

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

—Supplementary Materials—

A Implicit Bias of Discrepant Learning Rates in Linear Regression

In this part of the appendix, let us use a classical result for *underdetermined* linear regression problem to build up some intuitions behind the implicit bias of gradient descent for our problem formulation of robust learning problems. The high level message we aim to deliver through the simple example is that

- Gradient descent implicitly biases towards solutions with minimum ℓ_2 -norm.
- Discrepant learning rates lead to solutions with minimum *weighted* ℓ_2 -norm.

Underdetermined linear regression. Given observation $\mathbf{b} \in \mathbb{R}^{n_1}$ and wide data matrix $\mathbf{W} \in \mathbb{R}^{n_1 \times n_2}$ ($n_2 > n_1$), we want to find $\boldsymbol{\theta}$ which is a solution to

$$\min_{\boldsymbol{\theta} \in \mathbb{R}^{n_2}} \varphi(\boldsymbol{\theta}) = \frac{1}{2} \|\mathbf{b} - \mathbf{W}\boldsymbol{\theta}\|_2^2. \quad (16)$$

For $n_2 > n_1$ and full row-rank \mathbf{W} , the underdetermined problem (16) obviously has *infinite* many solutions, which forms a set

$$\mathcal{S} := \{ \boldsymbol{\theta}_{ln} + \mathbf{n} \mid \boldsymbol{\theta}_{ln} = \mathbf{W}^\dagger \mathbf{b}, \quad \mathbf{n} \in \mathcal{N}(\mathbf{W}) \},$$

where $\mathbf{W}^\dagger := \mathbf{W}^\top (\mathbf{W}\mathbf{W}^\top)^{-1}$ denotes the pseudo-inverse of \mathbf{W} , and

$$\mathcal{N}(\mathbf{W}) := \{ \mathbf{n} \mid \mathbf{W}\mathbf{n} = \mathbf{0} \}, \quad \mathcal{R}(\mathbf{W}) := \{ \mathbf{z} \mid \mathbf{z} = \mathbf{W}^\top \mathbf{v} \}$$

are the null space and row space of \mathbf{W} , respectively. Simple derivation shows that $\boldsymbol{\theta}_{ln}$ is a particular *least* ℓ_2 -norm solution to (16), that minimizes

$$\min_{\boldsymbol{\theta} \in \mathbb{R}^{n_2}} \frac{1}{2} \|\boldsymbol{\theta}\|_2^2, \quad \text{s.t. } \mathbf{W}\boldsymbol{\theta} = \mathbf{b}.$$

Gradient descent biases towards $\boldsymbol{\theta}_{ln}$. Starting from any initialization $\boldsymbol{\theta}_0$, gradient descent

$$\boldsymbol{\theta}_{k+1} = \boldsymbol{\theta}_k - \tau_{ls} \cdot \mathbf{W}^\top (\mathbf{W}\boldsymbol{\theta}_k - \mathbf{b}) \quad (17)$$

with a sufficiently small learning rate⁹ τ_{ls} always finds one of the global solutions for (16). Furthermore, it is now well-understood [59] that whenever the initialization $\boldsymbol{\theta}_0$ has zero component in $\mathcal{N}(\mathbf{W})$ (i.e., $\mathcal{P}_{\mathcal{N}(\mathbf{W})}(\boldsymbol{\theta}_0) = \mathbf{0}$), one interesting phenomenon is that the iterates $\boldsymbol{\theta}_\infty$ in (17) implicitly bias towards the minimum ℓ^2 -norm solution $\boldsymbol{\theta}_{ln}$. This happens because once initialized in $\mathcal{R}(\mathbf{W})$, gradient descent (17) implicitly biases towards iterates staying within $\mathcal{R}(\mathbf{W})$, such that

$$\mathcal{P}_{\mathcal{R}(\mathbf{W})}(\boldsymbol{\theta}_\infty) = \boldsymbol{\theta}_{ln}, \quad \mathcal{P}_{\mathcal{N}(\mathbf{W})}(\boldsymbol{\theta}_\infty) = \mathcal{P}_{\mathcal{N}(\mathbf{W})}(\boldsymbol{\theta}_0) = \mathbf{0}.$$

As we can see, a particular algorithm enforces specific regularization on the final solution.

Implicit bias of discrepant learning rates. The gradient update in (17) uses the same learning rate τ_{ls} for all coordinates of $\boldsymbol{\theta}$. If we use different learning rates for each coordinate (i.e., $\boldsymbol{\Lambda}$ is a diagonal matrix with positive diagonals)

$$\boldsymbol{\theta}_{k+1} = \boldsymbol{\theta}_k - \tau_{ls} \cdot \boldsymbol{\Lambda} \cdot \mathbf{W}^\top (\mathbf{W}\boldsymbol{\theta}_k - \mathbf{b}), \quad (18)$$

then by following a similar argument we conclude that the gradient update in (18) converges to a weighted regularized solution for

$$\min_{\boldsymbol{\theta} \in \mathbb{R}^{n_2}} \frac{1}{2} \left\| \boldsymbol{\Lambda}^{-1/2} \boldsymbol{\theta} \right\|_2^2, \quad \text{s.t. } \mathbf{W}\boldsymbol{\theta} = \mathbf{b}. \quad (19)$$

⁹This is because $\mathbf{W}\boldsymbol{\theta}_{k+1} - \mathbf{b} = (\mathbf{I} - \tau_{ls} \mathbf{W}\mathbf{W}^\top) (\mathbf{W}\boldsymbol{\theta}_k - \mathbf{b})$. If we choose $\tau_{ls} < \|\mathbf{W}\mathbf{W}^\top\|^{-1}$, then $\|\mathbf{W}\boldsymbol{\theta}_k - \mathbf{b}\|$ converges to 0 geometrically.

Remark 1. Let σ_i be the i -th diagonal of Λ and θ_i be the i -th element of θ . Then in (18), $\sigma_i \tau_{ls}$ is the learning rate for the variable θ_i , which varies for different variables. In words, the relation between (18) and (19) implies that for one particular optimization variable (e.g., θ_i) a large learning rate $\sigma_i \tau_{ls}$ in (18) leads to a small implicit regularization effect in (19). From a high-level perspective, this happens because a larger learning rate allows the optimization variable to move faster away from its initial point, resulting in a weaker regularization effect (which penalizes the distance of the variable to the initialization) on its solution path.

An alternative explanation of this is through a change of variable $\theta = \Lambda^{1/2} \tilde{\theta}$. Suppose we minimize

$$\min_{\tilde{\theta} \in \mathbb{R}^{n_2}} \tilde{\varphi}(\tilde{\theta}) = \frac{1}{2} \left\| \mathbf{b} - \mathbf{W} \Lambda^{1/2} \tilde{\theta} \right\|_2^2 \quad (20)$$

via standard gradient descent with a single learning rate

$$\tilde{\theta}_{k+1} = \tilde{\theta}_k - \tau_{ls} \cdot \Lambda^{1/2} \cdot \mathbf{W}^\top \left(\mathbf{W} \Lambda^{1/2} \tilde{\theta}_k - \mathbf{b} \right). \quad (21)$$

Thus, once initialized in $\mathcal{R}(\mathbf{W} \Lambda^{1/2})$, gradient descent (21) converges to the least ℓ_2 -norm solution to (20), i.e., the solution of the following problem

$$\min_{\tilde{\theta} \in \mathbb{R}^{n_2}} \frac{1}{2} \left\| \tilde{\theta} \right\|_2^2, \quad \text{s.t.} \quad \mathbf{W} \Lambda^{1/2} \tilde{\theta} = \mathbf{b}. \quad (22)$$

Finally, plugging $\tilde{\theta} = \Lambda^{-1/2} \theta$ into (21) and (22) gives (18) and (19), respectively, also indicating the gradient update (18) induces implicit weighted regularization towards the solution of (19).

B Proof of Theorem 1

In this part of the appendix, we provide the proof to our main technical result (i.e., Theorem 1) in Section 3. To make this part self-contained, we restate our result as follows.

Theorem 2. Assume that the measurement matrices $\mathbf{A}_1, \mathbf{A}_2, \dots, \mathbf{A}_m$ are symmetric and commutable, i.e.

$$\mathbf{A}_i \mathbf{A}_j = \mathbf{A}_j \mathbf{A}_i, \quad \forall 1 \leq i \neq j \leq m,$$

and the gradient flows of $\mathbf{U}_t(\gamma)$, $\mathbf{g}_t(\gamma)$, and $\mathbf{h}_t(\gamma)$ satisfy

$$\dot{\mathbf{U}}_t(\gamma) = \lim_{\tau \rightarrow 0} \frac{\mathbf{U}_{t+\tau}(\gamma) - \mathbf{U}_t(\gamma)}{\tau} = -\mathcal{A}^* (\mathbf{r}_t(\gamma)) \mathbf{U}_t(\gamma), \quad (23)$$

$$\begin{bmatrix} \dot{\mathbf{g}}_t(\gamma) \\ \dot{\mathbf{h}}_t(\gamma) \end{bmatrix} = \lim_{\tau \rightarrow 0} \left(\begin{bmatrix} \mathbf{g}_{t+\tau}(\gamma) \\ \mathbf{h}_{t+\tau}(\gamma) \end{bmatrix} - \begin{bmatrix} \mathbf{g}_t(\gamma) \\ \mathbf{h}_t(\gamma) \end{bmatrix} \right) / \tau = -\alpha \cdot \begin{bmatrix} \mathbf{r}_t(\gamma) \circ \mathbf{g}_t(\gamma) \\ -\mathbf{r}_t(\gamma) \circ \mathbf{h}_t(\gamma) \end{bmatrix}, \quad (24)$$

with $\mathbf{r}_t(\gamma) = \mathcal{A}(\mathbf{U}_t(\gamma) \mathbf{U}_t^\top(\gamma)) + \mathbf{g}_t(\gamma) \circ \mathbf{g}_t(\gamma) - \mathbf{h}_t(\gamma) \circ \mathbf{h}_t(\gamma) - \mathbf{y}$, and they are initialized by

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

Let $\mathbf{X}_t(\gamma) = \mathbf{U}_t(\gamma) \mathbf{U}_t^\top(\gamma)$, and let $\mathbf{X}_\infty(\gamma)$, $\mathbf{g}_\infty(\gamma)$, and $\mathbf{h}_\infty(\gamma)$ be the limit points defined as

$$\mathbf{X}_\infty(\gamma) := \lim_{t \rightarrow +\infty} \mathbf{X}_t(\gamma), \quad \mathbf{g}_\infty(\gamma) := \lim_{t \rightarrow +\infty} \mathbf{g}_t(\gamma), \quad \mathbf{h}_\infty(\gamma) := \lim_{t \rightarrow +\infty} \mathbf{h}_t(\gamma). \quad (25)$$

Suppose that our initialization is infinitesimally small such that

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

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

$$\min_{\mathbf{U} \in \mathbb{R}^{n \times n'}, \{\mathbf{g}, \mathbf{h}\} \subseteq \mathbb{R}^m} f(\mathbf{U}, \mathbf{g}, \mathbf{h}) := \frac{1}{4} \left\| \mathcal{A}(\mathbf{U} \mathbf{U}^\top) + (\mathbf{g} \circ \mathbf{g} - \mathbf{h} \circ \mathbf{h}) - \mathbf{y} \right\|_2^2, \quad (26)$$

with $\mathcal{A}(\hat{\mathbf{X}}) + \hat{\mathbf{s}} = \mathbf{y}$ and $\hat{\mathbf{s}} = \hat{\mathbf{g}} \circ \hat{\mathbf{g}} - \hat{\mathbf{h}} \circ \hat{\mathbf{h}}$. Then we have $\hat{\mathbf{g}} \circ \hat{\mathbf{h}} = \mathbf{0}$, and $(\hat{\mathbf{X}}, \hat{\mathbf{s}})$ is also a global optimal solution to

$$\min_{\mathbf{X} \in \mathbb{R}^{n \times n}, \mathbf{s} \in \mathbb{R}^m} \|\mathbf{X}\|_* + \lambda \cdot \|\mathbf{s}\|_1, \quad \text{s.t.} \quad \mathcal{A}(\mathbf{X}) + \mathbf{s} = \mathbf{y}, \quad \mathbf{X} \succeq \mathbf{0}. \quad (27)$$

with $\lambda = \alpha^{-1}$ and $\alpha > 0$ being the balancing parameter in (24).

Proof. From (23), we can derive the gradient flow for $\mathbf{X}_t(\gamma)$ via chain rule

$$\dot{\mathbf{X}}_t(\gamma) = \dot{\mathbf{U}}_t(\gamma)\mathbf{U}_t^\top(\gamma) + \mathbf{U}_t(\gamma)\dot{\mathbf{U}}_t^\top(\gamma) = -\mathcal{A}^*(\mathbf{r}_t(\gamma))\mathbf{X}_t(\gamma) - \mathbf{X}_t(\gamma)\mathcal{A}^*(\mathbf{r}_t(\gamma)). \quad (28)$$

We want to show that when the initialization is infinitesimally small (i.e., $\gamma \rightarrow 0$), the limit points of the gradient flows $\mathbf{X}_t(\gamma) = \mathbf{U}_t(\gamma)\mathbf{U}_t^\top(\gamma)$ and $\mathbf{s}_t(\gamma) = \mathbf{g}_t(\gamma) \circ \mathbf{g}_t(\gamma) - \mathbf{h}_t(\gamma) \circ \mathbf{h}_t(\gamma)$ are optimal solutions for (27) as $t \rightarrow +\infty$. Towards this goal, let us first look at the optimality condition for (27). From Lemma 1, we know that if $(\widehat{\mathbf{X}}, \widehat{\mathbf{s}})$ with

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

is an optimal solution for (27) then there exists a *dual certificate* $\boldsymbol{\nu}$ such that

$$\mathcal{A}(\widehat{\mathbf{X}}) + \widehat{\mathbf{s}} = \mathbf{y}, \quad (\mathbf{I} - \mathcal{A}^*(\boldsymbol{\nu})) \cdot \widehat{\mathbf{X}} = \mathbf{0}, \quad \mathcal{A}^*(\boldsymbol{\nu}) \preceq \mathbf{I}, \quad \boldsymbol{\nu} \in \lambda \cdot \text{sign}(\widehat{\mathbf{s}}), \quad \widehat{\mathbf{X}} \succeq \mathbf{0},$$

where $\text{sign}(\widehat{\mathbf{s}})$ is defined in (32), which is the subdifferential of $\|\cdot\|_1$. Thus, it suffices to construct a dual certificate $\boldsymbol{\nu}$ such that $(\widehat{\mathbf{X}}, \widehat{\mathbf{s}})$ satisfies the equation above.

Since $(\widehat{\mathbf{X}}, \widehat{\mathbf{g}}, \widehat{\mathbf{h}})$ is a global optimal solution to (26), we automatically have $\mathcal{A}(\widehat{\mathbf{X}}) + \widehat{\mathbf{s}} = \mathbf{y}$ and $\widehat{\mathbf{X}} \succeq \mathbf{0}$. On the other hand, given that $\{\mathbf{A}_i\}_{i=1}^m$ commutes and (28) and (24) hold for \mathbf{X}_t , \mathbf{g}_t and \mathbf{h}_t , by Lemma 2, we know that

$$\mathbf{X}_t(\gamma) = \exp(\mathcal{A}^*(\boldsymbol{\xi}_t(\gamma))) \cdot \mathbf{X}_0(\gamma) \cdot \exp(\mathcal{A}^*(\boldsymbol{\xi}_t(\gamma))), \quad (29)$$

$$\mathbf{g}_t(\gamma) = \mathbf{g}_0(\gamma) \circ \exp(\alpha \boldsymbol{\xi}_t(\gamma)), \quad \mathbf{h}_t(\gamma) = \mathbf{h}_0(\gamma) \circ \exp(-\alpha \boldsymbol{\xi}_t(\gamma)), \quad (30)$$

where $\boldsymbol{\xi}_t(\gamma) = -\int_0^t \mathbf{r}_\tau(\gamma) d\tau$. Let $\boldsymbol{\xi}_\infty(\gamma) := \lim_{t \rightarrow +\infty} \boldsymbol{\xi}_t(\gamma)$, by Lemma 3 and Lemma 4, we can construct

$$\boldsymbol{\nu}(\gamma) = \frac{\boldsymbol{\xi}_\infty(\gamma)}{\log(1/\gamma)},$$

such that

$$\lim_{\gamma \rightarrow 0} \mathcal{A}^*(\boldsymbol{\nu}(\gamma)) \preceq \mathbf{I}, \quad \lim_{\gamma \rightarrow 0} [\mathbf{I} - \mathcal{A}^*(\boldsymbol{\nu}(\gamma))] \cdot \widehat{\mathbf{X}} = \mathbf{0},$$

and

$$\lim_{\gamma \rightarrow 0} \boldsymbol{\nu}(\gamma) \in \alpha^{-1} \cdot \text{sign}(\widehat{\mathbf{s}}), \quad \lim_{\gamma \rightarrow 0} \mathbf{g}(\gamma) \circ \mathbf{h}(\gamma) = \mathbf{0}.$$

This shows the exists of the dual certificate $\boldsymbol{\nu}(\gamma)$ such that the optimality condition holds for $(\widehat{\mathbf{X}}, \widehat{\mathbf{s}})$. Hence, $(\widehat{\mathbf{X}}, \widehat{\mathbf{s}})$ is also a global optimal solution to (27). \square

Lemma 1. $(\widehat{\mathbf{X}}, \widehat{\mathbf{s}})$ is an optimal solution for (27) if there exists a dual certificate $\boldsymbol{\nu} \in \mathbb{R}^m$ such that the following conditions hold:

$$\mathcal{A}(\widehat{\mathbf{X}}) + \widehat{\mathbf{s}} = \mathbf{y}, \quad (\mathbf{I} - \mathcal{A}^*(\boldsymbol{\nu})) \cdot \widehat{\mathbf{X}} = \mathbf{0}, \quad \boldsymbol{\nu} \in \lambda \cdot \text{sign}(\widehat{\mathbf{s}}), \quad \mathbf{I} \succeq \mathcal{A}^*(\boldsymbol{\nu}), \quad \widehat{\mathbf{X}} \succeq \mathbf{0}, \quad (31)$$

where $\text{sign}(\mathbf{s})$ is the subdifferential of $\|\mathbf{s}\|_1$ with each entry

$$\text{sign}(s) := \begin{cases} s/|s| & s \neq 0, \\ [-1, 1] & s = 0. \end{cases} \quad (32)$$

Proof. The Lagrangian function of the problem can be written as

$$\mathcal{L}(\mathbf{X}, \mathbf{s}, \boldsymbol{\nu}, \boldsymbol{\Gamma}) = \text{trace}(\mathbf{X}) + \lambda \|\mathbf{s}\|_1 + \boldsymbol{\nu}^\top (\mathbf{y} - \mathcal{A}(\mathbf{X}) - \mathbf{s}) - \langle \mathbf{X}, \boldsymbol{\Gamma} \rangle,$$

with $\boldsymbol{\nu} \in \mathbb{R}^m$ and $\boldsymbol{\Gamma} \in \mathbb{R}^{n \times n}$ being the dual variables, where $\boldsymbol{\Gamma} \succeq \mathbf{0}$. Thus, we can derive the Karush-Kuhn-Tucker (KKT) optimality condition for (27) as

$$\begin{aligned} \mathbf{0} \in \partial \mathcal{L} : \quad & \mathbf{I} - \boldsymbol{\Gamma} - \mathcal{A}^*(\boldsymbol{\nu}) = \mathbf{0}, \quad \boldsymbol{\nu} \in \lambda \cdot \partial \|\mathbf{s}\|_1 = \lambda \cdot \text{sign}(\mathbf{s}), \\ \text{feasibility :} \quad & \mathcal{A}(\mathbf{X}) + \mathbf{s} = \mathbf{y}, \quad \mathbf{X} \succeq \mathbf{0}, \quad \boldsymbol{\Gamma} \succeq \mathbf{0}, \\ \text{complementary slackness :} \quad & \boldsymbol{\Gamma} \cdot \mathbf{X} = \mathbf{0}, \end{aligned}$$

where $\partial(\cdot)$ denotes the subdifferential operator and $\text{sign}(\mathbf{s})$ is the subdifferential of $\|\mathbf{s}\|_1$ with each entry

$$\text{sign}(s) = \begin{cases} s/|s| & s \neq 0, \\ [-1, 1] & s = 0. \end{cases}$$

Thus, we know that $(\widehat{\mathbf{X}}, \widehat{\mathbf{s}})$ is global solution to (27) as long as there exists a *dual certificate* $\boldsymbol{\nu}$ such that (31) holds, where we eliminated $\boldsymbol{\Gamma}$ by plugging in $\boldsymbol{\Gamma} = \mathbf{I} - \mathcal{A}^*(\boldsymbol{\nu})$. \square

Lemma 2. Suppose that $\{\mathbf{A}_i\}_{i=1}^m$ commutes. Suppose (14) and (24) hold for \mathbf{X}_t , \mathbf{g}_t and \mathbf{h}_t , then

$$\mathbf{X}_t = \exp(\mathcal{A}^*(\boldsymbol{\xi}_t)) \cdot \mathbf{X}_0 \cdot \exp(\mathcal{A}^*(\boldsymbol{\xi}_t)) \quad (33)$$

$$\mathbf{g}_t = \mathbf{g}_0 \circ \exp(\alpha \boldsymbol{\xi}_t), \quad \mathbf{h}_t = \mathbf{h}_0 \circ \exp(-\alpha \boldsymbol{\xi}_t), \quad (34)$$

where $\boldsymbol{\xi}_t = -\int_0^t \mathbf{r}_\tau d\tau$.

Proof. From (24), we know that

$$\frac{d\mathbf{g}_t}{dt} = -\alpha \mathbf{r}_t \circ \mathbf{g}_t,$$

where the differentiation $\frac{d\mathbf{g}_t}{dt}$ is entrywise for \mathbf{g}_t . Thus, we have

$$\int_0^t \frac{d\mathbf{g}_\tau}{\mathbf{g}_\tau} = -\alpha \int_0^t \mathbf{r}_\tau d\tau \implies \log \mathbf{g}_t - \log \mathbf{g}_0 = \alpha \boldsymbol{\xi}_t \implies \mathbf{g}_t = \mathbf{g}_0 \circ \exp(\alpha \boldsymbol{\xi}_t),$$

where all the operators are entrywise. Similarly, $\mathbf{h}_t = \mathbf{h}_0 \circ \exp(-\alpha \boldsymbol{\xi}_t)$ holds.

For (33), by using (28) and the fact that $\{\mathbf{A}_i\}_{i=1}^m$ commutes, we can derive it with an analogous argument. \square

Lemma 3. Under the settings of Theorem 2 and Lemma 2, for any $\gamma > 0$ there exists

$$\boldsymbol{\nu}(\gamma) = \frac{\boldsymbol{\xi}_\infty(\gamma)}{\log(1/\gamma)}, \quad (35)$$

such that

$$\lim_{\gamma \rightarrow 0} \mathcal{A}^*(\boldsymbol{\nu}(\gamma)) \preceq \mathbf{I}, \quad \lim_{\gamma \rightarrow 0} [\mathbf{I} - \mathcal{A}^*(\boldsymbol{\nu}(\gamma))] \cdot \widehat{\mathbf{X}} = \mathbf{0},$$

where $\boldsymbol{\xi}_\infty(\gamma) = \lim_{t \rightarrow \infty} \boldsymbol{\xi}_t(\gamma)$ with $\boldsymbol{\xi}_t(\gamma) = -\int_0^t \mathbf{r}_\tau(\gamma) d\tau$.

Proof. Given $\mathbf{U}_0 = \gamma \mathbf{I}$, we have $\mathbf{X}_0 = \mathbf{U}_0 \mathbf{U}_0^\top = \gamma^2 \mathbf{I}$. By the expression for \mathbf{X}_t in (29), we have

$$\mathbf{X}_\infty(\gamma) = \gamma^2 \cdot \exp(2\mathcal{A}^*(\boldsymbol{\xi}_\infty(\gamma))) \quad (36)$$

where $\boldsymbol{\xi}_\infty(\gamma) = \lim_{t \rightarrow \infty} \boldsymbol{\xi}_t(\gamma)$. Because $\{\mathbf{A}_i\}_{i=1}^m$ are symmetric and they commute, we know that they are simultaneously diagonalizable by an orthonormal basis $\boldsymbol{\Omega} = [\boldsymbol{\omega}_1, \dots, \boldsymbol{\omega}_n] \in \mathbb{R}^{n \times n}$, i.e.,

$$\boldsymbol{\Omega} \mathbf{A}_i \boldsymbol{\Omega}^\top = \boldsymbol{\Lambda}_i, \quad \boldsymbol{\Lambda}_i \text{ diagonal}, \quad \forall i = 1, 2, \dots, m,$$

and so is $\mathcal{A}^*(\mathbf{b})$ for any $\mathbf{b} \in \mathbb{R}^m$. Therefore, we have

$$\lambda_k(\mathbf{X}_\infty(\gamma)) = \gamma^2 \cdot \exp(2\lambda_k(\mathcal{A}^*(\boldsymbol{\xi}_\infty(\gamma)))) = \exp(2\lambda_k(\mathcal{A}^*(\boldsymbol{\xi}_\infty(\gamma))) + 2 \log \gamma), \quad (37)$$

where $\lambda_k(\cdot)$ denotes the k -th eigenvalue with respect to the k -th basis $\boldsymbol{\omega}_k$. Moreover $\mathbf{X}_\infty(\gamma)$ and its limit $\widehat{\mathbf{X}}$ have the same eigen-basis $\boldsymbol{\Omega}$. Since we have $\mathbf{X}_\infty(\gamma)$ converges to $\widehat{\mathbf{X}}$ when $\gamma \rightarrow 0$, then we have the eigenvalues

$$\lambda_k(\mathbf{X}_\infty(\gamma)) \rightarrow \lambda_k(\widehat{\mathbf{X}}), \quad \forall k = 1, 2, \dots, n, \quad (38)$$

whenever $\gamma \rightarrow 0$.

Case 1: $\lambda_k(\widehat{\mathbf{X}}) > 0$. For any k such that $\lambda_k(\widehat{\mathbf{X}}) > 0$, from (37) and (38), we have

$$\exp(2\lambda_k(\mathcal{A}^*(\boldsymbol{\xi}_\infty(\gamma))) + 2\log \gamma) \rightarrow \lambda_k(\widehat{\mathbf{X}}),$$

so that

$$2\lambda_k(\mathcal{A}^*(\boldsymbol{\xi}_\infty(\gamma))) + 2\log \gamma - \log \lambda_k(\widehat{\mathbf{X}}) \rightarrow 0,$$

which further implies that

$$\lambda_k\left(\mathcal{A}^*\left(\frac{\boldsymbol{\xi}_\infty(\gamma)}{\log(1/\gamma)}\right)\right) - 1 - \frac{\log \lambda_k(\widehat{\mathbf{X}})}{2\log(1/\gamma)} \rightarrow 0.$$

Now if we construct $\boldsymbol{\nu}(\gamma)$ as (35), so that we conclude

$$\lim_{\gamma \rightarrow 0} \lambda_k(\mathcal{A}^*(\boldsymbol{\nu}(\gamma))) = 1, \quad (39)$$

for any k such that $\lambda_k(\widehat{\mathbf{X}}) > 0$.

Case 2: $\lambda_k(\widehat{\mathbf{X}}) = 0$. On the other hand, for any k such that $\lambda_k(\widehat{\mathbf{X}}) = 0$, similarly from (37) and (38), we have

$$\exp(2\lambda_k(\mathcal{A}^*(\boldsymbol{\xi}_\infty(\gamma))) + 2\log \gamma) \rightarrow 0,$$

when $\gamma \rightarrow 0$. Thus, for any small $\epsilon \in (0, 1)$, there exists some $\gamma_0 \in (0, 1)$ such that for all $\gamma < \gamma_0$,

$$\exp(2\lambda_k(\mathcal{A}^*(\boldsymbol{\xi}_\infty(\gamma))) + 2\log \gamma) < \epsilon,$$

which implies that

$$\lambda_k\left(\mathcal{A}^*\left(\frac{\boldsymbol{\xi}_\infty(\gamma)}{\log(1/\gamma)}\right)\right) - 1 < \frac{\log \epsilon}{2\log(1/\gamma)} < 0.$$

Thus, given the construction of $\boldsymbol{\nu}(\gamma)$ in (35), we have

$$\lambda_k(\mathcal{A}^*(\boldsymbol{\nu}(\gamma))) < 1, \quad \forall \gamma < \gamma_0,$$

which further implies that for any k with $\lambda_k(\widehat{\mathbf{X}}) = 0$, we have

$$\lim_{\gamma \rightarrow 0} \lambda_k(\mathcal{A}^*(\boldsymbol{\nu}(\gamma))) < 1. \quad (40)$$

Putting things together. Combining our analysis in (39) and (40), we obtain

$$\lim_{\gamma \rightarrow 0} \mathcal{A}^*(\boldsymbol{\nu}(\gamma)) \preceq \mathbf{I}.$$

On the other hand, per our analysis, we know that there exists an orthogonal matrix $\boldsymbol{\Omega} \in \mathbb{R}^{n \times n}$, such that $\mathcal{A}^*(\boldsymbol{\nu}(\gamma))$ and $\widehat{\mathbf{X}}$ can be simultaneously diagonalized. Thus, we have

$$[\mathbf{I} - \mathcal{A}^*(\boldsymbol{\nu}(\gamma))] \cdot \widehat{\mathbf{X}} = \boldsymbol{\Omega} \cdot (\mathbf{I} - \boldsymbol{\Lambda}_{\mathcal{A}^*(\boldsymbol{\nu}(\gamma))}) \cdot \boldsymbol{\Lambda}_{\widehat{\mathbf{X}}} \cdot \boldsymbol{\Omega}^\top,$$

where $\boldsymbol{\Lambda}_{\mathcal{A}^*(\boldsymbol{\nu}(\gamma))}$ and $\boldsymbol{\Lambda}_{\widehat{\mathbf{X}}}$ are diagonal matrices, with entries being the eigenvalues of $\mathcal{A}^*(\boldsymbol{\nu}(\gamma))$ and $\boldsymbol{\Lambda}_{\widehat{\mathbf{X}}}$, respectively. From our analysis for Case 1 and Case 2, we know that $\lim_{\gamma \rightarrow 0} (\mathbf{I} - \boldsymbol{\Lambda}_{\mathcal{A}^*(\boldsymbol{\nu}(\gamma))}) \cdot \boldsymbol{\Lambda}_{\widehat{\mathbf{X}}} = \mathbf{0}$. Therefore, we also have

$$\lim_{\gamma \rightarrow 0} [\mathbf{I} - \mathcal{A}^*(\boldsymbol{\nu}(\gamma))] \cdot \widehat{\mathbf{X}} = \mathbf{0},$$

as desired. □

Lemma 4. Under the settings of Theorem 2 and Lemma 2, for any $\gamma > 0$ there exists

$$\boldsymbol{\nu}(\gamma) = \frac{\boldsymbol{\xi}_\infty(\gamma)}{\log(1/\gamma)}, \quad (41)$$

such that

$$\lim_{\gamma \rightarrow 0} \boldsymbol{\nu}(\gamma) \in \alpha^{-1} \cdot \text{sign}(\widehat{\mathbf{s}}), \quad \lim_{\gamma \rightarrow 0} \mathbf{g}(\gamma) \circ \mathbf{h}(\gamma) = \mathbf{0}, \quad (42)$$

where $\widehat{\mathbf{s}} = \widehat{\mathbf{g}} \circ \widehat{\mathbf{g}} - \widehat{\mathbf{h}} \circ \widehat{\mathbf{h}}$, and $\boldsymbol{\xi}_\infty(\gamma) = \lim_{t \rightarrow 0} \boldsymbol{\xi}_t(\gamma)$ with $\boldsymbol{\xi}_t(\gamma) = -\int_0^t \mathbf{r}_\tau(\gamma) d\tau$.

Proof. Let $g_\infty^i(\gamma)$ and $h_\infty^i(\gamma)$ be the i th coordinate of $\mathbf{g}_\infty(\gamma)$ and $\mathbf{h}_\infty(\gamma)$ defined in (25), respectively. It follows from (30) that

$$g_\infty^i(\gamma) = \gamma \cdot \exp(\alpha \cdot \xi_\infty^i(\gamma)), \quad h_\infty^i(\gamma) = \gamma \cdot \exp(-\alpha \cdot \xi_\infty^i(\gamma)), \quad \forall i = 1, 2, \dots, m.$$

When $\gamma \rightarrow 0$, we have

$$g_\infty^i(\gamma) \cdot h_\infty^i(\gamma) = \gamma^2 \rightarrow 0, \quad \forall i = 1, 2, \dots, m,$$

so that $\lim_{\gamma \rightarrow 0} \mathbf{g}(\gamma) \circ \mathbf{h}(\gamma) = \mathbf{0}$. This also implies that either $g_\infty^i(\gamma)$ or $h_\infty^i(\gamma)$ for any $i = 1, 2, \dots, m$.

On the other hand, let us define

$$\mathbf{s}_\infty(\gamma) = \mathbf{g}_\infty(\gamma) \circ \mathbf{g}_\infty(\gamma) - \mathbf{h}_\infty(\gamma) \circ \mathbf{h}_\infty(\gamma),$$

and let $s_\infty^i(\gamma)$ be the i th coordinate of $\mathbf{s}_\infty(\gamma)$ with

$$s_\infty^i(\gamma) = \gamma^2 \cdot \exp(2\alpha \cdot \xi_\infty^i(\gamma)) - \gamma^2 \cdot \exp(-2\alpha \cdot \xi_\infty^i(\gamma)). \quad (43)$$

Correspondingly, we know that $\widehat{\mathbf{s}} = \lim_{\gamma \rightarrow 0} \mathbf{s}_\infty(\gamma)$ and let \widehat{s}^i be the i th coordinate of $\widehat{\mathbf{s}} = \widehat{\mathbf{g}} \circ \widehat{\mathbf{g}} - \widehat{\mathbf{h}} \circ \widehat{\mathbf{h}}$. In the following, we leverage on these to show that our construction of $\nu(\gamma)$ satisfies (42). We classify the entries \widehat{s}_i of $\widehat{\mathbf{s}}$ ($i = 1, 2, \dots, m$) into three cases and analyze as follows.

- **Case 1:** $\widehat{s}_i > 0$. Since $\lim_{\gamma \rightarrow 0} s_\infty^i(\gamma) = \widehat{s}_i > 0$, from (43) we must have $\xi_\infty^i(\gamma) \rightarrow +\infty$ when $\gamma \rightarrow 0$, so that $\exp(2\alpha \cdot \xi_\infty^i(\gamma)) \rightarrow +\infty$ and $\exp(-2\alpha \cdot \xi_\infty^i(\gamma)) \rightarrow 0$. Therefore, when $\gamma \rightarrow 0$, we have

$$\begin{aligned} \gamma^2 \exp(2\alpha \cdot \xi_\infty^i(\gamma)) \rightarrow \widehat{s}_i &\implies 2\alpha \cdot \xi_\infty^i(\gamma) - 2 \log(1/\gamma) - \log \widehat{s}_i \rightarrow 0, \\ &\implies \nu_i(\gamma) = \frac{\xi_\infty^i(\gamma)}{\log(1/\gamma)} \rightarrow \frac{1}{\alpha}. \quad (\text{given } \log(1/\gamma) \rightarrow +\infty) \end{aligned}$$

- **Case 2:** $\widehat{s}_i < 0$. Since $\lim_{\gamma \rightarrow 0} s_\infty^i(\gamma) = \widehat{s}_i < 0$, from (43) we must have $\xi_\infty^i(\gamma) \rightarrow -\infty$ when $\gamma \rightarrow 0$, so that $\exp(2\alpha \cdot \xi_\infty^i(\gamma)) \rightarrow 0$ and $\exp(-2\alpha \cdot \xi_\infty^i(\gamma)) \rightarrow +\infty$. Therefore, when $\gamma \rightarrow 0$, we have

$$\begin{aligned} -\gamma^2 \exp(-2\alpha \cdot \xi_\infty^i(\gamma)) \rightarrow \widehat{s}_i &\implies -2\alpha \cdot \xi_\infty^i(\gamma) + 2 \log(1/\gamma) - \log \widehat{s}_i \rightarrow 0, \\ &\implies \nu_i(\gamma) = \frac{\xi_\infty^i(\gamma)}{\log(1/\gamma)} \rightarrow -\frac{1}{\alpha}. \end{aligned}$$

- **Case 3:** $\widehat{s}_i = 0$. Since $\lim_{\gamma \rightarrow 0} s_\infty^i(\gamma) = \widehat{s}_i = 0$, from (43) we must have $\gamma^2 \cdot \exp(2\alpha \cdot \xi_\infty^i(\gamma)) \rightarrow 0$ and $\gamma^2 \cdot \exp(-2\alpha \cdot \xi_\infty^i(\gamma)) \rightarrow 0$, when $\gamma \rightarrow 0$. Therefore, for any small $\epsilon \in (0, 1)$, there exists some $\gamma_0 > 0$, such that for all $\gamma \in (0, \gamma_0)$, we have

$$\begin{aligned} \gamma^2 \cdot \max\{\exp(2\alpha \cdot \xi_\infty^i(\gamma)), \exp(-2\alpha \cdot \xi_\infty^i(\gamma))\} &\leq \epsilon \\ \implies 2\alpha \cdot \max\left\{\frac{\xi_\infty^i(\gamma)}{\log(1/\gamma)}, -\frac{\xi_\infty^i(\gamma)}{\log(1/\gamma)}\right\} - 2 &< \frac{\log \epsilon}{\log(1/\gamma)} < 0, \end{aligned}$$

which further implies that

$$|\nu_i(\gamma)| = \max\left\{\frac{\xi_\infty^i(\gamma)}{\log(1/\gamma)}, -\frac{\xi_\infty^i(\gamma)}{\log(1/\gamma)}\right\} < \frac{1}{\alpha}.$$

Therefore, combining the results in the three cases above we obtain that

$$\lim_{\gamma \rightarrow 0} \nu_i(\gamma) = \frac{1}{\alpha} \text{sign}(\widehat{s}_i) = \begin{cases} \frac{\widehat{s}_i}{\alpha|\widehat{s}_i|} & \widehat{s}_i \neq 0, \\ [-1/\alpha, 1/\alpha] & \widehat{s}_i = 0, \end{cases}$$

so that we have (42) holds. \square

C Extra Experiments

Due to limited space in the main body, we here provide extra results for our experiments on robust image recovery presented in Section 4.2.

Varying corruption levels. In Figure 5, we display results of our method for robust image recovery with varying levels of salt-and-pepper corruption.

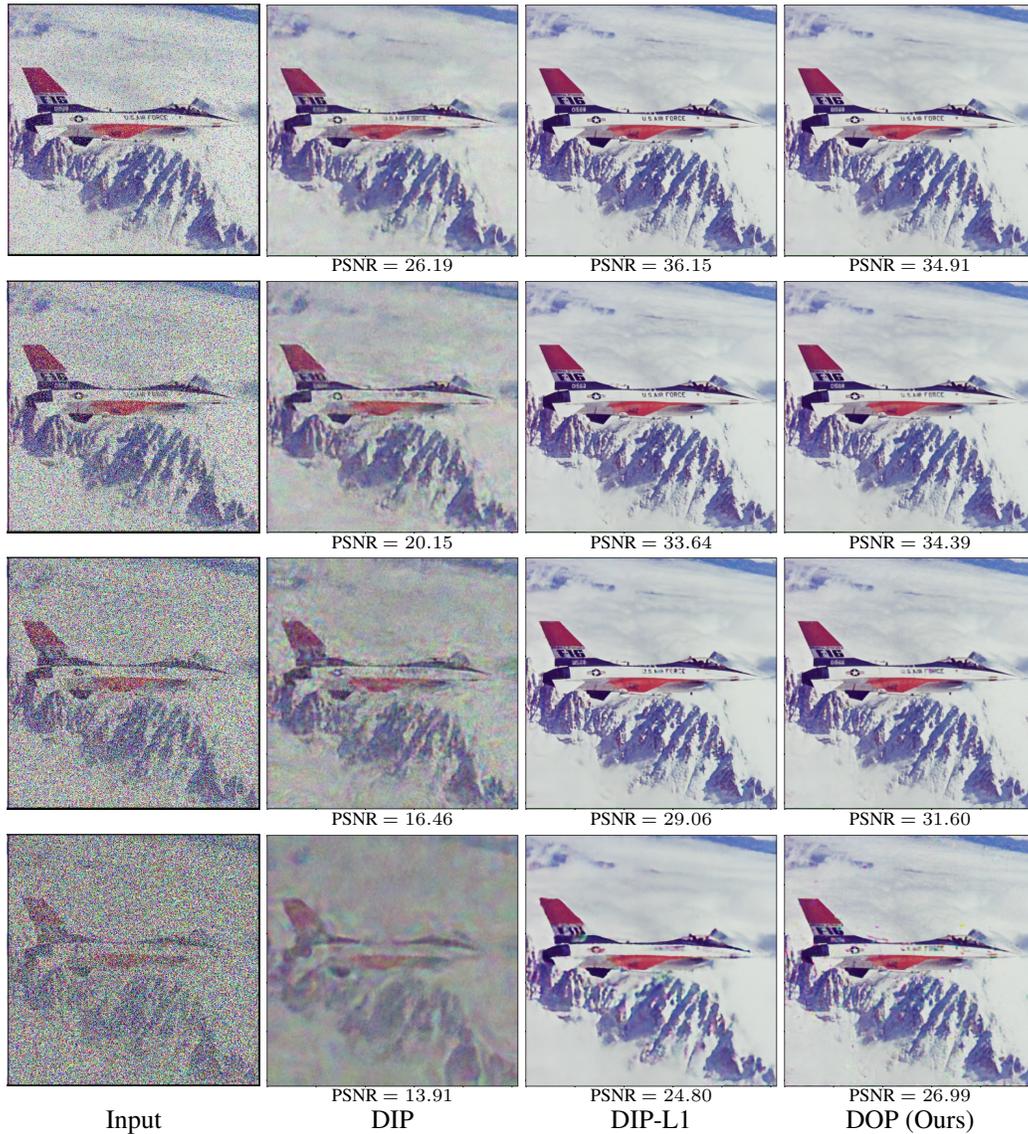


Figure 5: **Robust image recovery with 10%, 30%, 50%, and 70% salt-and-pepper noise.** PSNR of the results is displayed below the images. For our method, all cases use the same network width, learning rate, and termination condition. For DIP and DIP- ℓ_1 , case-dependent early stopping is used which is essential for their good performance. Despite that, our method achieves the highest PSNRs and best visual quality.