NSF grant #2008460. QQ acknowledges support from the Moore-Sloan fellowship and NSF DMS #2009752. We would like to thank Shuang Li for carefully proof-reading this manuscript, Yaodong Yu and Ryan Chan for the helpful discussions on connections between matrix factorization and deep image prior and on implicit bias with Hadamard product.

References

- [1] D. Soudry, E. Hoffer, M. S. Nacson, S. Gunasekar, and N. Srebro, “The implicit bias of gradient descent on separable data,” *The Journal of Machine Learning Research*, vol. 19, no. 1, pp. 2822–2878, 2018.
- [2] T. Vaskevicius, V. Kanade, and P. Rebeschini, “Implicit regularization for optimal sparse recovery,” in *Advances in Neural Information Processing Systems*, pp. 2968–2979, 2019.
- [3] P. Zhao, Y. Yang, and Q.-C. He, “Implicit regularization via hadamard product over-parametrization in high-dimensional linear regression,” *arXiv preprint arXiv:1903.09367*, 2019.
- [4] F. Wu and P. Rebeschini, “Hadamard wirtinger flow for sparse phase retrieval,” *arXiv preprint arXiv:2006.01065*, 2020.
- [5] S. Oymak and M. Soltanolkotabi, “Overparameterized nonlinear learning: Gradient descent takes the shortest path?,” in *International Conference on Machine Learning*, pp. 4951–4960, 2019.
- [6] S. Gunasekar, B. E. Woodworth, S. Bhojanapalli, B. Neyshabur, and N. Srebro, “Implicit regularization in matrix factorization,” in *Advances in Neural Information Processing Systems*, pp. 6151–6159, 2017.
- [7] Y. Li, T. Ma, and H. Zhang, “Algorithmic regularization in over-parameterized matrix sensing and neural networks with quadratic activations,” in *Conference On Learning Theory*, pp. 2–47, 2018.
- [8] S. Arora, N. Cohen, W. Hu, and Y. Luo, “Implicit regularization in deep matrix factorization,” in *Advances in Neural Information Processing Systems*, pp. 7411–7422, 2019.
- [9] M. A. Belabbas, “On implicit regularization: Morse functions and applications to matrix factorization,” *arXiv preprint arXiv:2001.04264*, 2020.
- [10] S. Gunasekar, J. D. Lee, D. Soudry, and N. Srebro, “Implicit bias of gradient descent on linear convolutional networks,” in *Advances in Neural Information Processing Systems*, pp. 9461–9471, 2018.
- [11] G. Gidel, F. Bach, and S. Lacoste-Julien, “Implicit regularization of discrete gradient dynamics in linear neural networks,” in *Advances in Neural Information Processing Systems*, pp. 3196–3206, 2019.
- [12] E. Alpaydin, *Introduction to machine learning*. 2020.
- [13] B. Woodworth, S. Gunasekar, J. D. Lee, E. Moroshko, P. Savarese, I. Golan, D. Soudry, and N. Srebro, “Kernel and rich regimes in overparametrized models,” *arXiv preprint arXiv:2002.09277*, 2020.
- [14] E. J. Candès, X. Li, Y. Ma, and J. Wright, “Robust principal component analysis?,” *Journal of the ACM (JACM)*, vol. 58, no. 3, pp. 1–37, 2011.
- [15] D. S. Weller, A. Pnueli, G. Divon, O. Radzyner, Y. C. Eldar, and J. A. Fessler, “Undersampled phase retrieval with outliers,” *IEEE transactions on computational imaging*, vol. 1, no. 4, pp. 247–258, 2015.
- [16] Y. Li, Y. Sun, and Y. Chi, “Low-rank positive semidefinite matrix recovery from corrupted rank-one measurements,” *IEEE Transactions on Signal Processing*, vol. 65, no. 2, pp. 397–408, 2016.
- [17] J. C. Duchi and F. Ruan, “Solving (most) of a set of quadratic equalities: Composite optimization for robust phase retrieval,” *Information and Inference: A Journal of the IMA*, vol. 8, no. 3, pp. 471–529, 2019.

- [18] Y. Chi, Y. M. Lu, and Y. Chen, “Nonconvex optimization meets low-rank matrix factorization: An overview,” *IEEE Transactions on Signal Processing*, vol. 67, no. 20, pp. 5239–5269, 2019.
- [19] X. Li, Z. Zhu, A. Man-Cho So, and R. Vidal, “Nonconvex robust low-rank matrix recovery,” *SIAM Journal on Optimization*, vol. 30, no. 1, pp. 660–686, 2020.
- [20] V. Chandrasekaran, S. Sanghavi, P. A. Parrilo, and A. S. Willsky, “Rank-sparsity incoherence for matrix decomposition,” *SIAM Journal on Optimization*, vol. 21, no. 2, pp. 572–596, 2011.
- [21] Z. Lin, M. Chen, and Y. Ma, “The augmented lagrange multiplier method for exact recovery of corrupted low-rank matrices,” *arXiv preprint arXiv:1009.5055*, 2010.
- [22] P. Jain and P. Kar, “Non-convex optimization for machine learning,” *Foundations and Trends® in Machine Learning*, vol. 10, no. 3-4, pp. 142–336, 2017.
- [23] S. Burer and R. D. Monteiro, “A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization,” *Mathematical Programming*, vol. 95, no. 2, pp. 329–357, 2003.
- [24] P. J. Rousseeuw and A. M. Leroy, *Robust regression and outlier detection*, vol. 589. John wiley & sons, 2005.
- [25] Z. Zhu, Y. Wang, D. Robinson, D. Naiman, R. Vidal, and M. Tsakiris, “Dual principal component pursuit: Improved analysis and efficient algorithms,” in *Advances in Neural Information Processing Systems 31*, pp. 2171–2181, 2018.
- [26] Q. Qu, Z. Zhu, X. Li, M. C. Tsakiris, J. Wright, and R. Vidal, “Finding the sparsest vectors in a subspace: Theory, algorithms, and applications,” 2020.
- [27] H. Jiang, D. P. Robinson, R. Vidal, and C. You, “A nonconvex formulation for low rank subspace clustering: algorithms and convergence analysis,” *Computational Optimization and Applications*, vol. 70, no. 2, pp. 395–418, 2018.
- [28] D. Davis, D. Drusvyatskiy, and C. Paquette, “The nonsmooth landscape of phase retrieval,” *arXiv preprint arXiv:1711.03247*, 2017.
- [29] S. Bhojanapalli, B. Neyshabur, and N. Srebro, “Global optimality of local search for low rank matrix recovery,” in *Advances in Neural Information Processing Systems*, pp. 3873–3881, 2016.
- [30] R. Ge, J. D. Lee, and T. Ma, “Matrix completion has no spurious local minimum,” in *Advances in Neural Information Processing Systems*, pp. 2973–2981, 2016.
- [31] Z. Zhu, Q. Li, G. Tang, and M. B. Wakin, “Global optimality in low-rank matrix optimization,” *IEEE Transactions on Signal Processing*, vol. 66, no. 13, pp. 3614–3628, 2018.
- [32] L. Prechelt, “Early stopping-but when?,” in *Neural Networks: Tricks of the trade*, pp. 55–69, Springer, 1998.
- [33] R. Heckel and M. Soltanolkotabi, “Denoising and regularization via exploiting the structural bias of convolutional generators,” *arXiv preprint arXiv:1910.14634*, 2019.
- [34] J. Mairal, F. Bach, J. Ponce, G. Sapiro, and A. Zisserman, “Non-local sparse models for image restoration,” in *2009 IEEE 12th international conference on computer vision*, pp. 2272–2279, IEEE, 2009.
- [35] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, “Image denoising by sparse 3-d transform-domain collaborative filtering,” *IEEE Transactions on image processing*, vol. 16, no. 8, pp. 2080–2095, 2007.
- [36] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, “Bm3d image denoising with shape-adaptive principal component analysis,” 2009.
- [37] S. Gu, L. Zhang, W. Zuo, and X. Feng, “Weighted nuclear norm minimization with application to image denoising,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2862–2869, 2014.
- [38] J. Xu, L. Zhang, and D. Zhang, “A trilateral weighted sparse coding scheme for real-world image denoising,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 20–36, 2018.
- [39] B. Lecouat, J. Ponce, and J. Mairal, “Fully trainable and interpretable non-local sparse models for image restoration,” 2020.

- [40] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, “Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising,” *IEEE Transactions on Image Processing*, vol. 26, no. 7, pp. 3142–3155, 2017.
- [41] Z. Yue, H. Yong, Q. Zhao, D. Meng, and L. Zhang, “Variational denoising network: Toward blind noise modeling and removal,” in *Advances in Neural Information Processing Systems*, pp. 1688–1699, 2019.
- [42] S. W. Zamir, A. Arora, S. Khan, M. Hayat, F. S. Khan, M.-H. Yang, and L. Shao, “Cycleisp: Real image restoration via improved data synthesis,” *arXiv preprint arXiv:2003.07761*, 2020.
- [43] D. Ulyanov, A. Vedaldi, and V. Lempitsky, “Deep image prior,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 9446–9454, 2018.
- [44] R. Heckel and P. Hand, “Deep decoder: Concise image representations from untrained non-convolutional networks,” in *International Conference on Learning Representations*, 2019.
- [45] Y. Zhang, Q. Qu, and J. Wright, “From symmetry to geometry: Tractable nonconvex problems,” *arXiv preprint arXiv:2007.06753*, 2020.
- [46] S. Tu, R. Boczar, M. Simchowitz, M. Soltanolkotabi, and B. Recht, “Low-rank solutions of linear matrix equations via procrustes flow,” in *International Conference on Machine Learning*, pp. 964–973, 2016.
- [47] D. Park, A. Kyrillidis, C. Carmanis, and S. Sanghavi, “Non-square matrix sensing without spurious local minima via the burer-monteiro approach,” in *Artificial Intelligence and Statistics*, pp. 65–74, 2017.
- [48] K. Geyer, A. Kyrillidis, and A. Kalev, “Low-rank regularization and solution uniqueness in over-parameterized matrix sensing,” in *International Conference on Artificial Intelligence and Statistics*, pp. 930–940, 2020.
- [49] E. J. Candès, J. Romberg, and T. Tao, “Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information,” *IEEE Transactions on information theory*, vol. 52, no. 2, pp. 489–509, 2006.
- [50] E. J. Candès and M. B. Wakin, “An introduction to compressive sampling,” *IEEE signal processing magazine*, vol. 25, no. 2, pp. 21–30, 2008.
- [51] E. J. Candès, “The restricted isometry property and its implications for compressed sensing,” *Comptes Rendus Mathématique*, vol. 346, no. 9-10, pp. 589–592, 2008.
- [52] J. Nocedal and S. Wright, *Numerical optimization*. Springer Science & Business Media, 2006.
- [53] Y. Li, C. Wei, and T. Ma, “Towards explaining the regularization effect of initial large learning rate in training neural networks,” in *Advances in Neural Information Processing Systems*, pp. 11669–11680, 2019.
- [54] K. You, M. Long, J. Wang, and M. Jordan, “How does learning rate decay help modern neural networks,” *arXiv preprint arXiv:1908.01878*, 2019.
- [55] P. Nakkiran, “Learning rate annealing can provably help generalization, even for convex problems,” *arXiv preprint arXiv:2005.07360*, 2020.
- [56] J. D. Lee, M. Simchowitz, M. I. Jordan, and B. Recht, “Gradient descent only converges to minimizers,” in *Conference on learning theory*, pp. 1246–1257, 2016.
- [57] M. Journée, F. Bach, P.-A. Absil, and R. Sepulchre, “Low-rank optimization on the cone of positive semidefinite matrices,” *SIAM Journal on Optimization*, vol. 20, no. 5, pp. 2327–2351, 2010.
- [58] Z. Lin, A. Ganesh, J. Wright, L. Wu, M. Chen, and Y. Ma, “Fast convex optimization algorithms for exact recovery of a corrupted low-rank matrix,” *Coordinated Science Laboratory Report no. UILU-ENG-09-2214, DC-246*, 2009.
- [59] S. Oymak and M. Soltanolkotabi, “Overparameterized nonlinear learning: Gradient descent takes the shortest path?,” *arXiv preprint arXiv:1812.10004*, 2018.