
Model Immunization from a Condition Number Perspective

Amber Yijia Zheng^{*1} Cedar Site Bai^{*1} Brian Bullins¹ Raymond A. Yeh¹

Abstract

Model immunization aims to pre-train models that are difficult to fine-tune on harmful tasks while retaining their utility on other non-harmful tasks. Though prior work has shown empirical evidence for immunizing text-to-image models, the key understanding of when immunization is possible and a precise definition of an immunized model remain unclear. In this work, we propose a framework, based on the condition number of a Hessian matrix, to analyze model immunization for linear models. Building on this framework, we design an algorithm with regularization terms to control the resulting condition numbers after pre-training. Empirical results on linear models and non-linear deep-nets demonstrate the effectiveness of the proposed algorithm on model immunization. The code is available at <https://github.com/amberyzheng/model-immunization-cond-num>.

1. Introduction

Model immunization, recently proposed by Zheng & Yeh (2024), studies how to pre-train a model that is more difficult to fine-tune on harmful content, but not others. The aim is to mitigate the risk of misuse (Brundage et al., 2018; Marchal et al., 2024) associated with open-sourced models by immunizing them before they are released to the public.

Zheng & Yeh (2024) focus on immunizing text-to-image models, where they formulate immunization as a bi-level optimization. Empirically, they show that pre-trained diffusion models that undergo immunization are more difficult to fine-tune on a given harmful concept dataset. To quantify this difficulty, they compare the generation quality of models with and without immunization after a fixed number of fine-tuning iterations. While the empirical results are promising,

^{*}Equal contribution ¹Department of Computer Science, Purdue University. Correspondence to: Raymond A. Yeh <rayyeh@purdue.edu>.

a definition of an immunized model and the circumstances that make immunization possible remain unclear.

To tackle this issue, we propose a framework to study model immunization using the condition number (Gloub & Van Loan, 1996). The effectiveness of immunization can be characterized by the condition number of the Hessian matrix. When using gradient-based methods during fine-tuning, a condition number closer to one indicates faster convergence (Boyd & Vandenberghe, 2004), *i.e.*, easier to fine-tune. With this perspective, we observe that the existence of an effective immunization for linear models is related to the angle between the singular vectors of the harmful fine-tuning dataset’s covariance matrix and the pre-training dataset’s covariance matrix.

From this condition number perspective, we propose an immunization algorithm to find such a model. In detail, we propose two additional terms to regularize the condition number during pre-training. Each of the introduced regularization terms can be shown to ensure a monotonic increase/decrease of the condition number under gradient updates.

Beyond the theoretical results, we empirically validate the proposed algorithm on linear models for regression and image classification tasks. Lastly, we conduct experiments using the proposed algorithm on non-linear models, *i.e.*, deep-nets. Despite the gap in theory, we observe that the proposed approach remains effective at model immunization across ResNet (He et al., 2016) and ViT (Dosovitskiy, 2021).

Our contributions are summarized as follows:

- We introduce a framework based on the condition number to study the task of model immunization. This framework leads to a concrete definition of an immunized model along with a novel experiment setup and evaluation metric to compare the quality of different immunization techniques.
- We propose regularizers to maximize/minimize the condition number, with a guaranteed monotonic increase/decrease when updated with the gradient-based method.
- Together with the task objective and regularizers, we demonstrate that the proposed algorithm effectively immunizes linear models and deep-nets on regression/image classification tasks.

2. Preliminaries

This section provides the background of the condition number and its connection to gradient descent. Additionally, we briefly review transfer learning (Zhuang et al., 2020), as it can be a technique for misusing open-source models.

Condition number and convergence of gradient descent. Given a general matrix \mathbf{S} , the condition number (Gloub & Van Loan, 1996) is defined as

$$\kappa(\mathbf{S}) \triangleq \|\mathbf{S}\|_2 \|\mathbf{S}^\dagger\|_2 = \sigma_{\mathbf{S}}^{\max} / \sigma_{\mathbf{S}}^{\min}, \quad (1)$$

where \dagger is the pseudoinverse and $\sigma_{\mathbf{S}}$ corresponds to the max/min singular value of \mathbf{S} . The condition number is related to the convergence rate of gradient-based algorithms.

Consider an optimization problem $\min_{\mathbf{w}} \mathcal{L}(\mathbf{w})$ where \mathcal{L} is strongly convex and has a Hessian $\nabla^2 \mathcal{L}$ with max/min singular values denoted as $\sigma^{\max/\min}$. In this case, the constant step-size steepest descent algorithm has a convergence rate (Bubeck, 2015) of the following

$$\|\mathbf{w}_t - \mathbf{w}^*\|^2 \leq \left(1 - \frac{\sigma^{\min}}{\sigma^{\max}}\right)^t \|\mathbf{w}_0 - \mathbf{w}^*\|^2, \quad (2)$$

where \mathbf{w}^* denotes the optimal solution, and \mathbf{w}^t denotes the steepest descent iterate at step t . We can observe that a larger condition number corresponds to a slower convergence.

Condition number regularization. Nenov et al. (2024) proposed a regularizer for minimizing the condition number of some general matrix \mathbf{S}

$$\mathcal{R}_{\text{we11}}(\mathbf{S}) = \frac{1}{2} \|\mathbf{S}\|_2^2 - \frac{1}{2p} \|\mathbf{S}\|_F^2, \quad (3)$$

in which p is the minimum dimension of \mathbf{S} , and the norms correspond to the spectral norm and Frobenius norm. They showed that $\mathcal{R}_{\text{we11}}(\mathbf{S})$ is a valid regularizer by proving its nonnegativity, and is an upper bound on $\log(\kappa(\mathbf{S}))$. In addition, they showed that $\mathcal{R}_{\text{we11}}(\mathbf{S})$ is differentiable under some mild conditions, and if updated with gradient descent, it is guaranteed to decrease the condition number monotonically. See Appendix A for the exact statements.

Different from Nenov et al. (2024), we propose a differentiable regularizer that is guaranteed to **increase** the condition number as an upper bound on $1/\log(\kappa(\mathbf{S}))$. For model immunization, instead of a general matrix \mathbf{S} , we need to consider the regularization of the Hessian of linear models composed of a feature extractor and a classifier, while preserving their differentiability and monotonicity guarantees during gradient updates to the feature extractor.

Transfer learning via linear probing. In this work, we focus on the transfer learning method of linear probing. Given a pre-trained feature extractor $f_\theta : \mathbb{R}^{D_{\text{in}}} \rightarrow \mathbb{R}^{D_{\text{hid}}}$, linear probing learns a linear classifier $h_{\mathbf{w}} : \mathbb{R}^{D_{\text{hid}}} \rightarrow$

$\mathbb{R}^{D_{\text{out}}}$ over the target dataset $\mathcal{D} = \{(\mathbf{x}, \mathbf{y})\}$ using the frozen feature extractor f_θ . This model learning is formulated as the following optimization problem

$$\min_{\mathbf{w}} \mathcal{L}(\mathcal{D}, \mathbf{w}, \theta) \triangleq \min_{\mathbf{w}} \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}} \ell(h_{\mathbf{w}} \circ f_\theta(\mathbf{x}), \mathbf{y}) \quad (4)$$

where ℓ denotes a suitable loss function, e.g., cross-entropy. By keeping θ fixed, the model leverages features learned from pre-training task and transfers them to the target task. This approach is effective when the target dataset is too small to train a model from scratch.

3. Immunization with Condition Number

The goal of model immunization is to learn a pre-trained model $g_\omega \circ f_{\theta^\dagger}$, consisting of a classifier g_ω and an immunized feature extractor f_{θ^\dagger} , such that fine-tuning f_{θ^\dagger} on a harmful task is difficult, but not for other tasks. The model should also maintain a good pre-training task performance. Specifically, we study the setting when a bad actor uses linear probing on a pre-trained linear feature extractor with gradient descent.

Immunization setting. We denote a pre-training dataset as $\mathcal{D}_{\text{P}} = \{(\mathbf{x}, \mathbf{y})\}$ and a harmful dataset as $\mathcal{D}_{\text{H}} = \{(\mathbf{x}, \tilde{\mathbf{y}})\}$ where $\mathbf{x} \in \mathbb{R}^{D_{\text{in}}}$. The bad actor performs linear probing using \mathcal{D}_{H} following Eq. (4) with an ℓ_2 loss. We will focus our analysis on linear pre-trained feature extractor without dimensionality reduction, i.e., $f_\theta \triangleq \mathbf{x}^\top \theta$ with $\theta \in \mathbb{R}^{D_{\text{in}} \times D_{\text{in}}}$.

Definition 3.1. Under this setting, a model is said to be *immunized* if it satisfies the following:

(a) It is more difficult to apply linear probing on the harmful task \mathcal{D}_{H} using the immunized feature extractor f_{θ^\dagger} than directly on the input data, i.e.,

$$\kappa(\nabla_{\mathbf{w}}^2 \mathcal{L}(\mathcal{D}_{\text{H}}, \mathbf{w}, \theta^\dagger)) \gg \kappa(\nabla_{\mathbf{w}}^2 \mathcal{L}(\mathcal{D}_{\text{H}}, \mathbf{w}, \mathbf{I})), \quad (5)$$

where \mathbf{I} denotes the identity matrix.

(b) It is not more difficult to apply linear probing on other tasks. As there is only one other task \mathcal{D}_{P} , an immunized feature extractor should have

$$\kappa(\nabla_{\omega}^2 \mathcal{L}(\mathcal{D}_{\text{P}}, \omega, \theta^\dagger)) \leq \kappa(\nabla_{\omega}^2 \mathcal{L}(\mathcal{D}_{\text{P}}, \omega, \mathbf{I})). \quad (6)$$

Note: we use ω to denote the classifier parameters of the pre-training task and \mathbf{w} for the harmful task.

(c) The immunized model should maintain a competitive task performance on the pre-training dataset \mathcal{D}_{P} , i.e.,

$$\min_{\omega, \theta} \mathcal{L}(\mathcal{D}_{\text{P}}, \omega, \theta) \approx \min_{\omega} \mathcal{L}(\mathcal{D}_{\text{P}}, \omega, \theta^\dagger). \quad (7)$$

For linear models, as long as θ^\dagger is invertible, exact equality can be achieved.

3.1. Analysis on Immunized Linear Models

To provide some intuition on how the feature extractor θ affects the convergence of linear probing, we study the analytical form of the singular values of the Hessian. For readability, we will rewrite linear probing in Eq. (4) by considering $f_\theta \triangleq \mathbf{x}^\top \theta$ and a ℓ_2 -loss.

Let $\mathbf{X}_H \in \mathbb{R}^{N \times D_{\text{in}}}$ and $\mathbf{Y}_H \in \mathbb{R}^{N \times D_{\text{out}}}$ denote data from \mathcal{D}_H stacked into matrices with $N \triangleq |\mathcal{D}_H|$. When using a ℓ_2 -loss, Eq. (4) can be written as

$$\min_{\mathbf{w}} \mathcal{L}(\mathcal{D}_H, \mathbf{w}, \theta) = \min_{\mathbf{w}} \|(\mathbf{X}_H \theta) \mathbf{w} - \mathbf{Y}\|_2^2. \quad (8)$$

In this case, the Hessian matrix

$$\mathbf{H}_H(\theta) \triangleq \nabla_{\mathbf{w}}^2 \mathcal{L}(\mathcal{D}_H, \mathbf{w}, \theta) = \theta^\top \mathbf{K}_H \theta, \quad (9)$$

where $\mathbf{K}_H \triangleq \mathbf{X}_H^\top \mathbf{X}_H$ is the data covariance matrix.

Proposition 3.2. *The singular values of the Hessian matrix in Eq. (9) are given by*

$$\sigma_i = \sum_{j=1}^{D_{\text{in}}} (\sigma_{\theta,i} (\mathbf{u}_{\theta,i}^\top \mathbf{q}_j) \sqrt{\gamma_j})^2, \quad \forall i \in \{1, \dots, D_{\text{in}}\}. \quad (10)$$

Here, $\sigma_{\theta,i}$ and $\mathbf{u}_{\theta,i}$ correspond to the i -th singular value and vector of θ . Next, γ_j and \mathbf{q}_j correspond to the j -th singular value and vector of the covariance \mathbf{K} .

Proof sketch. This result can be shown by using the fact that \mathbf{K}_H is a symmetric positive semi-definite matrix and decomposing via SVD. The complete proof is provided in Appendix B.1. \square

From Eq. (10), we can see that the singular value of the Hessian depends on the relative angle between the singular vectors between feature extractor θ and the covariance matrix of the data \mathbf{K}_H . As the feature extractor is shared between the pretrained \mathcal{D}_P and harmful \mathcal{D}_H datasets, the strength of the immunization depends on the relative angle between the singular vectors of \mathbf{K}_P and \mathbf{K}_H . For example, if the singular vectors (sorted by the singular values) are all perfectly aligned between the two, then no θ can simultaneously maximize $\kappa(\nabla_{\mathbf{w}}^2 \mathcal{L}(\mathcal{D}_H, \mathbf{w}, \theta))$ and minimize $\kappa(\nabla_{\omega}^2 \mathcal{L}(\mathcal{D}_P, \omega, \theta))$.

With a better understanding of the effect of the feature extractor θ on the condition number, we will next present an algorithm to immunize a model.

4. Algorithm for Immunizing a Model

We formulate model immunization as an optimization problem with the following objective:

$$\min_{\omega, \theta} \mathcal{R}_{\text{ill}}(\mathbf{H}_H(\theta)) + \mathcal{R}_{\text{well}}(\mathbf{H}_P(\theta)) + \mathcal{L}(\mathcal{D}_P, \omega, \theta), \quad (11)$$

where \mathcal{R}_{ill} , to be defined in Sec. 4.1, denotes our proposed regularizer to maximize the condition number, $\mathcal{R}_{\text{well}}$

Algorithm 1 Condition number regularized gradient descent for model immunization

input Primary task $\mathcal{D}_P = (\mathbf{X}_P, \mathbf{Y}_P)$, harmful task input \mathbf{X}_H , supervised loss \mathcal{L} , learning rate η , regularizing constants $\lambda_P, \lambda_H \in \mathbb{R}_+$, model initialization θ_0, ω_0

- 1: $\mathbf{K}_P = \mathbf{X}_P^\top \mathbf{X}_P$
- 2: $\mathbf{K}_H = \mathbf{X}_H^\top \mathbf{X}_H$
- 3: **for** $t = 0, 1, \dots, T - 1$ **do**
- 4: $\omega_{t+1} = \omega_t - \eta \nabla_{\omega} \mathcal{L}(\omega_t, \theta_t; \mathcal{D}_P)$
- 5: $\mathbf{H}_P(\theta_t) = \theta_t^\top \mathbf{K}_P \theta_t$, $\mathbf{H}_H(\theta_t) = \theta_t^\top \mathbf{K}_H \theta_t$
- 6: $\theta_{t+1} = \theta_t - \eta \nabla_{\theta} \mathcal{L}(\omega_t, \theta_t; \mathbf{X}_1)$
 $\quad - \eta \lambda_P \mathbf{K}_P^{-1} \nabla_{\theta} \mathcal{R}_{\text{well}}(\mathbf{H}_P(\theta_t))$
 $\quad - \eta \lambda_H \mathbf{K}_H^{-1} \nabla_{\theta} \mathcal{R}_{\text{ill}}(\mathbf{H}_H(\theta_t))$

7: **end for**

output Immunized feature extractor $\theta_I \triangleq \theta_T$.

in Eq. (3) denotes the regularizer to minimize the condition number, $\mathbf{H}_P(\theta) \triangleq \nabla_{\omega}^2 \mathcal{L}(\mathcal{D}_P, \omega, \theta) = \theta^\top \mathbf{K}_P \theta$ is the Hessian matrix of the pre-training task, and \mathcal{L} denotes the supervised loss.

Each of the terms encourages the model to satisfy the three immunization requirements in Definition 3.1. For readability, we have dropped the scalar hyperparameters balancing the terms. We propose to solve Eq. (11) using a gradient-based method as outlined in Alg. 1.

In the remainder of this section, we will first introduce the novel regularizer to *maximize* general matrices' condition number and their relevant properties (Sec. 4.1). We then show how to incorporate the regularizers \mathcal{R}_{ill} and $\mathcal{R}_{\text{well}}$ into the immunization setup (Sec. 4.2). Finally, we discuss the provable guarantees with respect to each of the regularizers (Sec. 4.3).

4.1. Regularizer for Maximizing the Condition Number

We analyze the condition number of a general matrix $\mathbf{S} \in \mathbb{R}^{p_r \times p_c}$, $p = \min\{p_r, p_c\}$, and $\text{rank}(\mathbf{S}) = k \leq p$. The compact SVD of \mathbf{S} is given by $\mathbf{S} = \mathbf{U} \text{Diag}(\boldsymbol{\sigma}) \mathbf{V}^\top$, in which $\boldsymbol{\sigma} = [\sigma_1, \dots, \sigma_k]^\top$ such that $\sigma_1^{\max} = \sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_k = \sigma_k^{\min} > 0$ and $\mathbf{u}_i, \mathbf{v}_i$ denotes the i^{th} column vector of \mathbf{U}, \mathbf{V} for $i \in [k]$.

Inspired by the regularizer for minimizing the condition number, we propose its counterpart for maximizing the condition number

$$\mathcal{R}_{\text{ill}}(\mathbf{S}) = \frac{1}{\frac{1}{2k} \|\mathbf{S}\|_F^2 - \frac{1}{2} (\sigma_k^{\min})^2}, \quad (12)$$

which satisfies the properties in the following theorem.

Theorem 4.1 (Properties of κ -maximizing regularizer $\mathcal{R}_{\text{ill}}(\mathbf{S})$).

- (1) [Nonnegativity] For any $\mathbf{S} \in \mathbb{R}^{p_r \times p_c}$, $\mathcal{R}_{\text{ill}}(\mathbf{S}) \geq 0$, and $\mathcal{R}_{\text{ill}}(\mathbf{S}) = 0$ if and only if $\kappa(\mathbf{S}) = \infty$.
- (2) [Upper Bound] $\frac{1}{\log(\kappa(\mathbf{S}))} \leq (\sigma_{\mathbf{S}}^{\max})^2 \mathcal{R}_{\text{ill}}(\mathbf{S})$, i.e., $\mathcal{R}_{\text{ill}}(\mathbf{S})$ upper bounds $\frac{1}{\log(\kappa(\mathbf{S}))}$ when $\sigma_{\mathbf{S}}^{\max}$ is reasonably away from ∞ .
- (3) [Differentiability] If $\sigma_{\mathbf{S}}^{\min} = \bar{\sigma}_k < \sigma_i$ for any $i < k$, i.e., $\sigma_{\mathbf{S}}^{\min}$ is unique, then $\mathcal{R}_{\text{ill}}(\mathbf{S})$ is differentiable and

$$\nabla_{\mathbf{S}} \mathcal{R}_{\text{ill}}(\mathbf{S}) = \frac{\sigma_k \mathbf{u}_k \mathbf{v}_k^{\top} - \frac{1}{k} \mathbf{S}}{\left(\frac{1}{2k} \|\mathbf{S}\|_F^2 - \frac{1}{2} (\sigma_{\mathbf{S}}^{\min})^2 \right)^2}. \quad (13)$$

- (4) [Monotonic Increase] If $\sigma_{\mathbf{S}}^{\min}$ is unique, update \mathbf{S} with $\nabla_{\mathbf{S}} \mathcal{R}_{\text{ill}}(\mathbf{S})$ such that $\mathbf{S}' = \mathbf{S} - \eta_2 \nabla_{\mathbf{S}} \mathcal{R}_{\text{ill}}(\mathbf{S})$ for $0 < \eta_2 < \frac{k}{k-1} \left(\frac{1}{2k} \|\mathbf{S}\|_F^2 - \frac{1}{2} (\sigma_{\mathbf{S}}^{\min})^2 \right)^2$, then $\kappa(\mathbf{S}') > \kappa(\mathbf{S})$.

Proof sketch. We provide some intuitive illustrations of the proof and defer the complete version to Appendix B.2.

For (1), as the squared Frobenius norm of a matrix equals the sum of the squares of its singular values, the denominator of $\mathcal{R}_{\text{ill}}(\mathbf{S})$ is the average of the squared singular values minus their minimum, ensuring it is nonnegative. It can be shown that $\mathcal{R}_{\text{ill}}(\mathbf{S})$ is inversely related to $\kappa(\mathbf{S})$, which indicates that $\mathcal{R}_{\text{ill}}(\mathbf{S}) = 0$ if and only if $\kappa(\mathbf{S}) = \infty$.

For (2) the upper bound holds by the design of $\mathcal{R}_{\text{ill}}(\mathbf{S})$ and applying the mean value inequality on

$$\log(\kappa(\mathbf{S})^2) = \log\left((\sigma_{\mathbf{S}}^{\max})^2\right) - \log\left((\sigma_{\mathbf{S}}^{\min})^2\right). \quad (14)$$

For (3), even though $\sigma_{\mathbf{S}}^{\min}$ is not differentiable since it involves taking the minimum of the singular values, its subdifferential is well-defined (Lewis, 1995). When $\sigma_{\mathbf{S}}^{\min}$ is unique, its subdifferential reduces to a singleton, i.e., its gradient, making $\mathcal{R}_{\text{ill}}(\mathbf{S})$ also differentiable.

For (4), one key observation is that the closed-form $\nabla_{\mathbf{S}} \mathcal{R}_{\text{ill}}(\mathbf{S})$ shares the same set of singular vectors as \mathbf{S} , so that the linear relation in gradient update can be passed on to singular values. By choosing a suitable step size, the increase in condition number can be guaranteed. \square

Theorem 4.1 demonstrates that the regularizer $\mathcal{R}_{\text{ill}}(\mathbf{S})$ introduced is a reasonable upper bound for maximizing condition numbers and indicates that under some mild condition, i.e., the minimum singular value is unique, simple first-order algorithms like gradient descent can be used to minimize the regularizer with guaranteed increase in condition number.

4.2. Incorporating Regularizers into Immunization

Given the immunization setup, we now analyze the regularizer \mathcal{R}_{ill} and $\mathcal{R}_{\text{we11}}$ for matrices with the specific structure

of feature covariance matrices, and propose the corresponding algorithm for model immunization.

As illustrated in the immunization setup, the feature extractor θ is the trainable parameter. For data $\mathbf{X} \in \mathbb{R}^{N \times D_{\text{in}}}$ of the feature extractor, we analyze the condition number of $\mathbf{H}(\theta) \triangleq \theta^{\top} \mathbf{K} \theta \in \mathbb{R}^{D_{\text{in}} \times D_{\text{in}}}$ with $\text{rank}(\mathbf{H}) = k$, and compact SVD $\mathbf{H} = \mathbf{U} \text{Diag}(\boldsymbol{\sigma}) \mathbf{V}^{\top}$. Recall, we define $\mathbf{K} = \mathbf{X}^{\top} \mathbf{X}$ to be the covariance matrix of the data.

In the following theorem, we show that under the same conditions, the introduced regularizers $\mathcal{R}_{\text{ill}}(\cdot)$ and $\mathcal{R}_{\text{we11}}(\cdot)$ are also differentiable w.r.t. θ when applied to $\theta^{\top} \mathbf{K} \theta$.

Theorem 4.2. For $\mathbf{H}(\theta) = \theta^{\top} \mathbf{K} \theta$, if its maximum and minimum singular values σ_1 and σ_k are unique, then

- (1) $\nabla_{\theta} \mathcal{R}_{\text{we11}}(\mathbf{H}(\theta)) = 2\mathbf{K}\theta \left(\sigma_1 \mathbf{v}_1 \mathbf{v}_1^{\top} - \frac{1}{D_{\text{in}}} \theta^{\top} \mathbf{K} \theta \right)$,
- (2) $\nabla_{\theta} \mathcal{R}_{\text{ill}}(\mathbf{H}(\theta)) = \frac{2\mathbf{K}\theta (\sigma_k \mathbf{v}_k \mathbf{v}_k^{\top} - \frac{1}{k} \theta^{\top} \mathbf{K} \theta)}{\left(\frac{1}{2k} \|\theta^{\top} \mathbf{K} \theta\|_F^2 - \frac{1}{2} \sigma_k^2 \right)^2}$.

Proof sketch. The differentiability follows from the same argument of Theorem 4.1 (3) under the condition that the maximum and minimum singular values are unique. The closed-form gradients are computed with the chain rule in matrix calculus defined by the Frobenius inner product. The complete proof can be found in Appendix B.3. \square

With the closed-form gradient of the regularizers w.r.t. θ , we propose our algorithm for model immunization in Alg. 1. Specifically, Alg. 1 employs the general gradient descent framework. Line 4 conducts standard updates for the classifier ω , minimizing the supervised loss \mathcal{L} . In lines 5 to 6, the regularizers \mathcal{R}_{ill} and $\mathcal{R}_{\text{we11}}$ are applied on the feature covariance $\mathbf{H}_{\text{H}}(\theta)$ of the harmful task and $\mathbf{H}_{\text{P}}(\theta)$ of the pre-training task. This is done by updating the feature extractor θ with the gradients $\nabla_{\theta} \mathcal{R}_{\text{ill}}(\mathbf{H}_{\text{H}})$ and $\nabla_{\theta} \mathcal{R}_{\text{we11}}(\mathbf{H}_{\text{P}})$ normalized by their input covariances and the gradient from the supervised loss $\nabla_{\theta} \mathcal{L}$.

4.3. Condition Number Guarantees

We show in the following theorem that the condition number decrease/increase guarantees introduced in Theorem A.1 (4) and Theorem 4.1 (4) are preserved for $\theta^{\top} \mathbf{K} \theta$ even when the gradient updates are taken in θ as in Alg. 1, instead of $\theta^{\top} \mathbf{K} \theta$.

Theorem 4.3. For the trainable feature extractor θ , feature covariance $\mathbf{H}_{\text{P}}(\theta) = \theta^{\top} \mathbf{K}_{\text{P}} \theta$ of the primary task and $\mathbf{H}_{\text{H}}(\theta) = \theta^{\top} \mathbf{K}_{\text{H}} \theta$ of the immunization task with $\text{rank}(\mathbf{H}_{\text{P}}) = k_{\text{P}}$, $\text{rank}(\mathbf{H}_{\text{H}}) = k_{\text{H}}$ and compact SVD $\mathbf{H}_{\text{P}}(\theta) = \mathbf{U}_{\text{P}} \text{Diag}(\boldsymbol{\sigma}_{\text{P}}) \mathbf{V}_{\text{P}}^{\top}$, $\mathbf{H}_{\text{H}}(\theta) = \mathbf{U}_{\text{H}} \text{Diag}(\boldsymbol{\sigma}_{\text{H}}) \mathbf{V}_{\text{H}}^{\top}$, for $\boldsymbol{\sigma}_{\text{P}} = [\sigma_{\text{P},1}, \dots, \sigma_{\text{P},k_{\text{P}}}]$, $\boldsymbol{\sigma}_{\text{H}} = [\sigma_{\text{H},1}, \dots, \sigma_{\text{H},k_{\text{H}}}]$,

- (1) if $\sigma_{\mathbf{H}_{\text{P}}}^{\max}$ is unique, i.e., $\sigma_{\mathbf{H}_{\text{P}}}^{\max} = \sigma_{\text{P},1} > \sigma_{\text{P},2}$, update θ such that $\theta' = \theta - \eta_{\text{P}} \mathbf{K}_{\text{P}}^{-1} \nabla_{\theta} \mathcal{R}_{\text{we11}}(\mathbf{H}_{\text{P}}(\theta))$ for $0 < \eta_{\text{P}} < \min \left\{ \frac{1}{(1 - \frac{1}{D_{\text{in}}}) \sigma_{\text{P},1}}, \frac{\sqrt{\sigma_{\text{P},1} \sigma_{\text{P},2} - \sigma_{\text{P},2}^2}}{\frac{2}{D_{\text{in}}} \sigma_{\text{P},2}^2} \right\}$, then

$$\kappa(\theta'^{\top} \mathbf{K}_P \theta') < \kappa(\theta^{\top} \mathbf{K}_P \theta),$$

- (2) if $\sigma_{\mathbf{H}_H}^{\min}$ is unique, i.e., $\sigma_{\mathbf{H}_H}^{\min} = \sigma_{H, k_H} < \sigma_{H, k_H-1}$, update θ such that $\theta' = \theta - \eta_H \mathbf{K}_H^{-1} \nabla_{\theta} \mathcal{R}_{i11}(\mathbf{H}_H(\theta))$ for $0 < \eta_H < \frac{1}{1 - 2\sigma_{\mathbf{H}_H}^{\min}/k_H} \left(\frac{1}{2k_H} \|\theta^{\top} \mathbf{K}_H \theta\|_F^2 - \frac{1}{2} (\sigma_{\mathbf{H}_H}^{\min})^2 \right)^2$, then $\kappa(\theta'^{\top} \mathbf{K}_H \theta') > \kappa(\theta^{\top} \mathbf{K}_H \theta)$.

Proof sketch. There is a mismatch between the gradient update on θ and the condition number update, which is observed for $\mathbf{H}(\theta)$. To address this, we carefully leverage the structure of the problem, noting that $\mathbf{H}(\theta)$, unlike a general matrix, is symmetric and positive semidefinite, with identical left and right singular vectors. Exploiting this property, along with our algorithm design, ensures that the linearity in singular value updates is preserved when expanding $\mathbf{H}(\theta')$ using the closed-form gradient in Theorem 4.2. Consequently, a monotonic increase or decrease in the condition number can be guaranteed by appropriately selecting the step size. The full proof is provided in Appendix B.4. \square

4.4. Additional Discussion

Implementation considerations. At a glance, it may seem that to implement Alg. 1 using automatic differentiation packages, e.g., Pytorch (Paszke et al., 2019), one would have to implement a custom optimizer and involve multiple update steps. Instead, we observe that by directly modifying the computation graph, it would only involve a single backward pass. This is done by introducing a “dummy layer” with an identified function as its forward pass and its backward pass multiplies the gradient by the inverse feature covariance matrix. The “dummy layer” implementation is inspired by prior works in gradient estimator (Bengio et al., 2013; Roeder et al., 2017). Pseudo-code is provided in Appendix C.3.

Limitations. The monotonicity guarantees in Theorem 4.3 serve as a theoretical justification for our proposed algorithm, albeit a partial reflection of the application setup. Note that the feature extractor is updated with the gradients of the two regularizers jointly together with that of the supervised loss and the guarantees may not linearly combine as such. In practice, maintaining the balance between $\kappa(\mathbf{H}_P(\theta))$ and $\kappa(\mathbf{H}_H(\theta))$ requires a proper choice of hyperparameters.

Next, the current framework we analyzed focuses on linear feature extractors and using linear probing for transfer learning. We are aware of the practical limitations of this setting. To address this, in the experiments, we empirically study the effect of the proposed method on non-linear models, i.e., deep-nets, and demonstrate our method’s potential despite the theoretical gap.

Table 1. Quantitative results of immunization in House Price dataset (Montoya & DataCanary, 2016), computed over 5 random seeds.

Method	Eq. (15) (i) \uparrow	Eq. (15) (ii) \downarrow	RIR \uparrow
\mathcal{R}_{i11} Only	90.02 ± 3.773	72.415 ± 3.545	1.244 ± 0.021
IMMA	7.053 ± 1.662	3.545 ± 0.880	2.001 ± 0.187
Opt κ	1.518 ± 0.027	0.016 ± 0.001	92.58 ± 4.492
Ours	18.92 ± 2.056	0.053 ± 0.002	356.20 ± 5.491

5. Experiments

We evaluate the proposed Alg. 1 on regression and image classification tasks using linear models, and also explored immunizing non-linear models, i.e., deep-nets. Experiment and implementation details are provided in Appendix C.

Evaluation metrics. We introduce the *relative immunization ratio (RIR)* to quantify the effectiveness of the immunization based on the ratio of the condition number of Hessian, defined as follows:

$$\text{RIR} \triangleq \frac{\underbrace{\left(\frac{\kappa(\mathbf{H}_H(\theta_I))}{\kappa(\mathbf{H}_H(\mathbf{I}))} \right)}_{(i)}}{\underbrace{\left(\frac{\kappa(\mathbf{H}_P(\theta_I))}{\kappa(\mathbf{H}_P(\mathbf{I}))} \right)}_{(ii)}} \quad (15)$$

where \mathbf{I} denotes the identity matrix. Each term here measures the ratio between condition numbers with and without the pre-trained feature extractor on the (i) harmful task or (ii) on the pre-training task.

A successful immunization is characterized by:

- (i) a large ratio $\frac{\kappa(\mathbf{H}_H(\theta_I))}{\kappa(\mathbf{H}_H(\mathbf{I}))}$, i.e., using the immunized feature extractor makes the optimization of linear probing more difficult on the harmful task.
- (ii) a small ratio $\frac{\kappa(\mathbf{H}_P(\theta_I))}{\kappa(\mathbf{H}_P(\mathbf{I}))}$, i.e., using the pre-trained extractor *do not* make optimization more difficult on the pre-training task.

To obtain a single metric, we compare (i) and (ii) relative to each other. In other words, an effective immunized model should have a relative immunization ratio $\text{RIR} \gg 1$.

Baselines. We consider three baselines for comparisons:

- \mathcal{R}_{i11} Only immunizes the model by minimizing only the regularizer $\mathcal{R}_{i11}(\mathbf{H}_H)$ as defined in Eq. (12) using gradient descent.
- IMMA (Zheng & Yeh, 2024) is formulated as a bi-level optimization program where both lower and upper tasks are solved via gradient descent. In the lower-level, it minimizes $\mathcal{L}(\mathcal{D}_H, \mathbf{w}, \theta)$ w.r.t. θ to obtain θ^* , and in the upper-level, it maximizes $\mathcal{L}(\mathcal{D}_H, \mathbf{w}, \theta^*) - \mathcal{L}(\mathcal{D}_P, \omega, \theta^*)$ w.r.t. θ by backpropagating through θ^* .
- Opt κ directly minimizes $\kappa(\mathbf{H}_P(\theta)) - \kappa(\mathbf{H}_H(\theta))$ w.r.t. θ

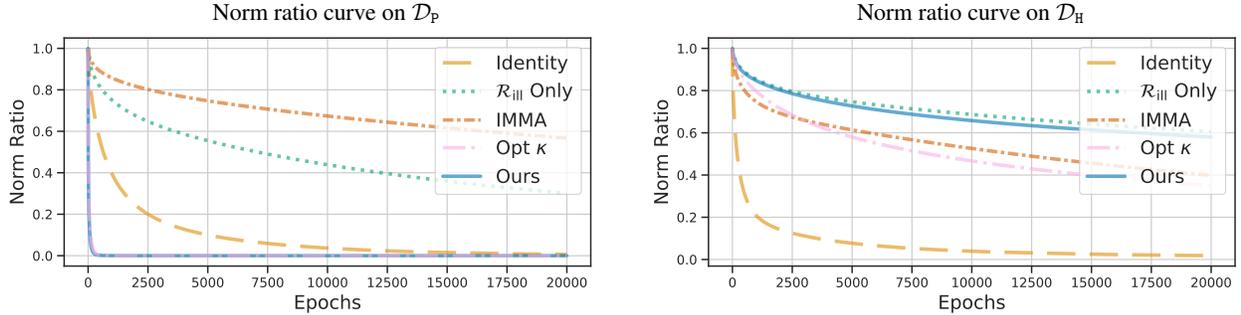


Figure 1. Norm ratio Eq. (16) vs. Epochs. We visualize the convergence of linear probing of different immunized models using gradient descent with an exact line search. Here, *Identity* corresponds to not using a feature extractor, *i.e.*, $\theta_I = \mathbf{I}$. Observe that Ours made the convergence faster on \mathcal{D}_P while slower in \mathcal{D}_H when compared to the other baselines; consistent with the results in Tab. 1.

via gradient descent instead of using our proposed regularizers.

5.1. Experiments on Immunizing Linear Models

Linear regression task. We use the regression task from the House prices dataset (Montoya & DataCanary, 2016). We split the data into \mathcal{D}_P and \mathcal{D}_H based on the feature MSZoning. For the pre-training task, we use the target of LotArea and for the harmful task we use the target of SalePrice. Both \mathcal{D}_P and \mathcal{D}_H contain input vectors of dimension 79. We immunized the model by running Alg. 1 for 100 epochs with $\eta = 0.005$. We choose λ_P and λ_H by balancing the gradient norm of $\mathcal{R}_{\text{well}}$ and \mathcal{R}_{ill} . The implementation details can be found in Appendix C.2.

In Tab. 1, we present the empirical results of immunizing a linear feature extractor θ . We observe that only *Opt κ* and our method successfully immunize the model achieving an RIR that’s much greater than 1. For *R_{III} Only* and *IMMA*, while they successfully made the harmful task more ill-conditioned, *i.e.*, Eq. (15) (i) went up, however, this is at the cost of making the other task ill-conditioned as well, *i.e.*, Eq. (15) (ii) went up.

Next, we demonstrate how a large condition number slows down the convergence of linear probing on the harmful task by analyzing the norm ratio defined as

$$\frac{\|\mathbf{w}_t - \mathbf{w}^*\|_2^2}{\|\mathbf{w}_0 - \mathbf{w}^*\|_2^2}, \quad (16)$$

which measures how the classifier weights \mathbf{w}_t at step t approach the optimal weights \mathbf{w}^* during fine-tuning. Note, naively choosing a step size will not reflect the difference in condition number. Hence, we use the exact line search (Boyd & Vandenberghe, 2004) which chooses the step size that minimizes the loss at each iteration.

As illustrated in Fig. 1, both our method and *Opt κ* slow down convergence in \mathcal{D}_H compared to *Identity* while accelerating convergence in \mathcal{D}_P . Furthermore, our method

Table 2. Quantitative results of immunization in MNIST (LeCun, 1998), computed over 3 random seeds and averaged over all digit pairs. Note that *Opt κ* has large STD in RIR, resulting in the deviation between RIR and the ratio of the averaged values.

Method	Eq. (15) (i) \uparrow	Eq. (15) (ii) \downarrow	RIR \uparrow
<i>R_{III} Only</i>	14.832 \pm 1.039	8.654 \pm 0.606	1.933 \pm 0.046
<i>IMMA</i>	4.522 \pm 0.139	2.774 \pm 0.094	1.774 \pm 0.041
<i>Opt κ</i>	3.196 \pm 1.225	0.756 \pm 1.171	69.73 \pm 54.00
Ours	6.345 \pm 0.188	0.149 \pm 0.009	70.04 \pm 3.280

achieves a stronger immunization effect than *Opt κ*. In contrast, *R_{III} Only* and *IMMA* slowed the convergence on both the harmful task \mathcal{D}_H and the pre-training task \mathcal{D}_P .

Image classification task. For image classification, we conduct experiments using MNIST (LeCun, 1998). The MNIST dataset consists of images over 10-digit classes, which can be formulated into 10 independent binary classification tasks. Across all pairs of tasks, we choose one to be the harmful task \mathcal{D}_H and the other the pre-training \mathcal{D}_P resulting in a total of 90 experiments. We ran Alg. 1 for 30 epochs with $\eta = 0.005$ for these experiments. The implementation details can be found in Appendix C.2.

In Tab. 2, we present the quantitative results on these binary task pairs. For each entry, the values are averaged over all 90 pairs. Based on the averaged results, we observe that our method effectively immunizes the linear feature extractor θ on \mathcal{D}_H without compromising performance on \mathcal{D}_P . Although *Opt κ* achieves comparable RIR with our method, the variances of the metric values are relatively large. This indicates that *Opt κ* is sensitive to random initialization while our method is robust.

In Fig. 2 we further analyze the results by visualizing the $\log(\text{RIR})$ for each digit pair. A blue block indicates successful immunization, while a red block indicates failure. It can

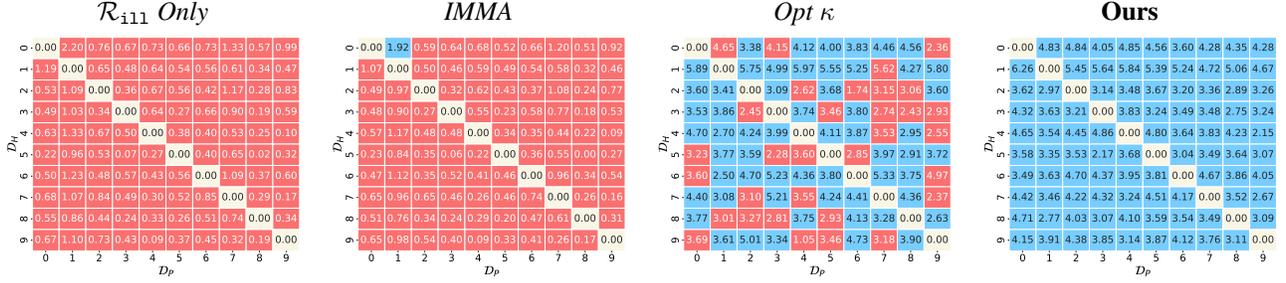


Figure 2. Visualization of $\log(\mathbf{RIR})$ of binary classification tasks created from MNIST. Each element in the figure corresponds to the $\log(\mathbf{RIR})$ of a model immunized against \mathcal{D}_H from the pre-training task of \mathcal{D}_P . We color the block blue if $\mathbf{RIR} \gg 1$, and red otherwise. Our method succeeds in immunizing the model across all digit pairs, while the baselines failed in most pairs.

be observed that \mathcal{R}_{111} Only fails for all digit pairs, $IMMA$ only succeeds in one pair, and $Opt \kappa$ fails for 32 out of 90 pairs. In contrast, our method achieves success across all digit pairs demonstrating its effectiveness for immunization.

Thus far, we have conducted experiments strictly following the immunization setting that we have proposed in Sec. 3. However, one limitation of the setting is that the feature extractor is assumed to be linear, which limits its real-world potential. To further study the practicality of our method, despite the theoretical gap, we conduct experiments with non-linear models, *i.e.*, deep-nets, on a larger-scale image classification dataset of ImageNet.

5.2. Experiments on Immunizing Deep-Nets

Immunization task. In this experiment, we consider a common setup of linear probing on models pre-trained on ImageNet (Deng et al., 2009), *i.e.*, ImageNet serves as \mathcal{D}_P . For \mathcal{D}_H we experiment with the Stanford Cars Dataset (Krause et al., 2013) and Country211 Dataset (Radford et al., 2021). These datasets have been previously used for studying transfer learning (Radford et al., 2021) for image classification. More dataset details are deferred to Appendix C.1.

Experiment setup. For non-linear models, we experiment with the architecture of ResNet18 (He et al., 2016) and ViT (Dosovitskiy, 2021). Here we study a practical setting where a given model with parameters θ_0 has already been trained on \mathcal{D}_P and would undergo immunization to obtain θ_I to be released to the public.

Note that as we are now using an initialization of θ_0 and a non-linear feature extractor f_θ , we extend the \mathbf{RIR} metric to consider those changes. Specifically, we propose

$$\mathbf{RIR}_{\theta_0} \triangleq \frac{\underbrace{\left(\frac{\kappa(\tilde{\mathbf{H}}_H(\theta_I))}{\kappa(\tilde{\mathbf{H}}_H(\theta_0))} \right)}_{(i)}}{\underbrace{\left(\frac{\kappa(\tilde{\mathbf{H}}_P(\theta_I))}{\kappa(\tilde{\mathbf{H}}_P(\theta_0))} \right)}_{(ii)}} \quad (17)$$

where we compare the immunized model θ_I relative to the

initialization model θ_0 . Here, $\tilde{\mathbf{H}}(\theta)$ denotes the Hessian for linear probing on \mathcal{D}_H with a non-linear f_θ , *i.e.*,

$$\tilde{\mathbf{H}}_H(\theta) = \nabla_{\mathbf{w}}^2 \mathcal{L}(\mathcal{D}_H, \mathbf{w}, \theta) = \tilde{\mathbf{X}}_H(\theta)^\top \tilde{\mathbf{X}}_H(\theta). \quad (18)$$

Here, $\tilde{\mathbf{X}}_H(\theta) \triangleq [f_\theta(\mathbf{x}); \forall \mathbf{x} \in \mathcal{D}_H] \in \mathbb{R}^{N \times D_{\text{hid}}}$ denotes the concatenation of the features, with dimensions D_{hid} , extracted from the input data. Due to memory constraints, we approximate Eq. (17) by randomly sampling 20 groups of training data, each containing 100 samples, and reporting the average values.

Finally, we also report the task performance after immunization. This is because, as the feature extractor is non-linear we are no longer guaranteed to retain the task performance. For ResNet18, we immunize only the last two convolutional blocks of the trained feature extractor and keep the rest of the parameters frozen as in θ_0 . For ViT, we only immunize the final transformer block. We optimize Eq. (11) using SGD with momentum, the default optimizer on ImageNet. Further details are provided in Appendix C.2.

Results. We present the quantitative results of immunizing deep-nets in Tab. 3. On both Cars and Country211 datasets, our method demonstrates strong performance when applied to ResNet18 and ViT, as indicated by $\mathbf{RIR}_{\theta_0} \gg 1$. In comparison, \mathcal{R}_{111} Only and $IMMA$ did not effectively immunize the models in all evaluated settings. Next, $Opt \kappa$ also succeeds in immunizing the models but our proposed method outperforms it in \mathbf{RIR}_{θ_0} .

Next, we report the test accuracy of the immunized models on \mathcal{D}_P , *i.e.*, ImageNet1K. On the ResNet18 architecture, we observe a reduction in test-accuracy from the initialization model θ_0 of 68.24% to 62.36% when \mathcal{D}_H is Cars and 65.01% when \mathcal{D}_H is Country211. Interestingly, on the ViT architecture the test-accuracy increased from 81.78% to 82.79% for Cars, and 83.17% for Country211. These results suggested that it is possible to immunize a non-linear model against the harmful task without losing the effectiveness of the other task.

Table 3. Quantitative results of immunization of model pre-trained on ImageNet (Deng et al., 2009), computed over 3 random seeds. The \mathcal{D}_p test accuracy for the off-the-shelf model initialization of θ_0 on ResNet18 is 68.24% and that of ViT is 81.78%. We report RIR_{θ_0} to measure the quality of immunization. Test accuracy of \mathcal{D}_p is reported to ensure the performance on the pre-training task is maintained.

\mathcal{D}_H	Method	ResNet18				ViT			
		Eq. (17) (i) \uparrow	Eq. (17) (ii) \downarrow	$\text{RIR}_{\theta_0} \uparrow$	\mathcal{D}_p Test Acc. (%) \uparrow	Eq. (17) (i) \uparrow	Eq. (17) (ii) \downarrow	$\text{RIR}_{\theta_0} \uparrow$	\mathcal{D}_p Test Acc. (%) \uparrow
	Init. θ_0	1.0	1.0	1.0	68.24	1.0	1.0	1.0	81.78
Cars	\mathcal{R}_{i11} Only	1.878 \pm 0.034	1.786 \pm 0.025	1.057 \pm 0.026	63.84 \pm 0.292	13.121 \pm 0.038	4.097 \pm 0.098	3.342 \pm 0.048	82.21 \pm 0.035
	IMMA	0.866 \pm 0.002	0.889 \pm 0.001	0.974 \pm 0.002	63.57 \pm 0.234	1.422 \pm 0.006	2.090 \pm 0.043	0.702 \pm 0.007	81.89 \pm 0.010
	Opt κ	1.217 \pm 0.021	0.798 \pm 0.005	1.527 \pm 0.019	63.65 \pm 0.148	3.598 \pm 0.510	0.171 \pm 0.033	26.369 \pm 2.814	82.51 \pm 0.085
	Ours	2.386 \pm 0.442	0.699 \pm 0.062	3.467 \pm 0.358	62.36 \pm 0.173	7.945 \pm 0.247	0.323 \pm 0.086	34.517 \pm 0.886	82.79 \pm 0.200
Country211	\mathcal{R}_{i11} Only	20.727 \pm 0.791	20.675 \pm 1.685	1.038 \pm 0.05	62.17 \pm 1.599	69.291 \pm 1.198	63.519 \pm 6.62	1.122 \pm 0.097	80.73 \pm 0.129
	IMMA	0.791 \pm 0.005	0.814 \pm 0.006	0.972 \pm 0.007	67.03 \pm 0.146	6.242 \pm 0.203	7.599 \pm 0.717	0.845 \pm 0.048	82.47 \pm 0.036
	Opt κ	1.538 \pm 0.155	1.053 \pm 0.091	1.472 \pm 0.043	66.81 \pm 0.115	4.589 \pm 0.079	0.300 \pm 0.106	16.498 \pm 5.183	82.79 \pm 0.023
	Ours	3.287 \pm 0.33	0.399 \pm 0.034	8.714 \pm 0.672	65.01 \pm 0.143	20.894 \pm 1.425	0.700 \pm 0.082	41.341 \pm 0.967	83.17 \pm 0.075

To further show a larger Eq. (17) (i) indicating that a model is better immunized, we report the linear probed (fine-tuned) results on different feature extractors and provide the test accuracy on \mathcal{D}_H , where \mathcal{D}_H is the Stanford Cars dataset. As shown in Fig. 3, our method exhibits the slowest convergence rate on both ResNet18 and ViT, indicated by the lowest test accuracy compared with baselines. In summary, our method remains effective on deep-nets, producing models that satisfy the requirements of an immunized model as in Definition 3.1.

6. Related Work

We briefly discuss related research on AI safety and the condition number.

AI safety, model un/re-learning, and immunization. AI safety has received attention lately, specifically in generative AI, due to the impressive progress. We refer the reader to Brundage et al. (2018); Marchal et al. (2024); Bengio et al. (2025) for a more in-depth discussion on this topic. In the following, we will discuss model unlearning, one of the ways to mitigate the potential of misuse, followed by model immunization, which protects a model against relearning.

Machine unlearning was first introduced by Cao & Yang (2015) to remove a user’s private information from a model. Approximate unlearning aims to achieve this by modifying the pre-trained model directly using the specific data samples to erase, without requiring full retraining (Nguyen et al., 2020; Wu et al., 2022; Guo et al., 2019; Sekhari et al., 2021; Neel et al., 2021). In the context of text-to-image models, several methods for concept erasure have been proposed. These include inference-time approaches (Brack et al., 2023; Schramowski et al., 2023), fine-tuning of diffusion models (Gandikota et al., 2023; Kim et al., 2023; Kumari et al., 2023), and direct model editing (Zhang et al., 2024; Gandikota et al., 2024).

While promising, these works still face potential risks of

the re-emergence/re-learning of harmful data (Zheng & Yeh, 2024; Zheng et al., 2024; Zhan et al., 2024; Bertran et al., 2024; Xu et al., 2025). To avoid relearning or further fine-tuning on harmful data, Zheng & Yeh (2024) propose to immunize the text-to-image models against malicious fine-tuning and Zheng & Yeh (2025) extend model immunization to multi-concept settings. Recent work highlights the importance of preventing re-finetuning or distillation on harmful tasks in language models (Huang et al., 2024; Savani et al., 2025) and encoder probing (Ding et al., 2025), which is closely related to our goal. While we also study the task of model immunization, different from Zheng & Yeh (2024) that primarily focuses on empirical applications on generative tasks, our work aims to provide a more principled understanding of model immunization by analyzing it through the lens of the condition number.

Minimizing Condition Number. Condition number has been a key factor in the convergence rates and accuracies of iterative methods, *e.g.*, Jacobi method (Arioli & Romani, 1985), steepest descent (Luenberger et al., 1984), conjugate gradient (Hestenes et al., 1952), for solving optimization problems from classic linear systems (Saad, 2003) to those with general nonlinear objectives (Nesterov, 2018) concerning modern machine learning applications. It is widely observed that a small condition number tends to speed up convergence and improve accuracy whereas a large condition number could lead to an unstable optimization procedure (Saarinen et al., 1993; Kress, 2012; Bengio et al., 2017; Guille-Escuret et al., 2021).

As a result, methods to minimize the condition number in various contexts have been proposed. Preconditioning (Evans, 1968), a technique that involves finding a matrix, *i.e.*, the preconditioner, to multiply with the original matrix, resulting in a new matrix with a significantly smaller condition number, is widely used for solving linear systems. The preconditioner can be constructed using methods such as semidefinite programming (Jambulapati et al., 2020; 2023; Qu et al., 2024) or matrix equilibration (Van der Sluis,

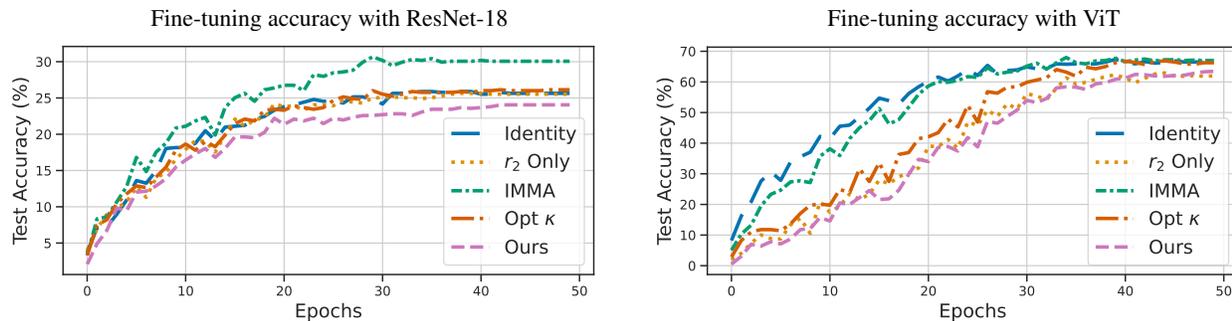


Figure 3. Test accuracy vs. Fine-tuning Epochs on \mathcal{D}_H . We visualize the test accuracy of linear probing on ImageNet of different immunized models using gradient descent. Here \mathcal{D}_H is the Stanford Cars dataset.

1969), and has recently found applications in deep learning (Saratchandran et al., 2024).

Most related to this work, Balazs et al. (2024) propose to regularize the condition number of weight matrices by directly adding the condition number term into the optimization objective and applying (sub)gradient descent. Observing that the condition number is discontinuous and nonconvex, Nenov et al. (2024) proposed a differentiable regularizer that minimizes the matrix condition number with a monotonic decrease guarantee if optimized with gradient descent. To the best of our knowledge, no notable effort has been made to increase or maximize the condition number.

7. Conclusion

We propose a framework for studying model immunization through the condition number of the Hessian matrix. We show that immunization can be achieved by increasing the condition number of harmful datasets while keeping it stable for the pre-training task. To achieve this, we introduce two differentiable regularizers and propose an algorithm that incorporates these regularizers into a gradient-based optimization algorithm. Empirical results on both linear and deep models demonstrate the effectiveness of our approach to model immunization. We believe that our proposed framework is a first step towards a more principled understanding of model immunization and will ultimately make open-sourced models safer.

Acknowledgements

This project is supported in part by an NSF Award #2420724 and the Ross-Lynn Research Scholar Grant.

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning and Optimization. While there are

many potential societal consequences of our work, we believe that the benefits outweigh the harms. Specifically, the topic of model immunization is towards making AI safer.

References

- Arioli, M. and Romani, F. Relations between condition numbers and the convergence of the jacobi method for real positive definite matrices. *Numerische Mathematik*, 1985.
- Balazs, P., Haider, D., Lostanlen, V., and Perfler, F. Trainable signal encoders that are robust against noise. In *INTER-NOISE and NOISE-CON Congress and Conference Proceedings*, 2024.
- Bengio, Y., Léonard, N., and Courville, A. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013.
- Bengio, Y., Goodfellow, I., and Courville, A. *Deep learning*. MIT press Cambridge, MA, USA, 2017.
- Bengio, Y., Mindermann, S., Privitera, D., Besiroglu, T., Bommasani, R., Casper, S., Choi, Y., Fox, P., Garfinkel, B., Goldfarb, D., Heidari, H., Ho, A., Kapoor, S., Khalatbari, L., Longpre, S., Manning, S., Mavroudis, V., Mazeika, M., Michael, J., Newman, J., Ng, K. Y., Okolo, C. T., Raji, D., Sastry, G., Seger, E., Skeadas, T., South, T., Strubell, E., Tramèr, F., Velasco, L., Wheeler, N., Acemoglu, D., Adekanmbi, O., Dalrymple, D., Dietterich, T. G., Felten, E. W., Fung, P., Gourinchas, P.-O., Heintz, F., Hinton, G., Jennings, N., Krause, A., Leavy, S., Liang, P., Ludermir, T., Marda, V., Margetts, H., McDermid, J., Munga, J., Narayanan, A., Nelson, A., Neppel, C., Oh, A., Ramchurn, G., Russell, S., Schaake, M., Schölkopf, B., Song, D., Soto, A., Tiedrich, L., Varoquaux, G., Yao, A., Zhang, Y.-Q., Ajala, O., Albalawi, F., Alserkal, M., Avrin, G., Busch, C., de Carvalho, A. C. P. d. L. F., Fox, B., Gill, A. S., Hatip, A. H., Heikkilä, J., Johnson, C.,

- Jolly, G., Katzir, Z., Khan, S. M., Kitano, H., Krüger, A., Lee, K. M., Ligot, D. V., López Portillo, J. R., Molchanovskiy, O., Monti, A., Mwamanzi, N., Nemer, M., Oliver, N., Pezoa Rivera, R., Ravindran, B., Riza, H., Rugege, C., Seoighe, C., Sheehan, J., Sheikh, H., Wong, D., and Zeng, Y. International AI safety report. Technical Report DSIT 2025/001, 2025. URL <https://www.gov.uk/government/publications/international-ai-safety-report-2025>.
- Bertran, M. A., Tang, S., Kearns, M., Morgenstern, J. H., Roth, A., and Wu, S. Reconstruction attacks on machine unlearning: Simple models are vulnerable. In *Proc. NeurIPS*, 2024.
- Boyd, S. and Vandenberghe, L. *Convex optimization*. Cambridge university press, 2004.
- Brack, M., Friedrich, F., Hintersdorf, D., Struppek, L., Schramowski, P., and Kersting, K. SEGA: Instructing text-to-image models using semantic guidance. In *Proc. NeurIPS*, 2023.
- Brundage, M., Avin, S., Clark, J., Toner, H., Eckersley, P., Garfinkel, B., Dafoe, A., Scharre, P., Zeitzoff, T., Filar, B., et al. The malicious use of artificial intelligence: Forecasting, prevention, and mitigation. *arXiv preprint arXiv:1802.07228*, 2018.
- Bubeck, S. Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning*, 2015.
- Cao, Y. and Yang, J. Towards making systems forget with machine unlearning. In *IEEE symposium on security and privacy*, 2015.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. ImageNet: A large-scale hierarchical image database. In *Proc. CVPR*, 2009.
- Ding, R., Zhou, T., Su, L., Ding, A. A., Xu, X., and Fei, Y. Probe-me-not: Protecting pre-trained encoders from malicious probing. In *Proc. NDSS*, 2025.
- Dosovitskiy, A. An image is worth 16x16 words: Transformers for image recognition at scale. In *Proc. ICLR*, 2021.
- Evans, D. J. The use of pre-conditioning in iterative methods for solving linear equations with symmetric positive definite matrices. *IMA Journal of Applied Mathematics*, 1968.
- Gandikota, R., Materzynska, J., Fiotto-Kaufman, J., and Bau, D. Erasing concepts from diffusion models. In *Proc. ICCV*, 2023.
- Gandikota, R., Orgad, H., Belinkov, Y., Materzyńska, J., and Bau, D. Unified concept editing in diffusion models. In *Proc. WACV*, 2024.
- Gloub, G. H. and Van Loan, C. F. Matrix computations. *Johns Hopkins University Press, 3rd edition*, 1996.
- Guille-Escuret, C., Girotti, M., Goujaud, B., and Mitliagkas, I. A study of condition numbers for first-order optimization. In *Proc. AISTATS*, 2021.
- Guo, C., Goldstein, T., Hannun, A., and Van Der Maaten, L. Certified data removal from machine learning models. *arXiv preprint arXiv:1911.03030*, 2019.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proc. CVPR*, 2016.
- Hestenes, M. R., Stiefel, E., et al. *Methods of conjugate gradients for solving linear systems*. NBS Washington, DC, 1952.
- Huang, T., Hu, S., Ilhan, F., Tekin, S. F., and Liu, L. Harmful fine-tuning attacks and defenses for large language models: A survey. *arXiv preprint arXiv:2409.18169*, 2024.
- Jambulapati, A., Li, J., Musco, C., Sidford, A., and Tian, K. Fast and near-optimal diagonal preconditioning. *arXiv preprint arXiv:2008.01722*, 2020.
- Jambulapati, A., Li, J., Musco, C., Shiragur, K., Sidford, A., and Tian, K. Structured semidefinite programming for recovering structured preconditioners. In *Proc. NeurIPS*, 2023.
- Kim, S., Jung, S., Kim, B., Choi, M., Shin, J., and Lee, J. Towards safe self-distillation of internet-scale text-to-image diffusion models. *arXiv preprint arXiv:2307.05977*, 2023.
- Kingma, D. P. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Krause, J., Stark, M., Deng, J., and Fei-Fei, L. 3d object representations for fine-grained categorization. In *Proc. ICCV Workshops*, 2013.
- Kress, R. *Numerical analysis*. Springer Science & Business Media, 2012.
- Kumari, N., Zhang, B., Wang, S.-Y., Shechtman, E., Zhang, R., and Zhu, J.-Y. Ablating concepts in text-to-image diffusion models. In *Proc. ICCV*, 2023.
- LeCun, Y. The MNIST database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998.
- Lewis, A. S. The convex analysis of unitarily invariant matrix functions. *Journal of Convex Analysis*, 1995.

- Luenberger, D. G., Ye, Y., et al. *Linear and nonlinear programming*, volume 2. Springer, 1984.
- Marchal, N., Xu, R., Elasmr, R., Gabriel, I., Goldberg, B., and Isaac, W. Generative AI misuse: A taxonomy of tactics and insights from real-world data. *arXiv preprint arXiv:2406.13843*, 2024.
- Montoya, A. and DataCanary. House prices - advanced regression techniques, 2016. Kaggle.
- Mordukhovich, B. S. *Variational analysis and applications*. Springer, 2018.
- Neel, S., Roth, A., and Sharifi-Malvajerdi, S. Descent-to-delete: Gradient-based methods for machine unlearning. In *Proc. ALT*, 2021.
- Nenov, R., Haider, D., and Balazs, P. (Almost) Smooth Sailing: Towards numerical stability of neural networks through differentiable regularization of the condition number. In *ICML Differentiable Almost Everything Workshop*, 2024.
- Nesterov, Y. *Lectures on convex optimization*, volume 137. Springer, 2018.
- Nguyen, Q. P., Low, B. K. H., and Jaillet, P. Variational bayesian unlearning. *Proc. NeurIPS*, 2020.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. Pytorch: An imperative style, high-performance deep learning library. In *Proc. NeurIPS*, 2019.
- Petersen, K. B., Pedersen, M. S., et al. The matrix cookbook. *Technical University of Denmark*, 7(15):510, 2008.
- Qu, Z., Gao, W., Hinder, O., Ye, Y., and Zhou, Z. Optimal diagonal preconditioning. *Operations Research*, 2024.
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al. Learning transferable visual models from natural language supervision. In *Proc. ICML*, 2021.
- Rockafellar, R. Convex analysis. *Princeton Mathematical Series*, 28, 1970.
- Roeder, G., Wu, Y., and Duvenaud, D. K. Sticking the landing: Simple, lower-variance gradient estimators for variational inference. In *Proc. NeurIPS*, 2017.
- Saad, Y. *Iterative methods for sparse linear systems*. SIAM, 2003.
- Saarinen, S., Bramley, R., and Cybenko, G. Ill-conditioning in neural network training problems. *SIAM Journal on Scientific Computing*, 1993.
- Saratchandran, H., Wang, T. X., and Lucey, S. Weight conditioning for smooth optimization of neural networks. In *Proc. ECCV*, 2024.
- Savani, Y., Trockman, A., Feng, Z., Schwarzschild, A., Robey, A., Finzi, M., and Kolter, J. Z. Antidistillation sampling. *arXiv preprint arXiv:2504.13146*, 2025.
- Schramowski, P., Brack, M., Deiseroth, B., and Kersting, K. Safe latent diffusion: Mitigating inappropriate degeneration in diffusion models. In *Proc. CVPR*, 2023.
- Sekhri, A., Acharya, J., Kamath, G., and Suresh, A. T. Remember what you want to forget: Algorithms for machine unlearning. In *Proc. NeurIPS*, 2021.
- Thomee, B., Shamma, D. A., Friedland, G., Elizalde, B., Ni, K., Poland, D., Borth, D., and Li, L.-J. YFCC100M: The new data in multimedia research. *Communications of the ACM*, 2016.
- Van der Sluis, A. Condition numbers and equilibration of matrices. *Numerische Mathematik*, 1969.
- Wightman, R. Pytorch image models. <https://github.com/rwightman/pytorch-image-models>, 2019.
- Wu, G., Hashemi, M., and Srinivasa, C. Puma: Performance unchanged model augmentation for training data removal. In *Proc. AAAI*, 2022.
- Xu, X., Yue, X., Liu, Y., Ye, Q., Hu, H., and Du, M. Unlearning isn't deletion: Investigating reversibility of machine unlearning in llms. *arXiv preprint arXiv:2505.16831*, 2025.
- Zhan, Q., Fang, R., Bindu, R., Gupta, A., Hashimoto, T., and Kang, D. Removing rlhf protections in gpt-4 via fine-tuning. In *Proc. NAACL*, 2024.
- Zhang, G., Wang, K., Xu, X., Wang, Z., and Shi, H. Forget-me-not: Learning to forget in text-to-image diffusion models. In *Proc. CVPR*, 2024.
- Zheng, A. Y. and Yeh, R. A. Imma: Immunizing text-to-image models against malicious adaptation. In *Proc. ECCV*, 2024.
- Zheng, A. Y. and Yeh, R. A. Multi-concept model immunization through differentiable model merging. In *Proc. AAAI*, 2025.
- Zheng, A. Y., Yang, C.-A., and Yeh, R. A. Learning to obstruct few-shot image classification over restricted classes. In *Proc. ECCV*, 2024.
- Zhuang, F., Qi, Z., Duan, K., Xi, D., Zhu, Y., Zhu, H., Xiong, H., and He, Q. A comprehensive survey on transfer learning. *Proceedings of the IEEE*, 2020.

Appendix

The appendix is organized as follows:

- In Sec. A, we provide the complete statements of the properties of $\mathcal{R}_{\text{well}}(\mathbf{S})$ for minimizing the condition number.
- In Sec. B, we provide the complete proof for the Theorems stated in the main paper.
- In Sec. C, we provide additional experiment details. The code will be open-sourced upon the acceptance of this paper.

A. Properties of the Condition Number Minimizing Regularizer

Theorem A.1 (Properties of κ -minimizing regularizer $\mathcal{R}_{\text{well}}(\mathbf{S})$, Theorem 2.1, 2.2, 3.1, 3.2 in [Nenov et al. \(2024\)](#)).

- (1) [Nonnegativity] $\forall \mathbf{S} \in \mathbb{R}^{p_r \times p_c}$, $\mathcal{R}_{\text{well}}(\mathbf{S}) \geq 0$. If $\mathbf{S} \neq \mathbf{0}$, $\mathcal{R}_{\text{well}}(\mathbf{S}) = 0$ if and only if \mathbf{S} has full rank and $\kappa(\mathbf{S}) = 1$.
- (2) [Upper Bound] $\kappa(\mathbf{S}) \leq e^{p(\sigma_{\mathbf{S}}^{\min})^{-2}\mathcal{R}_{\text{well}}(\mathbf{S})}$, i.e., $r(\mathbf{S})$ is an upper bound of $\log(\kappa(\mathbf{S}))$ as long as $\sigma_{\mathbf{S}}^{\min}$ is bounded away from 0.
- (3) [Differentiability] If $\sigma_{\mathbf{S}}^{\max} = \sigma_1 > \sigma_i$ for any $i > 1$, i.e., $\sigma_{\mathbf{S}}^{\max}$ is unique, then $\mathcal{R}_{\text{well}}(\mathbf{S})$ is differentiable and its gradient is given by $\nabla_{\mathbf{S}}\mathcal{R}_{\text{well}}(\mathbf{S}) = \sigma_1 \mathbf{u}_1 \mathbf{v}_1^\top - \frac{1}{p}\mathbf{S}$.
- (4) [Monotonic Decrease] If $\sigma_{\mathbf{S}}^{\max}$ is unique, update \mathbf{S} with $\nabla_{\mathbf{S}}\mathcal{R}_{\text{well}}(\mathbf{S})$ such that $\mathbf{S}' = \mathbf{S} - \eta_1 \nabla_{\mathbf{S}}\mathcal{R}_{\text{well}}(\mathbf{S})$ for $0 < \eta_1 < \frac{\kappa(\mathbf{S})-1}{(1-\frac{1}{p})\kappa(\mathbf{S})+\frac{1}{p}}$, then $\kappa(\mathbf{S}') < \kappa(\mathbf{S})$.

B. Proof of Propositions and Theorems

B.1. Proof of Proposition 3.2.

Proposition 3.2. *The singular values of the Hessian matrix in Eq. (9) are given by*

$$\sigma_i = \sum_{j=1}^{D_{\text{in}}} (\sigma_{\theta,i}(\mathbf{u}_{\theta,i}^\top \mathbf{q}_j) \sqrt{\gamma_j})^2, \quad \forall i \in \{1, \dots, D_{\text{in}}\}. \quad (10)$$

Here, $\sigma_{\theta,i}$ and $\mathbf{u}_{\theta,i}$ correspond to the i -th singular value and vector of θ . Next, γ_j and \mathbf{q}_j correspond to the j -th singular value and vector of the covariance \mathbf{K} .

Proof. Substitute the SVD of θ and the eigendecomposition of \mathbf{K} into $\theta^\top \mathbf{K} \theta$:

$$\theta^\top \mathbf{K} \theta = (\mathbf{U}_\theta \Sigma_\theta \mathbf{V}_\theta^\top)^\top (\mathbf{Q} \Gamma^2 \mathbf{Q}^\top) (\mathbf{U}_\theta \Sigma_\theta \mathbf{V}_\theta^\top).$$

Simplify the expression:

$$\theta^\top \mathbf{K} \theta = \mathbf{V}_\theta (\Sigma_\theta \mathbf{U}_\theta^\top \mathbf{Q} \Gamma^2 \mathbf{Q}^\top \mathbf{U}_\theta \Sigma_\theta) \mathbf{V}_\theta^\top.$$

Define $\mathbf{M} = \Sigma_\theta \mathbf{U}_\theta^\top \mathbf{Q} \Gamma$, so that:

$$\theta^\top \mathbf{K} \theta = \mathbf{V}_\theta (\mathbf{M} \mathbf{M}^\top) \mathbf{V}_\theta^\top.$$

The elements of \mathbf{M} are:

$$\mathbf{M}[i, j] = \sigma_{\theta,i}(\mathbf{u}_{\theta,i}^\top \mathbf{q}_j) \gamma_j,$$

where $\sigma_{\theta,i}$'s for $i \in [d]$ are the singular values of θ , γ_j 's for $i \in [d]$ are the diagonal entries of Γ , and $(\mathbf{u}_{\theta,i}^\top \mathbf{q}_j)$ measures the alignment between the i -th column of \mathbf{U}_θ and the j -th column of \mathbf{Q} .

We observe the following decomposition of M in to two matrices O and D :

$$\begin{aligned}
 M &= \begin{bmatrix} \ddots & \vdots & \ddots \\ \dots & \sigma_{\theta,i}(\mathbf{u}_{\theta,i}^\top \mathbf{q}_j) \gamma_j & \dots \\ \ddots & \vdots & \ddots \end{bmatrix} \\
 &= \begin{bmatrix} \ddots & \vdots & \ddots \\ \dots & \frac{\sigma_{\theta,i}(\mathbf{u}_{\theta,i}^\top \mathbf{q}_j) \gamma_j}{\sqrt{\sum_{j'} (\sigma_{\theta,i}(\mathbf{u}_{\theta,i}^\top \mathbf{q}_{j'}) \gamma_{j'})^2}} & \dots \\ \ddots & \vdots & \ddots \end{bmatrix} \begin{bmatrix} \ddots & 0 & 0 \\ 0 & \sqrt{\sum_{j'} (\sigma_{\theta,i}(\mathbf{u}_{\theta,i}^\top \mathbf{q}_{j'}) \gamma_{j'})^2} & 0 \\ 0 & 0 & \ddots \end{bmatrix} \\
 &= OD
 \end{aligned}$$

where O is an orthonormal matrix, i.e., $O^\top O = I$, and $D = \text{diag}(d_1, \dots, d_d)$ with $d_i = \sqrt{\sum_{j'} (\sigma_{\theta,i}(\mathbf{u}_{\theta,i}^\top \mathbf{q}_{j'}) \gamma_{j'})^2}$ is a diagonal matrix. As a result, diagonal entries of D^2 are:

$$d_i^2 = \sum_{j=1}^d (\sigma_{\theta,i}(\mathbf{u}_{\theta,i}^\top \mathbf{q}_j) \gamma_j)^2.$$

Thus, $MM^\top = (OD)(OD)^\top = OD^2O^\top$, and the eigenvalues of $\theta^\top K \theta$ are the diagonal entries of D^2 , given by:

$$\sigma_i = d_i^2 = \sum_{j=1}^d (\sigma_{\theta,i}(\mathbf{u}_{\theta,i}^\top \mathbf{q}_j) \gamma_j)^2, \quad i = 1, \dots, d.$$

□

B.2. Proof of Theorem 4.1

B.2.1. PROOF OF THEOREM 4.1 (1)

Theorem 4.1. (1) For any $\mathbf{S} \in \mathbb{R}^{p_r \times p_c}$, $\mathcal{R}_{i11}(\mathbf{S}) \geq 0$, and $\mathcal{R}_{i11}(\mathbf{S}) = 0$ if and only if $\kappa(\mathbf{S}) = \infty$.

Proof. By definition, $\mathcal{R}_{i11}(\mathbf{S}) = \frac{1}{\frac{1}{2k} \|\mathbf{S}\|_F^2 - \frac{1}{2} (\sigma_{\mathbf{S}}^{\min})^2}$. Denote $\mathcal{R}'_{i11}(\mathbf{S}) = \frac{1}{2} \left((\sigma_{\mathbf{S}}^{\min})^2 - \frac{1}{k} \|\mathbf{S}\|_F^2 \right)$, then we have $\mathcal{R}_{i11}(\mathbf{S}) = \frac{1}{-\mathcal{R}'_{i11}(\mathbf{S})}$, and

$$\begin{aligned}
 \mathcal{R}'_{i11}(\mathbf{S}) &= \frac{1}{2} \left(\sigma_k^2 - \frac{1}{k} \sum_{i=1}^k \sigma_i^2 \right) \\
 &= \frac{1}{2k} \sum_{i=1}^k (\sigma_k^2 - \sigma_i^2) \\
 &\leq 0,
 \end{aligned}$$

since $\forall i \in [k]$, $\sigma_{\mathbf{S}}^{\min} = \sigma_k \leq \sigma_i$. As a result, $-\mathcal{R}'_{i11}(\mathbf{S}) \geq 0$ and $\mathcal{R}_{i11}(\mathbf{S}) = \frac{1}{-\mathcal{R}'_{i11}(\mathbf{S})} \geq 0$, i.e., $\mathcal{R}_{i11}(\mathbf{S})$ is non-negative.

Also, by definition, $\sigma_1 = \kappa(\mathbf{S}) \sigma_k$. Therefore,

$$\begin{aligned} \mathcal{R}_{\text{ill}}(\mathbf{S}) &= \frac{2}{\frac{1}{k} \sum_{i=1}^k \sigma_i^2 - (\sigma_{\mathbf{S}}^{\min})^2} \\ &\leq \frac{2}{\frac{1}{k} \sigma_1^2 + \frac{k-1}{k} (\sigma_{\mathbf{S}}^{\min})^2 - (\sigma_{\mathbf{S}}^{\min})^2} \\ &= \frac{2}{\frac{1}{k} (\sigma_1^2 - (\sigma_{\mathbf{S}}^{\min})^2)} \\ &= \frac{2k}{(\kappa(\mathbf{S})^2 - 1) (\sigma_{\mathbf{S}}^{\min})^2}. \end{aligned}$$

If $\kappa(\mathbf{S}) = \infty$, $\mathcal{R}_{\text{ill}}(\mathbf{S}) \leq \frac{2k}{(\kappa(\mathbf{S})^2 - 1) (\sigma_{\mathbf{S}}^{\min})^2} = 0$ for $\sigma_{\mathbf{S}}^{\min} > 0$, which yields $\mathcal{R}_{\text{ill}}(\mathbf{S}) = 0$ given that $\mathcal{R}_{\text{ill}}(\mathbf{S}) \geq 0$.

Similarly, we have

$$\begin{aligned} \mathcal{R}_{\text{ill}}(\mathbf{S}) &= \frac{2}{\frac{1}{k} \sum_{i=1}^k \sigma_i^2 - (\sigma_{\mathbf{S}}^{\min})^2} \\ &\geq \frac{2}{\frac{k-1}{k} \sigma_1^2 + \frac{1}{k} (\sigma_{\mathbf{S}}^{\min})^2 - (\sigma_{\mathbf{S}}^{\min})^2} \\ &= \frac{2}{\frac{k-1}{k} (\sigma_1^2 - (\sigma_{\mathbf{S}}^{\min})^2)} \\ &= \frac{\frac{2k}{k-1}}{(\kappa(\mathbf{S})^2 - 1) (\sigma_{\mathbf{S}}^{\min})^2}. \end{aligned}$$

If $\mathcal{R}_{\text{ill}}(\mathbf{S}) = 0$, we have $\kappa(\mathbf{S}) \geq \sqrt{\frac{\frac{2k}{k-1}}{\mathcal{R}_{\text{ill}}(\mathbf{S}) (\sigma_{\mathbf{S}}^{\min})^2}} + 1 = \infty$ which yields $\kappa(\mathbf{S}) = \infty$. □

B.2.2. PROOF OF THEOREM 4.1 (2)

To prove Theorem 4.1 (2), we start by analyzing $\mathcal{R}'_{\text{ill}}(\mathbf{S}) = \frac{1}{2} \left((\sigma_{\mathbf{S}}^{\min})^2 - \frac{1}{k} \|\mathbf{S}\|_F^2 \right)$ with the following lemma.

Lemma B.1. For $\mathcal{R}'_{\text{ill}}(\mathbf{S}) = \frac{1}{2} \left((\sigma_{\mathbf{S}}^{\min})^2 - \frac{1}{k} \|\mathbf{S}\|_F^2 \right)$,

$$\frac{1}{\kappa(\mathbf{S})} \leq e^{\frac{k}{k-1} \sigma_1^{-2} \mathcal{R}'_{\text{ill}}(\mathbf{S})} \quad (19)$$

That is, $\mathcal{R}'_{\text{ill}}(\mathbf{S})$ is an upper bound of $\log \left(\frac{1}{\kappa(\mathbf{S})} \right)$, i.e., $-\log(\kappa(\mathbf{S}))$.

Proof. Similar to the proof of Theorem 3.2 in (Nenov et al., 2024),

$$\begin{aligned} 2\mathcal{R}'_{\text{ill}}(\mathbf{S}) &= (\sigma_{\mathbf{S}}^{\min})^2 - \frac{1}{k} \|\mathbf{S}\|_F^2 \\ &= (\sigma_{\mathbf{S}}^{\min})^2 - \frac{1}{k} \sum_{i=1}^k \sigma_i^2 \\ &\geq (\sigma_{\mathbf{S}}^{\min})^2 - \frac{1}{k} \left((k-1) \sigma_1^2 + (\sigma_{\mathbf{S}}^{\min})^2 \right) \\ &= \left(1 - \frac{1}{k} \right) \left((\sigma_{\mathbf{S}}^{\min})^2 - \sigma_1^2 \right) \end{aligned}$$

In the meantime,

$$\begin{aligned}
 2 \log \left(\frac{1}{\kappa(\mathbf{S})} \right) &= -\log(\kappa(\mathbf{S})^2) \\
 &= \log \left((\sigma_{\mathbf{S}}^{\min})^2 \right) - \log(\sigma_1^2) \\
 &\leq -\frac{1}{\sigma_1^2} \left(\sigma_1^2 - (\sigma_{\mathbf{S}}^{\min})^2 \right) \\
 &= \frac{1}{\sigma_1^2} \left((\sigma_{\mathbf{S}}^{\min})^2 - \sigma_1^2 \right)
 \end{aligned}$$

in which the inequality follows from the Mean Value Theorem. As a result,

$$\begin{aligned}
 \frac{1}{\kappa(\mathbf{S})} &\leq e^{\frac{1}{2\sigma_1^2} \left((\sigma_{\mathbf{S}}^{\min})^2 - \sigma_1^2 \right)} \\
 &\leq e^{\frac{1}{2\sigma_1^2} \left(\frac{k}{k-1} 2\mathcal{R}'_{\text{ill}}(\mathbf{S}) \right)} \\
 &= e^{\frac{k}{k-1} \sigma_1^{-2} \mathcal{R}'_{\text{ill}}(\mathbf{S})}
 \end{aligned}$$

□

Theorem 4.1. (2) $\frac{1}{\log(\kappa(\mathbf{S}))} \leq (\sigma_{\mathbf{S}}^{\max})^2 \mathcal{R}_{\text{ill}}(\mathbf{S})$, i.e., $\mathcal{R}_{\text{ill}}(\mathbf{S})$ upper bounds $\frac{1}{\log(\kappa(\mathbf{S}))}$ when $\sigma_{\mathbf{S}}^{\max}$ is reasonably away from ∞ .

Proof. Taking the logarithm of Lemma B.1, we have

$$-\log(\kappa(\mathbf{S})) \leq \frac{k}{k-1} \sigma_1^{-2} \mathcal{R}'_{\text{ill}}(\mathbf{S}).$$

Negating both sides,

$$\log(\kappa(\mathbf{S})) \geq -\frac{k}{k-1} \sigma_1^{-2} \mathcal{R}'_{\text{ill}}(\mathbf{S}).$$

Finally, taking the reciprocal,

$$\begin{aligned}
 \frac{1}{\log(\kappa(\mathbf{S}))} &\leq \frac{k-1}{k} \frac{\sigma_1^2}{-\mathcal{R}'_{\text{ill}}(\kappa(\mathbf{S}))} \\
 &= \frac{k-1}{k} \frac{\sigma_1^2}{\frac{1}{2} (\sigma_{\mathbf{S}}^{\min})^2 - \frac{1}{2k} \|\mathbf{S}\|_F^2} \\
 &\leq \sigma_1^2 \mathcal{R}_{\text{ill}}(\mathbf{S})
 \end{aligned}$$

□

B.2.3. PROOF OF THEOREM 4.1 (3)

To analyze the differentiability of $\mathcal{R}_{\text{ill}}(\mathbf{S}) = \frac{1}{\frac{1}{2k} \|\mathbf{S}\|_F^2 - \frac{1}{2} (\sigma_{\mathbf{S}}^{\min})^2}$, we start by analyzing the differentiability of $\mathcal{R}'_{\text{ill}}(\mathbf{S}) = \frac{1}{2} \left((\sigma_{\mathbf{S}}^{\min})^2 - \frac{1}{k} \|\mathbf{S}\|_F^2 \right)$, which needs the following lemma as a prerequisite.

Lemma B.2 (Theorem 3.1 in (Lewis, 1995) without Convexity). *If a function $f : \mathbb{R}^p \rightarrow \mathbb{R}$ is absolutely symmetric, that is, $\forall \mathbf{x} \in \mathbb{R}^p$ and any \mathbf{y} as a permutation of \mathbf{x} , $f(\mathbf{x}) = f(\mathbf{y})$, then $f \circ \sigma$ is differentiable at matrix $\mathbf{S} \in \mathbb{R}^{p_1 \times p_2}$ if and only if f is differentiable at $\boldsymbol{\sigma} = \boldsymbol{\sigma}(\mathbf{S})$. In this case, for the singular value decomposition $\mathbf{S} = \mathbf{U} \text{Diag}(\boldsymbol{\sigma}) \mathbf{V}^\top$,*

$$\nabla(f \circ \boldsymbol{\sigma})(\mathbf{S}) = \mathbf{U} \text{Diag}(\nabla f(\boldsymbol{\sigma})) \mathbf{V}^\top.$$

Proof. For the forward direction, by Corollary 2.5 in (Lewis, 1995), for $\mathbf{S} = \mathbf{U}\text{Diag}(\boldsymbol{\sigma})\mathbf{V}^\top$,

$$\partial(f \circ \boldsymbol{\sigma})(\mathbf{S}) = \left\{ \mathbf{U}\text{Diag}(\boldsymbol{\mu})\mathbf{V}^\top \mid \boldsymbol{\mu} \in \partial f(\boldsymbol{\sigma}) \right\}.$$

By Theorem 25.1 in (Rockafellar, 1970), since $f \circ \boldsymbol{\sigma}$ is differentiable at matrix $\mathbf{S} \in \mathbb{R}^{p_1 \times p_2}$, we know that its subgradient $\partial(f \circ \boldsymbol{\sigma})(\mathbf{S})$ is a singleton, meaning that $\mathbf{U}\text{Diag}(\boldsymbol{\mu})\mathbf{V}^\top$ is unique, and consequently, $\boldsymbol{\mu} \in \partial f(\boldsymbol{\sigma})$ is unique. As a result, $\partial f(\boldsymbol{\sigma})$ is also a singleton, which, again by Corollary 2.5 in (Lewis, 1995), indicates that f is differentiable at $\boldsymbol{\sigma}$. The reverse direction holds true following a similar argument. \square

Lemma B.3. For $\mathbf{S} = \mathbf{U}\text{Diag}(\boldsymbol{\sigma})\mathbf{V}^\top$, in which $\boldsymbol{\sigma} = [\sigma_1, \dots, \sigma_k]^\top$ such that $\sigma_{\mathbf{S}}^{\max} = \sigma_1 \geq \sigma_2 \geq \dots > \sigma_k = \sigma_{\mathbf{S}}^{\min}$, i.e., $\sigma_k < \sigma_i$ for any $i < k$, $\mathcal{R}'_{i11}(\mathbf{S}) = \frac{1}{2} \left((\sigma_{\mathbf{S}}^{\min})^2 - \frac{1}{k} \|\mathbf{S}\|_F^2 \right)$ is differentiable and for $\mathbf{u}_k, \mathbf{v}_k$ as the k^{th} column vector of \mathbf{U}, \mathbf{V} ,

$$\nabla \mathcal{R}'_{i11}(\mathbf{S}) = \sigma_{\mathbf{S}}^{\min} \mathbf{u}_k \mathbf{v}_k^\top - \frac{1}{k} \mathbf{S}. \quad (20)$$

Proof. For $\mathbf{x} \in \mathbb{R}^k$, denote

$$\mathcal{R}'_{i11,1}(\mathbf{x}) = \min_{i \in [k]} \frac{1}{2} x_i^2, \quad \mathcal{R}'_{i11,2}(\mathbf{x}) = \frac{1}{2k} \sum_{i=1}^k x_i^2.$$

With $\mathcal{R}'_{i11}(\mathbf{S}) = \frac{1}{2} (\sigma_{\mathbf{S}}^{\min})^2 - \frac{1}{2k} \|\mathbf{S}\|_F^2$, we first analyze $\frac{1}{2} (\sigma_{\mathbf{S}}^{\min})^2$. By the subdifferential of piecewise minimum given by Proposition 4.9 in (Mordukhovich, 2018), we have for $\mathbf{x} \in \mathbb{R}^k$,

$$\begin{aligned} \partial_{\mathbf{x}} \mathcal{R}'_{i11,1}(\mathbf{x}) &\subset \left\{ \partial_{\mathbf{x}} \left(\frac{1}{2} x_i^2 \right) \mid i \in \arg \min_{j \in [k]} \frac{1}{2} x_j^2 \right\} \\ &= \left\{ x_i \mathbf{e}_i \mid i \in \arg \min_{j \in [k]} \frac{1}{2} x_j^2 \right\} \\ &= \left\{ x_i \mathbf{e}_i \mid i \in \arg \min_{j \in [k]} |x_j| \right\} \end{aligned}$$

in which \mathbf{e}_i is the i^{th} vector from the k -dimensional standard basis. Therefore,

$$\partial_{\boldsymbol{\sigma}} \mathcal{R}'_{i11,1}(\boldsymbol{\sigma}) \subset \left\{ \sigma_i \mathbf{e}_i \mid i \in \arg \min_{j \in [k]} \sigma_j \right\}$$

Since for any $i < k$, $\sigma_k < \sigma_i$, i.e., the minimum non-zero singular value $\sigma_{\mathbf{S}}^{\min}$ is unique, we know that the subdifferential $\left\{ \sigma_i \mathbf{e}_i \mid i \in \arg \min_{j \in [k]} \sigma_j \right\} = \{\sigma_{\mathbf{S}}^{\min} \mathbf{e}_k\}$ is a singleton. Therefore, by Theorem 25.1 in (Rockafellar, 1970), we know $\mathcal{R}'_{i11,1}$ is differentiable with respect to $\boldsymbol{\sigma}$ and $\nabla_{\boldsymbol{\sigma}} \mathcal{R}'_{i11,1}(\boldsymbol{\sigma}) = \sigma_{\mathbf{S}}^{\min} \mathbf{e}_k$. Regarding $\boldsymbol{\sigma} = \boldsymbol{\sigma}(\mathbf{S})$ as a function of \mathbf{S} in which $\boldsymbol{\sigma}(\cdot)$ represents taking the singular values of a matrix, we have by Corollary 2.5 in (Lewis, 1995)

$$\begin{aligned} \partial_{\mathbf{S}} \left(\frac{1}{2} (\sigma_{\mathbf{S}}^{\min})^2 \right) &= \partial_{\mathbf{S}} (\mathcal{R}'_{i11,1} \circ \boldsymbol{\sigma})(\mathbf{S}) \\ &= \left\{ \mathbf{U}\text{Diag}(\boldsymbol{\mu})\mathbf{V}^\top \mid \boldsymbol{\mu} \in \partial_{\boldsymbol{\sigma}} \mathcal{R}'_{i11,1}(\boldsymbol{\sigma}) \right\} \end{aligned}$$

Given that $\mathcal{R}'_{i11,1}$ is differentiable and apparently also absolutely symmetric with respect to $\boldsymbol{\sigma}$, by Lemma B.2, we know $\frac{1}{2} (\sigma_{\mathbf{S}}^{\min})^2$ is also differentiable and

$$\begin{aligned} \nabla \left(\frac{1}{2} (\sigma_{\mathbf{S}}^{\min})^2 \right) &= \mathbf{U}\text{Diag}(\nabla_{\boldsymbol{\sigma}} \mathcal{R}'_{i11,1}(\boldsymbol{\sigma}))\mathbf{V}^\top \\ &= \mathbf{U}\text{Diag}(\sigma_{\mathbf{S}}^{\min} \mathbf{e}_p)\mathbf{V}^\top \\ &= \sigma_{\mathbf{S}}^{\min} \mathbf{u}_k \mathbf{v}_k^\top. \end{aligned}$$

In addition, we have

$$\begin{aligned}\partial_{\mathbf{S}} \left(\frac{1}{2k} \|\mathbf{S}\|_F^2 \right) &= \partial_{\mathbf{S}} \left(\frac{1}{2k} \sum_{i=1}^k \sigma(\mathbf{S})^2 \right) \\ &= \left\{ \mathbf{U} \text{Diag}(\boldsymbol{\mu}) \mathbf{V}^\top \mid \boldsymbol{\mu} \in \partial_{\boldsymbol{\sigma}} \mathcal{R}'_{i11,2}(\boldsymbol{\sigma}) \right\}\end{aligned}$$

by Corollary 2.5 in (Lewis, 1995). $\mathcal{R}'_{i11,2}$ is apparently differentiable with $\nabla \mathcal{R}'_{i11,2}(\mathbf{x}) = \frac{1}{k} \mathbf{x}$. Therefore, again by Lemma B.2,

$$\begin{aligned}\nabla \left(\frac{1}{2k} \|\mathbf{S}\|_F^2 \right) &= \mathbf{U} \text{Diag}(\nabla \mathcal{R}'_{i11,2}(\boldsymbol{\sigma}_S)) \mathbf{V}^\top \\ &= \frac{1}{k} \mathbf{U} \text{Diag}(\boldsymbol{\sigma}_S) \mathbf{V}^\top \\ &= \frac{1}{k} \mathbf{S}.\end{aligned}$$

By the linearity of gradients,

$$\begin{aligned}\nabla \mathcal{R}'_{i11}(\mathbf{S}) &= \nabla \left(\frac{1}{2} (\sigma_{\mathbf{S}^{\min}}^{\min})^2 \right) - \nabla \left(\frac{1}{2k} \|\mathbf{S}\|_F^2 \right) \\ &= \sigma_{\mathbf{S}^{\min}}^{\min} \mathbf{u}_k \mathbf{v}_k^\top - \frac{1}{k} \mathbf{S},\end{aligned}$$

which completes the proof. \square

Theorem 4.1. (3) If $\sigma_{\mathbf{S}^{\min}}^{\min} = \sigma_k < \sigma_i$ for any $i < k$, then $\mathcal{R}_{i11}(\mathbf{S})$ is differentiable and $\nabla_{\mathbf{S}} \mathcal{R}_{i11}(\mathbf{S}) = \frac{\sigma_k \mathbf{u}_k \mathbf{v}_k^\top - \frac{1}{k} \mathbf{S}}{\left(\frac{1}{2k} \|\mathbf{S}\|_F^2 - \frac{1}{2} (\sigma_{\mathbf{S}^{\min}}^{\min})^2 \right)^2}$.

Proof. Since $\mathcal{R}_{i11}(\mathbf{S}) = \frac{1}{\frac{1}{2k} \|\mathbf{S}\|_F^2 - \frac{1}{2} (\sigma_{\mathbf{S}^{\min}}^{\min})^2}$, we have

$$\begin{aligned}\partial \mathcal{R}_{i11}(\mathbf{S}) &= \frac{-\partial \left(\frac{1}{2k} \|\mathbf{S}\|_F^2 - \frac{1}{2} (\sigma_{\mathbf{S}^{\min}}^{\min})^2 \right)}{\left(\frac{1}{2k} \|\mathbf{S}\|_F^2 - \frac{1}{2} (\sigma_{\mathbf{S}^{\min}}^{\min})^2 \right)^2} \\ &= \frac{\partial \left(\frac{1}{2} (\sigma_{\mathbf{S}^{\min}}^{\min})^2 - \frac{1}{2k} \|\mathbf{S}\|_F^2 \right)}{\left(\frac{1}{2k} \|\mathbf{S}\|_F^2 - \frac{1}{2} (\sigma_{\mathbf{S}^{\min}}^{\min})^2 \right)^2} \\ &= \frac{\partial \mathcal{R}'_{i11}(\mathbf{S})}{\left(\frac{1}{2k} \|\mathbf{S}\|_F^2 - \frac{1}{2} (\sigma_{\mathbf{S}^{\min}}^{\min})^2 \right)^2}\end{aligned}$$

By Lemma B.3, we know that if $\sigma_{\mathbf{S}^{\min}}^{\min} = \sigma_k < \sigma_i$ for any $i < k$, $\mathcal{R}'_{i11}(\mathbf{S})$ is differentiable and $\nabla \mathcal{R}'_{i11}(\mathbf{S}) = \sigma_{\mathbf{S}^{\min}}^{\min} \mathbf{u}_k \mathbf{v}_k^\top - \frac{1}{k} \mathbf{S}$. Consequently, $\mathcal{R}_{i11}(\mathbf{S})$ is differentiable and

$$\nabla \mathcal{R}_{i11}(\mathbf{S}) = \frac{\sigma_{\mathbf{S}^{\min}}^{\min} \mathbf{u}_k \mathbf{v}_k^\top - \frac{1}{k} \mathbf{S}}{\left(\frac{1}{2k} \|\mathbf{S}\|_F^2 - \frac{1}{2} (\sigma_{\mathbf{S}^{\min}}^{\min})^2 \right)^2}.$$

\square

B.2.4. PROOF OF THEOREM 4.1 (4)

Theorem 4.1. (4) If $\sigma_{\mathbf{S}^{\min}}^{\min}$ is unique, update \mathbf{S} with $\nabla_{\mathbf{S}} \mathcal{R}_{i11}(\mathbf{S})$ such that $\mathbf{S}' = \mathbf{S} - \eta_2 \nabla_{\mathbf{S}} \mathcal{R}_{i11}(\mathbf{S})$ for $0 < \eta_2 < \frac{k}{k-1} \left(\frac{1}{2k} \|\mathbf{S}\|_F^2 - \frac{1}{2} (\sigma_{\mathbf{S}^{\min}}^{\min})^2 \right)^2$, then $\kappa(\mathbf{S}') > \kappa(\mathbf{S})$.

Proof. Given that $\mathbf{S}' = \mathbf{S} - \eta_2 \nabla \mathcal{R}_{i11}(\mathbf{S})$ and that $\nabla \mathcal{R}_{i11}(\mathbf{S}) = \frac{\sigma_{\mathbf{S}}^{\min} \mathbf{u}_k \mathbf{v}_k^\top - \frac{1}{k} \mathbf{S}}{\left(\frac{1}{2k} \|\mathbf{S}\|_F^2 - \frac{1}{2} (\sigma_{\mathbf{S}}^{\min})^2\right)^2} = \frac{1}{\mathcal{R}'_{i11}(\mathbf{S})^2} \left(\sigma_k \mathbf{u}_k \mathbf{v}_k^\top - \frac{1}{k} \mathbf{S}\right)$ for $\mathcal{R}'_{i11}(\mathbf{S}) = \frac{1}{2} \left(\sigma_k^2 - \frac{1}{k} \|\mathbf{S}\|_F^2\right)$,

$$\begin{aligned} \mathbf{S}' &= \mathbf{S} - \eta_2 \nabla \mathcal{R}_{i11}(\mathbf{S}) \\ &= \mathbf{S} - \frac{\eta_2}{\mathcal{R}'_{i11}(\mathbf{S})^2} \left(\sigma_k \mathbf{u}_k \mathbf{v}_k^\top - \frac{1}{k} \mathbf{S}\right) \\ &= \left(1 + \frac{\eta_2}{k \mathcal{R}'_{i11}(\mathbf{S})^2}\right) \mathbf{S} - \frac{\eta_2}{\mathcal{R}'_{i11}(\mathbf{S})^2} \sigma_k \mathbf{u}_k \mathbf{v}_k^\top \\ &= \left(1 + \frac{\eta_2}{k \mathcal{R}'_{i11}(\mathbf{S})^2}\right) \sum_{i=1}^k \sigma_i \mathbf{u}_i \mathbf{v}_i^\top - \frac{\eta_2}{\mathcal{R}'_{i11}(\mathbf{S})^2} \sigma_k \mathbf{u}_k \mathbf{v}_k^\top \\ &= \left(1 + \frac{\eta_2}{k \mathcal{R}'_{i11}(\mathbf{S})^2}\right) \sum_{i=1}^{k-1} \sigma_i \mathbf{u}_i \mathbf{v}_i^\top + \left(1 + \frac{\eta_2}{k \mathcal{R}'_{i11}(\mathbf{S})^2} - \frac{\eta_2}{\mathcal{R}'_{i11}(\mathbf{S})^2}\right) \sigma_k \mathbf{u}_k \mathbf{v}_k^\top \\ &= \mathbf{U} \text{Diag}(\boldsymbol{\sigma}_{\mathbf{S}'}) \mathbf{V}^\top. \end{aligned}$$

where $\boldsymbol{\sigma}_{\mathbf{S}'} = \left[\left(1 + \frac{\eta_2}{k \mathcal{R}'_{i11}(\mathbf{S})^2}\right) \sigma_1, \dots, \left(1 + \frac{\eta_2}{k \mathcal{R}'_{i11}(\mathbf{S})^2}\right) \sigma_{k-1}, \left(1 + \frac{\eta_2}{k \mathcal{R}'_{i11}(\mathbf{S})^2} - \frac{\eta_2}{\mathcal{R}'_{i11}(\mathbf{S})^2}\right) \sigma_k\right]^\top$ is the vector formed by the singular values of \mathbf{S}' but not necessarily in the decreasing order.

Now we argue that $\left(1 + \frac{\eta_2}{k \mathcal{R}'_{i11}(\mathbf{S})^2}\right) \sigma_1$ remains to be the maximum singular value while $\left(1 + \frac{\eta_2}{k \mathcal{R}'_{i11}(\mathbf{S})^2} - \frac{\eta_2}{\mathcal{R}'_{i11}(\mathbf{S})^2}\right) \sigma_k$ the minimum. Since $\sigma_k < \sigma_i$ for any $i < k$, i.e., $\sigma_{\mathbf{S}}^{\min} = \sigma_k$ is unique, we must have $0 < \beta < 1$ such that $\sigma_k = \beta \sigma_{k-1}$. Also, given that $\eta_2 < \frac{k}{k-1} \left(\frac{1}{2k} \|\mathbf{S}\|_F^2 - \frac{1}{2} \sigma_k^2\right)^2 = \frac{k \mathcal{R}'_{i11}(\mathbf{S})^2}{k-1}$, we have $1 + \frac{\eta_2}{k \mathcal{R}'_{i11}(\mathbf{S})^2} - \frac{\eta_2}{\mathcal{R}'_{i11}(\mathbf{S})^2} > 0$. Therefore,

$$\begin{aligned} &\left(1 + \frac{\eta_2}{k \mathcal{R}'_{i11}(\mathbf{S})^2} - \frac{\eta_2}{\mathcal{R}'_{i11}(\mathbf{S})^2}\right) \sigma_k \\ &= \left(1 + \frac{\eta_2}{k \mathcal{R}'_{i11}(\mathbf{S})^2}\right) \frac{1 + \frac{\eta_2}{k \mathcal{R}'_{i11}(\mathbf{S})^2} - \frac{\eta_2}{\mathcal{R}'_{i11}(\mathbf{S})^2}}{1 + \frac{\eta_2}{k \mathcal{R}'_{i11}(\mathbf{S})^2}} \sigma_k \\ &= \left(1 + \frac{\eta_2}{k \mathcal{R}'_{i11}(\mathbf{S})^2}\right) \left(1 - \frac{\frac{\eta_2}{\mathcal{R}'_{i11}(\mathbf{S})^2}}{1 + \frac{\eta_2}{k \mathcal{R}'_{i11}(\mathbf{S})^2}}\right) \sigma_k \\ &< \left(1 + \frac{\eta_2}{k \mathcal{R}'_{i11}(\mathbf{S})^2}\right) \sigma_k \\ &< \left(1 + \frac{\eta_2}{k \mathcal{R}'_{i11}(\mathbf{S})^2}\right) \frac{1}{\beta} \sigma_k \\ &= \left(1 + \frac{\eta_2}{k \mathcal{R}'_{i11}(\mathbf{S})^2}\right) \sigma_{k-1}. \end{aligned}$$

Since $\sigma_1 \geq \sigma_2(\mathbf{S}) \geq \dots \geq \sigma_{k-1} > \sigma_k$ and $1 + \frac{\eta_2}{k \mathcal{R}'_{i11}(\mathbf{S})^2} > 0$, we know that $\sigma_{\mathbf{S}'}^{\max} = \left(1 + \frac{\eta_2}{k \mathcal{R}'_{i11}(\mathbf{S})^2}\right) \sigma_1$ and $\sigma_{\mathbf{S}'}^{\min} = \left(1 + \frac{\eta_2}{k \mathcal{R}'_{i11}(\mathbf{S})^2} - \frac{\eta_2}{\mathcal{R}'_{i11}(\mathbf{S})^2}\right) \sigma_k$. Finally,

$$\begin{aligned} \kappa(\mathbf{S}') &= \frac{\sigma_{\mathbf{S}'}^{\max}}{\sigma_{\mathbf{S}'}^{\min}} \\ &= \frac{\left(1 + \frac{\eta_2}{k \mathcal{R}'_{i11}(\mathbf{S})^2}\right) \sigma_1}{\left(1 + \frac{\eta_2}{k \mathcal{R}'_{i11}(\mathbf{S})^2} - \frac{\eta_2}{\mathcal{R}'_{i11}(\mathbf{S})^2}\right) \sigma_k} \\ &= \frac{1 + \frac{\eta_2}{k \mathcal{R}'_{i11}(\mathbf{S})^2}}{1 + \frac{\eta_2}{k \mathcal{R}'_{i11}(\mathbf{S})^2} - \frac{\eta_2}{\mathcal{R}'_{i11}(\mathbf{S})^2}} \kappa(\mathbf{S}) \\ &> \kappa(\mathbf{S}). \end{aligned}$$

□

B.3. Proof of Theorem 4.2

Theorem 4.2. For $\mathbf{H}(\theta) = \theta^\top \mathbf{K}\theta$, if its maximum and minimum singular values σ_1 and σ_k are unique, then

$$(1) \nabla_\theta \mathcal{R}_{\text{well}}(\mathbf{H}(\theta)) = 2\mathbf{K}\theta \left(\sigma_1 \mathbf{v}_1 \mathbf{v}_1^\top - \frac{1}{D_{\text{in}}} \theta^\top \mathbf{K}\theta \right),$$

$$(2) \nabla_\theta \mathcal{R}_{\text{ill}}(\mathbf{H}(\theta)) = \frac{2\mathbf{K}\theta(\sigma_k \mathbf{u}_k \mathbf{v}_k^\top - \frac{1}{k} \theta^\top \mathbf{K}\theta)}{\left(\frac{1}{2k} \|\theta^\top \mathbf{K}\theta\|_F^2 - \frac{1}{2} \sigma_k^2 \right)^2}.$$

Proof. Given that $\mathbf{H}(\theta) = \theta^\top \mathbf{K}\theta$ for $\mathbf{K} = \mathbf{X}^\top \mathbf{X}$, we know $\mathbf{H}(\theta)$ is symmetric and positive semidefinite. Therefore, for compact SVD $\mathbf{H}(\theta) = \mathbf{U} \text{Diag}(\boldsymbol{\sigma}) \mathbf{V}^\top$, we have $\mathbf{U} = \mathbf{V}$.

(1) When the maximum singular value σ_1 of \mathbf{H} is unique, we know from Theorem A.1 (3) that $\mathcal{R}_{\text{well}}(\mathbf{H}(\theta))$ is differentiable with respect to \mathbf{H} , and $\nabla_{\mathbf{H}} \mathcal{R}_{\text{well}}(\mathbf{H}(\theta)) = \sigma_1 \mathbf{u}_1 \mathbf{v}_1^\top - \frac{1}{D_{\text{in}}} \mathbf{H} = \sigma_1 \mathbf{v}_1 \mathbf{v}_1^\top - \frac{1}{D_{\text{in}}} \mathbf{H}$.

Given the form $\mathbf{H}(\theta) = \theta^\top \mathbf{K}\theta$, we have $d\mathbf{H} = (d\theta)^\top \mathbf{K}\theta + \theta^\top \mathbf{K}(d\theta)$. Furthermore,

$$\begin{aligned} (d\mathcal{R}_{\text{well}})(\mathbf{H}(\theta)) &= \langle \nabla_{\mathbf{H}} \mathcal{R}_{\text{well}}(\mathbf{H}(\theta)), d\mathbf{H} \rangle_F \\ &= \text{Tr} \left(\left(\sigma_1 \mathbf{v}_1 \mathbf{v}_1^\top - \frac{1}{D_{\text{in}}} \mathbf{H} \right)^\top \left((d\theta)^\top \mathbf{K}\theta + \theta^\top \mathbf{K}(d\theta) \right) \right) \\ &= \text{Tr} \left(\left(\sigma_1 \mathbf{v}_1 \mathbf{v}_1^\top - \frac{1}{D_{\text{in}}} \mathbf{H} \right)^\top (d\theta)^\top \mathbf{K}\theta \right) + \text{Tr} \left(\left(\sigma_1 \mathbf{v}_1 \mathbf{v}_1^\top - \frac{1}{D_{\text{in}}} \mathbf{H} \right)^\top \theta^\top \mathbf{K}(d\theta) \right) \\ &= \text{Tr} \left(\mathbf{K}\theta \left(\sigma_1 \mathbf{v}_1 \mathbf{v}_1^\top - \frac{1}{D_{\text{in}}} \mathbf{H} \right)^\top (d\theta)^\top \right) + \text{Tr} \left((d\theta) \left(\sigma_1 \mathbf{v}_1 \mathbf{v}_1^\top - \frac{1}{D_{\text{in}}} \mathbf{H} \right)^\top \theta^\top \mathbf{K} \right), \end{aligned}$$

in which $\langle \cdot, \cdot \rangle_F$ denotes the Frobenius inner product, and that last equality follows from the cyclic property of trace. As a result, following the derivatives of traces as in Eq. (100) and Eq. (104) in Petersen et al. (2008),

$$\begin{aligned} \nabla_\theta \mathcal{R}_{\text{well}}(\mathbf{H}(\theta)) &= \frac{\partial \mathcal{R}_{\text{well}}(\mathbf{H}(\theta))}{\partial \theta} \\ &= \mathbf{K}\theta \left(\sigma_1 \mathbf{v}_1 \mathbf{v}_1^\top - \frac{1}{D_{\text{in}}} \mathbf{H} \right)^\top + \left(\left(\sigma_1 \mathbf{v}_1 \mathbf{v}_1^\top - \frac{1}{D_{\text{in}}} \mathbf{H} \right)^\top \theta^\top \mathbf{K} \right)^\top \\ &= \mathbf{K}\theta \left(\sigma_1 (\mathbf{v}_1 \mathbf{v}_1^\top)^\top - \frac{1}{D_{\text{in}}} \mathbf{H}^\top \right) + \mathbf{K}^\top \theta \left(\sigma_1 \mathbf{v}_1 \mathbf{v}_1^\top - \frac{1}{D_{\text{in}}} \mathbf{H} \right) \\ &= 2\mathbf{K}\theta \left(\sigma_1 \mathbf{v}_1 \mathbf{v}_1^\top - \frac{1}{D_{\text{in}}} \mathbf{H} \right). \end{aligned}$$

(2) When the minimum singular value σ_k of \mathbf{H} is unique, we know from Theorem 4.1 (3) that $\mathcal{R}_{\text{ill}}(\mathbf{H}(\theta))$ is differentiable with respect to \mathbf{H} , and $\nabla_{\mathbf{H}} \mathcal{R}_{\text{ill}}(\mathbf{H}(\theta)) = \frac{\sigma_k \mathbf{u}_k \mathbf{v}_k^\top - \frac{1}{k} \mathbf{H}}{\left(\frac{1}{2k} \|\mathbf{H}\|_F^2 - \frac{1}{2} \sigma_k^2 \right)^2}$. Following similar arguments as in (1), we have

$$\nabla_\theta \mathcal{R}_{\text{ill}}(\mathbf{H}(\theta)) = \frac{2\mathbf{K}\theta(\sigma_k \mathbf{u}_k \mathbf{v}_k^\top - \frac{1}{k} \mathbf{H})}{\left(\frac{1}{2k} \|\mathbf{H}\|_F^2 - \frac{1}{2} \sigma_k^2 \right)^2}.$$

□

B.4. Proof of Theorem 4.3

Theorem 4.3. For the trainable feature extractor θ , feature covariance $\mathbf{H}_P(\theta) = \theta^\top \mathbf{K}_P \theta$ of the primary task and $\mathbf{H}_H(\theta) = \theta^\top \mathbf{K}_H \theta$ of the immunization task with $\text{rank}(\mathbf{H}_P) = k_P$, $\text{rank}(\mathbf{H}_H) = k_H$ and compact SVD $\mathbf{H}_P(\theta) = \mathbf{U}_P \text{Diag}(\boldsymbol{\sigma}_P) \mathbf{V}_P^\top$, $\mathbf{H}_H(\theta) = \mathbf{U}_H \text{Diag}(\boldsymbol{\sigma}_H) \mathbf{V}_H^\top$, for $\boldsymbol{\sigma}_P = [\sigma_{P,1}, \dots, \sigma_{P,k_P}]$, $\boldsymbol{\sigma}_H = [\sigma_{H,1}, \dots, \sigma_{H,k_H}]$,

(1) if $\sigma_{H_P}^{\max}$ is unique, i.e., $\sigma_{H_P}^{\max} = \sigma_{P,1} > \sigma_{P,2}$, update θ such that $\theta' = \theta - \eta_P \mathbf{K}_P^{-1} \nabla_\theta \mathcal{R}_{\text{well}}(\mathbf{H}_P(\theta))$ for $0 < \eta_P <$

$$\min \left\{ \frac{1}{\left(1 - \frac{1}{D_{\text{in}}}\right) \sigma_{P,1}}, \frac{\sqrt{\sigma_{P,1} \sigma_{P,2} - \sigma_{P,2}^2}}{\frac{2}{D_{\text{in}}} \sigma_{P,2}^2} \right\}, \text{ then } \kappa(\theta'^\top \mathbf{K}_P \theta') < \kappa(\theta^\top \mathbf{K}_P \theta),$$

(2) if $\sigma_{\mathbf{H}_H}^{\min}$ is unique, i.e., $\sigma_{\mathbf{H}_H}^{\min} = \sigma_{H, k_H} < \sigma_{H, k_H-1}$, update θ such that $\theta' = \theta - \eta_H \mathbf{K}_H^{-1} \nabla_{\theta} \mathcal{R}_{\text{ill}}(\mathbf{H}_H(\theta))$ for $0 < \eta_H < \frac{1}{1-2\sigma_{\mathbf{H}_H}^{\min}/k_H} \left(\frac{1}{2k_H} \|\theta^{\top} \mathbf{K}_H \theta\|_F^2 - \frac{1}{2} (\sigma_{\mathbf{H}_H}^{\min})^2 \right)^2$, then $\kappa(\theta'^{\top} \mathbf{K}_H \theta') > \kappa(\theta^{\top} \mathbf{K}_H \theta)$.

Proof. (1) By Theorem 4.2 (1), we know $\nabla_{\theta} \mathcal{R}_{\text{well}}(\mathbf{H}_P(\theta)) = 2\mathbf{K}_P \theta \left(\sigma_{P,1} \mathbf{v}_{P,1} \mathbf{v}_{P,1}^{\top} - \frac{1}{D_{\text{in}}} \mathbf{H}_P \right)$. Since $\theta' = \theta - \eta_P \mathbf{K}_P^{-1} \nabla_{\theta} \mathcal{R}_{\text{well}}(\mathbf{H}_P(\theta))$, we have

$$\begin{aligned}
 \theta'^{\top} \mathbf{K}_P \theta' &= (\theta - \eta_P \mathbf{K}_P^{-1} \nabla_{\theta} \mathcal{R}_{\text{well}}(\mathbf{H}_P(\theta)))^{\top} \mathbf{K}_P (\theta - \eta_P \mathbf{K}_P^{-1} \nabla_{\theta} \mathcal{R}_{\text{well}}(\mathbf{H}_P(\theta))) \\
 &= \left(\theta - 2\eta_P \mathbf{K}_P^{-1} \mathbf{K}_P \theta \left(\sigma_{P,1} \mathbf{v}_{P,1} \mathbf{v}_{P,1}^{\top} - \frac{1}{D_{\text{in}}} \mathbf{H}_P \right) \right)^{\top} \mathbf{K}_P \left(\theta - 2\eta_P \mathbf{K}_P^{-1} \mathbf{K}_P \theta \left(\sigma_{P,1} \mathbf{v}_{P,1} \mathbf{v}_{P,1}^{\top} - \frac{1}{D_{\text{in}}} \mathbf{H}_P \right) \right) \\
 &= \left(\theta - 2\eta_P \theta \left(\sigma_{P,1} \mathbf{v}_{P,1} \mathbf{v}_{P,1}^{\top} - \frac{1}{D_{\text{in}}} \mathbf{H}_P \right) \right)^{\top} \mathbf{K}_P \left(\theta - 2\eta_P \theta \left(\sigma_{P,1} \mathbf{v}_{P,1} \mathbf{v}_{P,1}^{\top} - \frac{1}{D_{\text{in}}} \mathbf{H}_P \right) \right) \\
 &= \theta^{\top} \mathbf{K}_P \theta - 2\eta_P \left(\sigma_{P,1} \mathbf{v}_{P,1} \mathbf{v}_{P,1}^{\top} - \frac{1}{D_{\text{in}}} \mathbf{H}_P \right)^{\top} \theta^{\top} \mathbf{K}_P \theta - 2\eta_P \theta^{\top} \mathbf{K}_P \theta \left(\sigma_{P,1} \mathbf{v}_{P,1} \mathbf{v}_{P,1} - \frac{1}{D_{\text{in}}} \mathbf{H}_P \right) \\
 &\quad + 4\eta_P^2 \left(\sigma_{P,1} \mathbf{v}_{P,1} \mathbf{v}_{P,1}^{\top} - \frac{1}{D_{\text{in}}} \mathbf{H}_P \right)^{\top} \theta^{\top} \mathbf{K}_P \theta \left(\sigma_{P,1} \mathbf{v}_{P,1} \mathbf{v}_{P,1} - \frac{1}{D_{\text{in}}} \mathbf{H}_P \right) \\
 &= \mathbf{H}_P - 2\eta_P \left(\sigma_{P,1} \mathbf{v}_{P,1} \mathbf{v}_{P,1}^{\top} - \frac{1}{D_{\text{in}}} \mathbf{H}_P \right)^{\top} \mathbf{H}_P - 2\eta_P \mathbf{H}_P \left(\sigma_{P,1} \mathbf{v}_{P,1} \mathbf{v}_{P,1} - \frac{1}{D_{\text{in}}} \mathbf{H}_P \right) \\
 &\quad + 4\eta_P^2 \left(\sigma_{P,1} \mathbf{v}_{P,1} \mathbf{v}_{P,1}^{\top} - \frac{1}{D_{\text{in}}} \mathbf{H}_P \right)^{\top} \mathbf{H}_P \left(\sigma_{P,1} \mathbf{v}_{P,1} \mathbf{v}_{P,1} - \frac{1}{D_{\text{in}}} \mathbf{H}_P \right).
 \end{aligned}$$

Since $\mathbf{H}_P(\theta) = \theta^{\top} \mathbf{K}_P \theta$ for $\mathbf{K}_P = \mathbf{X}_P^{\top} \mathbf{X}_P$ is symmetric and positive semidefinite, we know for $\mathbf{H}_P(\theta) = \mathbf{U}_P \text{Diag}(\boldsymbol{\sigma}_P) \mathbf{V}_P^{\top}$, it holds that $\mathbf{U}_P = \mathbf{V}_P$. Furthermore,

$$\begin{aligned}
 \sigma_{P,1} \mathbf{v}_{P,1} \mathbf{v}_{P,1}^{\top} - \frac{1}{D_{\text{in}}} \mathbf{H}_P &= \sigma_{P,1} \mathbf{v}_{P,1} \mathbf{v}_{P,1}^{\top} - \frac{1}{D_{\text{in}}} \sum_{i=1}^{k_P} \sigma_{P,i} \mathbf{u}_{P,i} \mathbf{v}_{P,i}^{\top} \\
 &= \left(1 - \frac{1}{D_{\text{in}}} \right) \sigma_{P,1} \mathbf{u}_{P,1} \mathbf{v}_{P,1}^{\top} - \frac{1}{D_{\text{in}}} \sum_{i=2}^{k_P} \sigma_{P,i} \mathbf{u}_{P,i} \mathbf{v}_{P,i}^{\top} \\
 &= \mathbf{U}_P \text{Diag}(\tilde{\boldsymbol{\sigma}}_P) \mathbf{V}_P^{\top} \\
 &= \mathbf{V}_P \text{Diag}(\tilde{\boldsymbol{\sigma}}_P) \mathbf{V}_P^{\top}
 \end{aligned}$$

for $\text{Diag}(\tilde{\boldsymbol{\sigma}}_P) = \left[\left(1 - \frac{1}{D_{\text{in}}} \right) \sigma_{P,1}, -\frac{1}{D_{\text{in}}} \sigma_{P,2}, \dots, -\frac{1}{D_{\text{in}}} \sigma_{P,k_P} \right]$. Therefore, plugging this and the SVD of \mathbf{H}_P back in,

$$\begin{aligned}
 \theta'^{\top} \mathbf{K}_P \theta' &= \mathbf{V}_P \text{Diag}(\boldsymbol{\sigma}_P) \mathbf{V}_P^{\top} - 2\eta_P (\mathbf{V}_P \text{Diag}(\tilde{\boldsymbol{\sigma}}_P) \mathbf{V}_P^{\top})^{\top} \mathbf{V}_P \text{Diag}(\boldsymbol{\sigma}_P) \mathbf{V}_P^{\top} \\
 &\quad - 2\eta_P (\mathbf{V}_P \text{Diag}(\boldsymbol{\sigma}_P) \mathbf{V}_P^{\top})^{\top} \mathbf{V}_P \text{Diag}(\tilde{\boldsymbol{\sigma}}_P) \mathbf{V}_P^{\top} \\
 &\quad + 4\eta_P^2 (\mathbf{V}_P \text{Diag}(\tilde{\boldsymbol{\sigma}}_P) \mathbf{V}_P^{\top})^{\top} \mathbf{V}_P \text{Diag}(\boldsymbol{\sigma}_P) \mathbf{V}_P^{\top} \mathbf{V}_P \text{Diag}(\tilde{\boldsymbol{\sigma}}_P) \mathbf{V}_P^{\top} \\
 &= \mathbf{V}_P \text{Diag}(\boldsymbol{\sigma}_P) \mathbf{V}_P^{\top} - 2\eta_P \mathbf{V}_P \text{Diag}(\tilde{\boldsymbol{\sigma}}_P) \text{Diag}(\boldsymbol{\sigma}_P) \mathbf{V}_P^{\top} - 2\eta_P \mathbf{V}_P \text{Diag}(\boldsymbol{\sigma}_P) \text{Diag}(\tilde{\boldsymbol{\sigma}}_P) \mathbf{V}_P^{\top} \\
 &\quad + 4\eta_P^2 \mathbf{V}_P \text{Diag}(\tilde{\boldsymbol{\sigma}}_P) \text{Diag}(\boldsymbol{\sigma}_P) \text{Diag}(\tilde{\boldsymbol{\sigma}}_P) \mathbf{V}_P^{\top} \\
 &= \mathbf{V}_P \text{Diag}(\boldsymbol{\sigma}_P) \mathbf{V}_P^{\top} - 2\eta_P \mathbf{V}_P \text{Diag}(\tilde{\boldsymbol{\sigma}}_P \odot \boldsymbol{\sigma}_P) \mathbf{V}_P^{\top} - 2\eta_P \mathbf{V}_P \text{Diag}(\boldsymbol{\sigma}_P \odot \tilde{\boldsymbol{\sigma}}_P) \mathbf{V}_P^{\top} \\
 &\quad + 4\eta_P^2 \mathbf{V}_P \text{Diag}(\tilde{\boldsymbol{\sigma}}_P \odot \boldsymbol{\sigma}_P \odot \tilde{\boldsymbol{\sigma}}_P) \mathbf{V}_P^{\top} \\
 &= \mathbf{V}_P \text{Diag}(\boldsymbol{\sigma}_P - 4\eta_P \tilde{\boldsymbol{\sigma}}_P \odot \boldsymbol{\sigma}_P + 4\eta_P^2 \tilde{\boldsymbol{\sigma}}_P \odot \boldsymbol{\sigma}_P \odot \tilde{\boldsymbol{\sigma}}_P) \mathbf{V}_P^{\top} \\
 &= \mathbf{V}_P \text{Diag}(\boldsymbol{\sigma}'_P) \mathbf{V}_P^{\top},
 \end{aligned}$$

in which $\boldsymbol{\sigma}'_P = [\sigma'_{P,1}, \dots, \sigma'_{P,k_P}]^{\top}$ for $\sigma'_{P,i} = \begin{cases} \sigma_{P,1} - 4\eta_P \left(1 - \frac{1}{D_{\text{in}}} \right) \sigma_{P,1}^2 + 4\eta_P^2 \left(1 - \frac{1}{D_{\text{in}}} \right)^2 \sigma_{P,1}^3 & \text{if } i = 1 \\ \sigma_{P,i} + \frac{4\eta_P}{D_{\text{in}}} \sigma_{P,i}^2 + \frac{4\eta_P^2}{D_{\text{in}}^2} \sigma_{P,i}^3 & \text{if } i > 1 \end{cases}$, \odot denotes element-wise product and the second equality holds by the fact that \mathbf{V}_P is orthonormal, i.e., $\mathbf{V}_P^{\top} \mathbf{V}_P = \mathbf{I}$.

Since $\sigma_{H_p}^{\max}$ is unique, we know that $\exists \alpha > 1$ such that $\sigma_{P,1} = \alpha \sigma_{P,2}$. Therefore,

$$\begin{aligned}\sigma'_{P,2} &= \sigma_{P,2} + \frac{4\eta_P}{D_{\text{in}}} \sigma_{P,2}^2 + \frac{4\eta_P^2}{D_{\text{in}}^2} \sigma_{P,2}^3 \\ &= \left(1 + \frac{4\eta_P}{D_{\text{in}}} \sigma_{P,2} + \frac{4\eta_P^2}{D_{\text{in}}^2} \sigma_{P,2}^2\right) \sigma_{P,2} \\ &= \frac{1 + \frac{4\eta_P}{D_{\text{in}}} \sigma_{P,2} + \frac{4\eta_P^2}{D_{\text{in}}^2} \sigma_{P,2}^2}{\alpha} \sigma_{P,1}.\end{aligned}$$

With $\eta_P < \frac{\sqrt{\sigma_{P,1}\sigma_{P,2}} - \sigma_{P,2}}{\frac{2}{D_{\text{in}}}\sigma_{P,2}^2}$, we have $1 + \frac{4\eta_P}{D_{\text{in}}}\sigma_{P,2} + \frac{4\eta_P^2}{D_{\text{in}}^2}\sigma_{P,2}^2 < 1 - 4\eta_P\left(1 - \frac{1}{D_{\text{in}}}\right)\sigma_{P,1} + 4\eta_P^2\left(1 - \frac{1}{D_{\text{in}}}\right)^2\sigma_{P,1}^2$. As a result,

$$\frac{1 + \frac{4\eta_P}{D_{\text{in}}}\sigma_{P,2} + \frac{4\eta_P^2}{D_{\text{in}}^2}\sigma_{P,2}^2}{1 - 4\eta_P\left(1 - \frac{1}{D_{\text{in}}}\right)\sigma_{P,1} + 4\eta_P^2\left(1 - \frac{1}{D_{\text{in}}}\right)^2\sigma_{P,1}^2} < 1 < \alpha,$$

that is,

$$\frac{1 + \frac{4\eta_P}{D_{\text{in}}}\sigma_{P,2} + \frac{4\eta_P^2}{D_{\text{in}}^2}\sigma_{P,2}^2}{\alpha} < 1 - 4\eta_P\left(1 - \frac{1}{D_{\text{in}}}\right)\sigma_{P,1} + 4\eta_P^2\left(1 - \frac{1}{D_{\text{in}}}\right)^2\sigma_{P,1}^2.$$

Plugging this result back in,

$$\begin{aligned}\sigma'_{P,2} &= \frac{1 + \frac{4\eta_P}{D_{\text{in}}}\sigma_{P,2} + \frac{4\eta_P^2}{D_{\text{in}}^2}\sigma_{P,2}^2}{\alpha} \sigma_{P,1} \\ &< \left(1 - 4\eta_P\left(1 - \frac{1}{D_{\text{in}}}\right)\sigma_{P,1} + 4\eta_P^2\left(1 - \frac{1}{D_{\text{in}}}\right)^2\sigma_{P,1}^2\right) \sigma_{P,1} \\ &= \sigma'_{P,1}.\end{aligned}$$

In addition, $\sigma'_{P,2} = \sigma_{P,2} + \frac{4\eta_P}{D_{\text{in}}}\sigma_{P,2}^2 + \frac{4\eta_P^2}{D_{\text{in}}^2}\sigma_{P,2}^3 \geq \sigma_{P,i} + \frac{4\eta_P}{D_{\text{in}}}\sigma_{P,i}^2 + \frac{4\eta_P^2}{D_{\text{in}}^2}\sigma_{P,i}^3 = \sigma'_{P,i}$ for $i = 3, \dots, k_P$ since $\sigma_{P,2} \geq \sigma_{P,i}$ for $i = 3, \dots, k_P$ by definition. Therefore, $\sigma'_{P,1}$ remains to be the maximum singular value of $\theta'^{\top} \mathbf{K}_P \theta'$, and σ'_{P,k_P} the minimum. Finally,

$$\begin{aligned}\kappa\left(\theta'^{\top} \mathbf{K}_P \theta'\right) &= \frac{\sigma'_{P,1}}{\sigma'_{P,k_P}} \\ &= \frac{\sigma_{P,1} - 4\eta_P\left(1 - \frac{1}{D_{\text{in}}}\right)\sigma_{P,1}^2 + 4\eta_P^2\left(1 - \frac{1}{D_{\text{in}}}\right)^2\sigma_{P,1}^3}{\sigma_{P,k_P} + \frac{4\eta_P}{D_{\text{in}}}\sigma_{P,k_P}^2 + \frac{4\eta_P^2}{D_{\text{in}}^2}\sigma_{P,k_P}^3} \\ &< \frac{\sigma_{P,1} - 4\eta_P\left(1 - \frac{1}{D_{\text{in}}}\right)\sigma_{P,1}^2 + 4\eta_P^2\left(1 - \frac{1}{D_{\text{in}}}\right)^2\sigma_{P,1}^3}{\sigma_{P,k_P}} \\ &< \frac{\sigma_{P,1}}{\sigma_{P,k_P}} \\ &= \kappa\left(\theta^{\top} \mathbf{K}_P \theta\right)\end{aligned}$$

where the second inequality holds when $\eta_P < \frac{1}{\left(1 - \frac{1}{D_{\text{in}}}\right)\sigma_{P,1}}$ which indicates that $-4\eta_P\left(1 - \frac{1}{D_{\text{in}}}\right)\sigma_{P,1}^2 + 4\eta_P^2\left(1 - \frac{1}{D_{\text{in}}}\right)^2\sigma_{P,1}^3 < 0$.

(2) Denote $\mathcal{R}'_{i11}(\mathbf{H}_H) = \frac{1}{2}\sigma_{H,k_H}^2 - \frac{1}{2k_H}\|\mathbf{H}_H\|_F^2$, then by Theorem 4.2 (2), we know $\nabla_{\theta}\mathcal{R}_{i11}(\mathbf{H}_H(\theta)) = \frac{2\mathbf{K}_H\theta(\sigma_{H,k_H}\mathbf{u}_{H,k_H}\mathbf{v}_{H,k_H}^{\top} - \frac{1}{k_H}\mathbf{H}_H)}{\mathcal{R}'_{i11}(\mathbf{H}_H)^2}$. Since $\theta' = \theta - \eta_H\mathbf{K}_H^{-1}\nabla_{\theta}\mathcal{R}_{i11}(\mathbf{H}_H(\theta))$, we have

$$\begin{aligned}\theta'^{\top}\mathbf{K}_H\theta' &= (\theta - \eta_H\mathbf{K}_H^{-1}\nabla_{\theta}\mathcal{R}_{i11}(\mathbf{H}_H(\theta)))^{\top}\mathbf{K}_H(\theta - \eta_H\mathbf{K}_H^{-1}\nabla_{\theta}\mathcal{R}_{i11}(\mathbf{H}_H(\theta))) \\ &= \theta^{\top}\mathbf{K}_H\theta - \frac{2\eta_H}{\mathcal{R}'_{i11}(\mathbf{H}_H)^2}\left(\sigma_{H,k_H}\mathbf{u}_{H,k_H}\mathbf{v}_{H,k_H}^{\top} - \frac{1}{k_H}\mathbf{H}_H\right)^{\top}\theta^{\top}\mathbf{K}_H\theta - \frac{2\eta_H}{\mathcal{R}'_{i11}(\mathbf{H}_H)^2}\theta^{\top}\mathbf{K}_H\theta\left(\sigma_{H,k_H}\mathbf{u}_{H,k_H}\mathbf{v}_{H,k_H}^{\top} - \frac{1}{k_H}\mathbf{H}_H\right) \\ &\quad + \frac{4\eta_H^2}{\mathcal{R}'_{i11}(\mathbf{H}_H)^4}\left(\sigma_{H,k_H}\mathbf{u}_{H,k_H}\mathbf{v}_{H,k_H}^{\top} - \frac{1}{k_H}\mathbf{H}_H\right)^{\top}\theta^{\top}\mathbf{K}_H\theta\left(\sigma_{H,k_H}\mathbf{u}_{H,k_H}\mathbf{v}_{H,k_H}^{\top} - \frac{1}{k_H}\mathbf{H}_H\right) \\ &= \mathbf{H}_H - \frac{2\eta_H}{\mathcal{R}'_{i11}(\mathbf{H}_H)^2}\left(\sigma_{H,k_H}\mathbf{u}_{H,k_H}\mathbf{v}_{H,k_H}^{\top} - \frac{1}{k_H}\mathbf{H}_H\right)^{\top}\mathbf{H}_H - \frac{2\eta_H}{\mathcal{R}'_{i11}(\mathbf{H}_H)^2}\mathbf{H}_H\left(\sigma_{H,k_H}\mathbf{u}_{H,k_H}\mathbf{v}_{H,k_H}^{\top} - \frac{1}{k_H}\mathbf{H}_H\right) \\ &\quad + \frac{4\eta_H^2}{\mathcal{R}'_{i11}(\mathbf{H}_H)^4}\left(\sigma_{H,k_H}\mathbf{u}_{H,k_H}\mathbf{v}_{H,k_H}^{\top} - \frac{1}{k_H}\mathbf{H}_H\right)^{\top}\mathbf{H}_H\left(\sigma_{H,k_H}\mathbf{u}_{H,k_H}\mathbf{v}_{H,k_H}^{\top} - \frac{1}{k_H}\mathbf{H}_H\right).\end{aligned}$$

Since $\mathbf{H}_P(\theta) = \theta^{\top}\mathbf{K}_H\theta$ for $\mathbf{K}_H = \mathbf{X}_H^{\top}\mathbf{X}_H$ is also symmetric and positive semidefinite, we know for $\mathbf{H}_H(\theta) = \mathbf{U}_H\text{Diag}(\boldsymbol{\sigma}_H)\mathbf{V}_H^{\top}$, it holds that $\mathbf{U}_H = \mathbf{V}_H$. Following similar arguments as in (1),

$$\begin{aligned}\sigma_{H,k_H}\mathbf{u}_{H,k_H}\mathbf{v}_{H,k_H}^{\top} - \frac{1}{k_H}\mathbf{H}_H &= -\frac{1}{k_H}\sum_{i=1}^{k_H-1}\sigma_{H,i}\mathbf{u}_{H,i}\mathbf{v}_{H,i}^{\top} + \left(1 - \frac{1}{k_H}\right)\sigma_{H,k_H}\mathbf{u}_{H,k_H}\mathbf{v}_{H,k_H}^{\top} \\ &= \mathbf{V}_H\text{Diag}(\tilde{\boldsymbol{\sigma}}_H)\mathbf{V}_H^{\top}\end{aligned}$$

for $\text{Diag}(\tilde{\boldsymbol{\sigma}}_H) = \left[-\frac{1}{k_H}\sigma_{H,1}, \dots, -\frac{1}{k_H}\sigma_{H,k_H-1}, \left(1 - \frac{1}{k_H}\right)\sigma_{H,k_H}\right]$. Since \mathbf{V}_H is orthonormal, i.e., $\mathbf{V}_H^{\top}\mathbf{V}_H = \mathbf{I}$,

$$\begin{aligned}\theta'^{\top}\mathbf{K}_H\theta' &= \mathbf{V}_H\text{Diag}(\boldsymbol{\sigma}_H)\mathbf{V}_H^{\top} - \frac{2\eta_H}{\mathcal{R}'_{i11}(\mathbf{H}_H)^2}\left(\mathbf{V}_H\text{Diag}(\tilde{\boldsymbol{\sigma}}_H)\mathbf{V}_H^{\top}\right)^{\top}\mathbf{V}_H\text{Diag}(\boldsymbol{\sigma}_H)\mathbf{V}_H^{\top} \\ &\quad - \frac{2\eta_H}{\mathcal{R}'_{i11}(\mathbf{H}_H)^2}\left(\mathbf{V}_H\text{Diag}(\boldsymbol{\sigma}_H)\mathbf{V}_H^{\top}\right)^{\top}\mathbf{V}_H\text{Diag}(\tilde{\boldsymbol{\sigma}}_P)\mathbf{V}_P^{\top} \\ &\quad + \frac{4\eta_H^2}{\mathcal{R}'_{i11}(\mathbf{H}_H)^4}\left(\mathbf{V}_H\text{Diag}(\tilde{\boldsymbol{\sigma}}_H)\mathbf{V}_H^{\top}\right)^{\top}\mathbf{V}_H\text{Diag}(\boldsymbol{\sigma}_H)\mathbf{V}_H^{\top}\mathbf{V}_H\text{Diag}(\tilde{\boldsymbol{\sigma}}_H)\mathbf{V}_H^{\top} \\ &= \mathbf{V}_H\text{Diag}\left(\boldsymbol{\sigma}_H - \frac{4\eta_H}{\mathcal{R}'_{i11}(\mathbf{H}_H)^2}\tilde{\boldsymbol{\sigma}}_H \odot \boldsymbol{\sigma}_H + \frac{4\eta_H^2}{\mathcal{R}'_{i11}(\mathbf{H}_H)^4}\tilde{\boldsymbol{\sigma}}_H \odot \boldsymbol{\sigma}_H \odot \tilde{\boldsymbol{\sigma}}_H\right)\mathbf{V}_H^{\top} \\ &= \mathbf{V}_H\text{Diag}(\boldsymbol{\sigma}'_H)\mathbf{V}_H^{\top},\end{aligned}$$

$$\text{for } \boldsymbol{\sigma}'_H = \left[\sigma'_{H,1}, \dots, \sigma'_{H,k_H}\right]^{\top}, \sigma'_{H,i} = \begin{cases} \sigma_{H,i} + \frac{4\eta_H}{k_H\mathcal{R}'_{i11}(\mathbf{H}_H)^2}\sigma_{H,i}^2 + \frac{4\eta_H^2}{k_H^2\mathcal{R}'_{i11}(\mathbf{H}_H)^4}\sigma_{H,i}^3 & \text{if } i < k_H \\ \sigma_{H,k_H} - \frac{4\eta_H}{\mathcal{R}'_{i11}(\mathbf{H}_H)^2}\left(1 - \frac{1}{k_H}\right)\sigma_{H,k_H}^2 + \frac{4\eta_H^2}{\mathcal{R}'_{i11}(\mathbf{H}_H)^4}\left(1 - \frac{1}{k_H}\right)^2\sigma_{H,k_H}^3 & \text{if } i = k_H \end{cases},$$

and \odot denotes element-wise product.

Since $\sigma_{H_H}^{\min}$ is unique, we know that $\exists \beta \in (0, 1)$ such that $\sigma_{H,k_H} = \beta\sigma_{H,k_H-1}$. Then we have

$$\begin{aligned}\sigma'_{H,k_H} &= \sigma_{H,k_H} - \frac{4\eta_H}{\mathcal{R}'_{i11}(\mathbf{H}_H)^2}\left(1 - \frac{1}{k_H}\right)\sigma_{H,k_H}^2 + \frac{4\eta_H^2}{\mathcal{R}'_{i11}(\mathbf{H}_H)^4}\left(1 - \frac{1}{k_H}\right)^2\sigma_{H,k_H}^3 \\ &= \left(1 - \frac{4\eta_H}{\mathcal{R}'_{i11}(\mathbf{H}_H)^2}\left(1 - \frac{1}{k_H}\right)\sigma_{H,k_H} + \frac{4\eta_H^2}{\mathcal{R}'_{i11}(\mathbf{H}_H)^4}\left(1 - \frac{1}{k_H}\right)^2\sigma_{H,k_H}^2\right)\sigma_{H,k_H} \\ &= \left(1 - \frac{4\eta_H}{\mathcal{R}'_{i11}(\mathbf{H}_H)^2}\left(1 - \frac{1}{k_H}\right)\sigma_{H,k_H} + \frac{4\eta_H^2}{\mathcal{R}'_{i11}(\mathbf{H}_H)^4}\left(1 - \frac{1}{k_H}\right)^2\sigma_{H,k_H}^2\right)\beta\sigma_{H,k_H-1} \\ &= \left(1 + \frac{4\eta_H\sigma_{H,k_H}}{k_H\mathcal{R}'_{i11}(\mathbf{H}_H)^2} + \frac{4\eta_H^2\sigma_{H,k_H}^2}{k_H^2\mathcal{R}'_{i11}(\mathbf{H}_H)^4} - \frac{4\eta_H\sigma_{H,k_H}}{\mathcal{R}'_{i11}(\mathbf{H}_H)^2} + \frac{4\eta_H^2\sigma_{H,k_H}^2}{\mathcal{R}'_{i11}(\mathbf{H}_H)^4} - \frac{8\eta_H^2\sigma_{H,k_H}^2}{k_H\mathcal{R}'_{i11}(\mathbf{H}_H)^4}\right)\beta\sigma_{H,k_H-1}.\end{aligned}$$

Letting $0 < \eta_{\mathbb{H}} < \frac{\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^2}{1-2\sigma_{\mathbb{H},k_{\mathbb{H}}}/k_{\mathbb{H}}} = \frac{1}{1-2\sigma_{\mathbb{H},k_{\mathbb{H}}}/k_{\mathbb{H}}} \left(\frac{1}{2k_{\mathbb{H}}} \|\theta^{\top} \mathbf{K}_{\mathbb{H}} \theta\|_F^2 - \frac{1}{2} (\sigma_{\mathbb{H},k_{\mathbb{H}}}^{\min})^2 \right)^2$, we have

$$-\frac{4\eta_{\mathbb{H}}\sigma_{\mathbb{H},k_{\mathbb{H}}}}{\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^2} + \frac{4\eta_{\mathbb{H}}^2\sigma_{\mathbb{H},k_{\mathbb{H}}}^2}{\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^4} - \frac{8\eta_{\mathbb{H}}^2\sigma_{\mathbb{H},k_{\mathbb{H}}}^2}{k_{\mathbb{H}}\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^4} < 0. \quad (21)$$

Also, $1 - \frac{4\eta_{\mathbb{H}}}{\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^2} \left(1 - \frac{1}{k_{\mathbb{H}}}\right) \sigma_{\mathbb{H},k_{\mathbb{H}}} + \frac{4\eta_{\mathbb{H}}^2}{\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^4} \left(1 - \frac{1}{k_{\mathbb{H}}}\right)^2 \sigma_{\mathbb{H},k_{\mathbb{H}}}^2 = \left(1 - \frac{2\eta_{\mathbb{H}}}{\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^2} \left(1 - \frac{1}{k_{\mathbb{H}}}\right) \sigma_{\mathbb{H},k_{\mathbb{H}}}\right)^2 > 0$ for any $\eta_{\mathbb{H}} > 0$. Given that $\sigma_{\mathbb{H},k_{\mathbb{H}-1}} > \sigma_{\mathbb{H},k_{\mathbb{H}}}$ and Eq. (21),

$$\beta < 1$$

$$\begin{aligned} &< \frac{1 + \frac{4\eta_{\mathbb{H}}}{k_{\mathbb{H}}\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^2} \sigma_{\mathbb{H},k_{\mathbb{H}-1}} + \frac{4\eta_{\mathbb{H}}^2}{k_{\mathbb{H}}^2\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^4} \sigma_{\mathbb{H},k_{\mathbb{H}-1}^2}}{1 + \frac{4\eta_{\mathbb{H}}}{k_{\mathbb{H}}\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^2} \sigma_{\mathbb{H},k_{\mathbb{H}}} + \frac{4\eta_{\mathbb{H}}^2}{k_{\mathbb{H}}^2\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^4} \sigma_{\mathbb{H},k_{\mathbb{H}}}^2}} \\ &< \frac{1 + \frac{4\eta_{\mathbb{H}}}{k_{\mathbb{H}}\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^2} \sigma_{\mathbb{H},k_{\mathbb{H}-1}} + \frac{4\eta_{\mathbb{H}}^2}{k_{\mathbb{H}}^2\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^4} \sigma_{\mathbb{H},k_{\mathbb{H}-1}^2}}{1 + \frac{4\eta_{\mathbb{H}}}{\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^2} \sigma_{\mathbb{H},k_{\mathbb{H}}} + \frac{4\eta_{\mathbb{H}}^2}{k_{\mathbb{H}}^2\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^4} \sigma_{\mathbb{H},k_{\mathbb{H}}}^2 - \frac{4\eta_{\mathbb{H}}}{\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^2} \sigma_{\mathbb{H},k_{\mathbb{H}}} + \frac{4\eta_{\mathbb{H}}^2}{\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^4} \sigma_{\mathbb{H},k_{\mathbb{H}}}^2 - \frac{8\eta_{\mathbb{H}}^2}{k_{\mathbb{H}}\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^4} \sigma_{\mathbb{H},k_{\mathbb{H}}}^2}, \end{aligned}$$

indicating $\left(1 + \frac{4\eta_{\mathbb{H}}}{k_{\mathbb{H}}\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^2} \sigma_{\mathbb{H},k_{\mathbb{H}}} + \frac{4\eta_{\mathbb{H}}^2}{k_{\mathbb{H}}^2\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^4} \sigma_{\mathbb{H},k_{\mathbb{H}}}^2 - \frac{4\eta_{\mathbb{H}}}{\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^2} \sigma_{\mathbb{H},k_{\mathbb{H}}} + \frac{4\eta_{\mathbb{H}}^2}{\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^4} \sigma_{\mathbb{H},k_{\mathbb{H}}}^2 - \frac{8\eta_{\mathbb{H}}^2}{k_{\mathbb{H}}\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^4} \sigma_{\mathbb{H},k_{\mathbb{H}}}^2\right) \beta < 1 + \frac{4\eta_{\mathbb{H}}}{k_{\mathbb{H}}\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^2} \sigma_{\mathbb{H},k_{\mathbb{H}-1}} + \frac{4\eta_{\mathbb{H}}^2}{k_{\mathbb{H}}^2\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^4} \sigma_{\mathbb{H},k_{\mathbb{H}-1}^2}$. Therefore,

$$\begin{aligned} \sigma'_{\mathbb{H},k_{\mathbb{H}}} &= \left(1 + \frac{4\eta_{\mathbb{H}}\sigma_{\mathbb{H},k_{\mathbb{H}}}}{k_{\mathbb{H}}\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^2} + \frac{4\eta_{\mathbb{H}}^2\sigma_{\mathbb{H},k_{\mathbb{H}}}^2}{k_{\mathbb{H}}^2\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^4} - \frac{4\eta_{\mathbb{H}}\sigma_{\mathbb{H},k_{\mathbb{H}}}}{\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^2} + \frac{4\eta_{\mathbb{H}}^2\sigma_{\mathbb{H},k_{\mathbb{H}}}^2}{\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^4} - \frac{8\eta_{\mathbb{H}}^2\sigma_{\mathbb{H},k_{\mathbb{H}}}^2}{k_{\mathbb{H}}\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^4}\right) \beta \sigma_{\mathbb{H},k_{\mathbb{H}-1}} \\ &< \left(1 + \frac{4\eta_{\mathbb{H}}}{k_{\mathbb{H}}\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^2} \sigma_{\mathbb{H},k_{\mathbb{H}-1}} + \frac{4\eta_{\mathbb{H}}^2}{k_{\mathbb{H}}^2\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^4} \sigma_{\mathbb{H},k_{\mathbb{H}-1}^2}\right) \sigma_{\mathbb{H},k_{\mathbb{H}-1}} \\ &= \sigma'_{\mathbb{H},k_{\mathbb{H}-1}}. \end{aligned}$$

In addition, $\sigma'_{\mathbb{H},k_{\mathbb{H}-1}} = \sigma_{\mathbb{H},k_{\mathbb{H}-1}} + \frac{4\eta_{\mathbb{H}}}{k_{\mathbb{H}}\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^2} \sigma_{\mathbb{H},k_{\mathbb{H}-1}^2} + \frac{4\eta_{\mathbb{H}}^2}{k_{\mathbb{H}}^2\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^4} \sigma_{\mathbb{H},k_{\mathbb{H}-1}^3} \leq \sigma_{\mathbb{H},i} + \frac{4\eta_{\mathbb{H}}}{k_{\mathbb{H}}\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^2} \sigma_{\mathbb{H},i}^2 + \frac{4\eta_{\mathbb{H}}^2}{k_{\mathbb{H}}^2\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^4} \sigma_{\mathbb{H},i}^3 = \sigma'_{\mathbb{H},i}$ for $i = 1, \dots, k_{\mathbb{H}} - 2$ since $\sigma_{\mathbb{H},k_{\mathbb{H}-1}} \leq \sigma_{\mathbb{H},i}$ for $i = 2, \dots, k_{\mathbb{H}} - 1$ by definition. That is to say, $\sigma'_{\mathbb{H},k_{\mathbb{H}}}$ remains to be the minimum singular value of $\theta'^{\top} \mathbf{K}_{\mathbb{H}} \theta'$, and $\sigma'_{\mathbb{H},1}$ the maximum. Finally,

$$\begin{aligned} &\kappa(\theta'^{\top} \mathbf{K}_{\mathbb{H}} \theta') \\ &= \frac{\sigma'_{\mathbb{H},1}}{\sigma'_{\mathbb{H},k_{\mathbb{H}}}} \\ &= \frac{\left(1 + \frac{4\eta_{\mathbb{H}}}{k_{\mathbb{H}}\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^2} \sigma_{\mathbb{H},1} + \frac{4\eta_{\mathbb{H}}^2}{k_{\mathbb{H}}^2\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^4} \sigma_{\mathbb{H},1}^2\right) \sigma_{\mathbb{H},1}}{\left(1 + \frac{4\eta_{\mathbb{H}}}{k_{\mathbb{H}}\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^2} \sigma_{\mathbb{H},k_{\mathbb{H}}} + \frac{4\eta_{\mathbb{H}}^2}{k_{\mathbb{H}}^2\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^4} \sigma_{\mathbb{H},k_{\mathbb{H}}}^2 - \frac{4\eta_{\mathbb{H}}}{\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^2} \sigma_{\mathbb{H},k_{\mathbb{H}}} + \frac{4\eta_{\mathbb{H}}^2}{\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^4} \sigma_{\mathbb{H},k_{\mathbb{H}}}^2 - \frac{8\eta_{\mathbb{H}}^2}{k_{\mathbb{H}}\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^4} \sigma_{\mathbb{H},k_{\mathbb{H}}}^2\right) \sigma_{\mathbb{H},k_{\mathbb{H}}}} \\ &> \frac{\left(1 + \frac{4\eta_{\mathbb{H}}}{k_{\mathbb{H}}\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^2} \sigma_{\mathbb{H},1} + \frac{4\eta_{\mathbb{H}}^2}{k_{\mathbb{H}}^2\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^4} \sigma_{\mathbb{H},1}^2\right) \sigma_{\mathbb{H},1}}{\left(1 + \frac{4\eta_{\mathbb{H}}}{k_{\mathbb{H}}\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^2} \sigma_{\mathbb{H},k_{\mathbb{H}}} + \frac{4\eta_{\mathbb{H}}^2}{k_{\mathbb{H}}^2\mathcal{R}'_{111}(\mathbf{H}_{\mathbb{H}})^4} \sigma_{\mathbb{H},k_{\mathbb{H}}}^2\right) \sigma_{\mathbb{H},k_{\mathbb{H}}}} \\ &> \frac{\sigma_{\mathbb{H},1}}{\sigma_{\mathbb{H},k_{\mathbb{H}}}} \\ &= \kappa(\theta^{\top} \mathbf{K}_{\mathbb{H}} \theta) \end{aligned}$$

where the first inequality holds by Eq. (21) and the second by $\sigma_{\mathbb{H},1} > \sigma_{\mathbb{H},k_{\mathbb{H}}}$. \square

C. Detailed Experiment Setup

C.1. Datasets

Stanford Cars (Krause et al., 2013) contains 16,185 images of 196 car models and focuses on fine-grained image classification. Country211 (Radford et al., 2021) is a dataset used for country classification based on satellite images, comprising 211 country-level labels, each with 150 training images. This is a subset of the YFCC100M dataset (Thomee et al., 2016) providing user-generated photos and videos, used for domain adaptation evaluation.

C.2. Immunization training details

We summarize the hyper-parameters of training for model immunization in Tab. 4. We choose λ_P and λ_H by balancing the gradient norm of $\mathcal{R}_{\text{well}}$ and \mathcal{R}_{ill} . Specifically, we obtain the scale of λ_P and λ_H first and search over multiples of $\{1, 2, 3, 5\}$. For linear models, we search over the set of $\{0.0005, 0.001, 0.005, 0.01\}$ and report the best result. For ImageNet we followed the default learning rate $\eta = 1 \times 10^{-5}$. The number of epochs is based on early stopping using RIR and the test accuracy. All experiments are conducted using float64 precision to ensure numerical stability and reduce potential inaccuracies in computations.

Table 4. Hyperparameters for immunization training.

Dataset	Model	η	λ_P	λ_H	Epochs	\mathcal{L}
HousePrice	Linear	0.005	100	1×10^7	1000	Mean squared error
MNIST	Linear	0.001	1	5×10^7	30	Binary Cross-entropy (CE)
ImageNet vs. Stanford Cars	ResNet18	1×10^{-5}	5×10^{-5}	2×10^6	3	Label-smoothing CE
ImageNet vs. Country211	ResNet18	1×10^{-5}	1×10^{-4}	2×10^6	3	Label-smoothing CE
ImageNet vs. Stanford Cars	ViT	1×10^{-5}	3×10^{-6}	3×10^8	2	Label-smoothing CE
ImageNet vs. Country211	ViT	1×10^{-5}	1×10^{-6}	1×10^8	2	Label-smoothing CE

Details of immunizing linear models. For the regression task, the linear feature extractor $\theta \in \mathbb{R}^{79 \times 79}$ is a randomly initialized dummy linear layer, as discussed in Sec. 4.4. We handle missing values in the tabular data by filling NaNs with 0. Categorical features are converted into numerical values using `LabelEncoder`. Finally, the features and labels are normalized using their respective mean and standard deviation. To create \mathcal{D}_P and \mathcal{D}_H , we split the House prices dataset (Montoya & DataCanary, 2016) by the feature `MSZoning`. Specifically, all entries where `MSZoning = 'RL'` are assigned to \mathcal{D}_H , while the remaining entries form \mathcal{D}_P .

For the binary classification task on MNIST, the linear feature extractor $\theta \in \mathbb{R}^{784 \times 784}$ is also a randomly initialized dummy linear layer, and we construct a training dataset by selecting two specific target digits. The dataset is created using a custom `BinaryDataset` class, which filters the original MNIST dataset to include only the chosen digits and assigns new labels: one digit is mapped to label 0 and the other to label 1. To ensure balance in the dataset, we limit the number of samples for each digit to the smaller count between the two. For optimization, we use Adam (Kingma, 2014) with $\beta = (.9, 0.999)$ and $\epsilon = 1 \times 10^{-8}$ instead of the basic gradient descent in Alg. 1. For the linear model, we computed the Hessian inverse by solving a regularized least-squares system, where the Hessian is in the shape of $\mathbb{R}^{D_{\text{in}} \times D_{\text{in}}}$. Here $D_{\text{in}} = 79$ for the regression task and $D_{\text{in}} = 784$ for the image classification task.

Details of immunizing non-linear models. The pre-trained ResNet18 and ViT are loaded from Pytorch Image Models (Wightman, 2019) with the model name `resnet18` and `vit_base_patch16_clip_224`. We also create the dataset with the built-in function `create_dataset` from Wightman (2019). The feature embedding sizes for ResNet18 and ViT are 512 and 768, respectively. To facilitate balanced training when dataset sizes differ, we implement a `CombinedLoader`, which pairs batches from two data loaders. The longer dataset dictates training duration, while the shorter dataset cycles continuously using `itertools`. The number of epochs reported in Tab. 4 corresponds to the epochs of \mathcal{D}_H , *i.e.*, the shorter loader.

For optimization, we use SGD with Nesterov momentum to optimize Eq. (11), setting an initial learning rate of 1×10^{-5} with momentum 0.9. The trainable feature extractor parameters are optimized with zero weight decay, while the classifier parameters use a weight decay of 2×10^{-5} .

C.3. Pseudo-code of the dummy layer

We provide the Pseudo-code for implementing the dummy layer in Fig. 4 below. The `DummyLinear` layer extends `torch.nn.Linear` and incorporates an optional preconditioning mechanism in the backward pass using the inverse feature covariance matrix. The `LinearFunction` class defines the forward and backward computations, where the forward pass applies a standard linear transformation $\mathbf{X}\mathbf{W}^\top + \mathbf{b}$ and stores the input, weight, and bias for gradient computation. In the backward pass, the input gradient is computed normally, while the weight gradient is modified based on whether preconditioning is enabled (`use_precond=True`). If enabled, the weight gradient is adjusted by solving a *regularized least-squares system* using the inverse of the feature covariance matrix $\mathbf{X}^\top\mathbf{X} + \epsilon\mathbf{I}$, improving numerical stability.

```
class LinearFunction:
    @staticmethod
    def forward(ctx, input, weight, bias, lambda_reg, use_precond):
        # Save input tensors for backward
        ctx.save_for_backward(input, weight, bias)
        ctx.lambda_reg = lambda_reg
        ctx.use_precond = use_precond

        # Compute output
        output = input.mm(weight.t())
        if bias is not None:
            output += bias.unsqueeze(0).expand_as(output)
        return output

    @staticmethod
    def backward(ctx, grad_output):
        # Retrieve saved tensors
        input, weight, bias = ctx.saved_tensors
        lambda_reg = ctx.lambda_reg
        use_precond = ctx.use_precond

        # Initialize gradients
        grad_input = grad_weight = grad_bias = None

        if ctx.needs_input_grad[0]:
            grad_input = grad_output.mm(weight)

        if ctx.needs_input_grad[1]:
            base_grad_weight = grad_output.t().mm(input)
            if use_precond:
                XtX = input.t().mm(input)
                lambda_eye = lambda_reg * torch.eye(XtX.size(0), device=XtX.device)
                XtX_reg = XtX + lambda_eye
                grad_weight = torch.linalg.solve(XtX_reg, base_grad_weight)
            else:
                grad_weight = base_grad_weight

        if bias is not None and ctx.needs_input_grad[2]:
            grad_bias = grad_output.sum(0)

        return grad_input, grad_weight, grad_bias, None, None

class DummyLinear(nn.Linear):
    def forward(self, input, lambda_reg, use_precond):
        # Dynamically decide whether to use the covariance inversion as the preconditioner
        return LinearFunction.apply(input, self.weight, self.bias, lambda_reg, use_precond)
```

Figure 4. Dummy layer with selective inverse feature covariance matrix in backward function.