## Supplementary Material for: Adversarial Regression with Doubly Non-negative Weighting Matrices

In this supplementary, we give details for the proofs of our technical results in Section $A$. In Section $B$ we describe implementation details. In Section C, we illustrate additional empirical results.

We have released code for our proposed tools at:
https://github.com/lttam/Adversarial-Regression

## A Proofs and Additional Discussions

## A. 1 Proofs of Section 3

In the following, the symbol $\langle\cdot, \cdot\rangle$ will be used to represent both Frobenius norm of matrices and standard Euclidean norm of vectors. Also for a $p \times p$ matrix $A$, we let $\|A\|_{2}=\sqrt{\langle A, A\rangle}$ to denote its Frobenius norm and let $\lambda_{\max }(A)$ to represent its maximum eigenvalue. In order to prove Proposition 3.4 , we begin by computing the support function of the convex cone of symmetric positive definite (SPD) matrices.
Lemma A. 1 (Support function). For any matrices $A \in \mathbb{S}^{N+1}$ and $\Omega \in \mathbb{S}_{++}^{N+1}$, we have $\operatorname{Tr}[A \Omega] \leq$ $\lambda_{\max }(A) \operatorname{Tr}[\Omega]$, and $\operatorname{Tr}[A \Omega]<0$ if in addition $A \preceq 0$ and $A \neq 0$. Also for each $A \in \mathbb{S}^{N+1}$,

$$
\sup _{\Omega \succ 0} \operatorname{Tr}[A \Omega]= \begin{cases}+\infty & \text { if } \lambda_{\max }(A)>0 \\ 0 & \text { if } \lambda_{\max }(A) \leq 0\end{cases}
$$

Proof of Lemma A.1. Since $A$ is symmetric, we can decompose it as $D=Q^{\top} A Q$ with $D$ being a diagonal matrix formed by eigenvalues of $A$ and $Q$ being an orthogonal matrix whose columns are normalized eigenvectors of $A$. Then $Q^{\top} A=D Q^{\top}$ and hence

$$
\operatorname{Tr}[A \Omega]=\operatorname{Tr}\left[Q^{\top} A \Omega Q\right]=\operatorname{Tr}\left[D Q^{\top} \Omega Q\right] \leq \lambda_{\max }(D) \operatorname{Tr}\left[Q^{\top} \Omega Q\right]=\lambda_{\max }(A) \operatorname{Tr}[\Omega]
$$

where we have used the fact $Q^{\top} \Omega Q \succ 0$ to obtain the inequality. In case $A \preceq 0$ and $A \neq 0$, then as all diagonal elements of $D$ are nonpositive with at least one negative element we have $\operatorname{Tr}[A \Omega]=\operatorname{Tr}\left[D Q^{\top} \Omega Q\right]<0$. These give the first part of the lemma.
For the second part, let $v$ be an eigenvector of $A$ corresponding to eigenvalue $\lambda_{\max }(A)$. In case $\lambda_{\text {max }}(A)>0$, by taking $\Omega_{t}=I+t v v^{\top} \succ 0$ with $t \in(0,+\infty)$ and letting $t \rightarrow+\infty$ we see that $\operatorname{Tr}\left[A \Omega_{t}\right]=\operatorname{Tr}[A]+t \lambda_{\max }(A) \operatorname{Tr}\left[v v^{\top}\right]$ tends to $+\infty$. This implies that $\sup _{\Omega \succ 0} \operatorname{Tr}[A \Omega]=+\infty$. For the case $\lambda_{\max }(A) \leq 0$, we instead take $\Omega_{t}=t I+t v v^{\top} \succ 0$ with $t \in(0,1)$ and let $t \rightarrow 0^{+}$to conclude that $\sup _{\Omega \succ 0} \operatorname{Tr}[A \Omega] \geq 0$. On the other hand, due to the first part of the lemma we always have $\sup _{\Omega \succ 0} \operatorname{Tr}[A \Omega] \leq 0$. Therefore, the desired result follows.

## For Proposition 3.4

Proof of Proposition 3.4 The set $\mathcal{V}_{\mathrm{D}, \rho}(\widehat{\Omega})$ is bounded. Indeed, if $\rho=0$ then $\mathcal{V}_{\mathrm{D}, \rho}(\widehat{\Omega})=\{\widehat{\Omega}\}$ which is trivially bounded. If $\rho>0$, then any $\Omega \in \mathcal{V}_{\mathbb{D}, \rho}(\widehat{\Omega})$ satisfies $\operatorname{Tr}\left[\Omega \widehat{\Omega}^{-1}\right]-\log \operatorname{det}\left(\Omega \widehat{\Omega}^{-1}\right)-(N+$

1) $\leq \rho$. Since $\Omega \widehat{\Omega}^{-1}$ is similar to the matrix $\widehat{\Omega}^{-\frac{1}{2}} \Omega \widehat{\Omega}^{-\frac{1}{2}}$, we can rewrite this as

$$
\begin{equation*}
\operatorname{Tr}\left[\widehat{\Omega}^{-\frac{1}{2}} \Omega \widehat{\Omega}^{-\frac{1}{2}}\right]-\log \operatorname{det}\left(\widehat{\Omega}^{-\frac{1}{2}} \Omega \widehat{\Omega}^{-\frac{1}{2}}\right)-(N+1) \leq \rho \tag{A.1}
\end{equation*}
$$

We claim that this implies that $\Omega$ is bounded in the sense that there exist numbers $0<\underline{\sigma} \leq \bar{\sigma}<\infty$ depending only on $\rho, \widehat{\Omega}$, and $N$ such that $\underline{\sigma} I \preceq \Omega \preceq \bar{\sigma} I$. To see this, let $\sigma_{i}(i=1, \ldots, N+1)$ be the eigenvalues of $\widehat{\Omega}^{-\frac{1}{2}} \Omega \widehat{\Omega}^{-\frac{1}{2}}$. Then A.1 gives

$$
\sum_{i=1}^{N+1}\left[\sigma_{i}-\log \sigma_{i}-1\right] \leq \rho
$$

Because the function $\sigma \mapsto \sigma-\log \sigma-1$ is non-negative for every $\sigma>0$, we then find that $\left\{\sigma_{i}-\log \sigma_{i}-1\right\}_{i=1}^{N+1}$ is bounded. Therefore, $\sigma_{i}$ must be bounded between two positive constants depending only on $\rho$ and $N$. Now let $\lambda$ be an eigenvalue of $\Omega$ and let $v$ be its associated eigenvector. Then $\lambda\|v\|_{2}^{2}=\langle v, \Omega v\rangle=\left\langle\left(\widehat{\Omega}^{\frac{1}{2}} v\right), \widehat{\Omega}^{-\frac{1}{2}} \Omega \widehat{\Omega}^{-\frac{1}{2}}\left(\widehat{\Omega}^{\frac{1}{2}} v\right)\right\rangle$, and so

$$
\sigma_{*}\left\|\widehat{\Omega}^{\frac{1}{2}} v\right\|_{2}^{2} \leq \lambda\|v\|_{2}^{2} \leq \sigma^{*}\left\|\widehat{\Omega}^{\frac{1}{2}} v\right\|_{2}^{2}
$$

where $\sigma_{*}$ and $\sigma^{*}$ respectively denote the smallest and largest eigenvalues of $\widehat{\Omega}^{-\frac{1}{2}} \Omega \widehat{\Omega}^{-\frac{1}{2}}$. It follows that $\sigma_{*} \lambda_{\min }(\widehat{\Omega}) \leq \lambda \leq \sigma^{*} \lambda_{\max }(\widehat{\Omega})$. Thus all eigenvalues of $\Omega$ are bounded between two positive constants, and the claim is proved.
One now can rewrite $\mathcal{V}_{\mathrm{D}, \rho}(\widehat{\Omega})$ as

$$
\mathcal{V}_{\mathrm{D}, \rho}(\widehat{\Omega})=\left\{\Omega \in \mathbb{S}_{+}^{N+1}: \underline{\sigma} I \preceq \Omega \preceq \bar{\sigma} I, \quad \operatorname{Tr}\left[\Omega \widehat{\Omega}^{-1}\right]-\log \operatorname{det}\left(\Omega \widehat{\Omega}^{-1}\right)-(N+1) \leq \rho\right\} .
$$

The function $\Omega \mapsto \operatorname{Tr}\left[\Omega \widehat{\Omega}^{-1}\right]-\log \operatorname{det}\left(\Omega \widehat{\Omega}^{-1}\right)$ is convex and continuous on the set $\{\Omega: \underline{\sigma} I \preceq$ $\Omega \preceq \bar{\sigma} I\}$, thus the set $\mathcal{V}_{\mathrm{D}, \rho}(\widehat{\Omega})$ is convex and compact.

We now proceed to provide the support function of $\mathcal{V}_{\mathrm{D}, \rho}(\widehat{\Omega})$. One can verify that $\widehat{\Omega}$ is the Slater point of the convex set $\mathcal{V}_{\mathbb{D}, \rho}(\widehat{\Omega})$. Assume momentarily that $T \neq 0$, using a duality argument, we find

$$
\begin{align*}
& \sup _{\Omega \in \mathcal{V}_{\mathrm{D}, \rho}(\widehat{\Omega})} \operatorname{Tr}[T \Omega] \\
= & \sup _{\Omega \succ 0} \inf _{\gamma \geq 0} \operatorname{Tr}[T \Omega]+\gamma\left(\bar{\rho}-\operatorname{Tr}\left[\widehat{\Omega}^{-1} \Omega\right]+\log \operatorname{det} \Omega\right) \\
= & \inf _{\gamma \geq 0}\left\{\gamma \bar{\rho}+\sup _{\Omega \succ 0}\left\{\operatorname{Tr}\left[\left(T-\gamma \widehat{\Omega}^{-1}\right) \Omega\right]+\gamma \log \operatorname{det} \Omega\right\}\right\}, \tag{A.2}
\end{align*}
$$

where the last equality follows from strong duality [1, Proposition 5.3.1], and $\bar{\rho} \triangleq \rho+(N+1)-$ $\log \operatorname{det} \widehat{\Omega} \in \mathbb{R}$. Note that $\gamma=0$ is not an optimal solution to the minimization problem A.2). Indeed, if the maximum eigenvalue of $T$ is strictly positive, then the objective value of problem A.2) evaluated at $\gamma=0$ is $+\infty$ by Lemma A.1. However, because $\mathcal{V}_{\mathrm{D}, \rho}(\widehat{\Omega})$ is compact and bounded, we have $\sup _{\Omega \in \mathcal{V}_{\mathbb{D}_{. \rho}}(\widehat{\Omega})} \operatorname{Tr}[T \Omega]<\infty$. In case the maximum eigenvalue of $T$ is nonpositive, then from Lemma A.1 we see that the objective value of problem A.2 evaluated at $\gamma=0$ is 0 , while $\sup _{\Omega \in \mathcal{V}_{\mathrm{D}, \rho}(\widehat{\Omega})} \operatorname{Tr}[T \Omega]<0$. Thus, the infimum in problem A.2) can be restricted to $\gamma>0$.
If $T-\gamma \widehat{\Omega}^{-1} \nprec 0$, then the inner supremum problem in A.2 becomes unbounded according to Lemma A.1. If $T-\gamma \widehat{\Omega}^{-1} \prec 0$ then the inner supremum problem admits the unique optimal solution

$$
\begin{equation*}
\Omega^{\star}(\gamma)=\gamma\left(\gamma \widehat{\Omega}^{-1}-T\right)^{-1} \tag{A.3}
\end{equation*}
$$

which is obtained by solving the first-order optimality condition. By placing this optimal solution into the objective function and arranging terms together with using the definition of $\bar{\rho}$, we have

$$
\begin{equation*}
\sup _{\Omega \in \mathcal{V}_{\mathrm{D}, \rho}(\widehat{\Omega})} \operatorname{Tr}[T \Omega]=\inf _{\substack{\gamma>0 \\ \gamma \widehat{\Omega}^{-1} \succ T}} \gamma \rho-\gamma \log \operatorname{det}\left(I-\gamma^{-1} \widehat{\Omega}^{\frac{1}{2}} T \widehat{\Omega}^{\frac{1}{2}}\right) \tag{A.4}
\end{equation*}
$$

To complete the proof, we show that the reformulation A.4 holds even for $T=0$ or $\rho=0$. Notice that if $T=0$ or $\rho=0$, then the left-hand side of $A$. evaluates to 0 . If $T=0$, the infimum problem
on the right-hand side of A.4 also attains the optimal value of 0 asymptotically as $\gamma$ decreases to 0 . If $\rho=0$ and $T \neq 0$, then the infimum problem on the right-hand side of A.4 also attains the optimal value of 0 asymptotically as $\gamma$ increases to $+\infty$ by the l'Hopital rule

$$
\lim _{\gamma \uparrow+\infty}-\gamma \log \operatorname{det}\left(I-\widehat{\Omega}^{\frac{1}{2}} T \widehat{\Omega}^{\frac{1}{2}} / \gamma\right)=\operatorname{Tr}[T \widehat{\Omega}]
$$

This observation completes the proof.
For an $p \times p$ real matrix $A$, its spectral radius $\mathcal{R}(A)$ is defined as the largest absolute value of its eigenvalues. The following elementary fact is well known.
Lemma A. 2 (Nonnegativity of inverse). Let $A$ be an $p \times p$ real matrix such that $\mathcal{R}(A)<1$ and all its entries are nonnegative. Then the matrix $I-A$ is invertible and all entries of $(I-A)^{-1}$ are nonnegative.

Proof of Lemma A.2. For completeness, we include a proof here. By Gelfand's formula, we have

$$
\lim _{k \rightarrow+\infty}\left\|A^{k}\right\|_{2}^{\frac{1}{k}}=\mathcal{R}(A)<1
$$

Thus there exist constants $\tau \in(0,1)$ and $k_{0} \in \mathbb{N}$ such that $\left\|A^{k}\right\|_{2}<\tau^{k}$ for every $k \geq k_{0}$. Therefore, the Neumann series

$$
\sum_{k=0}^{\infty} A^{k}=(I-A)^{-1}
$$

converges. This together with the assumption about the nonnegativity of $A$ implies that all entries of $(I-A)^{-1}$ are nonnegative.

## For Theorem 3.2

Proof of Theorem 3.2. We find

$$
\begin{align*}
F_{\mathbb{D}, \rho}(\beta) & =\max _{\Omega \in \mathcal{U}_{\mathbb{D}, \rho}(\widehat{\Omega})}\langle\Omega, V(\beta)\rangle  \tag{A.5a}\\
& \leq \max _{\Omega \in \mathcal{V}_{\mathrm{D}, \rho}(\widehat{\Omega})}\langle\Omega, V(\beta)\rangle  \tag{A.5b}\\
& =\inf _{\gamma I \succ \widehat{\Omega}^{\frac{1}{2}} V(\beta) \widehat{\Omega}^{\frac{1}{2}}} \gamma \rho-\gamma \log \operatorname{det}\left(I-\gamma^{-1} \widehat{\Omega}^{\frac{1}{2}} V(\beta) \widehat{\Omega}^{\frac{1}{2}}\right), \tag{A.5c}
\end{align*}
$$

where equality $A .5 \mathrm{a}$ is the definition of $F_{\mathrm{D}, \rho}$ as in 7 , inequality A.5b) follows from the fact that $\mathcal{U}_{\mathrm{D}, \rho}(\widehat{\Omega}) \subseteq \mathcal{V}_{\mathrm{D}, \rho}(\widehat{\Omega})$, and equality A .5 c follows from Proposition 3.4. Notice that $V(\beta)$ has one nonnegative eigenvalue by virtue of Lemma 3.5 and thus the constraint $\gamma I \succ \widehat{\Omega}^{\frac{1}{2}} V(\beta) \widehat{\Omega}^{\frac{1}{2}}$ implies the condition $\gamma>0$.
The strictly positive radius condition $\rho \in(0,+\infty)$ implies the existence of a Slater point $\widehat{\Omega}$ of the set $\mathcal{U}_{\mathbb{D}, \rho}(\widehat{\Omega})$. By [1] Proposition 5.5.4], the existence of a solution $\gamma^{\star}>0$ that minimizes the dual problem A.5c is guaranteed. Moreover, the objective function of problem A.5c is strictly convex, and thus the solution $\gamma^{\star}$ is unique. By inspecting (A.3), the solution

$$
\Omega^{\star}\left(\gamma^{\star}\right)=\left(\widehat{\Omega}^{-1}-\frac{1}{\gamma^{\star}} V(\beta)\right)^{-1}
$$

thus solves the primal problem A.5b by [1, pp. 178].
Assumption 2.1 implies that both $\widehat{\Omega}$ and $V(\beta) / \gamma^{\star}$ are nonnegative matrices. Also the spectral radius of $\left(\gamma^{\star}\right)^{-1} \widehat{\Omega}^{\frac{1}{2}} V(\beta) \widehat{\Omega}^{\frac{1}{2}}$ is smaller than 1 by the feasibility of $\gamma^{\star}$ in problem A.5c). Hence the matrix $\left[I-\left(\gamma^{\star}\right)^{-1} \widehat{\Omega}^{\frac{1}{2}} V(\beta) \widehat{\Omega}^{\frac{1}{2}}\right]^{-1}$ is nonnegative by Lemma A.2. As $\Omega^{\star}\left(\gamma^{\star}\right)=\widehat{\Omega}^{\frac{1}{2}}[I-$ $\left.\left(\gamma^{\star}\right)^{-1} \widehat{\Omega}^{\frac{1}{2}} V(\beta) \widehat{\Omega}^{\frac{1}{2}}\right]^{-1} \widehat{\Omega}^{\frac{1}{2}}$, we conclude that $\Omega^{\star}\left(\gamma^{\star}\right)$ is a matrix with nonnegative entries. Moreover, $\Omega^{\star}\left(\gamma^{\star}\right)$ is also positive semidefinite. Thus $\Omega^{\star}\left(\gamma^{\star}\right)$ is doubly nonnegative. This observation completes the proof.

## For Lemma 3.3

Proof of Lemma 3.3. Because $\mathcal{U}_{\mathbb{D}, \rho}(\widehat{\Omega})$ is compact, by [2, Proposition A.22], the subdifferential of $F_{\mathrm{D}, \rho}$ is

$$
\partial F_{\mathrm{D}, \rho}(\beta)=\mathrm{ConvexHull}\left(2 \sum_{i=1}^{N} \Omega_{0 i}^{\star} \nabla_{\beta} \ell\left(\beta, \widehat{x}_{i}, \widehat{y}_{i}\right) \mid \Omega^{\star} \in \mathcal{O}^{\star}(\beta)\right),
$$

where $\mathcal{O}^{\star}(\beta)=\left\{\Omega \in \mathcal{U}_{\mathbb{D}, \rho}(\widehat{\Omega}):\langle\Omega, V(\beta)\rangle=F_{\mathrm{D}, \rho}(\beta)\right\}$ is the optimal solution set. By Theorem 3.2 the optimal solution set $\mathcal{O}^{\star}(\beta)$ is a singleton, which leads to the postulated result.

## For Lemma 3.5

Proof of Lemma 3.5. The symmetry of $V(\beta)$ follows from definition. The nonnegativity of $V(\beta)$ follows from Assumption 2.1. Let $\lambda$ be an eigenvalue of $V(\beta)$, then $\lambda$ solves the characteristic equation $\operatorname{det}(V(\beta)-\lambda I)=0$. By exploiting the form of $V(\beta)$ and by the determinant formula for the arrowhead matrix, the eigenvalue $\lambda$ then solves the algebraic equation $\lambda^{N-1}\left[\lambda^{2}-\sum_{i=1}^{N} \ell\left(\beta, \widehat{x}_{i}, \widehat{y}_{i}\right)^{2}\right]=0$. This completes the proof.

## A. 2 Proofs of Section 4

## For Theorem 4.2

Proof of Theorem 4.2. The statement regarding $\gamma^{\star}$ and the expression of $F_{\mathrm{W}, \rho}$ and $\Omega^{\star}$ follows from [3, Proposition A.4]. It remains to show that $\Omega^{\star}$ is nonnegative. Indeed, the constraint $\gamma^{\star} I \succ V(\beta)$ implies that the spectral radius of $\left(\gamma^{\star}\right)^{-1} V(\beta)$ is smaller than 1 . Therefore, the inverse matrix $\left[I-\left(\gamma^{\star}\right)^{-1} V(\beta)\right]^{-1}$ is nonnegative according to Lemma A. 2 . The matrix $\Omega^{\star}$ is thus the product of three nonnegative matrices and thus it is also nonnegative.

## For Lemma 4.3

Proof of Lemma 4.3. The proof of this lemma is similar to that of Lemma 3.3 with the optimal set is now given by $\mathcal{O}^{\star}(\beta)=\left\{\Omega \in \mathcal{U}_{\mathrm{W}, \rho}(\widehat{\Omega}):\langle\Omega, V(\beta)\rangle=F_{\mathrm{W}, \rho}(\beta)\right\}$. The singleton of this $\mathcal{O}^{\star}(\beta)$ is guaranteed by Theorem 4.2

## A. 3 Discussions on Assumption 2.1.

Assumption 2.11(i) is standard in the machine learning literature and it holds naturally in many classification and regression tasks. Regarding Assumption 2.1||ii) the non-negativity of $\widehat{\Omega}$ follows directly if the weighting function $\omega$ is also non-negative. As shown below, positive definiteness holds under mild conditions of the training data and the weighting kernel.
Lemma A. 3 (Positive definite nominal weighting matrix). Suppose that the weighting function $\omega$ can be represented by a strictly positive definite kernel $K: \mathcal{Z} \times \mathcal{Z} \rightarrow \mathbb{R}$ in the sense that $\omega\left(\widehat{z}_{i}\right)=K\left(z_{0}, \widehat{z}_{i}\right)$ for $i=1, \ldots, N$ and for some distinct covariates $\left(z_{0}, \widehat{z}_{1}, \ldots, \widehat{z}_{N}\right) \in \mathcal{Z}^{N+1}$. Then

$$
\widehat{\Omega} \triangleq\left[\begin{array}{cccc}
K\left(z_{0}, z_{0}\right) & K\left(z_{0}, \widehat{z}_{1}\right) & \cdots & K\left(z_{0}, \widehat{z}_{N}\right) \\
K\left(\widehat{z}_{1}, z_{0}\right) & K\left(\widehat{z}_{1}, \widehat{z}_{1}\right) & \cdots & K\left(\widehat{z}_{1}, \widehat{z}_{N}\right) \\
\vdots & \vdots & \ddots & \vdots \\
K\left(\widehat{z}_{N}, z_{0}\right) & K\left(\widehat{z}_{N}, \widehat{z}_{1}\right) & \cdots & K\left(\widehat{z}_{N}, \widehat{z}_{N}\right)
\end{array}\right]
$$

is positive definite with $\widehat{\Omega}_{0 i}=\widehat{\Omega}_{i 0}=\omega\left(\widehat{z}_{i}\right)$ for $i=1, \ldots, N$.
The proof of Lemma A.3 follows by noticing that $\widehat{\Omega}$ is the Gram matrix of a kernel $K$, and hence $\widehat{\Omega}$ is a positive definite matrix [4, pp. 2392].
The next result shows that given any weight $\omega\left(\widehat{z}_{i}\right)$, there exists a matrix $\widehat{\Omega}$ satisfying Assumption 2.1(ii). Notice that it is always possible to choose the numbers $d_{k}$ satisfying the specified conditions below.

Lemma $\mathbf{A . 4}$ (Existence of $\widehat{\Omega})$. Given any nonnegative weights $\omega\left(\widehat{z}_{i}\right)$ for $i=1, \ldots, N$, let $A$ be a symmetric matrix with all entries zero except for those on the first row and first column, and on the diagonal as follows

$$
A \triangleq\left[\begin{array}{cccc}
d_{0} & \omega\left(\widehat{z}_{1}\right) & \cdots & \omega\left(\widehat{z}_{N}\right) \\
\omega\left(\widehat{z}_{1}\right) & d_{1} & \cdots & 0 \\
\vdots & \vdots & \ddots & \vdots \\
\omega\left(\widehat{z}_{N}\right) & 0 & \cdots & d_{N}
\end{array}\right]
$$

Then $A$ is positive definite and nonnegative if $d_{k}$ are chosen such that $d_{0}>0, d_{1} \Delta_{1}>\omega\left(\widehat{z}_{1}\right)^{2}$, and $d_{k} \Delta_{k}>d_{1} \ldots d_{k-1} \omega\left(\widehat{z}_{k}\right)^{2}$ for $k=2, \ldots, N+1$, where $\Delta_{k}$ denotes the $k$ th leading principal minor of the matrix $A$ which is independent of $d_{k}, \ldots, d_{N+1}$.

Proof of Lemma A.4. It is clear that $A$ is symmetric and nonnegative. The conditions on $d_{k}$ also ensures that $\Delta_{k}>0$ for every $k=1, \ldots, N+1$. Thus $A$ is positive definite as well.

For the given non-negative weights $\omega\left(\widehat{z}_{i}\right)$, Lemma A.4 shows that there are infinitely many doubly non-negative matrices $\widehat{\Omega}$ that satisfy the condition $\widehat{\Omega}_{0 i}=\widehat{\Omega}_{i 0}=\omega\left(\widehat{z}_{i}\right)$ for $i=1, \ldots, N$.

In practice, the choice of $\widehat{\Omega}$ can impact the performance of our robust estimate, and fine-tuning the elements of $\widehat{\Omega}$ may improve the predictive power. For the scope of this paper, we aim to improve the robustness of NW and LLR estimators and therefore will mainly focus on the scenario described in Lemma A. 3 where $\widehat{\Omega}$ is given by a strictly positive definite kernel.

## B Implementation Details

## B. 1 Gradient Information

We provide here the gradient information for the convex problems (8) and (12). This information can be exploited to derive fast numerical routines to find the optimal dual solution $\gamma^{\star}$.

Denote by $g$ the objective function of problem (8). The gradient of $g$ is

$$
\nabla g(\gamma)=\rho-\log \operatorname{det}\left(I-\gamma^{-1} \widehat{\Omega}^{\frac{1}{2}} V(\beta) \widehat{\Omega}^{\frac{1}{2}}\right)-\frac{1}{\gamma}\left\langle\left(I-\gamma^{-1} \widehat{\Omega}^{\frac{1}{2}} V(\beta) \widehat{\Omega}^{\frac{1}{2}}\right)^{-1}, \widehat{\Omega}^{\frac{1}{2}} V(\beta) \widehat{\Omega}^{\frac{1}{2}}\right\rangle
$$

Let $h$ be the objective function of problem (12). The gradient of $h$ is

$$
\nabla h(\gamma)=\rho-\left\langle\widehat{\Omega},\left(I-\gamma^{\star}\left[\gamma^{\star} I-V(\beta)\right]^{-1}\right)^{2}\right\rangle
$$

## B. 2 Implementation

Following Lemma 3.5, for a fixed value of $\beta$, we can rewrite the low-rank matrix $V(\beta)$ using the eigenvalue decomposition $V(\beta)=Q \Lambda Q^{\top}$, where $\Lambda \in \mathbb{R}^{2 \times 2}$ is a diagonal matrix and $Q \in \mathbb{R}^{(N+1) \times 2}$ is an orthonormal matrix. Therefore, we can leverage the Woodbury matrix identity to implement the inverse in gradient computation (e.g., $\nabla g$ and $\nabla h$ ) efficiently. For examples, (i)

$$
(\gamma I-V(\beta))^{-1}=\gamma^{-1} I-\gamma^{-2} Q\left(-\Lambda^{-1}+\gamma^{-1} I_{2}\right)^{-1} Q^{\top}
$$

where we have exploited the fact that $Q^{\top} Q=I_{2}$ with $I_{2}$ being the $2 \times 2$ identity matrix; and (ii)

$$
\left(I-\gamma^{-1} \widehat{\Omega}^{\frac{1}{2}} V(\beta) \widehat{\Omega}^{\frac{1}{2}}\right)^{-1}=I-\widehat{\Omega}^{\frac{1}{2}} Q\left(-\gamma \Lambda^{-1}+Q^{\top} \widehat{\Omega} Q\right)^{-1} Q^{\top} \widehat{\Omega}^{\frac{1}{2}}
$$

Therefore, the inverse in gradient (e.g., $\nabla g$ and $\nabla h$ ) is computed on $2 \times 2$ matrices.

## C Additional Empirical Results

Details of datasets. Table A.1 lists the detailed statistical characteristics of the datasets used in our experiments.

Table A.1: Statistical characteristics of the datasets.

|  | Abalone | Bank | CPU | Elevators | KIN | POL | Puma | Slice |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| \#samples | 4177 | 8192 | 8192 | 16599 | 40000 | 15000 | 8192 | 53500 |
| \#features | 7 | 32 | 21 | 17 | 8 | 26 | 32 | 384 |

Further results for the ideal case: no sample perturbation. We illustrate further empirical results for all 8 datasets and with different nearest neighbor size $N$ (e.g., similar to Figure 1 with the nearest neighbor size $N=50$ ).

- For $N=50$, we illustrate results in Figure A. 1
- For $N=30$, we illustrate results in Figure A. 2
- For $N=20$, we illustrate results in Figure A. 3
- For $N=10$, we illustrate results in Figure A. 4


Figure A.1: Average RMSE for ideal case with no perturbation when $N=50$ for all 8 datasets.


Figure A.2: Average RMSE for ideal case with no perturbation when $N=30$ for all 8 datasets.


Figure A.3: Average RMSE for ideal case with no perturbation when $N=20$ for all 8 datasets.


Figure A.4: Average RMSE for ideal case with no perturbation when $N=10$ for all 8 datasets.

Further results for the cases when training samples are perturbed. We illustrate further empirical results for all 8 datasets with different nearest neighbor size $N$ (e.g., similar to Figure 2 where the nearest neighborr size $N=50$ ).

- For $N=50$, we illustrate results in Figure A. 5
- For $N=30$, we illustrate results in Figure A. 6
- For $N=20$, we illustrate results in Figure A. 7
- For $N=10$, we illustrate results in Figure A. 8


Figure A.5: RMSE for varying perturbation levels $\tau$ where $N=50$.


Figure A.6: RMSE for varying perturbation levels $\tau$ where $N=30$.


Figure A.7: RMSE for varying perturbation levels $\tau$ where $N=20$.


Figure A.8: RMSE for varying perturbation levels $\tau$ where $N=10$.

Further results for the effects of the uncertainty size $\rho$. We illustrate further empirical results for different nearest neighbor size $N$ in all 8 datasets (e.g., similar to Figure 3 where the nearest neighbor size $N=50$ and the perturbation $\tau=N$ ). Note that when $\rho=0$, the reweighting schemes are equivalent to the vanilla NW estimator.

- For $N=50$, we illustrate results for NW-LogDet and NW-BuresW in Figure A. 9 and Figure A. 10 respectively.
- For $N=30$, we illustrate results for NW-LogDet and NW-BuresW in Figure A. 11 and Figure A. 12 respectively.
- For $N=20$, we illustrate results for NW-LogDet and NW-BuresW in Figure A. 13 and Figure A. 14 respectively.
- For $N=10$, we illustrate results for NW-LogDet and NW-BuresW in Figure A. 15 and Figure A. 16 respectively.





$$
\begin{array}{|lll|}
\hline \\
\hline
\end{array} \rho=0.01 \quad \rho=0.1 \quad \ddagger \rho=1 \quad \rho=10
$$

Figure A.9: Effects of the uncertainty size $\rho$ on RMSE for NW-LogDet when $N=50$.


Figure A.10: Effects of the uncertainty size $\rho$ on RMSE for NW-BuresW when $N=50$.





\[

\]



Figure A.11: Effects of the uncertainty size $\rho$ on RMSE for NW-LogDet when $N=30$.







$$
\begin{array}{|ccc|}
\hline
\end{array} \rho=0.01 \quad \rho=0.1 \quad \rho=1 \quad \rho=10
$$

Figure A.12: Effects of the uncertainty size $\rho$ on RMSE for NW-BuresW when $N=30$.


Figure A.13: Effects of the uncertainty size $\rho$ on RMSE for NW-LogDet when $N=20$.







\[

\]

Figure A.14: Effects of the uncertainty size $\rho$ on RMSE for NW-BuresW when $N=20$.


Figure A.15: Effects of the uncertainty size $\rho$ on RMSE for NW-LogDet when $N=10$.


Figure A.16: Effects of the uncertainty size $\rho$ on RMSE for NW-BuresW when $N=10$.

Further results for the effects of the nearest neighbor size $N$. We illustrate further empirical results for different uncertainty size $\rho$ in all 8 datasets (e.g., similar to Figure 4 where $\rho=0.1$ in the KIN dataset).

- For $\rho=10$, we illustrate results for NW-LogDet and NW-BuresW in Figure A. 17 and Figure A. 18 respectively.
- For $\rho=1$, we illustrate results for NW-LogDet and NW-BuresW in Figure A. 19 and Figure A. 20 respectively.
- For $\rho=0.1$, we illustrate results for NW-LogDet and NW-BuresW in Figure A. 21 and Figure A. 22 respectively.
- For $\rho=0.01$, we illustrate results for NW-LogDet and NW-BuresW in Figure A. 23 and Figure A. 24 respectively.


Figure A.17: Effects of the nearest neighbor size $N$ on RMSE for NW-LogDet when $\rho=10$.


Figure A.18: Effects of the nearest neighbor size $N$ on RMSE for NW-BuresW when $\rho=10$.


Figure A.19: Effects of the nearest neighbor size $N$ on RMSE for NW-LogDet when $\rho=1$.


Figure A.20: Effects of the nearest neighbor size $N$ on RMSE for NW-BuresW when $\rho=1$.


Figure A.21: Effects of the nearest neighbor size $N$ on RMSE for NW-LogDet when $\rho=0.1$.


Figure A.22: Effects of the nearest neighbor size $N$ on RMSE for NW-BuresW when $\rho=0.1$.


Figure A.23: Effects of the nearest neighbor size $N$ on RMSE for NW-LogDet when $\rho=0.01$.


Figure A.24: Effects of the nearest neighbor size $N$ on RMSE for NW-BuresW when $\rho=0.01$.
Further results for varying shifting w.r.t. $\kappa$. We first illustrate corresponding results for NWBuresW in Figure A.25, similar as results in Figure 5 for NW-LogDet.

We next illustrate further empirical results for different uncertainty size $\rho$ and different perturbation $\tau$ in the KIN dataset (e.g., similar to Figure 5 and Figure A. 25 where $\rho=0.1$ for the left plots and $\tau=0.2 N$ for the right plots).

- For NW-LogDet, we illustrate results on effects of perturbation intensity $\kappa$ for different perturbation proportion $\tau$ in Figure A. 26 and Figure A. 27 respectively.
- For NW-BuresW, we illustrate results on effects of perturbation intensity $\kappa$ for different perturbation proportion $\tau$ in Figure A. 28 and Figure A. 29 respectively.


Figure A.25: Effects of perturbation intensity $\kappa$ for NW-BuresW estimate. Left plot: different perturbation $\tau$, right plot: different uncertainty size $\rho$.


Figure A.26: Effects of perturbation intensity $\kappa$ for NW-LogDet estimate for different perturbation proportion $\tau$.


Figure A.27: Effects of perturbation intensity $\kappa$ for NW-LogDet estimate for different uncertainty size $\rho$.


Figure A.28: Effects of perturbation intensity $\kappa$ for NW-BuresW estimate for different perturbation proportion $\tau$.


Figure A.29: Effects of perturbation intensity $\kappa$ for NW-BuresW estimate for different uncertainty size $\rho$.

Further results for effects of $N$ on computational time. We illustrate further effects of $N$ on computational time in all 8 datasets (e.g., similar to Figure 6 for the KIN dataset).

- For NW-LogDet, we illustrate further effects of $N$ on computational time in all 8 datasets in Figure A. 30
- For NW-BuresW, , we illustrate further effects of $N$ on computational time in all 8 datasets in Figure A.31.


Figure A.30: Effects of $N$ on computational time for NW-LogDet.


Figure A.31: Effects of $N$ on computational time for NW-BuresW.

## References

[1] Dimitri Bertsekas. Convex Optimization Theory. Athena Scientific, 2009.
[2] Dimitri P. Bertsekas. Control of Uncertain Systems with a Set-Membership Description of Uncertainty. PhD thesis, Massachusetts Institute of Technology, 1971.
[3] Viet Anh Nguyen, S. Shafieezadeh-Abadeh, D. Kuhn, and P. Mohajerin Esfahani. Bridging Bayesian and minimax mean square error estimation via Wasserstein distributionally robust optimization. arXiv preprint arxiv:1910.10583, 2019.
[4] Bharath K Sriperumbudur, Kenji Fukumizu, and Gert RG Lanckriet. Universality, characteristic kernels and RKHS embedding of measures. Journal of Machine Learning Research, 12(7), 2011.

