Minimaxity in Robust Nonparametric Regression: from Heavy-Tailed to Sub-Gaussian Assumptions

Yikun Li Matey Neykov

Department of Statistics and Data Science, Northwestern University

yikunli2028@u.northwestern.edu mneykov@northwestern.edu

Abstract

We study nonparametric regression over star-shaped function classes under strong adversarial contamination. The clean (original) observations $(\widetilde{X}_i, \widetilde{Y}_i)_{i=1}^n$ are i.i.d. and satisfy $\widetilde{Y}_i = f^*(\widetilde{X}_i) + \xi_i$, where $f^* \in \mathcal{F}$, $\widetilde{X}$ is distributed as $\mathbb{P}_X$, and the noise is conditionally centered. An adversary may inspect the clean sample and replace at most $\lfloor \epsilon n \rfloor$ observation pairs. We assume that $\mathcal{F} \subset L_2(\mathbb{P}_X)$ is star-shaped with finite diameter d.

Under a Bernstein-type moment condition on function increments and a bounded conditional second moment for the noise, we establish the matching lower and upper bounds of order $\left((\eta^*)^2 \vee \epsilon\right) \wedge d^2$, uniformly over admissible star-shaped function classes, therefore characterizing the minimax rate of estimation error, where

$$\eta^* := \sup\left\{\eta > 0 : n\eta^2 \le \log \mathcal{M}_{\mathcal{F}}^{\mathrm{loc}}(\eta, c)\right\}.$$

Here $\log \mathcal{M}_{\mathcal{F}}^{\mathrm{loc}}(\eta, c)$ denotes the local metric entropy of $\mathcal{F}$ in $L_2(\mathbb{P}_X)$, and c is a sufficiently large constant. Under a uniform relative sub-Gaussian condition on function increments and noise, the lower bound is adjusted to $\left((\eta^*)^2 \vee \epsilon^2\right) \wedge d^2$ while the upper bound improves to $\left((\eta^*)^2 \vee \epsilon^2 \log^2\left(\frac{e}{\epsilon}\right)\right) \wedge d^2$, yielding minimax optimality up to logarithmic factors.

Contents

1 Introduction **3**
- 1.1 Problem Formulation . . . 3
- 1.2 Related Literature . . . 5
- 1.3 Our Contributions . . . 7
- 1.4 Organization . . . 7
- 1.5 Notation and Definitions . . . 8

2 General Lower Bound **9**

3 Minimax rate under Bernstein-Type Function Class and Heavy-Tailed Noise **10**
- 3.1 Lower Bound . . . 12
- 3.2 Upper Bound . . . 13

4 Minimax Rate under Sub-Gaussian Assumptions **22**
- 4.1 Lower Bound . . . 24
- 4.2 Upper Bound . . . 24

5 Examples **26**
5.1 Sparse linear regression . . . 26
5.2 Low-rank trace regression . . . 28
5.3 Hölder smooth regression . . . 30

6 Discussion **34**
6.1 Summary . . . 34
6.2 Future Research Directions . . . 34

A Proof of Section 2 **38**
A.1 Proof of Lemma 2.2 . . . 38
A.2 Proof of Lemma 2.3 . . . 38

B Proof of Section 3 **39**
B.1 Proof of Lemma 3.1 . . . 39
B.2 Proof of Lemma 3.7 . . . 41
B.3 Proof of Corollary 3.8 . . . 42
B.4 Proof of Lemma 3.9 . . . 43
B.5 Proof of Corollary 3.10 . . . 43
B.6 Proof of Theorem 3.11 . . . 44
B.7 Proof of Corollary 3.12 . . . 46

C Proof of Section 4 **47**
C.1 Proof of Lemma 4.2 . . . 47
C.2 Proof of Lemma 4.3 . . . 48
C.3 Proof of Lemma 4.4 . . . 49

1 Introduction

1.1 Problem Formulation

Let $\mathcal{X} \subseteq \mathbb{R}^p$ be a measurable set, and let $\widetilde{X}_1, \ldots, \widetilde{X}_n$ be independent and identically distributed according to a probability distribution $\mathbb{P}_X$ on $\mathcal{X}$. The uncontaminated responses satisfy

$$\widetilde{Y}_i = f^*(\widetilde{X}_i) + \widetilde{\xi}_i, \qquad i = 1, \ldots, n, \tag{1.1}$$

where $f^* \in \mathcal{F}$ and $\mathcal{F}$ is a given star-shaped class of measurable functions from $\mathcal{X}$ to $\mathbb{R}$ (see Definition 1.1). We assume that $\mathcal{F} \subseteq L_2(\mathbb{P}_X)$ and that $\mathcal{F}$ has finite $L_2(\mathbb{P}_X)$-diameter (note that this does not necessarily imply that $\mathcal{F}$ is bounded in sup-norm). The noise-covariate pairs $(\widetilde{\xi}_i, \widetilde{X}_i), i = 1, 2, \ldots, n$ are independent, identically distributed, and have a common joint distribution $\mathbb{P}_{\xi,X}$. The noise $\widetilde{\xi}$ is allowed to depend on the covariate $\widetilde{X}$, but it always satisfies that $\mathbb{E}\left[\widetilde{\xi} \,\middle|\, \widetilde{X}\right] = 0$. We denote the family of such distributions as $\mathcal{P}_{\xi,X}$ and therefore $\mathbb{P}_{\xi,X} \in \mathcal{P}_{\xi,X}$. Such a setting includes the scenario that $\widetilde{\xi}$ is independent of $\widetilde{X}$ and $\mathbb{E}\widetilde{\xi} = 0$. In the following argument, we assume $\mathbb{P}_X$ is an arbitrary fixed distribution and write $\mathbb{P}_{\xi,X}$ and $\mathcal{P}_{\xi,X}$ as $\mathbb{P}_\xi$ and $\mathcal{P}_\xi$ for simplicity. The moment and tail conditions defining $\mathcal{P}_\xi$ will be specified for each result.

Write $\widetilde{D}_n := \left((\widetilde{X}_i, \widetilde{Y}_i)\right)_{i=1}^n$ for the uncontaminated sample. Before the sample is observed, an adaptive adversary may replace at most $\lfloor \epsilon n \rfloor$ observations by arbitrary elements of $\mathcal{X} \times \mathbb{R} \times \mathbb{R}$, where $\epsilon \in [0, 1/2)$ is a known upper bound on the contamination proportion. The adversary may use the entire uncontaminated sample, the distributions $\mathbb{P}_X$ and $\mathbb{P}_\xi$, and the regression function f^* when selecting the observations to replace and their replacement values.

More precisely, let $\mathfrak{C}_{n,\epsilon}$ denote the collection of admissible measurable, possibly randomized, contamination maps. For each $\mathcal{C} \in \mathfrak{C}_{n,\epsilon}$, the observed sample $D_n := ((X_i, Y_i))_{i=1}^n = \mathcal{C}(\widetilde{D}_n)$ satisfies

$$\sum_{i=1}^n \mathbf{1}_{\{(X_i,Y_i) \neq (\widetilde{X}_i, \widetilde{Y}_i)\}} \leq \lfloor \epsilon n \rfloor$$

almost surely. Thus, all covariates, responses, and the noise may be corrupted. This contamination mechanism is known as the *strong contamination model* Diakonikolas and Kane [2023], Oliveira and Resende [2025].

For $f, g \in L_2(\mathbb{P}_X)$, define

$$\|f - g\|_{L_2(\mathbb{P}_X)} := \left(\int_{\mathcal{X}} [f(x) - g(x)]^2 \, d\mathbb{P}_X(x) \right)^{1/2}. \tag{1.2}$$

The corresponding norm is $\|f\|_{L_2(\mathbb{P}_X)} := \|f - 0\|_{L_2(\mathbb{P}_X)}$. The expression in (1.2) is a pseudometric on functions and a metric modulo $\mathbb{P}_X$-almost-sure equality. Throughout the paper, estimation accuracy is measured by the squared $L_2(\mathbb{P}_X)$ loss.

Remark. We do not assume the full knowledge of $\mathbb{P}_X$ to exactly evaluate $L_2(\mathbb{P}_X)$-norm of $f \in \mathcal{F}$, as such exact evaluation is unnecessary. For the algorithms and results in this paper to hold, we only require an oracle to access the value of an equivalent norm to the $L_2(\mathbb{P}_X)$-norm. Here by equivalent norm we mean a norm $\|\cdot\|$, which agrees with $\|\cdot\|_{L_2(\mathbb{P}_X)}$ up to universal constants that are independent of $n, \mathcal{F}, \epsilon$. A common example is the ordinary linear regression where $\widetilde{Y}_i =$

$\widetilde{X}_i^\top \beta + \widetilde{\xi}_i$, $\widetilde{X}_i$ are independently, identically distributed with $\mathbb{E}\widetilde{X}_i = 0$ and $\mathrm{Cov}\left(\widetilde{X}_i\right) = \mathbf{\Sigma}$. In this scenario, having an oracle to $L_2(\mathbb{P}_X)$-norm is equivalent to knowing the covariance $\mathbf{\Sigma}$. Under the common assumption that we have nonzero lower and upper bounds κ, κ^{-1} such that $\kappa \leq \sigma_{\min}(\mathbf{\Sigma}) \leq \sigma_{\max}(\mathbf{\Sigma}) \leq \kappa^{-1}$ the $L_2(\mathbb{P}_X)$ norm is equivalent to the standard Euclidean norm. The full knowledge of $\mathbf{\Sigma}$ is hence unnecessary. Nonetheless, for notational simplicity, we assume for ease of exposition that $\mathbb{P}_X$ is known in the presentations of our results and proofs. They can be extended to the aforementioned scenario with minimal efforts.

Given the contaminated sample D_n, we seek to estimate f^* using a measurable estimator $\widehat{f}$ taking values in $\mathcal{F}$. For a specified noise family $\mathcal{P}_\xi$, define the minimax risk by

$$\mathcal{R}_{n,\epsilon}(\mathcal{F}; \mathcal{P}_\xi) := \inf_{\widehat{f}} \sup_{f^* \in \mathcal{F}} \sup_{\mathbb{P}_\xi \in \mathcal{P}_\xi} \sup_{\mathcal{C} \in \mathfrak{C}_{n,\epsilon}} \mathbb{E} \left\| \widehat{f}\left(\mathcal{C}(\widetilde{D}_n)\right) - f^* \right\|_{L_2(\mathbb{P}_X)}^2 . \tag{1.3}$$

The expectation is taken over the uncontaminated sample and any randomization used by the estimator or the adversary. The infimum ranges over estimators that may depend on the known quantities $\mathcal{F}$, $\mathbb{P}_X$, ϵ, and $\mathcal{P}_\xi$, but not on the unknown regression function, noise distribution, or contamination mechanism. Dependence on the fixed design distribution $\mathbb{P}_X$ is suppressed in the notation.

Structural constraints on a regression function provide a way to incorporate prior information into estimation. Sparsity, smoothness, monotonicity, and bounds on a reproducing kernel Hilbert space norm are examples of such constraints. The corresponding function classes are often star-shaped, including both convex classes and nonconvex classes arising from sparsity restrictions. At the same time, heavy-tailed noise and corrupted observations motivate the development of estimators whose guarantees remain valid beyond an uncontaminated model.

Sparse linear regression provides a motivating example of a nonconvex star-shaped class. For a sparsity level $1 \leq s < p$ and a radius $R > 0$, consider

$$\mathcal{F}_{s,R} := \left\{ f_\beta : x \mapsto x^\top \beta : \beta \in \mathbb{R}^p,\ \|\beta\|_0 \leq s,\ \|\beta\|_2 \leq R \right\}.$$

This class is star-shaped about zero because multiplication by a scalar in $[0, 1]$ preserves both constraints, although the class is not convex. Under a standard sub-Gaussian design ($\mathrm{Cov}\,(x) = \mathbf{I}_p$), $\|f_\beta - f_\gamma\|_{L_2(\mathbb{P}_X)} = \|\beta - \gamma\|_2$, so the coefficient constraint also ensures a finite $L_2\,(\mathbb{P}_X)$-diameter. Liu et al. [2020] study sparse linear regression under Gaussian designs when an adversary may arbitrarily corrupt a fraction of the covariate–response pairs. This example motivates allowing nonconvex structural constraints together with corruption of entire observations.

Smooth nonparametric regression provides another example. Hölder norm balls on a compact domain are convex and contain the zero function, and are therefore star-shaped. Such classes describe response surfaces whose variation is controlled by a smoothness constraint. Du et al. [2018] study robust nonparametric regression over Hölder and Sobolev classes under Huber's contamination model with a fixed-grid design. Their results illustrate the interaction between smoothness and contamination in determining estimation accuracy.

Shape-constrained regression provides a complementary setting in which prior information is expressed through qualitative restrictions. For example, the class of coordinatewise nondecreasing functions from $[0, 1]^p$ to $[-B, B]$, where $B > 0$, is convex and contains zero, hence is star-shaped.

Han et al. [2019] establish risk bounds for multivariate isotonic regression under both fixed and random designs in the uncontaminated setting. Studying such classes under contamination allows one to investigate how monotonicity constraints interact with the presence of unreliable observations.

Kernel regression supplies a further family of examples. Let $\mathcal{H}$ be a reproducing kernel Hilbert space with kernel k. The norm ball $\mathcal{F}_{\mathcal{H},R} := \{f \in \mathcal{H} : \|f\|_{\mathcal{H}} \leq R\}$ is convex and star-shaped about zero. If $\sup_{x\in\mathcal{X}} k(x,x) \leq \kappa^2$, the reproducing property gives $\sup_{x\in\mathcal{X}} |f(x)| \leq \kappa R$ for every $f \in \mathcal{F}_{\mathcal{H},R}$. Liu and Lian [2026] study conditional mean estimation in reproducing kernel Hilbert spaces with heavy-tailed errors and arbitrary contamination of responses, demonstrating the relevance of robustness to both distributional tails and corrupted data in nonlinear regression.

These examples motivate a unified study of regression over star-shaped function classes under explicit distributional assumptions. We consider a design distribution and a strong contamination mechanism under which an adversary may inspect the entire uncontaminated sample before replacing a prescribed fraction of observations. Our objective is to understand how the geometry of the function class, the noise assumptions, and the contamination proportion determine achievable squared $L_2(\mathbb{P}_X)$ risk bounds, and to establish their sharpness for each admissible star-shaped classes.

1.2 Related Literature

Our work connects the literature on local-entropy characterizations of minimax risk with robust regression under heavy-tailed observations and adversarial contamination.

Local entropy and constrained estimation. The information-theoretic bounds in Yang and Barron [1999] relate nonparametric minimax rates to metric entropy under suitable regularity conditions, with applications including regression. These results provide a classical foundation for using the complexity of a function class to characterize estimation difficulty. For convex sub-Gaussian function classes, Mendelson [2017] develops regression bounds governed by local covering parameters under appropriate assumptions on the noise. This work demonstrates the importance of local geometry in determining attainable prediction accuracy and motivates our use of local packing entropy. The minimax risk in the Gaussian sequence model under bounded convex constraints is characterized in Neykov [2023] through a critical radius defined by local metric entropy. Its multistage packing construction provides a geometric precursor to subsequent estimation procedures for constrained nonparametric regression. Under Gaussian noise and suitable boundedness or increment conditions, Prasadan and Neykov [2025] characterize the minimax squared $L_2(\mathbb{P}_X)$-risk over bounded star-shaped function classes through local entropy. This result provides the uncontaminated benchmark for the present problem, which additionally considers finite-variance noise and adaptive replacement of sample pairs.

Regression under heavy tails and contamination. The decision-theoretic framework in Chen et al. [2016] develops robust testing procedures and minimax lower bounds based on a total-variation modulus of continuity under Huber's contamination model. This framework explains why the contamination contribution depends on the statistical model and the loss, a consideration that also guides the formulation of lower bounds in our setting. The median-of-means tournaments in Lugosi and Mendelson [2020] attain optimal accuracy–confidence tradeoffs for regression from independent, uncontaminated observations under appropriate moment and norm-equivalence assumptions. Their use of robust pairwise comparisons provides a methodological foundation for estimation with

heavy-tailed observations. The median-of-means procedures in Lecué and Lerasle [2020] provide regression and regularization guarantees under weak moment assumptions in the presence of arbitrary outliers. Their bounds involve localized quadratic and multiplier complexities together with an outlier-dependent term, providing a useful comparison for our entropy-based analysis. Specializing the moment-dependent regression bound in Oliveira and Resende [2025] yields squared-error contamination terms of order ϵ under bounded second multiplier moments and $\epsilon^2 \log^2(e/\epsilon)$ under sub-exponential multiplier tails. Its quadratic-risk analysis treats convex classes satisfying a small-ball condition and uses localized empirical-process complexities, making the local-entropy formulation for star-shaped classes with known $\mathbb{P}_X$ a central point of comparison with our work.

Robust estimation under star-shaped constraints. For robust mean estimation under star-shaped constraints, Prasadan and Neykov [2026] combine local entropy with a squared-error contamination term of order ϵ^2 for Gaussian noise and $\epsilon^2 \log(1/\epsilon)$ for general unknown sub-Gaussian noise, with the noise scale held fixed. These results provide a closely related geometric characterization in a vector location model and motivate the corresponding investigation for random-design regression. Under a bounded covariance assumption and an additional sample-size condition involving the supremum of local entropy, Peng et al. [2026] characterize robust mean estimation over star-shaped constraints with a squared-error contamination term linear in ϵ. This heavy-tailed extension provides a useful comparison for how moment assumptions affect the contamination penalty, although its vector location model differs from the regression model considered here.

Structured regression models. Under fixed-design regression with Huber contamination, the local-median procedures in Du et al. [2018] achieve minimax rates over Hölder and Sobolev classes. These smoothness classes provide concrete nonparametric examples motivating a general theory for star-shaped regression classes. Regression-depth estimators in Gao [2020] attain minimax rates under Huber contamination for several regression models, including sparse linear and nonparametric regression. These results provide model-specific robust benchmarks for our study of general star-shaped classes under strong contamination. Computationally efficient procedures for sparse linear regression under adversarial corruption of both covariates and responses are developed in Liu et al. [2020], with guarantees for Gaussian designs. Since bounded sparse parameter classes are star-shaped and generally nonconvex, this setting illustrates the relevance of allowing star-shaped constraints. Covariate filtering in Pensia et al. [2025] yields computationally efficient robust linear-regression procedures under heavy-tailed covariates, heavy-tailed errors, and adversarial contamination. This work complements our geometric perspective by emphasizing tractable estimation within a specific regression family. The kernel-based procedures in Liu and Lian [2026] combine adaptive Huber loss with reproducing kernel Hilbert space methods to obtain estimation guarantees under heavy-tailed noise and response contamination. Kernel norm balls are convex and hence star-shaped, providing another nonparametric motivation for the class geometry considered here. The present work studies regression over bounded star-shaped classes under strong contamination and a known design distribution $\mathbb{P}_X$. Our focus is on local-entropy risk bounds under Bernstein-type increment conditions and finite-variance noise, together with sharper contamination dependence under relative sub-Gaussian increment assumptions and sub-Gaussian noise. In the heavy-tailed regime, the upper bound applies to every admissible class, while sharpness of the contamination term is understood for each star-shaped classes satisfying the assumptions.

1.3 Our Contributions

Building on the local-entropy characterization of uncontaminated regression in Prasadan and Neykov [2025], we investigate how strong contamination and distributional assumptions affect estimation over bounded star-shaped function classes. Applying the generalized Fano inequality to an admissible Gaussian noise submodel yields a lower bound for the uncontaminated problem, which remains valid under strong contamination because the adversary may leave the sample unchanged. This lower bound serves as a benchmark for the following contributions, with implicit constants depending on the moment and tail parameters specified in our assumptions.

- **Robust estimation under finite-variance noise.** Under a Bernstein-type moment condition on the increments of $\mathcal{F}$ and a bounded conditional second moment assumption on ξ, we construct the tree-based procedures in Algorithms 1 and 2 and establish the upper bound

$$\mathcal{R}_{n,\epsilon}(\mathcal{F}; \mathcal{P}_\xi) \lesssim \left((\eta^*)^2 \vee \epsilon\right) \wedge d^2.$$

  This bound holds for every star-shaped function class satisfying the stated assumptions. We further prove that the linear dependence on ϵ cannot be improved for each star-shaped class $\mathcal{F}$ uniformly over the admissible noise, therefore establishing the optimality of the contamination term.

- **Sharper bounds under sub-Gaussian assumptions.** Under relative sub-Gaussian increment and boundedness conditions on $\mathcal{F}$ and a sub-Gaussian assumption on ξ, we modify the procedures above to obtain the sharper upper bound

$$\mathcal{R}_{n,\epsilon}(\mathcal{F}; \mathcal{P}_\xi) \lesssim \left((\eta^*)^2 \vee \epsilon^2 \log^2(e/\epsilon)\right) \wedge d^2.$$

  We complement this result with the lower bound

$$\mathcal{R}_{n,\epsilon}(\mathcal{F}; \mathcal{P}_\xi) \gtrsim \left((\eta^*)^2 \vee \epsilon^2\right) \wedge d^2,$$

  thereby characterizing the minimax risk up to a logarithmic factor in the contamination term.

1.4 Organization

This paper is structured as follows. Section 2 establishes a general minimax lower bound in terms of the local metric entropy of the function class. Section 3 studies function classes satisfying a Bernstein-type condition with noise having bounded conditional second moments. We construct a robust tree-based estimator and establish matching minimax upper and lower bounds under strong adversarial corruption. Section 4 considers sub-Gaussian function increments and noise, obtaining a sharper dependence on the corruption proportion with upper and lower bounds that agree up to a logarithmic factor. Section 5 illustrates our results through sparse linear regression, low-rank trace regression, and Hölder smooth regression, including explicit calculations of their local entropy and resulting minimax rates. Section 6 summarizes the paper and introduces some possible future research avenues beyond this work. The proofs of the lemmas, theorems, and corollaries in Sections 2, 3, and 4 are deferred to Appendices A, B, and C, respectively.

1.5 Notation and Definitions

We introduce the notation and definitions used throughout this paper. Let n denote the sample size and p the dimension of the covariates $\widetilde{X}_i$. The parameter ϵ denotes the maximum fraction of observations that the adversary $\mathcal{C}$ may replace. We define $\widetilde{X} := (\widetilde{X}_1, \ldots, \widetilde{X}_n)$, $\widetilde{Y} := (\widetilde{Y}_1, \ldots, \widetilde{Y}_n)$, $\widetilde{\xi} = (\widetilde{\xi}_1, \ldots, \widetilde{\xi}_n)$ and $(X, Y) = \mathcal{C}(\widetilde{X}, \widetilde{Y})$. For a positive integer n, write $[n] := \{1, 2, \ldots, n\}$. For real numbers a and b, define $a \wedge b := \min\{a, b\}$ and $a \vee b := \max\{a, b\}$. For nonnegative quantities a and b, the notation $a \lesssim_{c,\ldots} b$ means that $a \leq Kb$ for some constant $K > 0$ depending only on the quantities indicated in the subscript. We write $a \gtrsim_{c,\ldots} b$ if $b \lesssim_{c,\ldots} a$.

For $f, g \in \mathcal{F}$, let $\|f - g\|_{L_2(\mathbb{P}_X)}$ denote their $L_2(\mathbb{P}_X)$-distance. Let $\mathbb{P}_n$ denote the empirical distribution of the observed covariates $X_1, \ldots, X_n$. The corresponding empirical L_2-distance is defined by

$$\|f - g\|_{L_2(\mathbb{P}_n)} := \left(\frac{1}{n} \sum_{i=1}^{n} \left[f(X_i) - g(X_i) \right]^2 \right)^{1/2}.$$

For $f \in L_2(\mathbb{P}_X)$ and $\delta \geq 0$, let $B_2(f, \delta)$ denote the closed $L_2(\mathbb{P}_X)$-ball centered at f with radius δ:

$$B_2(f, \delta) := \left\{ g \in L_2(\mathbb{P}_X) : \|f - g\|_{L_2(\mathbb{P}_X)} \leq \delta \right\}.$$

We denote the $L_2(\mathbb{P}_X)$-diameter of $\mathcal{F}$ by

$$d := \operatorname{diam}(\mathcal{F}) := \sup_{f,g \in \mathcal{F}} \|f - g\|_{L_2(\mathbb{P}_X)}.$$

For two probability distributions $\mathbb{P}$ and $\mathbb{Q}$, we use $\mathrm{KL}(\mathbb{P} \,\|\, \mathbb{Q}) := \mathbb{E}_{\mathbb{P}} \log\left(\frac{d\mathbb{P}}{d\mathbb{Q}}\right)$ to represent the Kullback-Leibler (KL) divergence and $\mathrm{TV}(\mathbb{P}, \mathbb{Q}) := \sup_{A \in \mathscr{F}} |\mathbb{P}(A) - \mathbb{Q}(A)|$ the total variation distance between $\mathbb{P}$ and $\mathbb{Q}$, where $\mathscr{F}$ is the corresponding sigma-algebra.

The following definitions of star-shapedness, packing numbers, and local entropy are fundamental throughout the analysis.

Definition 1.1 (Star-shaped set)**.** *A function class $\mathcal{F}$ is star-shaped with respect to $f_c \in \mathcal{F}$ if the line segment joining f_c to every $f \in \mathcal{F}$ is contained in $\mathcal{F}$; that is,*

$$(1 - t) f_c + t f \in \mathcal{F}, \qquad \forall f \in \mathcal{F},\ \forall t \in [0, 1].$$

Any such f_c is called a star center of $\mathcal{F}$. The class $\mathcal{F}$ is called star-shaped if it has at least one star center.

Definition 1.2 (δ-packing and packing number)**.** *Let $\delta > 0$. A finite subset $\{f_1, \ldots, f_M\} \subseteq \mathcal{F}$ is a δ-packing of $\mathcal{F}$ with respect to a norm or seminorm $\|\cdot\|$ if*

$$\|f_i - f_j\| \geq \delta, \qquad \forall i, j \in [M],\ i \neq j.$$

The packing number, denoted by $\mathcal{M}(\|\cdot\|, \delta, \mathcal{F})$, is the supremum of the cardinalities of all finite δ-packings of $\mathcal{F}$, with value $+\infty$ if these cardinalities are unbounded. When the norm or seminorm is clear from the context, we abbreviate this notation to $\mathcal{M}(\delta, \mathcal{F})$.

Definition 1.3 (Local packing number and local entropy). *For $\eta > 0$ and a fixed universal constant $c > 0$, define the local packing number of $\mathcal{F}$ by*

$$\mathcal{M}_{\mathcal{F}}^{\mathrm{loc}}(\eta, c) := \sup_{f \in \mathcal{F}} \mathcal{M}\left(\eta/c, B_2(f, \eta) \cap \mathcal{F}\right), \tag{1.4}$$

where the packing numbers are computed with respect to the $L_2(\mathbb{P}_X)$-distance. The quantity $\log \mathcal{M}_{\mathcal{F}}^{\mathrm{loc}}(\eta, c)$ is called the local metric entropy of $\mathcal{F}$. Throughout the paper, c is reserved for the constant in this definition and is taken to be sufficiently large.

An important property of the local packing number is that for any star-shaped function class $\mathcal{F}$, $\mathcal{M}_{\mathcal{F}}^{\mathrm{loc}}(\eta, c)$ is nonincreasing with respect to η, as verified by the following lemma.

Lemma 1.4 (Nonincreasingness of $\mathcal{M}(\eta, c)$, Lemma 1.4 from Prasadan and Neykov [2026]). *For any star-shaped class $\mathcal{F}$ and constant $c > 0$, $\mathcal{M}_{\mathcal{F}}^{loc}(\eta, c)$ is nonincreasing with respect to η.*

Lemma 1.4 guarantees the well-definedness of the critical radius, defined as follows.

Definition 1.5 (Critical radius). *Given the star-shaped $\mathcal{F}$ and n, the critical radius of $\mathcal{F}$ is defined as*

$$\eta^*(\mathcal{F}) := \sup\left\{\eta > 0 : n\eta^2 \le \log \mathcal{M}_{\mathcal{F}}^{\mathrm{loc}}(\eta, c)\right\}, \tag{1.5}$$

with the convention that $\eta^*(\mathcal{F}) = 0$ if the set is empty. We suppress the dependence on $\mathcal{F}$ when it is clear from the context.

2 General Lower Bound

In this section, we recall Fano's method to derive the general minimax lower bound of (1.3). The resulting lower bound for the uncontaminated problem serves as a benchmark for our analysis and remains valid under strong contamination because the adversary may leave the sample unchanged. Fano's method has been widely applied to statistical problems, including constrained mean estimation Prasadan and Neykov [2026], Neykov [2023], support recovery in sparse linear regression Wainwright [2009], Reeves and Gastpar [2008], and support recovery in sparse principal component analysis Amini and Wainwright [2009].

We state below the version of the generalized Fano inequality used in this paper. We omit the proof and refer the reader to Sections 15.3 and 15.4 of Wainwright [2019] for a detailed treatment and the necessary information-theoretic background.

Theorem 2.1 (Generalized Fano inequality). *Let $\mathcal{F} \subset L_2(\mathbb{P}_X)$, and let $\{f_1, \dots, f_m\} \subset \mathcal{F}$ be a finite δ-packing with $m \ge 2$ and $\delta > 0$, so that*

$$\|f_i - f_j\|_{L_2(\mathbb{P}_X)} \ge \delta, \qquad \text{for all distinct } i, j \in [m].$$

Fix an admissible noise distribution $\mathbb{P}_\xi$, and let J be uniformly distributed over $[m]$. Independently of J, draw $\widetilde{X}_1, \dots, \widetilde{X}_n$ independently from $\mathbb{P}_X$ and $\xi_1, \dots, \xi_n$ independently from $\mathbb{P}_\xi$, with the design variables and noise terms mutually independent. Conditional on $J = j$, generate the responses according to

$$\widetilde{Y}_i = f_j(\widetilde{X}_i) + \xi_i, \qquad i \in [n].$$

Then the minimax risk under squared $L_2(\mathbb{P}_X)$ loss satisfies

$$\mathcal{R}_{n,\epsilon}(\mathcal{F};\mathcal{P}_\xi) \geq \frac{\delta^2}{4}\left(1 - \frac{I(\widetilde{X},\widetilde{Y};J)+\log 2}{\log m}\right),$$

where $I(\widetilde{X},\widetilde{Y};J)$ denotes the mutual information between the full uncontaminated sample $(\widetilde{X},\widetilde{Y})$ and the index J under the joint distribution specified above.

By Fano's inequality above, to derive a nontrivial lower bound on $\mathcal{R}_{n,\epsilon}(\mathcal{F};\mathcal{P}_\xi)$, we need to control the mutual information $I(\widetilde{X},\widetilde{Y};J)$. We therefore restrict our attention to the uncontaminated Gaussian submodel with $\xi_i \sim \mathcal{N}(0,\sigma^2)$, where $\sigma > 0$ is chosen so that this noise distribution satisfies the stated assumptions. Since the minimax risk (1.3) takes a supremum over admissible noise distributions and corruption mechanisms, any lower bound for this submodel also applies to the full model. We also note that such inclusion of σ is compatible with our later additional assumptions introduced in Section 3 and Section 4.

The following lemma from Prasadan and Neykov [2025] provides a useful control given this additional assumption. We note that Gaussianity is used only in this lower-bound argument; it is not imposed on the noise distribution in our upper-bound results.

Lemma 2.2 (Lemma 2.2 of Prasadan and Neykov [2025])**.** *Given the same settings as Theorem 2.1 and additionally assuming that $\widetilde{\xi} \sim \mathcal{N}(0,\sigma^2)$ for some known σ, we have*

$$I(\widetilde{X},\widetilde{Y};J) \leq \frac{n}{2\sigma^2}\max_{j\in[m]}\|f_j - h\|^2_{L_2(\mathbb{P}_X)}, \qquad \forall h \in \mathcal{F}.$$

Combining Theorem 2.1 and Lemma 2.2, we come to the following lower bound on the minimax risk (1.3).

Lemma 2.3 (Lemma 2.3 of Prasadan and Neykov [2025])**.** *Given the conditions stated in Theorem 2.1 and Lemma 2.2, for any η satisfying $\log \mathcal{M}^{loc}_{\mathcal{F}}(\eta,c) > 4\left(\frac{n\eta^2}{2\sigma^2} \vee \log 2\right)$, we have*

$$\mathcal{R}_{n,\epsilon}(\mathcal{F};\mathcal{P}_\xi) \geq \frac{\eta^2}{8c^2}, \tag{2.1}$$

where we recall that c is the constant from Definition 1.3.

Lemma 2.2 follows from an inequality between the mutual information and the KL divergence and a direct calculation of KL divergence between Gaussian random variables. Lemma 2.3 is the result of a direct application of Theorem 2.1 and Lemma 2.2. For completeness, we include the proofs in Appendix A.1 and Appendix A.2.

3 Minimax rate under Bernstein-Type Function Class and Heavy-Tailed Noise

In this section, we study estimation under a Bernstein-type moment condition on the increments of $\mathcal{F}$ and a finite-variance assumption on the noise. Specifically, we assume that there exists a known constant $L > 0$ such that, for every integer $q \geq 2$ and every $f, g \in \mathcal{F}$,

$$\|f-g\|^{2q}_{L_{2q}(\mathbb{P}_X)} \leq \frac{q!}{2}L^{q-1}\|f-g\|^2_{L_2(\mathbb{P}_X)}. \tag{3.1}$$

Here, for $q \geq 1$, the $L_q(\mathbb{P}_X)$-norm is defined by

$$\|f - g\|_{L_q(\mathbb{P}_X)} := \left(\int_{\mathcal{X}} |f(x) - g(x)|^q \, d\mathbb{P}_X(x) \right)^{1/q}.$$

For the noise, we retain the assumptions that the conditional expectation of $\widetilde{\xi}$ given $\widetilde{X}$ is constantly zero, and additionally assume that there exists a known constant $\sigma > 0$ such that

$$\mathbb{E}\left[\widetilde{\xi}^2 \mid \widetilde{X}\right] \leq \sigma^2, \qquad \text{a.s.} \tag{3.2}$$

No higher-order moment or exponential-tail assumption is imposed on $\widetilde{\xi}$. Denote the collection of distributions of such $\widetilde{\xi}$ as $\mathcal{P}_\xi^{\mathrm{b}}$. We might wonder whether $\mathbb{E}\widetilde{\xi}^2 \leq \sigma^2$ for some σ is sufficient. However, as we remark later, this is insufficient and (3.2) is necessary for our argument.

We also note here that (3.1) actually implies an upper bound on the diameter of $\mathcal{F}$. By setting $q = 2$ in (3.1), we have

$$\|f - g\|_{L_4(\mathbb{P}_X)}^4 \leq L \, \|f - g\|_{L_2(\mathbb{P}_X)}^2 .$$

However, from Jensen's inequality, we also have

$$\|f - g\|_{L_2(\mathbb{P}_X)}^4 \leq \|f - g\|_{L_4(\mathbb{P}_X)}^4 .$$

Combining the two displays above, we know that $\|f - g\|_{L_2(\mathbb{P}_X)}^2 \leq L, \forall f, g \in \mathcal{F}$, which shows that $d^2 \leq L$.

Condition (3.1) is a joint assumption on the function class $\mathcal{F}$ and the design distribution $\mathbb{P}_X$. It controls the moments of the squared increments $(f(X) - g(X))^2$ and yields Bernstein-type concentration for their empirical averages, thereby relating empirical squared distances to their population counterparts; see Prasadan and Neykov [2025]. Under the standing bounded-diameter assumption, this condition is implied by either uniform sup-norm boundedness of $\mathcal{F}$ or a relative sub-Gaussian increment condition introduced later in Section 4. However, (3.1) does not require either of these stronger assumptions, as illustrated by the examples in Prasadan and Neykov [2025].

Although (3.1) is weaker than the relative sub-Gaussian increment condition, it still controls moments of all orders and implies exponential integrability of the squared increments. Its greater generality concerns the relationship between higher moments and the $L_2(\mathbb{P}_X)$-distance: it does not require the sub-Gaussian scale of an increment to be proportional to its $L_2(\mathbb{P}_X)$-norm. Accordingly, the term "heavy-tailed" in this subsection refers to the permitted noise distributions.

The (conditional) second-moment assumption on ξ allows distributions with polynomially decaying tails and, in particular, distributions with an infinite fourth moment. Control of the conditional second moment of ξ nevertheless ensures that, for every $f, g \in \mathcal{F}$,

$$\mathbb{E}\left[\widetilde{\xi}^2 \left(f(\widetilde{X}) - g(\widetilde{X})\right)^2\right] = \mathbb{E}\left[\left(f(\widetilde{X}) - g(\widetilde{X})\right)^2 \mathbb{E}\left[\widetilde{\xi}^2 \mid \widetilde{X}\right]\right] \leq \sigma^2 \, \|f - g\|_{L_2(\mathbb{P}_X)}^2 .$$

This provides second-moment control of the noise–function products that arise in regression comparisons without requiring higher-order noise moments.

3.1 Lower Bound

We now argue in the following lemma that $\epsilon\sigma^2 \wedge d^2$ is a valid lower bound of $\mathcal{R}_{n,\epsilon}(\mathcal{F}; \mathcal{P}_\xi)$ under the star-shapedness of $\mathcal{F}$, strong adversary $\mathcal{C}$, and Conditions (3.1), (3.2).

Lemma 3.1 (Lower bound regarding ϵ)**.** *For the nonparametric regression problem formulated in Section 1.1, with the presence of the strong adversary $\mathcal{C}$ and Conditions (3.1), (3.2), we have*

$$\mathcal{R}_{n,\epsilon}(\mathcal{F}; \mathcal{P}_\xi) \gtrsim \epsilon\sigma^2 \wedge d^2. \tag{3.3}$$

The proof, deferred to Appendix B.1, is based on a two-point construction. Star-shapedness allows us to select two functions in $\mathcal{F}$ whose squared $L_2(\mathbb{P}_X)$-distance is of order $\epsilon\sigma^2 \wedge d^2$. We then construct two regression models with these regression functions and a common conditional noise variance bounded by σ^2. The mixture probabilities depend on the covariate, ensuring that the conditional moment constraints hold while the two observations, under a suitable coupling, differ with marginal probability at most $\epsilon/2$. With probability bounded below by an absolute constant, the number of distinguishing observations does not exceed $\lfloor \epsilon n \rfloor$. The strong adversary can then replace their responses by common values, making the coupled corrupted samples identical. A two-point risk comparison therefore forces any estimator to incur squared error of order $\epsilon\sigma^2 \wedge d^2$ under at least one of the two models. This argument uses only star-shapedness and the conditional moment constraints on the noise.

Combining Lemma 3.1 with the general lower bound (2.1) yields

$$\mathcal{R}_{n,\epsilon}(\mathcal{F}; \mathcal{P}_\xi) \gtrsim_\sigma \left((\eta^*)^2 \vee \epsilon\right) \wedge d^2. \tag{3.4}$$

Recall in (1.3) that we define $\mathcal{R}_{n,\epsilon}(\mathcal{F}, \mathcal{P}_\xi)$ as the worst-case squared estimation error over all noise distributions that belong to $\mathcal{P}_\xi$. The following example demonstrates that, this general lower bound cannot be improved even when we restrict our scope to some good noise distributions. Or equivalently, we are free to expand our cover of the lower bound from some well-behaved noise distributions to the general distribution class $\mathcal{P}_\xi$ that satisfies the basic conditions $\mathbb{E}\left[\widetilde{\xi} \mid \widetilde{X}\right] = 0$.

Example 3.2 (Optimality of ϵ)**.** *Let $\widetilde{X}_i$ be distributed as a uniform distribution on $[0,1]$. The function class $\mathcal{F} = \left\{\theta \mathbf{1}_{\{x \in [0,t]\}} : \theta \in [0,1], t \in [0,1]\right\}$. Clearly $\mathcal{F}$ is a star-shaped function class. Also, since $|f - g| \le 1, \forall f, g \in \mathcal{F}$, we have*

$$\mathbb{E}\,|f-g|^{2p} \le \mathbb{E}\,|f-g|^2 \le \frac{p!}{2}\mathbb{E}\,|f-g|^2\,.$$

Therefore, the moment increment condition (3.1) holds for $\mathcal{F}$ with $L = 1$. Let $\widetilde{Y} = f^*(\widetilde{X}) + \widetilde{\xi}$ for some $f^* \in \mathcal{F}$ and $\xi \sim \mathcal{N}(0,1)$.

For a fixed and known contamination level ϵ, let $A = [0, \epsilon/2]$, $f_0(x) = 0, f_1(x) = \mathbf{1}_A(x)$. Then the squared $L_2(\mathbb{P}_X)$ distance between f_0 and f_1 is

$$\|f_1 - f_0\|^2_{L_2(\mathbb{P}_X)} = \mathbb{P}_X(A) = \frac{\epsilon}{2}.$$

Suppose our objective now is to distinguish between the cases when $f^* = f_0$ and $f^* = f_1$. Since $f_0(x) = f_1(x) = 0$ for $x \notin A$, any efforts must rely on the samples that are supported on A. However, consider the counting statistic

$$N_A = \sum_{i=1}^{n} \mathbf{1}_{\{\widetilde{X}_i \in A\}}.$$

By construction, $N_A \sim \mathrm{Bin}\,(n, \epsilon/2)$ *and consequently by the Chernoff bound, we have*

$$\mathbb{P}\,(N_A > \epsilon n) = \mathbb{P}\,(N_A > 2\mathbb{E}N_A) \le \exp\left\{-\frac{1}{6}\epsilon n\right\}.$$

Consider a simple strategy of $\mathcal{C}$ *that* $\mathcal{C}$ *replaces each sample on* A *with* 0 *when* $N_A \le \lfloor \epsilon N \rfloor$ *and otherwise leaves the samples unchanged. Denote* Q_0 *and* Q_1 *as the joint distributions of* (X, Y) *given* $f^* = f_0$ *and* $f^* = f_1$*, respectively. Then by the argument above, we know that on the event* $\{N_A \le \epsilon n\}$*, the behavior of samples is exactly the same under both hypotheses. Therefore, we have*

$$\mathrm{TV}(Q_0, Q_1) \le \exp\left\{-\frac{1}{6}\epsilon n\right\}.$$

Now from Le Cam's two-point lower bound (Donoho and Liu [1991]), for any estimator $\widehat{f}$*, we have*

$$\mathcal{R}_{n,\epsilon}(\mathcal{F}; \mathcal{P}_\xi) \ge \frac{\|f_1 - f_0\|^2_{L_2(\mathbb{P}_X)}}{8}\,(1 - \mathrm{TV}(Q_0, Q_1)) \ge \frac{\epsilon}{16}\left(1 - \exp\left\{-\frac{1}{6}\epsilon n\right\}\right) \gtrsim \epsilon,$$

provided n *is sufficiently large.*

In the next section, we establish a matching upper bound using tree-based Algorithms 1 and 2. This therefore determines the minimax rate of (1.3) up to constant factors under the Bernstein's condition (3.1) and finite conditional second moment (3.2) of ξ.

3.2 Upper Bound

To obtain an upper bound corresponding to the critical radius η^* appearing in Section 3.1, we construct a hierarchy of localized approximations to $\mathcal{F}$. Each successive level uses finer packing resolutions within smaller neighborhoods, allowing the estimation procedure to explore the function class at progressively decreasing scales. This multiscale strategy builds on the construction developed for constrained mean estimation in the Gaussian sequence model in Neykov [2023]. We first introduce the graph notation and then describe the construction.

Let $G = (V, E)$ be a directed graph whose vertices are labeled by functions in $\mathcal{F}$. For each vertex $u \in V$, define its parent set and offspring set by

$$\mathcal{P}(u) := \{v \in V : (v, u) \in E\}, \qquad \mathcal{O}(u) := \{v \in V : (u, v) \in E\}.$$

The vertices are organized into levels, with $\mathcal{L}(k)$ denoting the collection of vertices at level k. The root occupies level 1, and every edge connects a vertex at level k to a vertex at level $k+1$. Since the pruning step may assign several parents to the same vertex, the resulting structure is a layered directed acyclic graph, which we refer to as the pruned tree. When computing distances, we identify each vertex with its function label; distinct vertices may nevertheless carry the same label.

Choose an arbitrary root $u_1^1 \in \mathcal{F}$ and set $\mathcal{L}(1) = \{u_1^1\}$. Because d is the diameter of $\mathcal{F}$, we have $B_2(u_1^1, d) \cap \mathcal{F} = \mathcal{F}$. Construct a maximal (d/c)-packing of this set, take its elements as the vertices of $\mathcal{L}(2)$, and connect the root to each of them. The second level is retained without pruning.

For $k \ge 3$, suppose that the preceding level $\mathcal{L}(k-1)$ has already been finalized. For every $u \in \mathcal{L}(k-1)$, construct a maximal $d/(2^{k-1}c)$-packing of $B_2(u, d/2^{k-2}) \cap \mathcal{F}$. Each packing point generates a preliminary child of u. If the same function occurs in packings associated with different

parents, we initially represent these occurrences by separate vertices, so that every preliminary vertex has a unique parent.

Arrange the preliminary vertices according to a fixed ordering, write them as $u_1^k, \ldots, u_{N_k}^k$, and initialize the ordered collection of unprocessed vertices as $\mathcal{U}_k = \left[u_1^k, \ldots, u_{N_k}^k\right]$. The pruning procedure repeatedly retains the first unprocessed vertex and merges nearby candidates into it. Specifically, if u_i^k is the first element of $\mathcal{U}_k$, define

$$\mathcal{T}_k(u_i^k) := \left\{ u_j^k \in \mathcal{U}_k \;\middle|\; j \neq i,\ \left\| u_j^k - u_i^k \right\|_{L_2(\mathbb{P}_X)} \leq \frac{d}{2^{k-1}c} \right\}.$$

We retain u_i^k as the representative of this group. For every $u_j^k \in \mathcal{T}_k(u_i^k)$, let v_j be its unique parent. Replace the edge $v_j \to u_j^k$ with $v_j \to u_i^k$, adding the latter only if it is not already present, and discard u_j^k from the graph. Thus, the retained representative inherits the parent connections of all vertices merged into it.

After these edge reassignments, remove $\{u_i^k\} \cup \mathcal{T}_k(u_i^k)$ from $\mathcal{U}_k$, while keeping u_i^k in the graph. Continue until no unprocessed vertices remain. The surviving vertices form the finalized level

$$\mathcal{L}(k) = \bigcup_{u \in \mathcal{L}(k-1)} \mathcal{O}(u).$$

By construction, any two distinct vertices retained at this level are separated by more than $d/(2^{k-1}c)$, and every discarded candidate lies within this distance of its retained representative.

Applying the same construction recursively defines the complete hierarchy. Algorithm 1 summarizes this procedure, whose pruning mechanism follows Prasadan and Neykov [2026]. The subsequent Algorithm 2 uses this hierarchy with a different selection criterion tailored to the contaminated regression model we study.

The iterative construction and pruning procedures described above yield the following crucial lemmas from Prasadan and Neykov [2026], which establish the covering and packing properties at each level of the directed tree G. Specifically, the nodes at each level k constitute a valid covering and packing of the constraint set $\mathcal{F}$. Moreover, for any node u^{k-1} at level $k-1$, its offspring set $\mathcal{O}(u^{k-1})$ forms an effective cover for the localized neighborhood $\mathcal{B}(u^{k-1}, d/2^{k-2}) \cap \mathcal{F}$. These geometric structural properties play a foundational role in our subsequent analysis of the estimator's performance.

Lemma 3.3 (Lemma 3.1 of Prasadan and Neykov [2026]). *For any $J \geq 3$ and $u^{J-1} \in \mathcal{L}(J-1)$, we have:*
(1), $\mathcal{L}(J)$ forms a $d/(2^{J-2}c)$-covering of $\mathcal{F}$;
(2), $\mathcal{L}(J)$ forms a $d/(2^{J-1}c)$-packing of $\mathcal{F}$;
(3), the offspring $\mathcal{O}(u^{J-1})$ form a $d/(2^{J-2}c)$-covering of the set $\mathcal{B}(u^{J-1}, d/2^{J-2}) \cap \mathcal{F}$;
(4), the cardinality of $\mathcal{O}(u^{J-1})$ is upper bounded by $\mathcal{M}_{\mathcal{F}}^{\text{loc}}(d/2^{J-2}, 2c)$ for $J \geq 2$.

Lemma 3.4 (Lemma 3.2 of Prasadan and Neykov [2026]). *For any $f \in \mathcal{F}$ and $J \geq 2$, $\mathcal{L}(J-1) \cap B_2(f, d/2^{J-2})$ has a cardinality upper bounded by $\mathcal{M}_{\mathcal{F}}^{loc}(d/2^{J-2}, c) \leq \mathcal{M}_{\mathcal{F}}^{loc}(d/2^{J-2}, 2c)$.*

Lemma 3.5 (Lemma 3.3 of Prasadan and Neykov [2026]). *Let $[f_1, f_2, \ldots]$ be the nodes of an infinite path in the graph G, i.e., $f_{J+1} \in \mathcal{O}(f_J)$ for any $J \in \mathbb{N}$. Then for any integers $1 \leq J' \leq J$, we have $\|f_{J'} - f_J\|_{L_2(\mathbb{P}_X)} \leq \frac{d(2+4c)}{c2^{J'}}$.*

Algorithm 1: Construction of the pruned tree

Input: Function class $\mathcal{F}$ with diameter $d > 0$, root $f_1^1 \in \mathcal{F}$, and constant c

Output: The pruned graph G

Initialize G with the root f_1^1 and set $\mathcal{L}(1) \leftarrow \{f_1^1\}$;
Construct a maximal (d/c)-packing of $B_2(f_1^1, d) \cap \mathcal{F}$ and create its vertices as $\mathcal{L}(2)$;
Add these vertices to G and connect f_1^1 to each of them;
$k \leftarrow 3$;
while True **do**
  $\mathcal{U}_k \leftarrow \varnothing$;
  for $u \in \mathcal{L}(k-1)$ **do**
    Construct a maximal $d/(2^{k-1}c)$-packing of $B_2(u, d/2^{k-2}) \cap \mathcal{F}$;
    Create a separate vertex for each packing point and add these vertices to G and $\mathcal{U}_k$;
    Add an edge from u to each newly created vertex;
  Arrange $\mathcal{U}_k$ according to the fixed ordering;
  while $\mathcal{U}_k \neq \varnothing$ **do**
    Let $\widetilde{u}$ be the first vertex in $\mathcal{U}_k$;
    $\mathcal{T}_k(\widetilde{u}) \leftarrow \left\{ u \in \mathcal{U}_k \;\middle|\; u \neq \widetilde{u},\ \|u - \widetilde{u}\|_{L_2(\mathbb{P}_X)} \leq \frac{d}{2^{k-1}c} \right\}$;
    for $u' \in \mathcal{T}_k(\widetilde{u})$ **do**
      Let v be the unique parent of u';
      Remove the edge $v \to u'$;
      Add the edge $v \to \widetilde{u}$ if it is not already present;
      Delete u' from G;
    Remove $\{\widetilde{u}\} \cup \mathcal{T}_k(\widetilde{u})$ from $\mathcal{U}_k$;
  $\mathcal{L}(k) \leftarrow \bigcup_{u \in \mathcal{L}(k-1)} \mathcal{O}(u)$;
  $k \leftarrow k + 1$;

We refer the reader to the original source for the proofs. Algorithm 1 provides increasingly fine approximations of $\mathcal{F}$ and provides the hierarchy used by Algorithm 2 to approximate f^* at progressively finer scales.

Intuitively, for any candidate $g \in \mathcal{F}$ that is close enough to the true function f^*, we expect the mean squared error (MSE) of the estimation of Y using g to be small, at least compared with the function that is significantly farther. However, the adversary $\mathcal{C}$ breaks this intuition as it is known that the sample mean is not robust to the corruption. This motivates a more robust approach to estimate the MSE in the presence of the adversary—trimmed mean. The mathematical study of trimmed means dates back at least to Daniell [1920], which investigated the removal of extreme observations within a framework of weighted order statistics. More recently, finite-sample guarantees under heavy tails and adversarial contamination have been established for modified trimmed-mean procedures. Lugosi and Mendelson [2021] prove the effectiveness of such procedures in both the univariate and multivariate settings. These developments provide methodological foundations for using trimming in the present regression problem.

Lemma 3.6 (Trimmed mean, Lugosi and Mendelson [2021])**.** *Let X be a real-valued random vari-*

able with distribution $\mathbb{P}_X$, mean μ, and finite variance σ_X^2. Suppose that the observations $X_1, \dots, X_n$ are obtained from n independent copies of X after a strong adversary $\mathcal{C}$ replaces at most $\lfloor \epsilon n \rfloor$ observations. Assume that n is even, and let $\rho \in (0,1)$ satisfy $\rho \geq \exp\{-n/2\}/4$. Define

$$\bar{X} := X - \mu, \qquad \zeta := 8\epsilon + 24\frac{\log\left(\frac{4}{\rho}\right)}{n},$$

and assume that $\zeta < 1/4$. For $t \in (0,1)$, define the tail functional

$$\mathcal{E}(t, X) := \max\left\{\mathbb{E}\left[\left|\bar{X}\right| \mathbf{1}_{\left\{\bar{X} \leq Q_{t/2}(\bar{X})\right\}}\right], \mathbb{E}\left[\left|\bar{X}\right| \mathbf{1}_{\left\{\bar{X} \geq Q_{1-t/2}(\bar{X})\right\}}\right]\right\},$$

where $Q_u(\bar{X})$ denotes the u-quantile of $\bar{X}$.

Partition the observed sample into two halves, $\left(X_{1,1}, \dots, X_{1,n/2}\right)$ and $\left(X_{2,1}, \dots, X_{2,n/2}\right)$. Let $X_{1,1}^* \leq \cdots \leq X_{1,n/2}^*$ denote the order statistics of the first half, and set

$$\alpha := X^*_{1,\lceil \zeta n/2 \rceil}, \qquad \beta := X^*_{1,\lfloor (1-\zeta)n/2 \rfloor}.$$

Define the truncation function

$$\phi_{\alpha,\beta}(x) := \begin{cases} \alpha, & x \leq \alpha, \\ x, & \alpha < x \leq \beta, \\ \beta, & x > \beta, \end{cases}$$

and the modified trimmed mean estimator

$$\mathrm{Tm}(X) := \frac{2}{n}\sum_{i=1}^{n/2} \phi_{\alpha,\beta}\left(X_{2,i}\right).$$

Then the following bounds hold:

1. *With probability at least $1 - \rho$,*

$$|\mathrm{Tm}(X) - \mu| \leq 4\mathcal{E}(4\zeta, X) + 2\sigma_X\sqrt{\frac{2\log\left(\frac{4}{\rho}\right)}{n}}. \tag{3.5}$$

2. *With probability at least $1 - 4\exp\{-\zeta n/24\}$,*

$$|\mathrm{Tm}(X) - \mu| \leq 10\sqrt{\zeta}\sigma_X. \tag{3.6}$$

Remark. We will later use $\mathbb{P}\left(\bar{X} \geq Q_{1-t/2}(\bar{X})\right) \leq t/2$ to bound $\mathcal{E}(t, X)$. However, this is not true in general for random variables with atoms. As mentioned in Lugosi and Mendelson [2021], this issue can be amended by adding a sufficiently small absolutely continuous perturbation to X (for example, rG for sufficiently small r and G is a standard Gaussian random variable).

We now introduce the pairwise comparison statistic underlying Algorithm 2. For any $f, g \in \mathcal{F}$ satisfying $\|f - g\|_{L_2(\mathbb{P}_X)} > 0$, define the normalized squared-loss difference

$$Z_{f,g} := \frac{[Y - f(X)]^2 - [Y - g(X)]^2}{\|f - g\|_{L_2(\mathbb{P}_X)}}. \tag{3.7}$$

Let $\mathrm{Tm}_{Z_{f,g}}(X, Y)$ denote the estimator from Lemma 3.6 applied to the normalized loss differences computed from the observed sample $\{(X_i, Y_i)\}_{i=1}^n$. Define the pairwise comparison test

$$\psi_{f,g}(X, Y) := \mathbf{1}_{\left\{\mathrm{Tm}_{Z_{f,g}}(X,Y) > 0\right\}}. \tag{3.8}$$

We say that f is dominated by g, denoted by $f \prec g$ or $g \succ f$, whenever $\psi_{f,g}(X, Y) = 1$. This relation need not be transitive: $f \prec g$ and $g \prec h$ do not necessarily imply $f \prec h$. Moreover, for any two candidates $f, g \in \mathcal{F}$, (3.8) prevents the occurrence that both $f \prec g$ and $g \prec f$ happen simultaneously. To see this, let $z_{(1)} \le z_{(2)} \le \cdots \le z_{(n/2)}$ be the order statistics of the first half of $Z_{f,g}$ and $\omega_1, \omega_2, \ldots, \omega_{n/2}$ be the second half of $Z_{f,g}$. Let $k = \lceil \zeta n/2 \rceil$. By definition of the trimmed mean estimator, we have $\mathrm{Tm}(Z_{f,g}) = \frac{2}{n} \sum_{i=1}^{n/2} \phi_{z_{(k)}, z_{(n/2-k)}}(\omega_i)$. Since $Z_{f,g} = -Z_{g,f}$, we then have by the property of ϕ:

$$\mathrm{Tm}(Z_{g,f}) = \frac{2}{n} \sum_{i=1}^{n/2} \phi_{-z_{(n/2+1-k)}, -z_{(k+1)}}(-\omega_i) = -\frac{2}{n} \sum_{i=1}^{n/2} \phi_{z_{(k+1)}, z_{(n/2+1-k)}}(\omega_i) \le -\mathrm{Tm}(Z_{f,g}).$$

Hence, $\mathrm{Tm}(Z_{f,g}) > 0$ implies that $\mathrm{Tm}(Z_{g,f}) < 0$.

Applying Lemma 3.6 to the normalized loss differences yields the following concentration bound, which is a key ingredient in establishing the upper bound for the estimator constructed by Algorithm 2.

Lemma 3.7. *Suppose that $\mathcal{F}$ satisfies the Bernstein-type moment condition (3.1) and that the noise $\widetilde{\xi}$ satisfies Condition (3.2). Let $f, g \in \mathcal{F}$ be fixed functions, and suppose that, for some $\delta > 0$,*

$$\|g - f^*\|_{L_2(\mathbb{P}_X)} \le \delta, \qquad \|f - g\|_{L_2(\mathbb{P}_X)} \ge C\delta,$$

where $C := c/2 - 1 \ge 3$ is a universal constant. Let ρ and the corresponding trimming parameter ζ satisfy the conditions of Lemma 3.6. Then, with probability at least $1 - \rho$,

$$\left|\mathrm{Tm}_{Z_{f,g}}(X, Y) - \mathbb{E}Z_{f,g}\right| \lesssim \sqrt{L + \sigma^2} \left[\sqrt{\epsilon} + \sqrt{\log\left(\frac{4}{\rho}\right)/n}\right], \tag{3.9}$$

where the expectation is taken under the uncontaminated distribution and the implicit constant is universal.

By bounding $\mathbb{E}Z_{f,g}$ and properly selecting the constants in Lemma 3.7, we obtain the following corollary, which is analogous to Lemma 3.8 of Prasadan and Neykov [2025].

Corollary 3.8. *Under the same settings of Lemma 3.7, additionally assume that* $\epsilon \le \frac{1}{64}$ *and* $a \le \frac{1}{192}$. *If* $\delta^2 \gtrsim_{L,\sigma} \epsilon$, *we have*

$$\mathbb{P}\left(\psi_{f,g}(X,Y)=0\right) \le 4\exp\left\{-\frac{an\delta^2}{L+\sigma^2}\right\}.$$

Equivalently, the pairwise comparison test (3.8) *correctly selects* g *over* f *with probability greater than* $1-4\exp\left\{-an\delta^2/(L+\sigma^2)\right\}$

The proofs of Lemma 3.7 and Corollary 3.8 are provided in Appendix B.2 and Appendix B.3, respectively.

To describe the selection procedure on the pruned tree constructed in Algorithm 1, we first introduce some notation. For a nonempty finite set $S \subset \mathcal{F}$, a scale $\delta > 0$, and a function $f \in S$, define the set of admissible competitors of f by

$$S_f := \left\{g \in S : g \succ f,\ \|f-g\|_{L_2(\mathbb{P}_X)} \ge C\delta\right\},$$

where C is the constant appearing in Lemma 3.7. The dependence of S_f on S and δ is suppressed for convenience. We define the *discrepancy* of f relative to S at scale δ as

$$H(\delta, f, S) := \begin{cases} \max\limits_{g\in S_f} \|f-g\|_{L_2(\mathbb{P}_X)}, & \text{if } S_f \neq \varnothing, \\ 0, & \text{if } S_f = \varnothing. \end{cases} \tag{3.10}$$

Combining Corollary 3.8 with the discrepancy criterion (3.10) yields the following selection guarantee. It ensures that the probability of incorrect selection is small when at least one function from S is close to the true function f^*.

Lemma 3.9. *Suppose that the conditions of Corollary 3.8 hold at scale* $\delta > 0$. *Let* $S = \{f_1, \dots, f_M\} \subset \mathcal{F}$ *be a fixed finite set satisfying*

$$\min_{i\in[M]} \|f_i - f^*\|_{L_2(\mathbb{P}_X)} \le \delta.$$

Select

$$f_S \in \arg\min_{f\in S} H(\delta, f, S),$$

with ties resolved according to a fixed ordering of S. *Then*

$$\mathbb{P}\left(\|f_S - f^*\|_{L_2(\mathbb{P}_X)} \ge (C+1)\,\delta\right) \le 4M\exp\left\{-\frac{an\delta^2}{L+\sigma^2}\right\},$$

where $a > 0$ *is the constant appearing in Corollary 3.8.*

The proof follows from Corollary 3.8, the triangle inequality, and a union bound and is given in Appendix B.4.

We now describe the selection procedure on the pruned tree constructed in Algorithm 1. Although the tree may have infinitely many levels, the procedure traverses only finitely many of them. Starting from the root f_1^1, we set $f_1^* := f_1^1$. At each iteration $k \ge 1$, we select f_{k+1}^* from the offspring set $\mathcal{O}(f_k^*)$ by minimizing the discrepancy $H\left(\delta = \frac{d}{2^k(C+1)}, f, \mathcal{O}(f_k^*)\right)$ over $f \in \mathcal{O}(f_k^*)$. Ties are resolved by selecting the first minimizer according to a fixed ordering of the nodes. The procedure is summarized in Algorithm 2, where the stopping index J^* will be specified later.

Algorithm 2: Sequential selection on the pruned tree

Input: Function class $\mathcal{F}$, observations $(X_1, Y_1), \ldots, (X_n, Y_n)$, the pruned tree G from Algorithm 1, and a stopping index J^*

Output: The sequence $\mathcal{F}^* = [f_1^*, \ldots, f_{J^*}^*]$

$f_1^* \leftarrow f_1^1$;
$k \leftarrow 1$;
$\mathcal{F}^* \leftarrow [f_1^*]$;
while $k < J^*$ **do**
  Select $f_{k+1}^* \in \underset{f \in \mathcal{O}(f_k^*)}{\arg\min}\, H\left(\frac{d}{2^k(C+1)}, f, \mathcal{O}(f_k^*)\right)$; `/* Choose the first minimizer in the fixed node ordering */`
  Append f_{k+1}^* to $\mathcal{F}^*$;
  $k \leftarrow k + 1$;
return $\mathcal{F}^* = [f_1^*, \ldots, f_{J^*}^*]$;

Corollary 3.10 (One-step error bound along the tree traversal)**.** *Under the assumptions of the preceding tournament selection lemma, denote by f_j^* the node selected at level j by Algorithm 2, with f_1^* being the root. Thus, f_j^* is selected after $j-1$ iterations. For each $j \geq 1$, define*

$$\delta_j := \frac{d}{2^{j-2}c} = \frac{d}{2^{j-1}(C+1)}, \qquad \eta_j := \sqrt{\frac{a}{L+\sigma^2}}\,\delta_j,$$

and, for each level $k \geq 1$, define the failure event

$$A_k := \left\{ \|f_k^* - f^*\|_{L_2(\mathbb{P}_X)} > \frac{d}{2^{k-1}} = (C+1)\delta_k \right\}. \tag{3.11}$$

Suppose that $\epsilon \leq b\delta_j^2/(L+\sigma^2)$, where $b > 0$ is a sufficiently small constant determined by the preceding testing bound and independent of d. Then

$$\mathbb{P}\left(A_{j-1}^{\complement} \cap A_j\right) \leq 4\left[\mathcal{M}_{\mathcal{F}}^{\mathrm{loc}}\left(\frac{d}{2^{j-2}}, 2c\right)\right]^2 \exp\left\{-n\eta_j^2\right\}.$$

We note that Corollary 3.10 is a direct deduction from Lemma 3.4, Algorithm 2, and Corollary 3.8. We provide its proof in Appendix B.5.

With the help of Corollary 3.10, we can "telescope" the chain $\mathcal{F}^*$ to analyse the increment of the traversal. As a consequence, we are finally prepared to present the main theorem induced by Algorithm 2, which also serves as an upper bound for the minimax rate of (1.3) uniformly over the admissible star-shaped function class $\mathcal{F}$ and adversary $\mathcal{C}$.

Theorem 3.11. *Let J^* be the maximal integer such that*

$$n\eta_j^2 > 4\log\left[\mathcal{M}_{\mathcal{F}}^{loc}\left(\frac{d}{2^{j-2}}, 2c\right)\right] \vee \log 2. \tag{3.12}$$

If this condition never holds, we let $J^* = 1$.

For the nonparametric regression problem formulated in Section 1.1, suppose that $\epsilon \le \frac{1}{64}$ as required in Corollary 3.8. Let $f^*_{J^*}$ be the J^*th element of the sequence $\mathcal{F}^*$ after J^*-1 iterations of Algorithm 2. Then we have

$$\sup_{f^*\in\mathcal{F}} \sup_{\xi:\mathbb{P}_\xi\in\mathcal{P}_\xi^b} \sup_{\mathcal{C}\in\mathfrak{C}_{n,\epsilon}} \mathbb{E}_{\widetilde{X},\widetilde{Y}} \left\| f^*_{J^*}(\mathcal{C}(\widetilde{X},\widetilde{Y})) - f^* \right\|^2_{L_2(\mathbb{P}_X)} \lesssim \left(\eta^2_{J^*} \vee \epsilon\right) \wedge d^2. \tag{3.13}$$

Here a constant that only depends on L, σ, c is omitted.

Remark (Finiteness of the tree construction)**.** Boundedness of $\mathcal{F}$ in $L_2(\mathbb{P}_X)$ does not, by itself, guarantee that its packing numbers are finite. For example, the unit ball of an infinite-dimensional Hilbert space is bounded but contains an infinite collection of points with pairwise distances bounded away from zero. Thus, boundedness alone would not justify constructing a finite tree to an arbitrary depth.

For our estimator, however, it suffices to construct the tree through the stopping level J^*. When $J^* \ge 2$, the local entropy condition defining J^* guarantees that the local packing number controlling this level is finite. By the monotonicity of local packing numbers under star-shapedness established in Lemma 1.4, the corresponding packing numbers at all preceding levels are finite as well. Consequently, each node has finitely many offspring, and induction over the levels shows that the entire truncated tree is finite. When $J^* = 1$, the estimator simply returns the root, and no packing construction is needed. Therefore, the construction does not require an additional assumption that $\mathcal{F}$ is totally bounded. This observation ensures that the truncated construction is finite, although its computational cost can still be prohibitive.

Theorem 3.11 implies that $\mathcal{R}_{n,\epsilon}(\mathcal{F};\mathcal{P}_\xi) \lesssim_{L,\sigma} \left(\eta^2_{J^*} \vee \epsilon\right) \wedge d^2$. Recalling the definition of η^*, we obtain the following corollary that the upper bound (3.13) matches the lower bound (3.4), characterizing the exact minimax rate of $\mathcal{R}_{n,\epsilon}(\mathcal{F};\mathcal{P}_\xi)$ under (3.1) and (3.2).

Corollary 3.12. *For the nonparametric regression problem formulated in Section 1.1, under Conditions (3.1) and (3.2), the minimax rate of $\mathcal{R}_{n,\epsilon}(\mathcal{F};\mathcal{P}_\xi)$ is*

$$\mathcal{R}_{n,\epsilon}(\mathcal{F};\mathcal{P}_\xi) \asymp_{L,\sigma} \left((\eta^*)^2 \vee \epsilon\right) \wedge d^2. \tag{3.14}$$

Corollary 3.12 originates from Theorem 3.11 and the definition of η^*. We defer the proofs of Theorem 3.11 and Corollary 3.12 to Appendices B.6 and B.7, respectively.

Remark. Recall that we previously argued that the conditional second-moment assumption (3.2) is essential to the minimax rate. Intuitively, such a condition prevents the noise variance from concentrating on regions where competing regression functions differ. The following example shows that an unconditional second-moment bound alone is insufficient for both the variance estimate used in the proof and the resulting minimax guarantee.

Example 3.13 (Insufficiency of unconditional variance control)**.** *Let $X \sim \mathrm{U}[0,1]$ and consider the fixed function class previously discussed in Example 3.2:*

$$\mathcal{F} := \left\{ f_{\theta,t}(x) = \theta \mathbf{1}_{[0,t]}(x) : \theta, t \in [0,1] \right\}.$$

We know that this class is star-shaped about zero and has diameter $d = 1$. Moreover, the Bernstein condition holds with $L = 1$. Also, in this example, we consider the uncontaminated case, where $\epsilon = 0$ and $\left(\widetilde{X}_i, \widetilde{Y}_i, \widetilde{\xi}_i\right) = (X_i, Y_i, \xi_i), i = 1, 2, \ldots, n$.

For $q \in (0, 1/2)$, write $h_q := \mathbf{1}_{[0,q]}$ and let $G \sim \mathcal{N}(0,1)$ be independent of X. Define

$$\xi_q := \frac{\sigma}{\sqrt{2q}} h_q(X) G.$$

Then

$$\mathbb{E}[\xi_q \mid X] = 0, \qquad \mathbb{E}\xi_q^2 = \frac{\sigma^2}{2} < \sigma^2, \qquad \mathbb{E}[\xi_q^2 \mid X] = \frac{\sigma^2}{2q} h_q(X).$$

Thus, the unconditional second moment is uniformly bounded, while the conditional variance on $[0, q]$ diverges as $q \downarrow 0$.

Take $f^* = g = 0$ and $f = h_q$. Since $\|h_q\|_{L_2(\mathbb{P}_X)} = \sqrt{q}$, the pairwise comparison statistic satisfies

$$Z_{h_q,0} := \frac{(Y - h_q(X))^2 - Y^2}{\|h_q\|_{L_2(\mathbb{P}_X)}} = \frac{h_q(X)}{\sqrt{q}}(1 - 2\xi_q).$$

Consequently,

$$\mathbb{E}Z_{h_q,0} = \sqrt{q}, \qquad \mathrm{Var}(Z_{h_q,0}) = 1 - q + \frac{2\sigma^2}{q}.$$

Hence the variance cannot be bounded uniformly by a constant depending only on L and σ, although the function class and its diameter remain fixed.

The same construction also demonstrates failure of the critical-radius minimax guarantee. Let $\widetilde{\mathcal{P}}_\xi$ denote the noise class obtained by requiring only $\mathbb{E}[\xi \mid X] = 0$ and $\mathbb{E}\xi^2 \le \sigma^2$. Let $\mathbb{P}_0$ and $\mathbb{P}_1$ be the laws of $(X, Y) = (\widetilde{X}, \widetilde{Y})$ under $f_0 = 0$ and $f_1 = h_q$, respectively, using the same noise ξ_q under both hypotheses. The two conditional response distributions coincide outside $[0, q]$ and are Gaussian with common variance $\sigma^2/(2q)$ inside $[0, q]$. Therefore,

$$\|f_1 - f_0\|_{L_2(\mathbb{P}_X)}^2 = q, \qquad \mathrm{KL}\left(\mathbb{P}_0^{\otimes n} \,\middle\|\, \mathbb{P}_1^{\otimes n}\right) = nq \frac{1}{2(\sigma^2/(2q))} = \frac{nq^2}{\sigma^2}.$$

Choose $q = \frac{1}{4} \wedge \frac{\sigma}{4\sqrt{n}}$. The preceding divergence is at most $1/16$, so Pinsker's inequality and the two-point testing bound give

$$\mathcal{R}_{n,0}(\mathcal{F}; \widetilde{\mathcal{P}}_\xi) \ge \frac{q}{8}\left(1 - \mathrm{TV}\left(P_0^{\otimes n}, P_1^{\otimes n}\right)\right) \ge \frac{q}{16} \gtrsim 1 \wedge \frac{\sigma}{\sqrt{n}}.$$

The dependence of q on n is permitted because the minimax risk takes a supremum over the noise class for each sample size; $\mathcal{F}$ and $\mathbb{P}_X$ remain fixed.

On the other hand, this function class has small metric entropy. Indeed, suppose $t' \le t$. Then we have

$$\left\|f_{\theta,t} - f_{\theta',t'}\right\|_{L_2(\mathbb{P}_X)}^2 = t'(\theta - \theta')^2 + \theta^2(t - t') \le \left|\theta - \theta'\right|^2 + \left|t - t'\right|.$$

Discretizing θ and t with mesh sizes $\rho/4$ and $\rho^2/16$, respectively, gives a covering with radius strictly smaller than $\rho/2$ and at most a constant multiple of ρ^{-3} elements (think about the t-θ grid). Consequently,

$$\mathcal{M}(\rho, \mathcal{F}) \lesssim \rho^{-3}, \qquad 0 < \rho \le 1,$$

We argue the following inequality:

$$(\eta^*)^2 \lesssim_c \frac{\log(en)}{n}.$$

To see this, take any $r > 0$ satisfying $nr^2 \le \log \mathcal{M}_{\mathcal{F}}^{loc}(r, c)$. We split the region of r as follows:
(i), $\boldsymbol{r < n^{-1/2}}$: in this case, we trivially have $r^2 \le \frac{1}{n} \le \frac{\log(en)}{n}$;
(ii), $\boldsymbol{r \ge n^{-1/2}}$: in this case, every r/c-packing of $B_2(f, r) \cap \mathcal{F}$ is also a $1/c\sqrt{n}$-packing of the whole class: its points belong to $\mathcal{F}$, and their pairwise distances are at least $r/c \ge 1/(c\sqrt{n})$. Therefore,

$$\mathcal{M}_{\mathcal{F}}^{\mathrm{loc}}(r, c) \le \mathcal{M}\left(\frac{1}{c\sqrt{n}}, \mathcal{F}\right) \lesssim c^3 n^{3/2}.$$

Taking logarithms gives

$$nr^2 \lesssim \log(c^3) + \frac{3}{2}\log n \lesssim \log(en).$$

Since r is any value that satisfies $nr^2 \le \log \mathcal{M}_{\mathcal{F}}^{loc}(r, c)$, taking the limit yields the desired inequality.

For every fixed $\sigma > 0$, the minimax lower bound above is asymptotically larger than the minimax rate determined by the critical radius. Therefore, replacing conditional variance control by an unconditional second-moment bound can invalidate the claimed minimax rate even when $\epsilon = 0$.

4 Minimax Rate under Sub-Gaussian Assumptions

In this section, we study the minimax rate under the sub-Gaussian conditions on ξ and the increment of functions in $\mathcal{F}$. Namely, we assume that there exist (known) constants L_{sg} and L_ξ such that for each integer $q \ge 1$:

$$\begin{aligned} \|f - g\|_{L_q(\mathbb{P}_X)} &\le L_{\mathrm{sg}}\sqrt{q}\,\|f - g\|_{L_2(\mathbb{P}_X)}\,, \\ \|\xi\|_q &\le L_\xi\sqrt{q}\,\|\xi\|_2\,, \end{aligned} \tag{4.1}$$

Recall that we previously assumed an upper bound on the conditional variance of ξ and demonstrated that it is necessary for the derivations of minimax rate (3.14) under the Bernstein's condition (3.1). However, we find that, under the stronger sub-Gaussian type conditions (4.1), only the more general marginal control on the second moment of ξ is required. Namely, in this section, σ^2 is the variance proxy of ξ as we assume

$$\mathbb{E}\xi^2 \le \sigma^2. \tag{4.2}$$

We denote the distribution family of the noise satisfying (4.1) and (4.2) as $\mathcal{P}_\xi^{\mathrm{sg}}$.

Though we say that (4.1) is stronger compared with (3.1), the implication from (4.1) to (3.1) requires some care regarding the dependence of the constants on the diameter. Assume without loss of generality that $L_{\mathrm{sg}} \ge 1$. The first inequality in (4.1) gives, for every integer $q \ge 1$,

$$\|f - g\|_{L_{2q}(\mathbb{P}_X)}^{2q} \le (2q)^q L_{\mathrm{sg}}^{2q}\,\|f - g\|_{L_2(\mathbb{P}_X)}^{2q} \le q!\,(2eL_{\mathrm{sg}}^2)^q d^{2q-2}\,\|f - g\|_{L_2(\mathbb{P}_X)}^2\,,$$

where we used $\|f - g\|_{L_2(\mathbb{P}_X)} \le d$ and $q^q \le e^q q!$. Setting $B = 2eL_{\mathrm{sg}}^2$, we have $B \ge 1$ and $2B^q \le (2B^2)^{q-1}$ for $q \ge 2$. Consequently,

$$\|f - g\|_{L_{2q}(\mathbb{P}_X)}^{2q} \le \frac{q!}{2} L^{q-1} \|f - g\|_{L_2(\mathbb{P}_X)}^2, \qquad L = 8e^2 L_{\mathrm{sg}}^4 d^2.$$

Thus, for a class of finite diameter, (4.1) implies (3.1), but the resulting moment constant is allowed to depend on d. The noise assumption is not needed for this implication.

This dependence reflects an essential difference between the two conditions. Recall that (3.1) implies $d^2 \le L$. Therefore, fixing L automatically bounds the diameter. In contrast, the relative

sub-Gaussian condition is invariant under multiplication of every function by the same positive scalar, so fixing L_{sg} does not control d, as shown by the following example.

Example 4.1. *Consider an uncontaminated scenario. Let $X \sim \mathrm{Bernoulli}(1/2)$ and, for any $d > 0$, consider the convex, complete, star-shaped class*

$$\mathcal{F}_d = \{f_\theta : \theta \in [0, d]\}, \qquad f_\theta(x) = \sqrt{2}\,\theta x.$$

For any $\theta, \theta' \in [0, d]$,

$$\|f_\theta - f_{\theta'}\|_{L_2(\mathbb{P}_X)} = |\theta - \theta'|,$$

so the diameter of $\mathcal{F}_d$ is exactly d. Moreover, for every $p \geq 2$,

$$\|f_\theta - f_{\theta'}\|_{L_p(\mathbb{P}_X)} = 2^{1/2-1/p}|\theta - \theta'| \leq \sqrt{p}\,\|f_\theta - f_{\theta'}\|_{L_2(\mathbb{P}_X)}.$$

Hence the relative sub-Gaussian condition holds with $L_{\mathrm{sg}} = 1$, uniformly over all $d > 0$. However, applying the general moment condition with $p = 2$ to the endpoints f_0 and f_d requires

$$2d^4 = \|f_d - f_0\|^4_{L_4(\mathbb{P}_X)} \leq L\|f_d - f_0\|^2_{L_2(\mathbb{P}_X)} = Ld^2.$$

Thus any admissible moment constant must satisfy $L \geq 2d^2$, even though L_{sg} remains equal to one.

This example also illustrates why an additional diameter bound is relevant to a uniform statistical guarantee. Consider the clean regression model $Y = f_\theta(X) + \xi$, with $\xi \sim \mathcal{N}(0, 1)$ independent of X. On the event

$$A_n = \{X_1 = \cdots = X_n = 0\}, \qquad \mathbb{P}(A_n) = 2^{-n},$$

the conditional distribution of the observations is identical for all θ. In particular, the endpoints f_0 and f_d cannot be distinguished on this event. Using

$$\frac{\|u - f_0\|_2^2 + \|u - f_d\|_2^2}{2} = \left\|u - \frac{f_0 + f_d}{2}\right\|_2^2 + \frac{d^2}{4},$$

we obtain, even without contamination,

$$\inf_{\widehat{f}} \sup_{\theta \in [0,d]} \mathbb{E}_\theta \|\widehat{f} - f_\theta\|_2^2 \geq \frac{d^2}{4}\, 2^{-n}.$$

Meanwhile, $\mathcal{F}_d$ is isometric to the interval $[0, d]$, whose local packing entropy is bounded by a constant for a fixed packing ratio. Consequently, $(\eta^*)^2 \asymp n^{-1}$ whenever $d \geq n^{-1/2}$. Allowing d to grow arbitrarily therefore rules out an upper bound of order $(\eta^*)^2$ with a multiplicative constant depending only on the fixed sub-Gaussian and noise parameters.

Accordingly, in the sub-Gaussian setting we impose a common, known bound

$$d \leq M,$$

where M is fixed independently of n, p and of the particular class under consideration. Constants may depend on M, but remain uniform over the actual diameters $d \leq M$. This is stronger than merely requiring each individual class to have finite diameter. Under this assumption, the conversion above is valid with the uniform choice $L = 8e^2 L_{\mathrm{sg}}^4 M^2$.

4.1 Lower Bound

We now derive the lower bound of the minimax rate of (1.3) relevant to the corruption rate ϵ in the following lemma, complementing Lemma 2.3 under the sub-Gaussian setting, where the admissible distributions include $\mathcal{N}(0, \sigma^2)$

Lemma 4.2 (Lower bound in the sub-Gaussian case)**.** *For the nonparametric regression problem formulated in Section 1, suppose that the admissible noise distributions include $\mathcal{N}(0, \sigma^2)$. Then*

$$\mathcal{R}_{n,\epsilon}(\mathcal{F}; \mathcal{P}_\xi^{\mathrm{sg}}) \geq \frac{1}{96}\big(\sigma^2\epsilon^2 \wedge d^2\big).$$

This lower bound already holds when the noise distribution is fixed to $\mathcal{N}(0, \sigma^2)$.

Similar to Lemma 3.1, the proof is based on a two-point construction. This time, we select two functions in $\mathcal{F}$ whose $L_2(\mathbb{P}_X)$-distance is of order $\sigma\epsilon \wedge d$. By Pinsker's inequality and a coupling argument, we argue that the probability that the number of different samples from the two hypotheses does not exceed the corruption budget $\lfloor \epsilon n \rfloor$ is at least $3/4$, in which case the adversary can tamper with the original dataset to make them identical, yielding the results.

Again, combining Lemma 4.2 with the general lower bound (2.1) in the uncontaminated case, we conclude that the minimax rate is lower bounded by

$$\mathcal{R}_{n,\epsilon}(\mathcal{F}; \mathcal{P}_\xi^{\mathrm{sg}}) \gtrsim \big[(\eta^*)^2 \vee \epsilon^2\big] \wedge d^2, \tag{4.3}$$

where the implicit constant is independent of n, ϵ, and the actual diameter d.

Lemma 4.2 suggests a minimax rate in the order of ϵ^2. We note here that Example 3.2, with a minimax lower bound ϵ, does not meet the sub-Gaussian requirement (4.1). Indeed, consider $h_t = \mathbf{1}_{\{x \in [0,t]\}}$ for some $t \in [0, 1]$. For $p > 2$, we have

$$\|h_t\|_{L_p(\mathbb{P}_X)} = t^{\frac{1}{p}}, \|h_t\|_{L_2(\mathbb{P}_X)} = \sqrt{t}.$$

Therefore the ratio

$$\frac{\|h_t\|_{L_p(\mathbb{P}_X)}}{\sqrt{p}\,\|h_t\|_{L_2(\mathbb{P}_X)}} = \frac{t^{\frac{1}{p}-\frac{1}{2}}}{\sqrt{p}}$$

diverges as $t \to 0$, which implies that there is no universal L_{sg} for which (4.1) holds uniformly.

4.2 Upper Bound

The constructive design required for the upper bound is also the tree-based packings given in Algorithm 1 and the selection criterion given in Algorithm 2, described in Section 3.2. However, as shown in the following lemmas, the sub-Gaussian conditions (4.1) allow us to deduce stronger properties for the pairwise comparison statistic $Z_{f,g}$, which consequently leads to the wider region of feasibility of Algorithm 2.

The following lemma indicates that under (4.1), $Z_{f,g}$ behaves as a sub-exponential random variable for each fixed pair of f, g. We defer the proof to Appendix C.2. The reader can verify that the marginal second moment control $\mathbb{E}\xi^2 \leq \sigma^2$ is sufficient in the proof as we leverage Hölder's inequality.

Lemma 4.3. *For the nonparametric regression problem formulated in Section 1.1 and the pairwise comparison statistic $Z_{f,g}$ defined in (3.7), under Conditions (4.1) and (4.2), and*

$$\|g - f^*\|_{L_2(\mathbb{P}_X)} \le \delta, \qquad \|f - g\|_{L_2(\mathbb{P}_X)} \ge C\delta$$

for some $\delta > 0, C > 3$, we have

$$\|Z_{f,g} - \mathbb{E}Z_{f,g}\|_{L_p(\mathbb{P}_{\xi,X})} \le K(\|f - g\|_{L_2(\mathbb{P}_X)} + \sigma)p \tag{4.4}$$

for any $p \ge 2$, where K is a constant depending on L_{sg}, L_ξ, and C.

Recall that the effectiveness of trimmed mean is fundamental for $Z_{f,g}$ to correctly select between two candidates with one close to f^* while the other is sufficiently far. As remarked in Lugosi and Mendelson [2021], their result (Lemma 3.6) can be reinforced when the underlying distribution satisfies stronger moment control. The following lemma establishes one such reinforcement under the sub-Gaussian conditions (4.1).

Lemma 4.4. *For two functions $f, g \in \mathcal{F}$, assume that g is close to f^* while f is far from f^* with respect to $\delta > 0$:*

$$\|g - f^*\|_{L_2(\mathbb{P}_X)} \le \delta, \qquad \|f - g\|_{L_2(\mathbb{P}_X)} \ge C\delta.$$

Suppose additionally we have (4.1), (4.2), and $\epsilon \le \min\{1/128, \epsilon_0\}$ for some ϵ_0 depending only on L_{sg}, L_ξ, C (will be determined later). Then for $t_\delta := a\min\left\{\delta^2/\sigma^2, 1\right\}$, $a \le \min\left\{\frac{1}{384}, \frac{1}{1024K^2e^2}\right\}$, $K = 8L_{sg}(L_{sg} + L_\xi)$, as long as $\delta \ge \frac{3456Ke}{C}\sigma\epsilon\log\left(\frac{e}{\epsilon}\right)$, the test $\psi_{f,g}(X, Y)$ from (3.8) can correctly select g over f with probability higher than $1 - 4\exp\{-nt_\delta\}$.

The proof of Lemma 4.4 relies on sharply bounding $\mathcal{E}(4\zeta, Z_{f,g})$ upon (4.1), using the $L_p(\mathbb{P}_{X,\xi})$-norm control derived in Lemma 4.3. We defer the proof to Appendix C.3 for the reader's reference.

Analogous to Corollary 3.8, we conclude the error probability of one-step incorrect selection along the trajectory of the tree traversal. The following argument is straightforward since any nonvacuous comparison must have $\delta \le d \le M$ and consequently we have $t_\delta = a\min\left\{\delta^2/\sigma^2, 1\right\} \ge a\delta^2/(M^2 + \sigma^2)$.

Corollary 4.5. *Under the same conditions and assumptions as Lemma 4.4, we have*

$$\sup_{f^*\in\mathcal{F}} \sup_{\|g-f^*\|_{L_2(\mathbb{P}_X)}\le\delta, \|f-g\|_{L_2(\mathbb{P}_X)}\ge C\delta} \mathbb{P}\left(\psi_{f,g}(X, Y) = 0\right) \le 4\exp\left\{-\frac{an\delta^2}{M^2 + \sigma^2}\right\}. \tag{4.5}$$

Equivalently, the probability of incorrect selection scales with $n\delta^2$. When $4\exp\left\{-an\delta^2/(M^2 + \sigma^2)\right\} \ge 1$, this inequality trivially holds.

We comment here that once we obtain Lemma 4.4 and Corollary 4.5, the rest of the argument parallels that derived in Section 3. Specifically, an alternative version of Lemma 3.9, Corollary 3.10, Theorem 3.11, and Corollary 3.12 and their proofs continue to hold under the sub-Gaussian case. The only, but also the crucial difference is that the feasible region of Lemma 4.4 is now $\delta \ge \frac{3456K}{C}\sigma\epsilon\log\left(\frac{e}{\epsilon}\right)$ compared with $\delta^2 \gtrsim_{L,\sigma} \epsilon$ required in Corollary 3.8. Such difference directly changes the performance of Algorithm 2, where the term regarding ϵ changes from ϵ to $\epsilon^2\log^2\left(\frac{e}{\epsilon}\right)$. Formally, we have the following Theorem for the sub-Gaussian case.

Theorem 4.6. *Define J^* in the same way as in Theorem 3.11, where δ_j, η_j are similarly defined as:*

$$\delta_j = \frac{d}{2^{j-1}(C+1)}, \qquad \eta_j := \sqrt{\frac{a}{M^2+\sigma^2}}\delta_j, \qquad a \le \min\left\{\frac{1}{384}, \frac{1}{1024K^2e^2}\right\}.$$

For the nonparametric regression problem defined in Section 1.1, suppose that the function class $\mathcal{F}$ and the noise ξ satisfy the sub-Gaussian conditions (4.1) and the marginal control (4.2). For the output of Algorithm 2 $f^*_{J^*}$ after J^*-1 iterations, we have

$$\sup_{f^*\in\mathcal{F}}\sup_{\xi:\mathbb{P}_\xi\in\mathcal{P}_\xi^{sg}}\sup_{\mathcal{C}\in\mathfrak{C}_{n,\epsilon}} \mathbb{E}_{\widetilde{X},\widetilde{Y}}\left\|f^*_{J^*}(\mathcal{C}(\widetilde{X},\widetilde{Y})) - f^*\right\|^2_{L_2(\mathbb{P}_X)} \lesssim \left(\eta^2_{J^*} \vee \epsilon^2\log^2\left(\frac{e}{\epsilon}\right)\right) \wedge d^2. \tag{4.6}$$

Combining with (4.3), we conclude that the minimax rate of (1.3) is characterized by (4.3) and (4.6) up to a logarithmic factor $\log^2\left(\frac{e}{\epsilon}\right)$.

5 Examples

5.1 Sparse linear regression

Sparse linear regression provides a classical example of a star-shaped function class that need not be convex. Its uncontaminated minimax rates have been studied extensively; see, for example, Raskutti et al. [2011]. Robust estimation under Huber's contamination model is considered in Gao [2020]. We show below how our general results specialize to sparse linear regression with Gaussian design under strong adversarial contamination.

Let $\widetilde{X} \sim \mathcal{N}(0, \mathbf{I}_p)$ be a Gaussian random vector (when $\widetilde{X}$ follows a sub-Gaussian distribution with zero mean and isotropic covariance matrix, the following arguments hold with minor modification), and let s, with $1 \le s \le p$, denote the sparsity level of the model, and fix $0 < R \le M/2$. Consider the parameter set and associated function class

$$\Theta_{s,R} := \{\beta \in \mathbb{R}^p : \|\beta\|_0 \le s,\ \|\beta\|_2 \le R\}, \qquad \mathcal{F}_{s,R} := \left\{f_\beta : f_\beta(x) = x^\top\beta,\ \beta \in \Theta_{s,R}\right\}.$$

The class $\mathcal{F}_{s,R}$ is star-shaped about the zero function and has diameter $d = 2R \le M$. Moreover, for any $\beta, \gamma \in \Theta_{s,R}$, Gaussianity and isotropy of X give

$$f_\beta(X) - f_\gamma(X) \sim \mathcal{N}\left(0, \|\beta-\gamma\|_2^2\right).$$

Therefore, we have a general expression for $L_q(\mathbb{P}_X)$-norm for any integer $q \ge 2$:

$$\|f_\beta - f_\gamma\|_{L_2(\mathbb{P}_X)} = \|\beta-\gamma\|_2, \qquad \|f_\beta - f_\gamma\|_{L_q(\mathbb{P}_X)} = \|\beta-\gamma\|_2 \|G\|_q,$$

where $\|G\|_q = \left(\mathbb{E}|G|^q\right)^{1/q}, G \sim \mathcal{N}(0,1)$. From integration by parts and Hölder's inequality:

$$\mathbb{E}|G|^q = (q-1)\mathbb{E}|G|^{q-2} \le (q-1)\left(\mathbb{E}|G|^q\right)^{(q-2)/q},$$

where we conclude $\|G\|_q^2 \le q-1 < q (q \ge 2)$. When $q=1$, we directly have $\|G\|_1 \le \|G\|_2$ Hence, we have $\|f_\beta - f_\gamma\|_{L_q(\mathbb{P}_X)} \le \sqrt{q}\|f_\beta - f_\gamma\|_{L_2(\mathbb{P}_X)}$, which means the first condition in (4.1) holds with $L_{\text{sg}} = 1$. By the moment comparison in Section 4 and the bound $d \le M$, the Bernstein condition (3.1) also holds with $L = 8e^2M^2$.

It remains to calculate the local entropy. Write $k := s \log(ep/s)$. Fix a center $\beta_0 \in \Theta_{s,R}$ and a radius $t > 0$. For each coordinate set $S \subseteq [p]$ of cardinality s, the vectors supported on S and lying within distance t of β_0 are contained in an s-dimensional Euclidean ball of radius t, centered at the orthogonal projection of β_0 onto that coordinate subspace. The volumetric packing bound therefore gives

$$\mathcal{M}^{\mathrm{loc}}_{\mathcal{F}_{s,R}}(t, c) \le \binom{p}{s} (1 + 2c)^s .$$

Taking logarithms yields, uniformly over $t > 0$,

$$\log \mathcal{M}^{\mathrm{loc}}_{\mathcal{F}_{s,R}}(t, c) \le s \log\left(\frac{ep}{s}\right) + s \log(1 + 2c) \le C_c k,$$

where $C_c > 0$ depends only on the fixed packing constant c.

For the reverse inequality, the standard sparse packing construction, combined with a Euclidean packing on a fixed support when $s > p/2$, provides universal constants $a, b > 0$ and a finite set

$$\mathcal{V} \subseteq \{v \in \mathbb{R}^p : \|v\|_0 \le s,\ \|v\|_2 \le 1\}$$

such that

$$\log|\mathcal{V}| \ge ak, \qquad \|v - w\|_2 \ge b \quad \text{for all distinct } v, w \in \mathcal{V};$$

see the sparse packing arguments in Raskutti et al. [2011]. For $0 < t \le R$, the functions $\{f_{tv} : v \in \mathcal{V}\}$ belong to $B_2(0, t) \cap \mathcal{F}_{s,R}$ and have pairwise distances at least bt. Taking the universal constant c sufficiently large that $1/c \le b$, we obtain

$$\log \mathcal{M}^{\mathrm{loc}}_{\mathcal{F}_{s,R}}(t, c) \ge ak, \qquad 0 < t \le R.$$

These bounds identify the critical radius, including the regime where the diameter determines the risk. Indeed, the upper entropy bound implies $(\eta^*)^2 \le C_c k/n$. Conversely, setting $t_0 := R \wedge \sqrt{ak/n}$ gives $nt_0^2 \le ak \le \log \mathcal{M}^{\mathrm{loc}}_{\mathcal{F}_{s,R}}(t_0, c)$, so the definition of η^* implies $\eta^* \ge t_0$. Since $d = 2R$, it follows that

$$(\eta^*)^2 \wedge d^2 \asymp \frac{s \log(ep/s)}{n} \wedge d^2.$$

Therefore, for fixed $\sigma > 0$ and the strong contamination model from Section 1.1, under the conditional moment assumptions (3.2), we have

$$\mathcal{R}_{n,\epsilon}\left(\mathcal{F}_{s,R}; \mathcal{P}^{\mathrm{b}}_{\xi}\right) \asymp \left[\frac{s \log(ep/s)}{n} \vee \epsilon\right] \wedge d^2.$$

Under the sub-Gaussian noise condition in (4.1), together with $\mathbb{E}[\xi \mid X] = 0$ and $\mathbb{E}\xi^2 \le \sigma^2$, we instead obtain

$$\begin{aligned}\left[\frac{s \log(ep/s)}{n} \vee \epsilon^2\right] \wedge d^2 &\lesssim \mathcal{R}_{n,\epsilon}\left(\mathcal{F}_{s,R}; \mathcal{P}^{\mathrm{sg}}_{\xi}\right) \\ &\lesssim \left[\frac{s \log(ep/s)}{n} \vee \epsilon^2 \log^2\left(\frac{e}{\epsilon}\right)\right] \wedge d^2.\end{aligned}$$

The sub-Gaussian bounds hold for $0 \le \epsilon \le \epsilon_0$, where $\epsilon_0 > 0$ is a sufficiently small constant as required in Section 4; the logarithmic term is interpreted as zero when $\epsilon = 0$. The implicit constants depend only on M, σ, c and, in the sub-Gaussian case, L_ξ. In particular, they are independent of n, p, s, ϵ and the actual diameter $d = 2R$.

The heavy-tailed conclusion concerns the supremum over the full conditional-moment noise family $\mathcal{P}_\xi^{\mathrm{b}}$, which allows dependence between ξ and X. In the sub-Gaussian setting, the bounds retain a logarithmic gap in their dependence on ϵ. For comparison, Gao [2020] obtains an ϵ^2 contamination term for sparse regression with Gaussian errors under Huber's contamination model while we consider a more general class of noise distributions and the strong contamination model.

5.2 Low-rank trace regression

Low-rank trace regression is a matrix analogue of linear regression in which the unknown coefficient matrix has a low-rank structure. Classical results on estimation and minimax rates in this model can be found in Rohde and Tsybakov [2011], Candès and Plan [2011]. We consider Gaussian design matrices, for which the prediction distance coincides with the Frobenius distance between the coefficient matrices.

Let $m_1, m_2 \geq 1$, $1 \leq r \leq \min\{m_1, m_2\}$, and $0 < R \leq M/2$, where M is a fixed constant. Suppose that the entries of $X \in \mathbb{R}^{m_1 \times m_2}$ are independent standard Gaussian random variables, and consider the regression model

$$Y = \operatorname{tr}\left(X^\top A^*\right) + \xi, \qquad A^* \in \Theta_{r,R} := \left\{A \in \mathbb{R}^{m_1 \times m_2} : \operatorname{rank}(A) \leq r,\ \|A\|_{\mathrm{F}} \leq R\right\}.$$

Define the corresponding function class and its effective dimension by

$$\mathcal{F}_{r,R} := \left\{f_A : f_A(X) = \operatorname{tr}\left(X^\top A\right),\ A \in \Theta_{r,R}\right\}, \qquad k_r := r(m_1 + m_2 - r).$$

Here $\|A\|_{\mathrm{F}}^2 = \sum_{j=1}^{m_1}\sum_{k=1}^{m_2} A_{jk}^2$. The class $\mathcal{F}_{r,R}$ is star-shaped around the zero function and has diameter $d = 2R \leq M$, as witnessed by the two rank-one matrices $Re_1e_1^\top$ and $-Re_1e_1^\top$, with the two coordinate vectors taken in their respective ambient spaces. Moreover, for any $A, B \in \Theta_{r,R}$, independence and standardization of the Gaussian entries give

$$\|f_A - f_B\|_{L_2(\mathbb{P}_X)}^2 = \mathbb{E}\left[\operatorname{tr}\left(X^\top(A-B)\right)^2\right] = \|A - B\|_{\mathrm{F}}^2.$$

Using the same argument as in the preceding example, $f_A(X) - f_B(X)$ is a centered Gaussian random variable with variance $\|A - B\|_{\mathrm{F}}^2$. Hence, for every $q \geq 3$,

$$\|f_A - f_B\|_{L_q(\mathbb{P}_X)} \leq \sqrt{q}\,\|A - B\|_{\mathrm{F}} = \sqrt{q}\,\|f_A - f_B\|_{L_2(\mathbb{P}_X)}.$$

We now calculate the local entropy. Fix the constant $c \geq 2$ appearing in the definition of η^*, and write

$$H_c(t) := \log \mathcal{M}^{\mathrm{loc}}_{\mathcal{F}_{r,R}}(t, c).$$

We will show that there are positive constants a_c, b_c, depending only on c, such that

$$H_c(t) \leq b_c k_r \quad \text{for every } t > 0, \qquad H_c(t) \geq a_c k_r \quad \text{for } 0 < t \leq R.$$

For the upper bound, let $\mathcal{N}(u, \mathcal{A})$ denote the covering number of a matrix set $\mathcal{A}$ by Frobenius balls of radius u. For $1 \leq k \leq \min\{m_1, m_2\}$, define

$$\mathcal{S}_k := \{A : \operatorname{rank}(A) \leq k,\ \|A\|_{\mathrm{F}} = 1\}, \qquad \mathcal{K}_k := \{A : \operatorname{rank}(A) \leq k,\ \|A\|_{\mathrm{F}} \leq 1\}.$$

The low-rank covering bound of Candès and Plan [2011] states that, for $0 < u < 1$,

$$\mathcal{N}(u, \mathcal{S}_k) \leq \left(\frac{9}{u}\right)^{(m_1+m_2+1)k}.$$

Their result concerns the unit sphere. To obtain the corresponding bound for the unit ball, take a $u/2$-net of $\mathcal{S}_k$ and a grid of $[0,1]$ with mesh at most $u/2$ and cardinality at most $4/u$. Every nonzero matrix in $\mathcal{K}_k$ can be written as aU, where $a \in [0,1]$ and $U \in \mathcal{S}_k$. Choosing a grid point b and a net point V with $|a-b| \leq u/2$ and $\|U-V\|_{\mathrm{F}} \leq u/2$ yields

$$\|aU - bV\|_{\mathrm{F}} \leq a\|U-V\|_{\mathrm{F}} + |a-b|\|V\|_{\mathrm{F}} \leq u.$$

Including zero in the radial grid therefore gives

$$\mathcal{N}(u, \mathcal{K}_k) \leq \frac{4}{u}\left(\frac{18}{u}\right)^{(m_1+m_2+1)k} \leq \left(\frac{18}{u}\right)^{(m_1+m_2+1)k+1}.$$

Fix any center $A_0 \in \Theta_{r,R}$ and radius $t > 0$. Differences between two matrices of rank at most r have rank at most $2r$. Thus, setting $r_0 := \min\{2r, m_1, m_2\}$, we have

$$\{A - A_0 : A \in \Theta_{r,R},\ \|A - A_0\|_{\mathrm{F}} \leq t\} \subseteq t\mathcal{K}_{r_0}.$$

Applying the preceding covering bound with $u = 1/(3c)$ shows that $t\mathcal{K}_{r_0}$ can be covered by at most $(54c)^{(m_1+m_2+1)r_0+1}$ Frobenius balls of radius $t/(3c)$. Each such ball contains at most one point of a t/c-separated set, because its diameter is strictly smaller than t/c. Taking the supremum over A_0 and using the isometry established above yields

$$H_c(t) \leq [(m_1+m_2+1)r_0 + 1]\log(54c) \leq 7k_r \log(54c).$$

For the last inequality, we used $r_0 \leq 2r$, $m_1+m_2+1 \leq 3(m_1+m_2)/2$, and $k_r = r(m_1+m_2-r) \geq \frac{r(m_1+m_2)}{2} \geq 1$.

This proves the upper bound for every $t > 0$. In particular, the covering resolution after rescaling is the fixed number $1/(3c)$, which explains why no factor involving $\log(R/t)$ appears in the local entropy.

For the lower bound, suppose without loss of generality that $m_1 \geq m_2$; otherwise, transpose all matrices. Consider the linear subspace

$$\mathcal{V} := \left\{A = [B \ \ 0] : B \in \mathbb{R}^{m_1 \times r}\right\}.$$

Every matrix in $\mathcal{V}$ has rank at most r, and $\mathcal{V}$ has dimension $D = m_1 r$. Therefore, for $0 < t \leq R$, the set

$$\{A \in \mathcal{V} : \|A\|_{\mathrm{F}} \leq t\}$$

is contained in the local parameter set centered at zero and is isometric to a Euclidean ball of radius t in $\mathbb{R}^D$.

Take a maximal t/c-separated subset of this Euclidean ball, and denote its cardinality by N. Maximality implies that balls of radius t/c around these points cover the radius-t ball. Comparing D-dimensional volumes gives

$$N \geq \frac{\mathrm{vol}_D\left(t\mathbb{B}_2^D\right)}{\mathrm{vol}_D\left((t/c)\,\mathbb{B}_2^D\right)} = c^D,$$

where $\mathbb{B}_2^D$ is the Euclidean unit ball. Since the zero function belongs to $\mathcal{F}_{r,R}$, this packing supplies a lower bound for its local packing number. Consequently,

$$H_c(t) \geq D \log c \geq \frac{k_r}{2} \log c, \qquad 0 < t \leq R,$$

where we used $k_r \leq r(m_1 + m_2) \leq 2m_1 r = 2D$. Combining the two bounds proves

$$\log \mathcal{M}^{\mathrm{loc}}_{\mathcal{F}_{r,R}}(t, c) \asymp_c r(m_1 + m_2 - r), \qquad 0 < t \leq R.$$

It remains to calculate the critical radius. Set $a_c = (\log c)/2$ and $b_c = 7 \log(54c)$. By the definition of η^* and the entropy upper bound valid at every radius, we have $(\eta^*)^2 \leq \frac{b_c k_r}{n}$.

For the reverse direction, take $t_0 := R \wedge \sqrt{\frac{a_c k_r}{n}}$. Since $0 < t_0 \leq R$, the entropy lower bound gives $n t_0^2 \leq a_c k_r \leq H_c(t_0)$. Thus t_0 belongs to the set defining η^*, and $\eta^* \geq t_0$. Recalling that $d = 2R$, we conclude that

$$(\eta^*)^2 \wedge d^2 \asymp_c \frac{r(m_1 + m_2 - r)}{n} \wedge d^2.$$

Substituting this critical radius into Corollary 3.12 proves the rate of (1.3) for the conditional second-moment noise family. Specifically, under the strong corruption model, for sufficiently small ϵ, if the noise satisfies Condition (3.2), then

$$\mathcal{R}_{n,\epsilon}\left(\mathcal{F}_{r,R}; \mathcal{P}_\xi^{\mathrm{b}}\right) \asymp \left[\frac{k_r}{n} \vee \epsilon\right] \wedge d^2.$$

Applying the upper and lower bounds from Section 4 controls the rate under the sub-Gaussian condition (4.1):

$$\left[\frac{k_r}{n} \vee \epsilon^2\right] \wedge d^2 \lesssim \mathcal{R}_{n,\epsilon}\left(\mathcal{F}_{r,R}; \mathcal{P}_\xi^{\mathrm{sg}}\right) \lesssim \left[\frac{k_r}{n} \vee \epsilon^2 \log^2\left(\frac{e}{\epsilon}\right)\right] \wedge d^2.$$

The corruption term is interpreted as zero when $\epsilon = 0$. The implicit constants may depend on the fixed parameters M, $\sigma > 0$, and the sub-Gaussian noise constant, but do not depend on n, m_1, m_2, r, R, d, or ϵ.

5.3 Hölder smooth regression

We next consider regression over a Hölder ball on a fixed-dimensional domain. Throughout this example, the smoothness parameter $\alpha > 0$ and the dimension $m \geq 1$ are fixed. Constants may depend on α and m, but remain independent of the sample size, corruption proportion, and actual diameter of the function class.

Let $k := \lceil \alpha \rceil - 1$ and $\beta := \alpha - k \in (0, 1]$. For a function with continuous partial derivatives through order k on $[0,1]^m$, define

$$\|f\|_{C^\alpha} := \max_{|\nu| \leq k} \|D^\nu f\|_\infty + \max_{|\nu| = k} \sup_{\substack{x, y \in [0,1]^m \\ x \neq y}} \frac{|D^\nu f(x) - D^\nu f(y)|}{\|x - y\|_2^\beta},$$

where ν is a multi-index and $|\nu|$ is its total order. This convention includes integer smoothness parameters by requiring the derivatives of order $\alpha - 1$ to be Lipschitz continuous.

Suppose that X is uniformly distributed on $[0,1]^m$ and

$$Y = f^*(X) + \xi, \qquad f^* \in \mathcal{F}_{\alpha,R} := \{f : \|f\|_{C^\alpha} \le R\}, \qquad 0 < R \le M/2,$$

where M is a fixed constant. Assume that the noise belongs to the family $\mathcal{P}_\xi^{\mathrm{b}}$ satisfying

$$\mathbb{E}[\xi \mid X] = 0, \qquad \mathbb{E}\left[\xi^2 \mid X\right] \le \sigma^2.$$

The class $\mathcal{F}_{\alpha,R}$ is a norm ball and contains the zero function, so it is convex and star-shaped. Since $\|f\|_\infty \le R$ for every $f \in \mathcal{F}_{\alpha,R}$, its $L_2(\mathbb{P}_X)$-diameter is at most $2R$. Equality follows by considering the two constant functions $f \equiv R$ and $g \equiv -R$.

Moreover, for every integer $q \ge 2$ and $f, g \in \mathcal{F}_{\alpha,R}$,

$$\begin{aligned}
\|f-g\|_{L_{2q}(\mathbb{P}_X)}^{2q} &\le \|f-g\|_\infty^{2q-2} \|f-g\|_{L_2(\mathbb{P}_X)}^2 \\
&\le M^{2q-2} \|f-g\|_{L_2(\mathbb{P}_X)}^2 \\
&\le \frac{q!}{2} \left(M^2\right)^{q-1} \|f-g\|_{L_2(\mathbb{P}_X)}^2 .
\end{aligned}$$

Thus Condition (3.1) holds with $L = M^2$, independently of the actual diameter d.

We next calculate the local entropy. Fix the sufficiently large constant c used in the definition of η^*, taking $c \ge 4$, and write

$$H_c(t) := \log \mathcal{M}^{\mathrm{loc}}_{\mathcal{F}_{\alpha,R}}(t, c).$$

We will establish

$$H_c(t) \asymp_{\alpha,m,c} \left(\frac{R}{t}\right)^{m/\alpha}, \qquad 0 < t \le R.$$

For the upper bound, let $\mathcal{N}(u, \mathcal{F}, \|\cdot\|_\infty)$ denote the covering number in the uniform norm. Theorem 2.7.1 of van der Vaart and Wellner [1996], applied to the unit Hölder ball on $[0,1]^m$, gives

$$\log \mathcal{N}(u, \mathcal{F}_{\alpha,1}, \|\cdot\|_\infty) \le A_{\alpha,m} u^{-m/\alpha}, \qquad 0 < u \le 1,$$

where $A_{\alpha,m}$ depends only on α and m. Scaling the function class by R yields

$$\log \mathcal{N}(u, \mathcal{F}_{\alpha,R}, \|\cdot\|_\infty) \le A_{\alpha,m} \left(\frac{R}{u}\right)^{m/\alpha}, \qquad 0 < u \le R.$$

For $u \ge R$, the class is covered by the single uniform-norm ball centered at zero with radius u, so the same upper bound remains valid.

Because $\mathbb{P}_X$ is a probability measure, $\|f-g\|_{L_2(\mathbb{P}_X)} \le \|f-g\|_\infty$. Therefore, a uniform-norm ball of radius $t/(3c)$ contains at most one point of a t/c-separated set in $L_2(\mathbb{P}_X)$. Every local packing is also a packing of the entire class, and hence

$$\begin{aligned}
H_c(t) &\le \log \mathcal{N}\left(\frac{t}{3c}, \mathcal{F}_{\alpha,R}, \|\cdot\|_\infty\right) \\
&\le A_{\alpha,m} \left(\frac{3cR}{t}\right)^{m/\alpha} = b_{\alpha,m,c} \left(\frac{R}{t}\right)^{m/\alpha}, \qquad t > 0.
\end{aligned}$$

For the lower bound, we construct a packing inside the single local ball centered at zero. Choose a smooth function ψ compactly supported in $(1/4, 3/4)^m$, normalized so that

$$\int_{\mathbb{R}^m} \psi(x)^2 \, dx = 1.$$

For an integer $K \geq 1$, define the index set $\mathcal{J}_K := \{0, \ldots, K-1\}^m$ and let $N_K := K^m$. For each binary vector $\theta = (\theta_j)_{j \in \mathcal{J}_K} \in \{0,1\}^{N_K}$, set

$$g_\theta(x) := t \sum_{j \in \mathcal{J}_K} \theta_j \psi(Kx - j).$$

The bumps have disjoint supports, each lying strictly inside one of the grid cubes of side length K^{-1}. Differentiation contributes a factor $K^{|\nu|}$, and the β-Hölder seminorm of a derivative of order k contributes an additional factor K^β. The separation between supports controls differences involving distinct bumps. Consequently, there exists a constant $C_0 \geq 1$, depending only on α, m and the fixed bump ψ, such that $\|g_\theta\|_{C^\alpha} \leq C_0 t K^\alpha$ uniformly over K and θ. Since ψ is fixed once α, m are fixed, we regard C_0 as depending only on α, m.

Disjointness of the supports and the normalization of ψ also give the exact identities

$$\|g_\theta\|^2_{L_2(\mathbb{P}_X)} = t^2 K^{-m} \sum_{j \in \mathcal{J}_K} \theta_j^2 \leq t^2,$$

and

$$\|g_\theta - g_{\theta'}\|^2_{L_2(\mathbb{P}_X)} = t^2 K^{-m} \rho_{\mathrm{H}}(\theta, \theta'),$$

where ρ_{H} denotes Hamming distance.

First suppose that $0 < t \leq R/C_0$, and choose $K := \left\lfloor \left(\frac{R}{C_0 t}\right)^{1/\alpha} \right\rfloor$. Then $K \geq 1$ and $C_0 t K^\alpha \leq R$, so every g_θ belongs to $\mathcal{F}_{\alpha,R} \cap B_2(0,t)$. The Varshamov–Gilbert bound, as stated in Lemma 2.9 of Tsybakov [2009], supplies a subset $\Theta_K \subseteq \{0,1\}^{N_K}$ satisfying

$$\log |\Theta_K| \geq \frac{\log 2}{8} N_K, \qquad \rho_{\mathrm{H}}(\theta, \theta') \geq \frac{N_K}{8} \quad \text{for distinct } \theta, \theta' \in \Theta_K.$$

The cited lemma applies when $N_K \geq 8$; when $N_K < 8$, the all-zero and all-one vectors give these same inequalities. It follows that $\|g_\theta - g_{\theta'}\|_{L_2(\mathbb{P}_X)} \geq \frac{t}{\sqrt{8}} > \frac{t}{c}$. Thus these functions form an admissible local packing. Using $\lfloor x \rfloor \geq x/2$ for $x \geq 1$, we obtain

$$\begin{aligned} H_c(t) &\geq \frac{\log 2}{8} K^m \\ &\geq \frac{\log 2}{8 \cdot 2^m C_0^{m/\alpha}} \left(\frac{R}{t}\right)^{m/\alpha}, \qquad 0 < t \leq R/C_0. \end{aligned}$$

For the remaining radii $R/C_0 < t \leq R$, the constant functions $f \equiv t$ and $g \equiv -t$ both belong to $\mathcal{F}_{\alpha,R} \cap B_2(0,t)$ and are $2t$ apart. Hence $H_c(t) \geq \log 2$. Since $(R/t)^{m/\alpha} \leq C_0^{m/\alpha}$ in this range, this also gives a lower bound of the required order. Combining the two ranges, we have

$$a_{\alpha,m} \left(\frac{R}{t}\right)^{m/\alpha} \leq H_c(t) \leq b_{\alpha,m,c} \left(\frac{R}{t}\right)^{m/\alpha}, \qquad 0 < t \leq R,$$

for positive constants independent of R and t. The upper bound continues to hold for every $t > 0$.

We now calculate the critical radius. Define $v_n := R^{\frac{2m}{2\alpha+m}} n^{-\frac{2\alpha}{2\alpha+m}}$. Whenever $nt^2 \le H_c(t)$, the entropy upper bound implies $nt^{2+m/\alpha} \le b_{\alpha,m,c} R^{m/\alpha}$, and therefore $(\eta^*)^2 \le b_{\alpha,m,c}^{\frac{2\alpha}{2\alpha+m}} v_n$. Conversely, set $t_0 := R \wedge \left(a_{\alpha,m} R^{m/\alpha}/n\right)^{\frac{\alpha}{2\alpha+m}}$. Then $0 < t_0 \le R$ and

$$nt_0^2 \le a_{\alpha,m} \left(\frac{R}{t_0}\right)^{m/\alpha} \le H_c(t_0).$$

Thus $\eta^* \ge t_0$. Recalling $d = 2R$, we conclude that

$$(\eta^*)^2 \wedge d^2 \asymp_{\alpha,m,c} R^{\frac{2m}{2\alpha+m}} n^{-\frac{2\alpha}{2\alpha+m}} \wedge d^2.$$

Finally, by Corollary 3.12, under the strong corruption model, for sufficiently small ϵ, we have:

$$\mathcal{R}_{n,\epsilon}\left(\mathcal{F}_{\alpha,R}; \mathcal{P}_\xi^{\mathrm{b}}\right) \asymp_{\alpha,m,M,\sigma} \left[R^{\frac{2m}{2\alpha+m}} n^{-\frac{2\alpha}{2\alpha+m}} \vee \epsilon\right] \wedge d^2.$$

In particular, for fixed $R > 0$, the uncontaminated component has the classical nonparametric rate $n^{-2\alpha/(2\alpha+m)}$. The asserted minimax rate now follows from Corollary 3.12.

Remark. Although the functions in $\mathcal{F}_{\alpha,R}$ are uniformly bounded, the class does not satisfy the relative sub-Gaussian increment condition (4.1). Indeed, using the bump ψ in the argument above, let

$$h_K(x) := \frac{R}{C_0} K^{-\alpha} \psi(Kx), \qquad K \ge 1.$$

The same scaling argument shows that $h_K \in \mathcal{F}_{\alpha,R}$, whereas

$$\frac{\|h_K\|_{L_4(\mathbb{P}_X)}}{\|h_K\|_{L_2(\mathbb{P}_X)}} = K^{m/4} \frac{\|\psi\|_{L_4(\mathbb{R}^m)}}{\|\psi\|_{L_2(\mathbb{R}^m)}} \longrightarrow \infty.$$

Since the zero function also belongs to the class, no finite L_{sg} can make the required increment inequality hold uniformly, even at $q = 4$. Consequently, the sub-Gaussian minimax theorem cannot be applied to this full Hölder ball merely by imposing sub-Gaussian noise.

Remark. Section 3.1 of Gao [2020] studies a related regression problem with uniform design on $[0,1]$, independent standard Gaussian noise, and a Sobolev ball defined through the Fourier coefficients of the regression function. Under Huber's ϵ-contamination model, Theorem 3.1 therein reports a high-probability bound of order $n^{-2\alpha/(2\alpha+1)} \vee \epsilon^2$ for the integrated squared error of a regression depth estimator. For fixed $R > 0$, the uncontaminated term for our Hölder ball is of order $n^{-2\alpha/(2\alpha+m)}$, so the two exponents agree when $m = 1$. The contamination terms differ: in our setting, the minimax risk has a term of order ϵ, subject to the diameter cap d^2. This linear dependence is a worst-case statement over conditionally centered noise laws satisfying $\mathbb{E}\left[\xi^2 \mid X\right] \le \sigma^2$ and adversaries from the strong contamination model. It does not assert that the same linear dependence is necessary when the noise law is fixed to be Gaussian. We also acknowledge one limitation of our approach that we are not able to obtain the ϵ^2 term on the contamination even when the noise is Gaussian, as what is achieved by Gao [2020], since by the last remark the function class $\mathcal{F}_{\alpha,R}$ in general does not satisfy the sub-Gaussian increment conditions (4.1).

6 Discussion

6.1 Summary

We study nonparametric regression over star-shaped function classes under strong adversarial contamination, allowing the adversary to replace both covariates and responses. Our results relate the minimax squared $L_2(\mathbb{P}_X)$ risk to the local metric entropy of the function class and the contamination proportion. Under the Bernstein-type increment condition and a bounded conditional second moment of the noise, we establish matching upper and lower bounds of order $\left[(\eta^*)^2 \vee \epsilon\right] \wedge d^2$. Under relative sub-Gaussian increment and noise conditions, together with a common diameter bound, the contamination contribution lies between ϵ^2 and $\epsilon^2 \log^2(e/\epsilon)$. The upper bounds follow from a hierarchy of local packings combined with robust pairwise comparisons based on modified trimmed means. The examples of sparse linear regression, low-rank trace regression, and Hölder smooth regression illustrate how the local-entropy characterization yields explicit rates. Our analysis also distinguishes the roles of conditional and marginal noise assumptions: conditional second-moment control is essential for the general guarantee under the Bernstein-type condition, whereas marginal second-moment control suffices under the stronger sub-Gaussian assumptions.

6.2 Future Research Directions

An immediate question is the exact dependence on ϵ in the sub-Gaussian setting. The present upper bound uses sub-exponential moment control of normalized squared-loss differences, while the contamination lower bound is obtained from a Gaussian noise submodel. These arguments leave open whether either of the logarithmic factors can be removed. Resolving this question may require sharper comparisons that exploit conditional centering more directly, or stronger lower bounds using the full family of admissible sub-Gaussian noise distributions. The total-variation modulus approach developed by Chen et al. [2016] provides a useful perspective on how contamination sensitivity depends on the statistical model. Applying this perspective here would require accounting for the fixed design distribution, the restrictions on the conditional noise laws, and the finite replacement budget of the strong adversary.

A second direction is computational efficiency. The packing construction establishes statistical achievability, but constructing the candidate sets and performing all pairwise comparisons can be computationally prohibitive. It is therefore natural to seek efficient procedures for structured subclasses, including the sparse and low-rank models considered in Section 5. Robust iterative hard thresholding in Liu et al. [2020] and covariate filtering in Pensia et al. [2025] demonstrate how additional structure can support tractable robust regression methods. Li et al. [2026] study estimation of the natural parameter in uncontaminated exponential-family models under star-shaped constraints and develop a computationally efficient estimator under additional geometric assumptions. In the related robust signal-detection problems, Li and Neykov [2026b,a] develop efficient robust testing procedures with matching minimax upper bound up to a logarithmic factor. A corresponding goal in our setting is to replace explicit packing searches by suitable optimization procedures while retaining the localization properties and comparison guarantees used in the analysis. Identifying the geometric or computational assumptions needed for such a replacement would help distinguish statistical achievability from efficient achievability.

Another direction concerns unknown design geometry and tuning parameters. Our construction assumes access to the $L_2(\mathbb{P}_X)$ geometry, or to an equivalent norm with controlled comparison

constants. When such an oracle is unavailable, the geometry must be estimated from covariates that may themselves be corrupted. Developing a joint analysis of robust distance estimation and function selection would extend the scope of the present results. It is also of interest to construct procedures that adapt to an unknown contamination proportion or unknown moment bounds. Such adaptation requires calibrating the pairwise comparisons without knowing the scale at which contamination prevents further reliable refinement. An important question is whether this can be achieved without changing the local-entropy rate or introducing additional logarithmic losses.

Finally, the two distributional regimes considered here suggest studying intermediate moment assumptions and more restricted noise families. Bounds on conditional moments of order greater than two may yield contamination penalties between the finite-variance and sub-Gaussian regimes. Requiring independence between the noise and the covariates, or restricting contamination to the responses, may also change the least favorable distributions and the attainable rates. Characterizing these changes would clarify which features of the contamination penalty are determined by the geometry of the function class and which arise from the permitted dependence and tail behavior of the noise.

AI Usage Disclosure

Throughout this study and the write-up of this manuscript, ChatGPT 6 and 6.1 were used mainly to find and correct the grammar mistakes, improve the word usage, and polish the language. They were also leveraged to search and provide some insights on the studies that share similar background with the current work where the authors are unaware of. Some constants in proofs of some lemmas, theorems, and corollaries are also validated by AI for correctness. The authors are responsible for the novelty of the idea of this work, the main framework of, and the derivations of the lemmas, theorems, corollaries, and algorithms in the manuscript.

References

A. A. Amini and M. J. Wainwright. High-dimensional analysis of semidefinite relaxations for sparse principal components. *The Annals of Statistics*, 37(5B):2877 – 2921, 2009. doi: 10.1214/08-AOS664. URL https://doi.org/10.1214/08-AOS664.

E. J. Candès and Y. Plan. Tight oracle inequalities for low-rank matrix recovery from a minimal number of noisy random measurements. *IEEE Transactions on Information Theory*, 57(4):2342–2359, 2011. doi: 10.1109/TIT.2011.2111771.

M. Chen, C. Gao, and Z. Ren. A general decision theory for Huber's ϵ-contamination model. *Electronic Journal of Statistics*, 10(2):3752–3774, 2016. doi: 10.1214/16-EJS1216.

P. J. Daniell. Observations weighted according to order. *American Journal of Mathematics*, 42(4): 222–236, 1920. doi: 10.2307/2370465.

I. Diakonikolas and D. M. Kane. *Algorithmic High-Dimensional Robust Statistics*. Cambridge University Press, 2023. doi: 10.1017/9781108943161.

D. L. Donoho and R. C. Liu. Geometrizing Rates of Convergence, II. *The Annals of Statistics*, 19(2):633 – 667, 1991. doi: 10.1214/aos/1176348114. URL https://doi.org/10.1214/aos/1176348114.

S. S. Du, Y. Wang, S. Balakrishnan, P. Ravikumar, and A. Singh. Robust nonparametric regression under Huber's ϵ-contamination model, 2018.

C. Gao. Robust regression via multivariate regression depth. *Bernoulli*, 26(2):1139–1170, 2020. doi: 10.3150/19-BEJ1144.

Q. Han, T. Wang, S. Chatterjee, and R. J. Samworth. Isotonic regression in general dimensions. *The Annals of Statistics*, 47(5):2440–2471, 2019. doi: 10.1214/18-AOS1753.

G. Lecué and M. Lerasle. Robust machine learning by median-of-means: Theory and practice. *The Annals of Statistics*, 48(2):906–931, 2020. doi: 10.1214/19-AOS1828.

Y. Li and M. Neykov. Efficient robust constrained signal detection via kolmogorov width approximations, 2026a. URL https://arxiv.org/abs/2605.11238.

Y. Li and M. Neykov. Robust signal detection with quadratically convex orthosymmetric constraints, 2026b. URL https://arxiv.org/abs/2308.13036.

Y. Li, G. Yi, and M. Neykov. Minimaxity and efficiency in exponential family regression: From star-shaped to convex constraints, 2026. URL https://arxiv.org/abs/2503.10794.

J. Liu and H. Lian. Kernel-based adaptive Huber mean regression with heavy tails and contamination. *Journal of the American Statistical Association*, 2026. doi: 10.1080/01621459.2026.2696486. Advance online publication.

L. Liu, Y. Shen, T. Li, and C. Caramanis. High dimensional robust sparse regression. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 411–421. PMLR, 2020. URL https://proceedings.mlr.press/v108/liu20b.html.

G. Lugosi and S. Mendelson. Risk minimization by median-of-means tournaments. *Journal of the European Mathematical Society*, 22(3):925–965, 2020. doi: 10.4171/JEMS/937.

G. Lugosi and S. Mendelson. Robust multivariate mean estimation: The optimality of trimmed mean. *The Annals of Statistics*, 49(1):393 – 410, 2021. doi: 10.1214/20-AOS1961. URL https://doi.org/10.1214/20-AOS1961.

S. Mendelson. "local" vs. "global" parameters—breaking the Gaussian complexity barrier. *The Annals of Statistics*, 45(5):1835–1862, 2017. doi: 10.1214/16-AOS1510.

M. Neykov. On the minimax rate of the Gaussian sequence model under bounded convex constraints. *IEEE Transactions on Information Theory*, 69(2):1244–1260, 2023. doi: 10.1109/TIT.2022.3213141.

R. I. Oliveira and L. Resende. Trimmed sample means for robust uniform mean estimation and regression. *The Annals of Statistics*, 53(5):2153–2178, 2025. doi: 10.1214/25-AOS2536.

T. Peng, A. Prasadan, and M. Neykov. Robust mean estimation under star-shaped constraints with heavy-tailed noise, 2026. URL https://arxiv.org/abs/2604.05063.

A. Pensia, V. Jog, and P.-L. Loh. Robust regression with covariate filtering: Heavy tails and adversarial contamination. *Journal of the American Statistical Association*, 120(550):1002–1013, 2025. doi: 10.1080/01621459.2024.2392906.

A. Prasadan and M. Neykov. Characterizing the minimax rate of nonparametric regression under bounded star-shaped constraints. *Electronic Journal of Statistics*, 19(2):3449 – 3488, 2025. doi: 10.1214/25-EJS2419. URL https://doi.org/10.1214/25-EJS2419.

A. Prasadan and M. Neykov. Information theoretic limits of robust sub-Gaussian mean estimation under star-shaped constraints. *The Annals of Statistics*, 54(1):490 – 515, 2026. doi: 10.1214/25-AOS2576. URL https://doi.org/10.1214/25-AOS2576.

G. Raskutti, M. J. Wainwright, and B. Yu. Minimax rates of estimation for high-dimensional linear regression over ℓ_q-balls. *IEEE Transactions on Information Theory*, 57(10):6976–6994, 2011. doi: 10.1109/TIT.2011.2165799.

G. Reeves and M. Gastpar. Sampling bounds for sparse support recovery in the presence of noise. In *2008 IEEE International Symposium on Information Theory*, pages 2187–2191, 2008. doi: 10.1109/ISIT.2008.4595378.

A. Rohde and A. B. Tsybakov. Estimation of high-dimensional low-rank matrices. *The Annals of Statistics*, 39(2):887–930, 2011. doi: 10.1214/10-AOS860.

A. B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer Series in Statistics. Springer, 2009. doi: 10.1007/b13794.

A. W. van der Vaart and J. A. Wellner. *Weak Convergence and Empirical Processes: With Applications to Statistics*. Springer Series in Statistics. Springer, 1996. doi: 10.1007/978-1-4757-2545-2.

M. J. Wainwright. Information-theoretic limits on sparsity recovery in the high-dimensional and noisy setting. *IEEE Transactions on Information Theory*, 55:5728–5741, 2009. URL https://api.semanticscholar.org/CorpusID:1446.

M. J. Wainwright. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2019.

Y. Yang and A. Barron. Information-theoretic determination of minimax rates of convergence. *The Annals of Statistics*, 27(5):1564–1599, 1999. doi: 10.1214/aos/1017939142.

A Proof of Section 2

A.1 Proof of Lemma 2.2

Proof. Denote $\mathbb{P}_{\widetilde{X};j}, \mathbb{P}_{\widetilde{Y}|\widetilde{X};j}, \mathbb{P}_{\widetilde{X},\widetilde{Y};j}$ as the law of $\widetilde{X}$, $\widetilde{Y}$ conditioned on $\widetilde{X}$, and $(\widetilde{X},\widetilde{Y})$ when $J=j$ (i.e., $f^*=f_j$). Let $\bar{\mathbb{P}}_{\widetilde{X},\widetilde{Y}} := \frac{1}{m}\sum_{j=1}^{m}\mathbb{P}_{\widetilde{X},\widetilde{Y};j}$ be the mixture distribution of $(\widetilde{X},\widetilde{Y})$ over the index J. From the definition above, we know $\mathbb{P}_{\widetilde{X},\widetilde{Y};j} = \mathbb{P}_{\widetilde{X};j}\times\mathbb{P}_{\widetilde{Y}|\widetilde{X};j}$. By the definition of mutual information, we have

$$I(\widetilde{X},\widetilde{Y};J) = \frac{1}{m}\sum_{j=1}^{m}\mathrm{KL}(\mathbb{P}_{\widetilde{X},\widetilde{Y};j}\,\|\,\bar{\mathbb{P}}_{\widetilde{X},\widetilde{Y}}) \overset{\text{(i)}}{\le} \frac{1}{m}\sum_{j=1}^{m}\mathrm{KL}(\mathbb{P}_{\widetilde{X},\widetilde{Y};j}\,\|\,\mathbb{P}^{h}_{\widetilde{X},\widetilde{Y}}).$$

Here $\mathbb{P}^{h}_{\widetilde{X},\widetilde{Y}}$ represents the joint probability distribution of $(\widetilde{X},\widetilde{Y})$ when $f^*=h\in\mathcal{F}$, and (i) leverages the property that $\bar{\mathbb{P}}_{\widetilde{X},\widetilde{Y}}$ minimizes the objective $\frac{1}{m}\sum_{j=1}^{m}\mathrm{KL}(\mathbb{P}_{\widetilde{X},\widetilde{Y};j}\,\|\,\mathbb{Q})$ over arbitrary probability distribution $\mathbb{Q}$. Now by the definition of KL divergence, we further have

$$\begin{aligned}
I(\widetilde{X},\widetilde{Y};J) &\le \frac{1}{m}\sum_{j=1}^{m}\mathbb{E}_{\widetilde{X},\widetilde{Y};j}\left(\log\frac{d\mathbb{P}_{\widetilde{X},\widetilde{Y};j}}{d\mathbb{P}^{h}_{\widetilde{X},\widetilde{Y}}}\right)\\
&= \frac{1}{m}\sum_{j=1}^{m}\mathbb{E}_{\widetilde{X}}\mathbb{E}_{\widetilde{Y}|\widetilde{X};j}\left(\log\frac{d\mathbb{P}_{\widetilde{Y}|\widetilde{X};j}}{d\mathbb{P}^{h}_{\widetilde{Y}|\widetilde{X}}}\right)\\
&= \frac{1}{m}\sum_{j=1}^{m}\sum_{i=1}^{n}\mathbb{E}_{\widetilde{X}_i}\mathbb{E}_{\widetilde{Y}_i|\widetilde{X}_i;j}\left(\log\frac{d\mathbb{P}_{\widetilde{Y}_i|\widetilde{X}_i;j}}{d\mathbb{P}^{h}_{\widetilde{Y}_i|\widetilde{X}_i}}\right)\\
&\overset{\text{(i)}}{\le} \frac{1}{2m\sigma^2}\sum_{j=1}^{m}\sum_{i=1}^{n}\|f_j-h\|^2_{L_2(\mathbb{P}_X)}\\
&\le \frac{n}{2\sigma^2}\max_j\|f_j-h\|^2_{L_2(\mathbb{P}_X)},
\end{aligned}$$

where (i) uses the property that given $\widetilde{X}_i$, $\widetilde{Y}_i$ is distributed as $\mathcal{N}(f^*(\widetilde{X}_i),\sigma^2)$ and the KL divergence between the distributions $\mathcal{N}\left(f_j(\widetilde{X}_i),\sigma^2\right)$ and $\mathcal{N}\left(h(\widetilde{X}_i),\sigma^2\right)$ is $\frac{\left[f^*(\widetilde{X}_i)-h(\widetilde{X}_i)\right]^2}{2\sigma^2}$, followed by the definition of $L_2(\mathbb{P}_X)$-distance. □

A.2 Proof of Lemma 2.3

Proof. For any η satisfying $\log\mathcal{M}^{\mathrm{loc}}_{\mathcal{F}}(\eta,c)\ge 4\left(\frac{n\eta^2}{2\sigma^2}\vee\log 2\right)$ (in the case that $\log\mathcal{M}^{\mathrm{loc}}_{\mathcal{F}}(\eta,c)$ is infinite, choose any packing set that is finite and satisfies the inequality), choose some $f\in\mathcal{F}$ admitting a η/c-packing set $\left\{f_1,\ldots,f_{\mathcal{M}(\eta/c,B_2(f,\eta)\cap\mathcal{F})}\right\}$ of $B_2(f,\eta)\cap\mathcal{F}$. (If the local packing number is infinite, choose any finite subset of the packing set whose cardinality satisfies the condition above.) Therefore, from Theorem 2.1 and Lemma 2.2, we know that

$$\mathcal{R}_{n,\epsilon}(\mathcal{F};\mathcal{P}_\xi)\ge\frac{\eta^2}{4c^2}\left(1-\frac{I(\widetilde{X},\widetilde{Y};J)+\log 2}{\log\mathcal{M}^{\mathrm{loc}}_{\mathcal{F}}(\eta,c)}\right)\ge\frac{\eta^2}{4c^2}\left(1-\frac{\frac{n\eta^2}{2\sigma^2}+\log 2}{\log\mathcal{M}^{\mathrm{loc}}_{\mathcal{F}}(\eta,c)}\right)\ge\frac{\eta^2}{8c^2}.$$

The proof is completed. □

B Proof of Section 3

B.1 Proof of Lemma 3.1

Proof. The conclusion is trivial if $\epsilon = 0$, $\sigma = 0$, or $d = 0$. We therefore assume that these quantities are positive. Let $f_c \in \mathcal{F}$ be a star center of $\mathcal{F}$. By the definition of the diameter and the triangle inequality, there exists $\bar{f} \in \mathcal{F}$ such that $\left\|\bar{f} - f_c\right\|_{L_2(\mathbb{P}_X)} > d/3$ (also see Lemma 1.3 of Prasadan and Neykov [2026]). Set

$$r^2 := \epsilon\sigma^2 \wedge \frac{d^2}{9}, \qquad f_0 := f_c, \qquad f_1 := f_c + \frac{r}{\left\|\bar{f} - f_c\right\|_{L_2(\mathbb{P}_X)}} \left(\bar{f} - f_c\right).$$

Star-shapedness ensures that $f_0, f_1 \in \mathcal{F}$, and $\|f_1 - f_0\|_{L_2(\mathbb{P}_X)} = r$. Define

$$q(x) := \frac{[f_1(x) - f_0(x)]^2}{2\sigma^2 + [f_1(x) - f_0(x)]^2}.$$

We construct two admissible regression models with regression functions f_0 and f_1. Let $X \sim \mathbb{P}_X$, draw $B \mid X \sim \text{Bernoulli}(q(X))$, and let R be a Rademacher random variable independent of (X, B). For $j \in \{0, 1\}$, define

$$Y^{(j)} := \begin{cases} \frac{f_0(X)+f_1(X)}{2} + (2j-1)\dfrac{f_1(X) - f_0(X)}{2q(X)} B, & f_0(X) \neq f_1(X), \\ \frac{f_0(X)+f_1(X)}{2} + \frac{\sigma}{\sqrt{2}} R, & f_0(X) = f_1(X). \end{cases}$$

Denote the law of $(X, Y^{(j)})$ by $\mathbb{P}_j$ and set $\xi_j := Y^{(j)} - f_j(X)$.

When $f_0(X) \neq f_1(X)$, we have

$$\mathbb{E}\left[Y^{(j)} \mid X\right] = \frac{f_0(X) + f_1(X)}{2} + (2j-1)\frac{f_1(X) - f_0(X)}{2} = f_j(X),$$

and

$$\mathbb{E}\left[\xi_j^2 \mid X\right] = \frac{[f_1(X) - f_0(X)]^2}{4q(X)^2} q(X)\left(1 - q(X)\right) = \frac{\sigma^2}{2}.$$

Similar conditional mean and variance identities of ξ_j hold on $\{f_1(X) = f_0(X)\}$ by the definition of R. Consequently, both models satisfy

$$\mathbb{E}\left[\xi_j \mid X\right] = 0, \qquad \mathbb{E}\left[\xi_j^2 \mid X\right] = \frac{\sigma^2}{2} \quad \text{a.s.}$$

Thus both conditional noise laws are admissible. Moreover, the unconditional probability of the event $\{B = 1\}$ satisfies

$$\bar{q} := \mathbb{E}\left[q(X)\right] \leq \frac{\|f_1 - f_0\|_{L_2(\mathbb{P}_X)}^2}{2\sigma^2} = \frac{r^2}{2\sigma^2} \leq \frac{\epsilon}{2}.$$

To construct the samples, take independent copies $(\widetilde{X}_i, B_i, R_i)_{i=1}^n$ of (X, B, R) and use the same copies to generate both

$$\widetilde{\mathcal{D}}_j := \left\{ \left(\widetilde{X}_i, \widetilde{Y}_i^{(j)} \right) \right\}_{i=1}^n, \qquad j \in \{0, 1\}.$$

Each marginal original sample has law $\mathbb{P}_j^{\otimes n}$.

Let $m := \lfloor \epsilon n \rfloor$. Consider the following deterministic adversary $\mathcal{C}_0$. Given a sample $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^n$, define

$$I(\mathcal{D}) := \left\{ i : f_0(x_i) \neq f_1(x_i), y_i \neq \frac{f_0(x_i) + f_1(x_i)}{2} \right\}.$$

If $|I(\mathcal{D})| \leq m$, replace y_i by $\frac{f_0(x_i)+f_1(x_i)}{2}$ for every $i \in I(\mathcal{D})$. Otherwise, leave the sample unchanged. This adversary changes at most m observations for every input and is therefore admissible. Under either constructed model,

$$I\left(\widetilde{\mathcal{D}}_j\right) = \{i : B_i = 1\} \quad \text{a.s.}$$

Set $N := \sum_{i=1}^n B_i$ and $\mathcal{E} := \{N \leq m\}$. Since $\mathbb{E}N = n\bar{q}$, Markov's inequality gives

$$\mathbb{P}\left(\mathcal{E}^{\mathrm{c}}\right) = \mathbb{P}\left(N \geq m + 1\right) \leq \frac{n\bar{q}}{m+1} \leq \frac{n\epsilon}{2(m+1)} < \frac{1}{2}.$$

In particular, $\mathbb{P}(\mathcal{E}) \geq 1/2$, including when $\lfloor \epsilon n \rfloor = 0$.

Define the corrupted samples by $\mathcal{D}_j := \mathcal{C}_0(\widetilde{\mathcal{D}}_j)$. On $\mathcal{E}$, all observations corresponding to $B_i = 1$ are replaced by the common response $\frac{f_0(X_i)+f_1(X_i)}{2}$. All remaining observations already coincide under the coupling. Hence

$$\mathcal{D}_0 = \mathcal{D}_1 \qquad \text{on } \mathcal{E}.$$

Fix any estimator $\widehat{f}$. If the estimator is randomized, couple its auxiliary randomness identically under the two models. On $\mathcal{E}$, its outputs therefore coincide. For every possible common output $\widehat{f}$, we have

$$\left\| \widehat{f} - f_0 \right\|_{L_2(\mathbb{P}_X)}^2 + \left\| \widehat{f} - f_1 \right\|_{L_2(\mathbb{P}_X)}^2 = 2 \left\| \widehat{f} - \frac{f_0 + f_1}{2} \right\|_{L_2(\mathbb{P}_X)}^2 + \frac{1}{2} \left\| f_1 - f_0 \right\|_{L_2(\mathbb{P}_X)}^2 \geq \frac{r^2}{2}.$$

Writing $\mathbb{E}_{\mathrm{cpl}}$ for expectation under the coupling, we obtain

$$\max_{j \in \{0,1\}} \mathbb{E}_{P_j^{\otimes n}} \left[\left\| \widehat{f}\left(\mathcal{C}_0(\widetilde{\mathcal{D}}_j) \right) - f_j \right\|_2^2 \right] \geq \frac{1}{2} \mathbb{E}_{\mathrm{cpl}} \left[\left(\| \widehat{f}(\mathcal{D}_0) - f_0 \|_2^2 + \| \widehat{f}(\mathcal{D}_1) - f_1 \|_2^2 \right) \mathbf{1}_{\mathcal{E}} \right] \geq \frac{r^2}{8}.$$

Both regression models and the adversary $\mathcal{C}_0$ are included in the supremum defining the minimax risk. Taking the infimum over $\widehat{f}$ therefore gives

$$\mathcal{R}_{n,\epsilon}\left(\mathcal{F}; \mathcal{P}_\xi\right) \geq \frac{1}{8} \left(\epsilon\sigma^2 \wedge \frac{d^2}{9} \right) \geq \frac{1}{72} \left(\epsilon\sigma^2 \wedge d^2 \right).$$

This proves the claim with an absolute constant independent of d, n, ϵ, and σ. □

B.2 Proof of Lemma 3.7

Proof. We compute the variance of $Z_{f,g}$ (under the uncontaminated samples, hence $\left(\left(\widetilde{X}_i, \widetilde{Y}_i, \widetilde{\xi}_i\right)\right)_{i=1}^n = ((X_i, Y_i, \xi_i))_{i=1}^n$). By rewriting $Z_{f,g}$, we have

$$\begin{aligned}
Z_{f,g} &= \frac{[Y - f(X)]^2 - [Y - g(X)]^2}{\|f - g\|_{L_2(\mathbb{P}_X)}} \\
&= \frac{[g(X) - f(X)]\,[2Y - f(X) - g(X)]}{\|f - g\|_{L_2(\mathbb{P}_X)}} \\
&= \frac{[g(X) - f(X)]}{\|f - g\|_{L_2(\mathbb{P}_X)}} \cdot [2f^*(X) - f(X) - g(X)] + \frac{2\,[g(X) - f(X)]}{\|f - g\|_{L_2(\mathbb{P}_X)}} \cdot \xi.
\end{aligned}$$

For the variance of $Z_{f,g}$, we have

$$\begin{aligned}
\mathrm{Var}\,(Z_{f,g}) &\le \mathbb{E} Z_{f,g}^2 \\
&\overset{\text{(i)}}{=} \mathbb{E}\left[\frac{[g(X) - f(X)]^2}{\|f - g\|_{L_2(\mathbb{P}_X)}^2} \cdot [2f^*(X) - f(X) - g(X)]^2 + \frac{4\,[g(X) - f(X)]^2}{\|f - g\|_{L_2(\mathbb{P}_X)}^2} \cdot \xi^2\right] \\
&\overset{\text{(ii)}}{\le} \mathbb{E}\left[2 \cdot \frac{[g(X) - f(X)]^2}{\|f - g\|_{L_2(\mathbb{P}_X)}^2} \cdot \left([f^*(X) - f(X)]^2 + [f^*(X) - g(X)]^2\right) + \frac{4\,[g(X) - f(X)]^2}{\|f - g\|_{L_2(\mathbb{P}_X)}^2} \cdot \xi^2\right] \\
&\overset{\text{(iii)}}{\le} 2 \cdot \frac{\|f - g\|_{L_4(\mathbb{P}_X)}^2}{\|f - g\|_{L_2(\mathbb{P}_X)}^2} \left(\|f^* - f\|_{L_4(\mathbb{P}_X)}^2 + \|f^* - g\|_{L_4(\mathbb{P}_X)}^2\right) + 4\sigma^2,
\end{aligned}$$

where (i) follows from the fact that $\mathbb{E}[\xi \mid X] = 0$; (ii) is from the inequality $(a+b)^2 \le 2a^2 + 2b^2, \forall a, b \in \mathbb{R}$, and (iii) is from $\mathbb{E}\left[\xi^2 \mid X\right] \le \sigma^2$ and the Cauchy–Schwarz inequality.

From the moment condition (3.1), we have

$$\begin{aligned}
\|f - g\|_{L_4(\mathbb{P}_X)}^4 &\le L\,\|f - g\|_{L_2(\mathbb{P}_X)}^2 \\
\|f^* - g\|_{L_4(\mathbb{P}_X)}^4 &\le L\,\|f^* - g\|_{L_2(\mathbb{P}_X)}^2 \\
\|f^* - f\|_{L_4(\mathbb{P}_X)}^4 &\le L\,\|f^* - f\|_{L_2(\mathbb{P}_X)}^2.
\end{aligned}$$

Therefore, we have

$$\begin{aligned}
\mathrm{Var}\,(Z_{f,g}) &\le 2L \cdot \frac{\|f^* - f\|_{L_2(\mathbb{P}_X)} + \|f^* - g\|_{L_2(\mathbb{P}_X)}}{\|f - g\|_{L_2(\mathbb{P}_X)}} + 4\sigma^2 \\
&\overset{\text{(i)}}{\le} 2L \cdot \frac{\|f - g\|_{L_2(\mathbb{P}_X)} + 2\,\|f^* - g\|_{L_2(\mathbb{P}_X)}}{\|f - g\|_{L_{2(\mathbb{P}_X)}}} + 4\sigma^2,
\end{aligned}$$

where (i) is from the triangle inequality. Note that we assume $\|f^* - g\|_{L_2(\mathbb{P}_X)} \le \delta$ while $\|f - g\|_{L_2(\mathbb{P}_X)} \ge C\delta$. Therefore, we have

$$\mathrm{Var}\,(Z_{f,g}) \le 2L\left(1 + \frac{2}{C}\right) + 4\sigma^2.$$

By (3.6) in Lemma 3.6, we have

$$\left|\mathrm{Tm}_{Z_{f,g}}(X,Y) - \mathbb{E}Z_{f,g}\right| \le 10\sqrt{2L(1+\frac{2}{C})+4\sigma^2}\cdot\sqrt{8\epsilon + 24\frac{\log\left(\frac{4}{\rho}\right)}{n}}.$$

The proof is completed. □

B.3 Proof of Corollary 3.8

Proof. We first verify the requirements of Lemma 3.7. The inequality trivially holds if $\rho = 4\exp\left\{-an\delta^2/(L+\sigma^2)\right\} \ge 1$. Otherwise, we have

$$4\exp\left\{-\frac{an\delta^2}{L+\sigma^2}\right\} \overset{\text{(i)}}{\ge} 4\exp\left\{-\frac{1}{2}n\right\} > \frac{1}{4}\exp\left\{-\frac{1}{2}n\right\},$$

where (i) is from the boundedness that $d^2 \le L$, $a \le \frac{1}{192}$, and that every nonvacuous comparison must have $\delta^2 \le d^2 \le L$. Also, the quantile ζ satisfies that

$$\zeta = 8\epsilon + 24\frac{\log\left(\frac{4}{\rho}\right)}{n} = 8\epsilon + 24\frac{a\delta^2}{L+\sigma^2} \le 8\epsilon + 24a \le \frac{1}{4},$$

given that $\epsilon \le \frac{1}{64}$ and $a \le \frac{1}{192}$. Second, we calculate $\mathbb{E}Z_{f,g}$. For $f, g \in \mathcal{F}$ such that

$$\|f-g\|_{L_2(\mathbb{P}_X)} \ge C\delta, \qquad \|f^*-g\|_{L_{2(\mathbb{P}_X)}} \le \delta,$$

the trimmed mean test (3.8) correctly selects g if $\mathbb{E}Z_{f,g}$ dominates the variation term (RHS) in (3.9). By the proof of Lemma 3.7, it suffices that

$$\mathbb{E}Z_{f,g} \ge 10\sqrt{2L(1+\frac{2}{C})+4\sigma^2}\cdot\sqrt{8\epsilon + 24\frac{\log\left(\frac{4}{\rho}\right)}{n}} \asymp \sqrt{L+\sigma^2}\left[\sqrt{\epsilon} + \sqrt{\log\left(\frac{4}{\rho}\right)/n}\right].$$

For $\mathbb{E}Z_{f,g}$, we have

$$\begin{aligned}
\mathbb{E}Z_{f,g} &= \mathbb{E}\left[\frac{[\xi + f^*(X) - f(X)]^2 - [\xi + f^*(X) - g(X)]^2}{\|f-g\|_{L_2(\mathbb{P}_X)}}\right] \\
&= \frac{\left(\|f^*-f\|_{L_2(\mathbb{P}_X)} - \|f^*-g\|_{L_2(\mathbb{P}_X)}\right)\left(\|f^*-f\|_{L_2(\mathbb{P}_X)} + \|f^*-g\|_{L_2(\mathbb{P}_X)}\right)}{\|f-g\|_{L_2(\mathbb{P}_X)}} \\
&\overset{\text{(i)}}{\ge} \|f-g\|_{L_2(\mathbb{P}_X)} - 2\|f^*-g\|_{L_2(\mathbb{P}_X)} \\
&\ge (C-2)\delta,
\end{aligned}$$

where (i) is again from the triangle inequality. Therefore, to let the variation term be dominated by the true signal, we require

$$\begin{cases} (C-2)\delta \gtrsim \sqrt{L+\sigma^2}\sqrt{\epsilon} \Leftrightarrow \delta^2 \gtrsim_{L,\sigma} \epsilon \text{ (the desired condition in Corollary 3.8)}, \\ (C-2)\delta \gtrsim \sqrt{L+\sigma^2}\sqrt{\log(4/\rho)/n} \Leftrightarrow (C-2)\delta \gtrsim \sqrt{a}\delta \text{ (already achieved by our choice of } a). \end{cases}$$

□

B.4 Proof of Lemma 3.9

Proof. The proof follows from an application of the union bound. Denote $f_{i^*} = \arg\min_{i\in[M]} \|f_i - f^*\|_{L_2(\mathbb{P}_X)}$, then by the assumption of the lemma, we have $\|f_{i^*} - f^*\|_{L_2(\mathbb{P}_X)} \le \delta$. Consider any $f_i \in S$ such that $\|f_i - f^*\|_{L_2(\mathbb{P}_X)} \ge (C+1)\delta$, then by the triangle inequality, we know $\|f_{i^*} - f_i\|_{L_2(\mathbb{P}_X)} \ge C\delta$. By Corollary 3.8, we have

$$\mathbb{P}(f_{i^*} \text{ fails to dominate } f_i) \le 4\exp\left\{-\frac{an\delta^2}{L+\sigma^2}\right\}.$$

Consequently, we have by the union bound

$$\mathbb{P}(f_{i^*} \succ f_i, \forall f_i, \|f_{i^*} - f_i\|_{L_2(\mathbb{P}_X)} \ge C\delta) \ge 1 - 4M\exp\left\{-\frac{an\delta^2}{L+\sigma^2}\right\}.$$

However, on the event $\left\{f_{i^*} \succ f_i, \forall f_i, \|f_{i^*} - f_i\|_{L_2(\mathbb{P}_X)} \ge C\delta\right\}$, we have $H(\delta, f_i, S) > 0$ for all f_i such that $\|f_i - f^*\|_{L_2(\mathbb{P}_X)} \ge (C+1)\delta$ and $H(\delta, f_{i^*}, S) = 0$ (by definition of H and the triangle inequality), which further implies that $\|f_S - f^*\|_{L_2(\mathbb{P}_X)} < (C+1)\delta$. Therefore, we have

$$\mathbb{P}\left(\|f_S - f^*\|_{L_2(\mathbb{P}_X)} \ge (C+1)\delta\right) = 1 - \mathbb{P}\left(\|f_S - f^*\|_{L_2(\mathbb{P}_X)} < (C+1)\delta\right) \le 4M\exp\left\{-\frac{an\delta^2}{L+\sigma^2}\right\}.$$

The proof is completed. □

B.5 Proof of Corollary 3.10

Proof. Let's directly evaluate $\mathbb{P}\left(A_{j-1}^{\complement} \cap A_j\right)$ for $j \ge 3$:

$$\begin{aligned}
\mathbb{P}\left(A_{j-1}^{\complement} \cap A_j\right) &= \mathbb{P}\left(\left\|f_{j-1}^* - f^*\right\|_{L_2(\mathbb{P}_X)} \le \frac{d}{2^{j-2}}, \left\|f_j^* - f^*\right\|_{L_2(\mathbb{P}_X)} > \frac{d}{2^{j-1}}\right) \\
&= \sum_{f\in\mathcal{L}(j-1)\cap B_2(f^*, \frac{d}{2^{j-2}})} \mathbb{P}\left(f_{j-1}^* = f, \left\|f_j^* - f^*\right\|_{L_2(\mathbb{P}_X)} > \frac{d}{2^{j-1}}\right) \\
&\le \sum_{f\in\mathcal{L}(j-1)\cap B_2(f^*, \frac{d}{2^{j-2}})} \mathbb{P}\left(\left\|f_j^* - f^*\right\|_{L_2(\mathbb{P}_X)} > \frac{d}{2^{j-1}}\right) \\
&\overset{\text{(i)}}{\le} \sum_{f\in\mathcal{L}(j-1)\cap B_2(f^*, \frac{d}{2^{j-2}})} 4\mathcal{M}_{\mathcal{F}}^{\text{loc}}\left(\frac{d}{2^{j-2}}, 2c\right)\exp\left\{-n\eta_j^2\right\} \\
&\overset{\text{(ii)}}{\le} 4\left[\mathcal{M}_{\mathcal{F}}^{\text{loc}}\left(\frac{d}{2^{j-2}}, 2c\right)\right]^2\exp\left\{-n\eta_j^2\right\}.
\end{aligned}$$

The term $\mathcal{M}_{\mathcal{F}}^{\text{loc}}$ in (i) is from (4) of Lemma 3.3 and exponential term is from Lemma 3.9 by observing $d/2^{j-1} = (C+1)\delta_j$. (ii) originates from Lemma 3.4. When $i = 2$, since $\mathcal{L}(1)$ only consists of a single point and $d > 0$, we have

$$\mathbb{P}\left(A_1^{\complement} \cap A_2\right) \le 4\mathcal{M}_{\mathcal{F}}^{\text{loc}}(d, c)\exp\left\{-n\eta_2^2\right\} \le 4\left[\mathcal{M}_{\mathcal{F}}^{\text{loc}}(d, 2c)\right]^2\exp\left\{-n\eta_2^2\right\}.$$

For $i = 1$, we trivially have by the definition of diameter:

$$\mathbb{P}(A_1) = \mathbb{P}(\|f_1^* - f^*\|_{L_2(\mathbb{P}_X)} > d) = 0.$$

The proof is completed. □

B.6 Proof of Theorem 3.11

Proof. If $d = 0$ or the sample size n is overly small, then (3.12) might never hold, in which case $J^* = 1$ and $\eta_{J^*} = \eta_1 \asymp d$. Therefore, the upper bound trivially holds. Assume in the following that $d > 0$ and (3.12) holds for some $J \geq 2$.

For any $J \geq 2$ and a sequence of measurable sets $A_1, \ldots, A_J$, we have

$$\begin{aligned}
A_J &= (A_J \cap A_1) \cup \left(A_J \cap A_1^{\complement}\right) \\
&= (A_J \cap A_1) \cup \left(A_J \cap A_1^{\complement} \cap A_2\right) \cup \left(A_J \cap A_1^{\complement} \cap A_2^{\complement}\right) \\
&= \ldots \\
&= \left[\bigcup_{i=1}^{J} \left(A_J \cap \left(\bigcap_{j=1}^{i-1} A_j^{\complement}\right) \cap A_i\right)\right] \cup \left(A_J \cap \left(\bigcap_{j=1}^{J} A_i^{\complement}\right)\right) \\
&= \bigcup_{i=1}^{J} \left(A_J \cap \left(\bigcap_{j=1}^{i-1} A_j^{\complement}\right) \cap A_i\right) \\
&\subseteq \bigcup_{i=1}^{J} \left(A_{i-1}^{\complement} \cap A_i\right),
\end{aligned}$$

where we define $A_0 = \varnothing$. Therefore, we have

$$\mathbb{P}(A_J) \leq \sum_{i=1}^{J} \mathbb{P}\left(A_{i-1}^{\complement} \cap A_i\right). \tag{B.1}$$

Applying Corollary 3.10 yields that for $J \leq J^*$:

$$\begin{aligned}
\mathbb{P}\left(A_J\right) &\leq \sum_{i=1}^{J} \mathbb{P}\left(A_{i-1}^{\complement} \cap A_i\right) \\
&\leq \sum_{j=2}^{J} 4\left[\mathcal{M}_{\mathcal{F}}^{\mathrm{loc}}\left(\frac{d}{2^{j-2}}, 2c\right)\right]^2 \exp\left\{-n\eta_j^2\right\} \\
&\overset{\text{(i)}}{\leq} 4\left[\mathcal{M}_{\mathcal{F}}^{\mathrm{loc}}\left(\frac{d}{2^{J-2}}, 2c\right)\right]^2 \sum_{j=2}^{J} \exp\left\{-n\eta_j^2\right\} \\
&\leq 4\left[\mathcal{M}_{\mathcal{F}}^{\mathrm{loc}}\left(\frac{d}{2^{J-2}}, 2c\right)\right]^2 \frac{\exp\left\{-n\eta_J^2\right\}}{1 - \exp\left\{-n\eta_J^2\right\}} \\
&\overset{\text{(ii)}}{\leq} 4\exp\left\{\frac{n\eta_J^2}{2}\right\} \cdot \frac{\exp\left\{-n\eta_J^2\right\}}{1 - \exp\left\{-n\eta_J^2\right\}} \\
&\overset{\text{(iii)}}{\leq} 8\exp\left\{-\frac{n\eta_J^2}{2}\right\},
\end{aligned} \tag{B.2}$$

where (i) is from Lemma 1.4; (ii) is from the fact that $n\eta_J^2 > 4\log\left[\mathcal{M}_{\mathcal{F}}^{\text{loc}}\left(\frac{d}{2^{J-2}}, 2c\right)\right]$ and (iii) is from $n\eta_J^2 > \log 2$, both originate from Condition (3.12) for $J \leq J^*$.

From Algorithm 2 and the argument above, we know that the closeness between f_J^* and f^* (on the rate of η_J) depends on the correct selection for f_{J-1}^*, and hence on all previous trajectories $f_1^*, \ldots, f_{J-1}^*$. This, surely, is not guaranteed to happen. However, we can control the error probability for large-scale deviation given the control on $\mathbb{P}(A_J)$ from (B.2) and the property that the sequence $f_1^*, \ldots, f_{J^*}^*, \ldots$ is a Cauchy sequence (Lemma 3.5). Consider the event $\left\{\|f_{J^*}^* - f^*\|_{L_2(\mathbb{P}_X)} \geq \frac{d(1+3c)}{2^{J-1}c}\right\}$ for some $J \leq J^*$, we have

$$\begin{aligned}
\mathbb{P}\left(\|f_{J^*}^* - f^*\|_{L_2(\mathbb{P}_X)} \geq \frac{d(1+3c)}{2^{J-1}c}\right) &\leq \mathbb{P}\left(\|f_{J^*}^* - f_J^*\|_{L_2(\mathbb{P}_X)} + \|f_J^* - f^*\|_{L_2(\mathbb{P}_X)} \geq \frac{d(1+3c)}{2^{J-1}c}\right) \\
&\overset{\text{(i)}}{\leq} \mathbb{P}\left(\|f_J^* - f^*\|_{L_2(\mathbb{P}_X)} + \frac{d(2+4c)}{2^J \cdot c} \geq \frac{d(1+3c)}{2^{J-1}c}\right) \\
&\leq 8\exp\left\{-\frac{n\eta_J^2}{2}\right\},
\end{aligned}$$

where (i) is from Lemma 3.5. Note that we can extend this inequality to the region $J \leq 0$, where the LHS is trivially zero followed by $\frac{d(1+3c)}{2^{J-1}c} \geq 6d, J \leq 0$. Define $\kappa = \frac{1+3c}{2}\sqrt{\frac{L+\sigma^2}{a}}$, then $\frac{d(1+3c)}{2^{J-1}c} = \kappa\eta_J$. For any $x \geq \eta_{J^*}$, there exists an integer $J \leq J^*$ such that $\eta_J \leq x \leq \eta_{J-1} = 2\eta_J$. Therefore, we have

$$\mathbb{P}\left(\|f_{J^*}^* - f^*\|_{L_2(\mathbb{P}_X)} \geq \kappa x\right) \leq \mathbb{P}\left(\|f_{J^*}^* - f^*\|_{L_2(\mathbb{P}_X)} \geq \kappa\eta_J\right) \leq 8\exp\left\{-\frac{n\eta_J^2}{2}\right\} \leq 8\exp\left\{-\frac{nx^2}{8}\right\}. \tag{B.3}$$

Based on (B.3), we can well bound the expectation of the squared loss.

$$\begin{aligned}
\mathbb{E}_{\widetilde{X},\widetilde{Y}}\left\|f_{J^*}^*(\mathcal{C}(\widetilde{X},\widetilde{Y})) - f^*\right\|_{L_2(\mathbb{P}_X)}^2 &= 2\int_0^\infty x \cdot \mathbb{P}\left(\|f_{J^*}^* - f^*\|_{L_2(\mathbb{P}_X)} \geq x\right) dx \\
&= 2\kappa^2 \int_0^\infty t \cdot \mathbb{P}\left(\|f_{J^*}^* - f^*\|_{L_2(\mathbb{P}_X)} \geq \kappa t\right) dt \\
&\leq 2\kappa^2 \int_0^{\eta_{J^*}} t dt + 2\kappa^2 \int_{\eta_{J^*}}^\infty t \cdot \mathbb{P}\left(\|f_{J^*}^* - f^*\|_{L_2(\mathbb{P}_X)} \geq \kappa t\right) dt \\
&\leq \kappa^2\eta_{J^*}^2 + 2\kappa^2 \int_{\eta_{J^*}}^\infty 8t\exp\left\{-\frac{nt^2}{8}\right\} dt \\
&= \kappa^2\eta_{J^*}^2 + \frac{64\kappa^2}{n}\exp\left\{-\frac{n\eta_{J^*}^2}{8}\right\} \\
&\overset{\text{(i)}}{\leq} \kappa^2\left(1 + \frac{64}{\log 2}\right)\eta_{J^*}^2 \\
&= \frac{L+\sigma^2}{a}\left(\frac{1+3c}{2}\right)^2\left(1 + \frac{64}{\log 2}\right)\eta_{J^*}^2,
\end{aligned} \tag{B.4}$$

where (i) is from the definition of J^* (3.12). As an important note, we should be aware that the bound on the squared loss (B.4) depends on the condition that $\epsilon \lesssim \eta_{J^*}^2$. When $\epsilon \gtrsim \eta_{J^*}^2$, this crucial basis – Lemma 3.7, Lemma 3.9, and Corollary 3.10 – of the probability control on A_J can fail before

J increases to J^* and η_J decreases to η_{J^*}. In this case, we can alternatively define

$$J_\epsilon := \max\left(\{1\} \cup \left\{2 \le j \le J^* : \epsilon \le \frac{b\delta_j^2}{L+\sigma^2}\right\}\right)$$

and parallel the argument above by replacing J^* with J_ϵ since in this case $\eta_{J_\epsilon}^2 \ge \eta_{J^*}^2$. Consequently, a counterpart inequality of (B.4) on the squared loss holds:

$$\mathbb{E}_{\widetilde{X},\widetilde{Y}} \left\| f_{J^*}^*(\mathcal{C}(\widetilde{X},\widetilde{Y})) - f^* \right\|_{L_2(\mathbb{P}_X)}^2 \le \frac{L+\sigma^2}{a}\left(\frac{1+3c}{2}\right)^2\left(1+\frac{64}{\log 2}\right)\epsilon. \tag{B.5}$$

Combining (B.4) and (B.5), we know that $\eta_{J^*}^2 \vee \epsilon$ is an upper bound for the minimax rate (1.3). Since d^2 is always a valid upper bound and the RHS of (B.4) and (B.5) do not depend on $\mathcal{C}$, f^*, and ξ, we conclude that

$$\mathcal{R}_{n,\epsilon}(\mathcal{F}; \mathcal{P}_\xi^b) \lesssim (\eta_{J^*}^2 \vee \epsilon) \wedge d^2,$$

where a constant that depends only on L, σ, c is omitted. The proof is completed. □

B.7 Proof of Corollary 3.12

Proof. The lower bound follows from (3.4), so it remains to prove the upper bound. The case $d = 0$ is trivial. Suppose $d > 0$, and write

$$\kappa := \sqrt{\frac{a}{L+\sigma^2}}, \qquad H_c(r) := \log \mathcal{M}_{\mathcal{F}}^{\mathrm{loc}}(r, c), \qquad \eta_j = \frac{\kappa d}{2^{j-2}c}, \quad j \ge 1.$$

First, star-shapedness and the definition of the diameter imply that $\mathcal{F}$ contains a line segment of length $d/3$. Consequently, $H_c(r) \ge \log 2$ for every $0 < r \le d/3$. Applying the definition of η^* at $r = (d/3) \wedge \sqrt{(\log 2)/n}$ gives

$$(\eta^*)^2 \ge \frac{d^2}{9} \wedge \frac{\log 2}{n}, \qquad \frac{1}{n} \wedge d^2 \lesssim (\eta^*)^2 \wedge d^2.$$

We next compare the packing constants c and $2c$. Take any finite $r/(2c)$-packing of $B_2(f, r) \cap \mathcal{F}$ and partition it using a maximal r/c-packing of $B_2(f, r) \cap \mathcal{F}$. There are at most $\mathcal{M}_{\mathcal{F}}^{\mathrm{loc}}(r, c)$ groups by maximality, each contained in a ball of radius r/c, and each group has at most $\mathcal{M}_{\mathcal{F}}^{\mathrm{loc}}(r/c, 2)$ points (recall these packing points are $r/2c$-separated, and $(r/c)/(r/2c) = 2$). Taking logarithms and using $c \ge 2$ and Lemma 1.4, we obtain

$$H_{2c}(r) \le H_c(r) + H_2(r/c) \le 2H_c(r/c).$$

Let $s := \eta_{J^*+1} = \eta_{J^*}/2$. By the definition of J^*, Condition (3.12) fails at $J^* + 1$, including when the convention $J^* = 1$ applies. The preceding packing comparison therefore yields

$$ns^2 \le 4H_{2c}\left(\frac{cs}{\kappa}\right) \vee \log 2 \le 8H_c\left(\frac{s}{\kappa}\right) \vee \log 2.$$

Since $2\eta^* > \eta^*$, the definition of the critical radius gives $H_c(2\eta^*) < 4n(\eta^*)^2$. If $s \ge 2\kappa\eta^*$, monotonicity of H_c consequently implies $ns^2 \le 32n(\eta^*)^2 \vee \log 2$. Otherwise, $s^2 < 4\kappa^2(\eta^*)^2$. Thus, in either case,

$$\eta_{J^*}^2 = 4s^2 \lesssim_{L,\sigma} (\eta^*)^2 \vee \frac{1}{n}.$$

Combining this with the critical-radius lower bound established above gives

$$\eta_{J^*}^2 \wedge d^2 \lesssim_{L,\sigma} \left((\eta^*)^2 \vee \frac{1}{n}\right) \wedge d^2 \lesssim_{L,\sigma} (\eta^*)^2 \wedge d^2.$$

Finally, Theorem (3.11) and the identity $(u \vee v) \wedge w = (u \wedge w) \vee (v \wedge w)$ yield

$$\mathcal{R}_{n,\epsilon}(\mathcal{F}; \mathcal{P}_\xi) \lesssim_{L,\sigma} \left(\eta_{J^*}^2 \vee \epsilon\right) \wedge d^2 \lesssim_{L,\sigma} \left((\eta^*)^2 \vee \epsilon\right) \wedge d^2.$$

Together with (3.4), this proves the corollary. Here a and c are fixed universal constants, and none of the implicit constants depends on d. □

C Proof of Section 4

C.1 Proof of Lemma 4.2

Proof. Throughout the proof, write $\|\cdot\|_2 = \|\cdot\|_{L_2(\mathbb{P}_X)}$. The case $d = 0$ or $\epsilon = 0$ is immediate. Hence, suppose that $d > 0$ and $\epsilon > 0$. Since Gaussian noise is an admissible submodel, it suffices to establish the contamination-dependent lower bound under $\xi \sim \mathcal{N}\left(0, \sigma^2\right)$, independently of X.

Let f_c be a star center of $\mathcal{F}$. The triangle inequality gives

$$d \le 2 \sup_{f \in \mathcal{F}} \|f - f_c\|_2 .$$

Thus, there exists $u \in \mathcal{F}$ with $\|u - f_c\|_2 \ge d/3$. By star-shapedness, the entire segment joining f_c and u belongs to $\mathcal{F}$. Setting

$$D := \min\left\{\frac{d}{3}, \frac{\sigma\epsilon}{2}\right\},$$

we can therefore choose $f_0, f_1 \in \mathcal{F}$ such that $\|f_1 - f_0\|_2 = D$.

For $j \in \{0, 1\}$, let P_j denote the joint distribution of one clean observation under

$$\widetilde{X} \sim \mathbb{P}_X, \qquad \widetilde{Y} = f_j\left(\widetilde{X}\right) + \xi, \qquad \xi \sim \mathcal{N}\left(0, \sigma^2\right), \qquad \xi \perp \widetilde{X}.$$

Because the design distribution is the same under both hypotheses, the conditional Gaussian likelihoods yield

$$\mathrm{KL}\left(P_0 \,\|\, P_1\right) = \frac{1}{2\sigma^2}\mathbb{E}_{\widetilde{X}}\left[\left(f_0\left(\widetilde{X}\right) - f_1\left(\widetilde{X}\right)\right)^2\right] = \frac{D^2}{2\sigma^2}.$$

By Pinsker's inequality,

$$q := \mathrm{TV}\left(P_0, P_1\right) \le \sqrt{\frac{1}{2}\mathrm{KL}\left(P_0 \,\|\, P_1\right)} = \frac{D}{2\sigma} \le \frac{\epsilon}{4}.$$

We next construct an admissible corruption procedure that makes the two experiments difficult to distinguish. By maximal coupling, there exists a pair (U, V) with marginal distributions P_0 and P_1 such that

$$\mathbb{P}\left(U \ne V\right) = q.$$

Take n independent copies (U_i, V_i) of this coupling, and define

$$\widetilde{\mathcal{D}}_0 := \left(U_1, \ldots, U_n\right), \qquad \widetilde{\mathcal{D}}_1 := \left(V_1, \ldots, V_n\right).$$

Then $\widetilde{\mathcal{D}}_j \sim P_j^{\otimes n}$, and the number of differing observations satisfies

$$W := \sum_{i=1}^{n} \mathbf{1}\{U_i \neq V_i\} \sim \mathrm{Bin}(n, q).$$

Under f_0, let $\mathcal{C}_0$ be the identity corruption map. Under f_1, the adversary first samples a companion dataset $\widetilde{\mathcal{D}}_0$ from its conditional distribution given the observed $\widetilde{\mathcal{D}}_1$ under the coupling above. If $W \leq \epsilon n$, it replaces precisely the differing observations and outputs $\widetilde{\mathcal{D}}_0$. If $W > \epsilon n$, it leaves $\widetilde{\mathcal{D}}_1$ unchanged. Denote this procedure by $\mathcal{C}_1$.

Since W is integer-valued, $\mathcal{C}_1$ always replaces at most $\lfloor \epsilon n \rfloor$ observations. Hence, it is an admissible strong corruption procedure. Its use of auxiliary randomness is harmless for a minimax lower bound: the supremum over deterministic corruption maps is at least the risk averaged over these random seeds.

Let

$$Q_0 := P_0^{\otimes n}, \qquad Q_1 := \mathcal{L}\left(\mathcal{C}_1\left(\widetilde{\mathcal{D}}_1\right)\right).$$

Under the coupling just constructed, the output $\mathcal{C}_1\left(\widetilde{\mathcal{D}}_1\right)$ agrees with $\widetilde{\mathcal{D}}_0$ whenever $W \leq \epsilon n$. Therefore, the coupling inequality and Markov's inequality give

$$\mathrm{TV}(Q_0, Q_1) \leq \mathbb{P}(W > \epsilon n) \leq \frac{\mathbb{E}[W]}{\epsilon n} = \frac{q}{\epsilon} \leq \frac{1}{4}.$$

This comparison uses the unconditional law of $\widetilde{\mathcal{D}}_0$; it does not require that conditioning on $W \leq \epsilon n$ preserves that law.

For completeness, given any estimator $\widehat{f}$, construct a test by selecting whichever of f_0 and f_1 is closer to $\widehat{f}$ in $L_2(\mathbb{P}_X)$, with ties resolved arbitrarily. Whenever this test is incorrect under hypothesis j, the triangle inequality implies $\left\|\widehat{f} - f_j\right\|_2 \geq D/2$. The sum of the two testing error probabilities is at least $1 - \mathrm{TV}(Q_0, Q_1)$. Consequently,

$$\max_{j \in \{0,1\}} \mathbb{E}_{Q_j}\left[\left\|\widehat{f} - f_j\right\|_2^2\right] \geq \frac{D^2}{8}\left[1 - \mathrm{TV}(Q_0, Q_1)\right] \geq \frac{3D^2}{32}.$$

Taking the infimum over $\widehat{f}$, and recalling that these two functions and corruption procedures are admissible, yields

$$\mathcal{R}_{n,\epsilon}(\mathcal{F}; \mathcal{P}_\xi^{\mathrm{sg}}) \geq \frac{3}{32} \min\left\{\frac{d^2}{9}, \frac{\sigma^2\epsilon^2}{4}\right\} \geq \frac{1}{96}\left(\sigma^2\epsilon^2 \wedge d^2\right).$$

The proof is completed. □

C.2 Proof of Lemma 4.3

Proof. Recall that

$$Z_{f,g} := \frac{[Y - f(X)]^2 - [Y - g(X)]^2}{\|f - g\|_{L_2(\mathbb{P}_X)}} = \frac{g(X) - f(X)}{\|f - g\|_{L_2(\mathbb{P}_X)}}\left[[2f^*(X) - f(X) - g(X)] + 2\xi\right].$$

First, by a direct application of (4.1), we know

$$\left\| \frac{g-f}{\|f-g\|_{L_2(\mathbb{P}_X)}} \right\|_{L_{2p}(\mathbb{P}_X)} \leq L_{\mathrm{sg}}\sqrt{2p}.$$

Similarly, we have

$$\|2f^*-f-g\|_{L_{2p}(\mathbb{P}_X)} \leq L_{\mathrm{sg}}\sqrt{2p}\left(\|f^*-f\|_{L_2(\mathbb{P}_X)}+\|f^*-g\|_{L_2(\mathbb{P}_X)}\right) \leq L_{\mathrm{sg}}\sqrt{2p}\left(\|f-g\|_{L_2(\mathbb{P}_X)}+2\delta\right),$$

where the latter step uses the assumption that $\|f^*-g\|_{L_2(\mathbb{P}_X)} \leq \delta$ and the triangle inequality.

Combining the last two arguments with Hölder's inequality and the assumption that $\|f-g\|_{L_2(\mathbb{P}_X)} \geq C\delta$, we come to

$$\left\| \frac{g-f}{\|f-g\|_{L_2(\mathbb{P}_X)}}(2f^*-f-g) \right\|_{L_p(\mathbb{P}_X)} \leq 2L_{\mathrm{sg}}^2(\|f-g\|_{L_2(\mathbb{P}_X)}+2\delta)p \leq 4L_{\mathrm{sg}}^2\|f-g\|_{L_2(\mathbb{P}_X)}p.$$

By a parallel argument, we have

$$\left\| 2\frac{g-f}{\|f-g\|_{L_2(\mathbb{P}_X)}}\xi \right\|_{L_p(\mathbb{P}_{\xi,X})} \leq 2\cdot L_{\mathrm{sg}}\sqrt{2p}\cdot L_\xi\sqrt{2p}\sigma = 4L_{\mathrm{sg}}L_\xi\sigma p.$$

Therefore, we have

$$\|Z_{f,g}-\mathbb{E}Z_{f,g}\|_{L_p(\mathbb{P}_{\xi,X})} \leq 2\|Z_{f,g}\|_{L_p(\mathbb{P}_{\xi,X})} \leq Kp(\|f-g\|_{L_2(\mathbb{P}_X)}+\sigma),$$

where $K = 8L_{\mathrm{sg}}(L_{\mathrm{sg}}+L_\xi)$ depends only on L_{sg} and L_ξ. □

C.3 Proof of Lemma 4.4

Proof. We first evaluate $\mathcal{E}(\eta, Z_{f,g})$ (recall its definition from Lemma 3.6). By Hölder's inequality and Lemma 4.3, we have

$$\mathbb{E}\left[\left|\bar{Z}_{f,g}\right|\mathbf{1}_{\{\bar{Z}_{f,g}\geq Q_{1-\eta/2}(\bar{Z})\}}\right] \leq \left(\mathbb{E}\left|\bar{Z}_{f,g}\right|^p\right)^{\frac{1}{p}}\left(\frac{\eta}{2}\right)^{1-1/p} \leq K\left(\|f-g\|_{L_2(\mathbb{P}_X)}+\sigma\right)p\left(\frac{\eta}{2}\right)^{1-\frac{1}{p}}$$

Select $p = 2\vee\log\left(\frac{2e}{\eta}\right)$. We have:

$$\left(\frac{\eta}{2}\right)^{-\frac{1}{p}} = \begin{cases} \exp\left\{\frac{1}{1-1/\log(\eta/2)}\right\} & ,\ (\eta \leq \frac{2}{e}), \\ \sqrt{\frac{2}{\eta}} & ,\ (\frac{2}{e} < \eta < 1) \end{cases}$$

One can easily verify that $\sqrt{2} \leq \left(\frac{\eta}{2}\right)^{-1/p} \leq e$. Therefore, we have:

$$\begin{aligned}\mathcal{E}(\eta, Z_{f,g}) &\leq \frac{Kpe\eta}{2}\left(\|f-g\|_{L_2(\mathbb{P}_X)}+\sigma\right) \\ &\leq Ke\eta\log\left(\frac{2e}{\eta}\right)\left(\|f-g\|_{L_2(\mathbb{P}_X)}+\sigma\right).\end{aligned}$$

We have $\eta = 8\epsilon + 24t_\delta \le 8\epsilon + 24a \le \frac{1}{8}$ by our selection of ϵ and a. Also, $4\exp\{-nt_\delta\} \ge 4\exp\{-an\} \ge \exp\{-n/2\}/4$. Applying (3.5) in Lemma 3.6 with X replaced by $Z_{f,g}$ and Lemma 4.3 with $p = 2$, we know that with probability at least $1 - 4\exp\{-nt_\delta\}$, we have

$$\begin{aligned}
|\mathrm{Tm}(Z_{f,g}) - \mathbb{E}Z_{f,g}| &\le 16Ke\eta \log\left(\frac{e}{2\eta}\right)\left(\|f-g\|_{L_2(\mathbb{P}_X)} + \sigma\right) + 4\sqrt{2}K\left(\|f-g\|_{L_2(\mathbb{P}_X)} + \sigma\right)\sqrt{t_\delta} \\
&\le 16Ke\left(\|f-g\|_{L_2(\mathbb{P}_X)} + \sigma\right)\left(\eta\log\left(\frac{e}{\eta}\right) + \sqrt{t_\delta}\right) \\
&\overset{\text{(i)}}{\le} 16Ke\left(\|f-g\|_{L_2(\mathbb{P}_X)} + \sigma\right)\left(8\epsilon\log\left(\frac{e}{\epsilon}\right) + 24t_\delta\log\left(\frac{e}{t_\delta}\right) + \sqrt{t_\delta}\right) \\
&\overset{\text{(ii)}}{\le} 128Ke\underbrace{\left(\|f-g\|_{L_2(\mathbb{P}_X)} + \sigma\right)\epsilon\log\left(\frac{e}{\epsilon}\right)}_{\text{contamination term}} + 384Ke\underbrace{\left(\|f-g\|_{L_2(\mathbb{P}_X)} + \sigma\right)\sqrt{t_\delta}}_{\text{confidence term}},
\end{aligned}$$

where in (i) we use $(a+b)\log(e/(a+b)) \le a\log(e/a) + b\log(e/b)$ and (ii) follows from the fact that $x\log(e/x) \le \sqrt{x}$ for small x (equivalently, sufficiently small a).

For the confidence term, observe that

$$\begin{aligned}
(\|f-g\|_{L_2(\mathbb{P}_X)} + \sigma)\sqrt{t_\delta} &= \sqrt{a}\left(\|f-g\|_{L_2(\mathbb{P}_X)} + \sigma\right)\min\left\{\frac{\delta}{\sigma}, 1\right\} \\
&\le \sqrt{a}\left(\|f-g\|_{L_2(\mathbb{P}_X)} + \delta\right) \\
&\le \sqrt{a}\left(1 + \frac{1}{C}\right)\|f-g\|_{L_2(\mathbb{P}_X)}.
\end{aligned}$$

The middle inequality holds in both regimes: for $\delta \le \sigma$, expand the product; for $\delta > \sigma$, use $\sigma < \delta$. The last inequality follows from $\|f-g\|_{L_2(\mathbb{P}_X)} \ge C\delta$.

For the contamination term, suppose $\delta \ge \frac{3456Ke}{C}\sigma\epsilon\log\left(\frac{e}{\epsilon}\right)$. Then

$$\begin{aligned}
(\|f-g\|_{L_2(\mathbb{P}_X)} + \sigma)\epsilon\log\left(\frac{e}{\epsilon}\right) &\le \|f-g\|_{L_2(\mathbb{P}_X)}\epsilon\log\left(\frac{e}{\epsilon}\right) + \frac{C\delta}{3456Ke} \\
&\le \|f-g\|_{L_2(\mathbb{P}_X)}\left(\epsilon\log\left(\frac{e}{\epsilon}\right) + \frac{1}{3456Ke}\right).
\end{aligned}$$

Combining these estimates,

$$|\,\mathrm{TM}(Z_{f,g}) - \mathbb{E}Z_{f,g}| \le 384Ke\,\|f-g\|_{L_2(\mathbb{P}_X)}\left[\epsilon\log\left(\frac{e}{\epsilon}\right) + \frac{1}{3456Ke} + \sqrt{a}\left(1 + \frac{1}{C}\right)\right].$$

Take ϵ_0 small enough that $384Ke\epsilon_0\log\left(\frac{e}{\epsilon_0}\right) \le 1/9$. The bracketed expression, including K, is then at most $1/3$. Recall that $\mathbb{E}Z_{f,g} \ge \|f-g\|_{L_2(\mathbb{P}_X)} - 2\delta \ge \left(1 - \frac{2}{C}\right)\|f-g\|_{L_2(\mathbb{P}_X)} \ge 1/3\,\|f-g\|_{L_2(\mathbb{P}_X)}$, which implies that the clean signal dominates the variation. Consequently, we correctly select g over f. The proof is completed. □