



Available online at www.sciencedirect.com

ScienceDirect

Journal of Approximation Theory 273 (2022) 105660

JOURNAL OF
Approximation
Theory

www.elsevier.com/locate/jat

Full Length Article

A statistical learning assessment of Huber regression

Yunlong Feng^{a,*}, Qiang Wu^b

^a Department of Mathematics and Statistics, University at Albany, United States of America

^b Department of Mathematical Sciences, Middle Tennessee State University, United States of America

Received 19 February 2021; received in revised form 15 August 2021; accepted 8 October 2021

Available online 22 October 2021

Communicated by Ding-Xuan Zhou

Abstract

As one of the triumphs and milestones of robust statistics, Huber regression plays an important role in robust inference and estimation. It has also been finding a great variety of applications in machine learning. In a parametric setup, it has been extensively studied. However, in the statistical learning context where a function is typically learned in a nonparametric way, there is still a lack of theoretical understanding of how Huber regression estimators learn the conditional mean function and why it works in the absence of light-tailed noise assumptions. To address these fundamental questions, this paper conducts an assessment of Huber regression from a statistical learning viewpoint. First, we show that the usual risk consistency property of Huber regression estimators, which is usually pursued in machine learning, cannot guarantee their learnability in mean regression. Second, we argue that Huber regression should be implemented in an adaptive way to perform mean regression, implying that one needs to tune the scale parameter in accordance with the sample size and the moment condition of the noise. Third, with an adaptive choice of the scale parameter, we demonstrate that Huber regression estimators can be asymptotic mean regression calibrated under $(1 + \epsilon)$ -moment conditions ($\epsilon > 0$) on the conditional distribution. Last but not least, under the same moment conditions, we establish almost sure convergence rates for Huber regression estimators. Note that the $(1 + \epsilon)$ -moment conditions accommodate the special case where the response variable possesses infinite variance and so the established convergence rates justify the robustness feature of Huber regression estimators. In the above senses, the present study provides a systematic statistical learning assessment of Huber regression estimators and justifies their merits in terms of robustness from a theoretical viewpoint.

© 2021 Elsevier Inc. All rights reserved.

Keywords: Learning theory; Huber regression; Robustness; Convergence rates; Weak moment condition

* Corresponding author.

E-mail address: ylfeng@albany.edu (Y. Feng).

1. Introduction and motivation

In this paper, we are concerned with the robust regression problem where one aims at seeking a functional relation between input and output when the response variable may be heavy-tailed [11,15,18,22]. In such scenarios, the traditionally frequently used least squares regression paradigms may not work well due to the amplification of the least squares loss to large residuals. As an alternative, the Huber loss was proposed in the seminal work [13] in the context of robust estimation of location parameters. The Huber loss and the theoretical findings in location parameter estimation then applied and carried over to robust regression problems. The regression paradigm that is associated with the Huber loss is termed as *Huber regression* and the resulting estimator is termed as the *Huber regression estimator*. The introduction of Huber regression led to the development of various subsequent M-estimators and fostered the development of robust statistics into a discipline.

Denoting X as the input variable that takes values in a compact metric space $\mathcal{X} \subset \mathbb{R}^d$ and Y the response variable taking values in $\mathcal{Y} \subset \mathbb{R}$, given i.i.d observations $\mathbf{z} = \{(x_i, y_i)\}_{i=1}^n$, in the context of parametric regression, the Huber regression estimator $f_{\mathbf{z}, \sigma}(X) = X^\top \hat{\beta}$ is learned from the following empirical risk minimization (ERM) scheme

$$f_{\mathbf{z}, \sigma} := \arg \min_{f \in \mathcal{H}} \frac{1}{n} \sum_{i=1}^n \ell_\sigma(y_i - f(x_i)), \quad (1)$$

where \mathcal{H} is the function space from \mathcal{X} to \mathbb{R} consisting of linear functions of the form $f(x) = x^\top \beta$ and ℓ_σ is the well-known Huber loss defined by

$$\ell_\sigma(t) = \begin{cases} t^2, & \text{if } |t| \leq \sigma, \\ 2\sigma|t| - \sigma^2, & \text{otherwise.} \end{cases} \quad (2)$$

Assuming that the conditional mean function $f^*(X) = \mathbb{E}(Y|X)$ can be parametrically represented as $f^*(X) = X^\top \beta^*$ and the noise $Y - f^*(X)$ is zero mean when conditioned on X , asymptotic properties of $\hat{\beta}$ and its convergence to β^* have been extensively studied in the literature of parametric statistics. An incomplete list of related literature includes [11,12,14–16,18,20,22,28] and many references therein. Note that in the aforementioned studies, the scale parameter σ in the Huber loss is set to be fixed and chosen according to the 95% asymptotic efficiency rule. In a high-dimensional setting, Huber regression with a fixed scale parameter, however, may not be able to learn β^* when the noise is asymmetric, as argued recently in [9,25]. There the authors proposed to choose the scale parameter by relating it to the dimension of the input space, the moment condition of the noise distribution, and the sample size so that one may debias the resulting regression estimator and the scale parameter can play a trade-off role between bias and robustness.

In a nonparametric statistical learning context where functions in \mathcal{H} , in general, do not admit parametric representations, theoretical investigations of Huber regression estimators are still sparse though they have been applied extensively into various applications where robustness is a concern. To proceed with our discussion, denote \mathcal{H} as a compact subset of the space $C(\mathcal{X})$ of continuous functions on \mathcal{X} , ρ the underlying unknown distribution over $\mathcal{X} \times \mathcal{Y}$, and $\mathcal{R}^\sigma(f)$ the *generalization error* of $f : \mathcal{X} \rightarrow \mathbb{R}$ defined by

$$\mathcal{R}^\sigma(f) = \mathbb{E} \ell_\sigma(Y - f(X)),$$

where the expectation is taken jointly with respect to X and Y . Recall that the objective is to learn the conditional mean function $f^*(X) = \mathbb{E}(Y|X)$ robustly. Existing studies in the literature

of statistical learning theory remind that Huber regression estimators are \mathcal{R}^σ -risk consistent, i.e., $\mathcal{R}^\sigma(f_{\mathbf{z},\sigma}) \rightarrow \min_{f \in \mathcal{H}} \mathcal{R}^\sigma(f)$ as $n \rightarrow \infty$. To see this, one simply notes that the Huber loss is Lipschitz continuous on \mathbb{R} and so the usual learning theory arguments may apply [2,6,7,10,21]. However, such a risk consistency property neither says anything about the role that the scale parameter σ plays nor implies the convergence of the regression estimator $f_{\mathbf{z},\sigma}$ to the mean regression function f^* .

In the present study, we aim to conduct a statistical learning assessment of the Huber regression estimator $f_{\mathbf{z},\sigma}$. More specifically, we pursue answers to the following fundamental questions:

- **Question 1:** Whether \mathcal{R}^σ -risk consistency implies the convergence of $f_{\mathbf{z},\sigma}$ to f^* ?
- **Question 2:** What is the role that σ plays when learning f^* through ERM (1)?
- **Question 3:** How to develop exponential-type fast convergence rates of $f_{\mathbf{z},\sigma}$ to f^* ?
- **Question 4:** How to justify the learnability of $f_{\mathbf{z},\sigma}$ in the absence of light-tail noise?

Answers to these questions represent our main contributions. In particular, if \mathcal{R}^σ -risk consistency implies the convergence of $f_{\mathbf{z},\sigma}$ to f^* , we say that Huber regression (1) is mean regression calibrated. We show that Huber regression is generally not mean regression calibrated for any fixed scale parameter σ . Instead, it should be implemented in an adaptive way in order to perform mean regression, where the adaptiveness refers to the dependence of the scale parameter on the sample size and the moment condition. We also show that the scale parameter needs to diverge in accordance with the sample size to ensure that the Huber regression estimator $f_{\mathbf{z},\sigma}$ learns the mean regression function f^* , which we term as the asymptotic mean regression calibration property. Furthermore, such an asymptotic mean regression calibration property can be established under $(1 + \epsilon)$ -moment conditions ($\epsilon > 0$) on the conditional distribution. This is a rather weak condition as it admits the case where the conditional distribution possesses infinite variance. To develop fast exponential-type convergence rates, we establish a relaxed Bernstein condition. The idea is to bound the second moment of associated random variables by using their first moment and an additional bias term that diminishes towards 0 when the sample size tends to infinity. These preparations allow us to establish fast exponential-type convergence rates for $f_{\mathbf{z},\sigma}$. Interestingly, but not surprisingly, it is shown that σ plays a trade-off role between bias and learnability, and the convergence rates of $f_{\mathbf{z},\sigma}$ depend on the order of the imposed moment conditions.

The rest of this paper is organized as follows. In Section 2, we argue that risk consistency is insufficient in guaranteeing learnability and so does not necessarily imply the convergence of Huber regression estimators to the mean regression function. In Section 3, we demonstrate that Huber regression is asymptotically mean regression calibrated if the scale parameter is chosen in a diverging manner in accordance to the sample size and the moment condition. Some efforts are then made in Section 4 to develop fast exponential-type convergence rates by relaxing the standard Bernstein condition in learning theory. In Section 5, we establish fast convergence rates for Huber regression estimators under weak moment conditions. Proofs of Theorems are collected in Section 6. The paper is concluded in Section 7.

Notation and Convention. Throughout this paper, we assume that f^* is bounded and $\mathcal{H} \subset C(\mathcal{X})$ is uniformly bounded and denote $M = \max\{\|f^*\|_\infty, \sup_{f \in \mathcal{H}} \|f\|_\infty\}$. Denoting $\rho_{\mathcal{X}}$ as the marginal distribution of ρ on \mathcal{X} , then $\|\cdot\|_{2,\rho}$ defines the L_2 -norm induced by $\rho_{\mathcal{X}}$. The notation $a \lesssim b$ denotes the fact that there exists a positive constant c such that $a \leq cb$. For any $t \in \mathbb{R}$, let $t_+ = \max(0, t)$.

2. Risk consistency is insufficient in guaranteeing learnability

In this section, we shall make efforts to answer Question 1 listed in the introduction by arguing that risk consistency is insufficient in guaranteeing learnability of the Huber regression estimator, where by *risk consistency* we refer to the convergence of $\mathcal{R}^\sigma(f_{z,\sigma})$ to $\min_{f \in \mathcal{H}} \mathcal{R}^\sigma(f)$ while *learnability* refers to the convergence of $f_{z,\sigma}$ to f^* .

Following existing studies on empirical risk minimization schemes induced by convex loss functions, it is easy to deduce that $f_{z,\sigma}$ is \mathcal{R}^σ -risk consistent for any fixed σ value. Moreover, under certain mild assumptions, probabilistic convergence rates may also be established. To see this, note that the deduction of the risk consistency property of Huber regression estimators as well as their convergence rates involves the following set of random variables

$$\mathcal{G}_\mathcal{H} = \left\{ \xi_f \mid \xi_f := \ell_\sigma(Y - f(X)) - \ell_\sigma(Y - f_{\mathcal{H},\sigma}(X)), f \in \mathcal{H} \right\},$$

where $f_{\mathcal{H},\sigma} = \arg \min_{f \in \mathcal{H}} \mathcal{R}^\sigma(f)$ is the population version of $f_{z,\sigma}$. Notice that the Huber loss ℓ_σ in (2) is Lipschitz continuous on \mathbb{R} with Lipschitz constant 2σ . Therefore, the random variables in $\mathcal{G}_\mathcal{H}$ and their variances can be uniformly upper bounded by constants involving σ . Applying learning theory arguments and concentration inequalities to $\mathcal{G}_\mathcal{H}$, under mild assumptions, convergence rates can be derived. However, due to the dependence of $f_{z,\sigma}$ on the scale parameter σ , it may possess much flexibility and can be quite different with different choices of the σ values. Consequently, the \mathcal{R}^σ -risk consistency property as well as the convergence rates of $\mathcal{R}^\sigma(f_{z,\sigma}) - \min_{f \in \mathcal{H}} \mathcal{R}^\sigma(f)$ may not be informative and may not indicate the learnability of $f_{z,\sigma}$ in learning f^* even if \mathcal{H} is perfectly chosen such that $f^* \in \mathcal{H}$.

To illustrate this phenomenon numerically, consider a toy example with the model

$$Y = 2 \sin(\pi X) + (1 + 2X)\varepsilon,$$

where X follows a uniform distribution on $[0, 1]$ and $\varepsilon \sim 0.5N(0, 2.5^2) + 0.5N(0, 0.5^2)$. It is apparent that the noise distribution admits zero mean and is skewed. Simple calculation shows that for this regression model, the conditional mean function $f^*(X) = 2 \sin(\pi X)$. In this experiment, we visualize the function $f_{z,\sigma}$ and compare it with f^* . We choose the hypothesis space \mathcal{H} as a ball of the reproducing kernel Hilbert space associated with the Gaussian kernel $K(x_i, x_j) = \exp\{-\|x_i - x_j\|^2/h^2\}$. Both the kernel bandwidth h and the radius of the ball are tuned via cross-validation under the least absolute deviation error criterion while the scale parameter σ in the Huber loss is set to be fixed with $\sigma = 0.01$. A set of independent observations are sampled from the above regression model and are used as the training data. Then $f_{z,\sigma}$ is plotted in Fig. 1. The conditional mean function is also plotted for comparison. As discussed earlier, due to the Lipschitz continuity of the Huber loss, with the choice of $\sigma = 0.01$, the risk consistency can be guaranteed. However, from the plots in Fig. 1, clearly, $f_{z,\sigma}$ does not approach the conditional mean function.

The fact that \mathcal{R}^σ -risk consistency of Huber regression estimators cannot guarantee their ability to learn the conditional mean function can be further justified through the following example. Let \mathcal{M} be the space of measurable functions from \mathcal{X} to \mathbb{R} and define

$$f_\sigma := \arg \min_{f \in \mathcal{M}} \mathcal{R}^\sigma(f). \quad (3)$$

Intuitively, f_σ can be regarded as the best Huber regression estimator learned in an ideal case where infinite observations are available and the hypothesis space is perfectly selected.

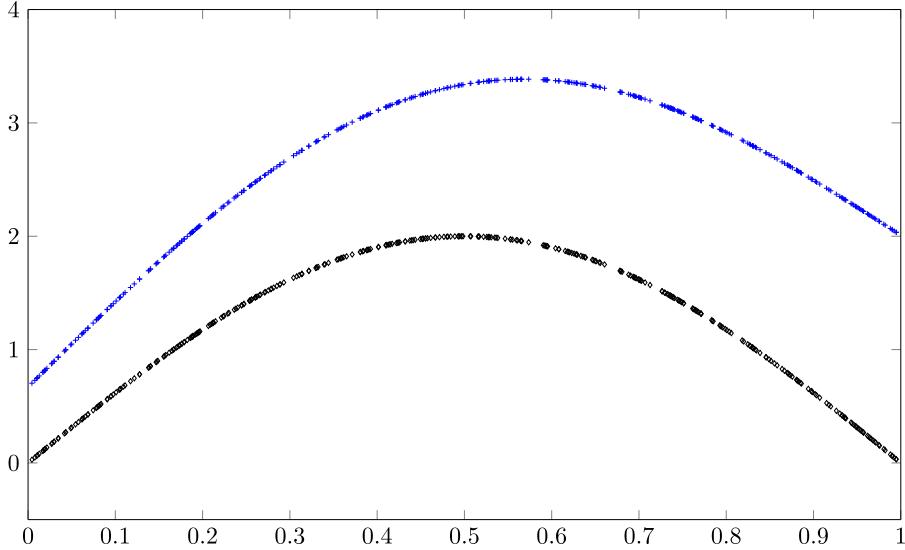


Fig. 1. The bottom black curve gives the conditional mean function. The top blue curve represents the learned Huber regression estimator with $\sigma = 0.01$.

Example 1. Consider the Huber regression problem where one aims to learn the conditional mean function f^* from the following homoscedastic regression model

$$Y = f^*(X) + \varepsilon,$$

where ε is the zero-mean noise variable with density

$$p_\varepsilon(t) = \begin{cases} \frac{1}{2}e^{-(t+\frac{1}{4})}, & \text{if } t \geq -\frac{1}{4}, \\ e^{2(t+\frac{1}{4})}, & \text{if } t < -\frac{1}{4}. \end{cases}$$

Then there exists a constant c with $c \neq 0$ such that $f_\sigma(x) = f^*(x) + c$ for all $x \in \mathcal{X}$. As a result, if $f_\sigma \in \mathcal{H}$ and $\mathcal{R}^\sigma(f_{\mathbf{z},\sigma})$ converges to $\min_{f \in \mathcal{H}} \mathcal{R}^\sigma(f)$ as $n \rightarrow \infty$ with large probability, then $f_{\mathbf{z},\sigma}$ does not converge to f^* with large probability.

Proof. Recalling the definition of f_σ in (3), for any $x \in \mathcal{X}$, we can re-express it as follows

$$\begin{aligned} f_\sigma(x) &= \arg \min_{\nu \in \mathbb{R}} \int_{\mathbb{R}} \ell_\sigma(t - \nu) p_{Y|X=x}(t) dt \\ &= \arg \min_{\nu \in \mathbb{R}} \int_{\mathbb{R}} \ell_\sigma(t - \nu) p_\varepsilon(t - f^*(x)) dt \\ &= \arg \min_{\nu \in \mathbb{R}} \int_{\mathbb{R}} \ell_\sigma(u - (\nu - f^*(x))) p_\varepsilon(u) du. \end{aligned}$$

Therefore, for any $x \in \mathcal{X}$, we have

$$f_\sigma(x) - f^*(x) = \arg \min_{\nu \in \mathbb{R}} \int_{\mathbb{R}} \ell_\sigma(u - \nu) p_\varepsilon(u) du.$$

The assumption that the noise ε is independent of x tells us that

$$\arg \min_{v \in \mathbb{R}} \int_{\mathbb{R}} \ell_{\sigma}(u - v) p_{\varepsilon}(u) du \quad (4)$$

is a unique constant for all $x \in \mathcal{X}$. To prove the first part of the assertion, we only need to show that 0 is not a solution to the above minimization problem.

From the definition of the Huber loss (2), we know that

$$\begin{aligned} \int_{\mathbb{R}} \ell_{\sigma}(u - v) p_{\varepsilon}(u) du &= \int_{\mathbb{R}} \ell_{\sigma}(u - v) p_{\varepsilon}(u) du + \int_{\mathbb{R}} (u - v)^2 p_{\varepsilon}(u) du - \int_{\mathbb{R}} (u - v)^2 p_{\varepsilon}(u) du \\ &= \int_{\mathbb{R}} (u - v)^2 p_{\varepsilon}(u) du + \int_{|u-v| \geq \sigma} (2\sigma |u - v| - \sigma^2) p_{\varepsilon}(u) du \\ &\quad - \int_{|u-v| \geq \sigma} (u - v)^2 p_{\varepsilon}(u) du \\ &= \int_{\mathbb{R}} (u - v)^2 p_{\varepsilon}(u) du - \int_{|u-v| \geq \sigma} (|u - v| - \sigma)^2 p_{\varepsilon}(u) du \\ &= \int_{\mathbb{R}} (u - v)^2 p_{\varepsilon}(u) du - \int_{u-v \geq \sigma} (u - v - \sigma)^2 p_{\varepsilon}(u) du \\ &\quad - \int_{u-v \leq -\sigma} (u - v + \sigma)^2 p_{\varepsilon}(u) du. \end{aligned}$$

Therefore, we have

$$\begin{aligned} \frac{d(\int_{\mathbb{R}} \ell_{\sigma}(u - v) p_{\varepsilon}(u) du)}{dv} &= -2 \int_{\mathbb{R}} (u - v) p_{\varepsilon}(u) du + 2 \int_{v+\sigma}^{+\infty} (u - v - \sigma) p_{\varepsilon}(u) du \\ &\quad + 2 \int_{-\infty}^{v-\sigma} (u - v + \sigma) p_{\varepsilon}(u) du. \end{aligned}$$

The zero-mean noise assumption tells us that

$$\frac{d(\int_{\mathbb{R}} \ell_{\sigma}(u - v) p_{\varepsilon}(u) du)}{dv} \Big|_{v=0} = 2 \int_{\sigma}^{+\infty} (u - \sigma) p_{\varepsilon}(u) du + 2 \int_{-\infty}^{-\sigma} (u + \sigma) p_{\varepsilon}(u) du.$$

If $\sigma \geq \frac{1}{4}$, then we have

$$\begin{aligned} &2 \int_{\sigma}^{+\infty} (u - \sigma) p_{\varepsilon}(u) du + 2 \int_{-\infty}^{-\sigma} (u + \sigma) p_{\varepsilon}(u) du \\ &= \int_{\sigma}^{+\infty} (u - \sigma) e^{-(u+\frac{1}{4})} du + 2 \int_{-\infty}^{-\sigma} (u + \sigma) e^{2(u+\frac{1}{4})} du \\ &= e^{-\sigma-\frac{1}{4}} - \frac{1}{2} e^{\frac{1}{2}-2\sigma} > 0, \end{aligned}$$

where we used the fact that $g(a) = a - \frac{e}{2}a^2$ is positive for $a \in (0, \frac{2}{e})$ and $0 < e^{-\frac{1}{4}-\sigma} \leq e^{-\frac{1}{2}} < \frac{2}{e}$.

If $0 < \sigma < \frac{1}{4}$, then we have

$$\begin{aligned} &2 \int_{\sigma}^{+\infty} (u - \sigma) p_{\varepsilon}(u) du + 2 \int_{-\infty}^{-\sigma} (u + \sigma) p_{\varepsilon}(u) du \\ &= \int_{\sigma}^{+\infty} (u - \sigma) e^{-(u+\frac{1}{4})} du + \int_{-\frac{1}{4}}^{-\sigma} (u + \sigma) e^{-(u+\frac{1}{4})} du + 2 \int_{-\infty}^{-\frac{1}{4}} (u + \sigma) e^{2(u+\frac{1}{4})} du \\ &= 2\sigma + e^{-\sigma-\frac{1}{4}} - e^{\sigma-\frac{1}{4}} > 0, \end{aligned}$$

where we used the fact that the function $g(\sigma) = 2\sigma + e^{-\sigma-\frac{1}{4}} - e^{\sigma-\frac{1}{4}}$ has $g(0) = 0$ and is strictly increasing.

Therefore, $\frac{d(\int_{\mathbb{R}} \ell_{\sigma}(u-v) p_{\varepsilon}(u) du)}{dv} \Big|_{v=0}$ is non-zero for fixed $\sigma > 0$, which implies that 0 is not a solution to the minimization problem (4). This proves the first claim in [Example 1](#).

To prove the divergence of $f_{z,\sigma}$ to f^* for any fixed σ , first note that, by the convexity of the Huber loss, f_{σ} is the unique minimizer of $\mathcal{R}^{\sigma}(f)$. Since $f_{\sigma} = f^* + c \neq f^*$, we have $\mathcal{R}^{\sigma}(f_{\sigma}) \neq \mathcal{R}^{\sigma}(f^*)$. If $f_{\sigma} \in \mathcal{H}$ and $\mathcal{R}^{\sigma}(f_{z,\sigma})$ converges to $\min_{f \in \mathcal{H}} \mathcal{R}^{\sigma}(f) = \mathcal{R}^{\sigma}(f_{\sigma})$ with large probability, then there exists some large enough N such that for any $n > N$,

$$|\mathcal{R}^{\sigma}(f_{z,\sigma}) - \mathcal{R}^{\sigma}(f_{\sigma})| \leq \frac{1}{2} |\mathcal{R}^{\sigma}(f_{\sigma}) - \mathcal{R}^{\sigma}(f^*)|$$

holds with large probability. By the Lipschitz property of the Huber loss, we have

$$\begin{aligned} \|f_{z,\sigma} - f^*\|_{2,\rho} &\geq \frac{1}{2\sigma} |\mathcal{R}^{\sigma}(f_{z,\sigma}) - \mathcal{R}^{\sigma}(f^*)| \\ &\geq \frac{1}{2\sigma} (|\mathcal{R}^{\sigma}(f_{\sigma}) - \mathcal{R}^{\sigma}(f^*)| - |\mathcal{R}^{\sigma}(f_{z,\sigma}) - \mathcal{R}^{\sigma}(f_{\sigma})|) \\ &\geq \frac{1}{4\sigma} |\mathcal{R}^{\sigma}(f_{\sigma}) - \mathcal{R}^{\sigma}(f^*)|. \end{aligned}$$

This proves that $f_{z,\sigma}$ does not converge to f^* for any fixed σ . \square

[Example 1](#) tells us that in some scenarios, $f_{z,\sigma}$ may not converge to the conditional mean function which one aims to learn even if infinite samples are given and the hypothesis space is perfectly chosen. This is due to the inherent bias brought by the integrated scale parameter in the Huber loss when pursuing robustness. Continuing our discussion at the beginning of this section, the general answer to Question 1 is that, \mathcal{R}^{σ} -risk consistency cannot guarantee its learnability as the gain in robustness may entail a bias. That is, in general, the Huber regression scheme (1) may not be mean regression calibrated. To address this problem and to learn the conditional mean function f^* through Huber regression, one needs to tune the scale parameter σ to reduce the bias and learn in an adaptive way, as argued in the next section.

3. Huber regression is asymptotically mean regression calibrated

In this section, we shall show that, in a distribution-free setup, with properly selected scale parameter σ , Huber regression can be asymptotically mean regression calibrated, meaning that risk consistency implies the convergence of $f_{z,\sigma}$ to the conditional mean function f^* when $\sigma \rightarrow \infty$.

3.1. Mean regression calibration property

Recall that in the context of regression learning, one of the central concerns is the convergence of the learned empirical target function to the unknown truth function of interest, that is, the conditional mean function f^* in this study. While the distance between the empirical target function and f^* is not directly accessible, one settles for bounding the excess generalization error. The underlying philosophy is that the generalization error of a learning machine can be approximated by using its empirical counterpart and the excess generalization error can be bounded via learning theory arguments. As mentioned in the introduction, the regression estimator is called *mean regression calibrated* if the convergence of the excess

generalization error towards 0 implies the convergence of the empirical target function to the conditional mean function [23].

Translating into the context of Huber regression, one is concerned with whether the convergence of the excess generalization error $\mathcal{R}^\sigma(f_{z,\sigma}) - \mathcal{R}^\sigma(f_\sigma)$ towards 0 implies the convergence of $f_{z,\sigma}$ to f^* . However, the numerical experiment and the counterexample in the preceding section suggest a negative answer and demonstrate that the desired mean regression calibration property may, in general, not be true. This conflicts with our intuition and common understanding that the Huber loss can serve as a robust alternate of the least squares loss when the scale parameter is chosen sufficiently large. To bypass this problem, in what follows, noticing our interest in learning the conditional mean function f^* , we turn to investigate the relation between the convergence of $\mathcal{R}^\sigma(f_{z,\sigma})$ to $\mathcal{R}^\sigma(f^*)$ and the convergence of $f_{z,\sigma}$ to f^* . More specifically, we shall show that under mild conditions, the convergence of $\mathcal{R}^\sigma(f_{z,\sigma})$ to $\mathcal{R}^\sigma(f^*)$ does imply the convergence of $f_{z,\sigma}$ to f^* when $\sigma \rightarrow \infty$ and $n \rightarrow \infty$. This justifies the mean regression calibration property in an asymptotic sense.

3.2. A comparison theorem

We now look into the asymptotic mean calibration property by establishing a comparison theorem. For regression estimators that are produced by empirical risk minimization schemes with convex loss functions, some efforts on investigating their mean regression calibration properties have been made in the literature; see e.g., [23]. For Huber regression estimators, it was concluded that they are mean regression calibrated if the response variable is upper bounded or the conditional noise variable $\varepsilon|X$ admits a symmetric probability density function. Recall that one of the most prominent merits of Huber regression estimators lies in that they can perform mean regression in the absence of light-noise assumptions. In this sense, boundedness or symmetry constraints on the noise variable should be considered as stringent ones. In this study, we are seeking to assess the Huber regression estimator $f_{z,\sigma}$ and investigate its mean regression calibration properties without resorting to light-tail distributional assumptions on the conditional distribution or on the noise. To this end, we introduce the following weak moment condition.

Assumption 1. There exists a constant $\epsilon > 0$ such that $\mathbb{E}|Y|^{1+\epsilon} < +\infty$.

The moment condition in [Assumption 1](#) is rather weak in the sense that it admits the case where the response variable Y possesses infinite variance. The same comment condition also applies to the distributions of the conditional random variable $Y|X$ and the conditional noise variable $\varepsilon|X$ under the additive data generating model, implying that heavy-tailed noise is allowed.

As discussed earlier, without further distributional assumptions on the noise variable, $f_{z,\sigma}$ is, in general, biased and its population version f_σ may be different from f^* almost everywhere on \mathcal{X} . However, such a bias can be upper bounded and may decrease with the increase of the σ values. Results in this regard are stated in the following theorem under the above $(1+\epsilon)$ -moment condition.

Theorem 1. Let $\sigma > \max\{2M, 1\}$. Under [Assumption 1](#), there exists a constant $c_\epsilon > 0$ independent of σ such that for any measurable function $f : \mathcal{X} \rightarrow \mathbb{R}$ with $\|f\|_\infty \leq M$, we have

$$\left| [\mathcal{R}^\sigma(f) - \mathcal{R}^\sigma(f^*)] - \|f - f^*\|_{2,\rho}^2 \right| \leq \frac{c_\epsilon}{\sigma^\epsilon}. \quad (5)$$

Theorem 1 states that for any bounded measurable function f , under the $(1 + \epsilon)$ -moment conditions, the gap between $\mathcal{R}^\sigma(f) - \mathcal{R}^\sigma(f^*)$ and $\|f - f^*\|_{2,\rho}^2$ is up to $\mathcal{O}(\sigma^{-\epsilon})$. With a sufficiently large σ value or sufficiently light-tailed noise, this gap could be sufficiently small. This is called the asymptotic regression calibration property. It characterizes the relationship between the loss function ℓ_σ and the L^2 -distance and is independent of the sample. When applying the property specifically to $f_{z,\sigma}$, the sample \mathbf{z} is involved and one may choose $\sigma = \sigma(n)$ so that $\sigma \rightarrow \infty$ as $n \rightarrow \infty$. As a special case, let us consider the presence of Gaussian or sub-Gaussian noise where the moment condition holds for arbitrarily large ϵ values. In this scenario, the gap between the above two quantities can be arbitrarily small. These findings remind us that in order to debias the Huber regression estimator, one may relate the σ value to the sample size n . In other words, from an asymptotic viewpoint, with diverging σ values, according to **Theorem 1**, Huber regression is asymptotically mean regression calibrated. Following this spirit, we shall proceed with the assessment based on diverging σ values by deriving their convergence rates to the conditional mean function f^* .

4. Relaxing the Bernstein condition for assessing Huber regression

From our previous discussions, in order to derive convergence rates for $f_{z,\sigma}$, one needs to bound the excess generalization error $\mathcal{R}^\sigma(f_{z,\sigma}) - \mathcal{R}^\sigma(f^*)$ which essentially requires us to deal with the following set of random variables

$$\mathcal{F}_H := \left\{ \xi \mid \xi(X, Y) = \ell_\sigma(Y - f(X)) - \ell_\sigma(Y - f^*(X)), \quad f \in \mathcal{H} \right\}.$$

The existing studies in learning theory remind us that it is crucial to establish the so-called *Bernstein condition*, i.e., bounding the second moment of $\xi \in \mathcal{F}_H$ by using its first moment. We will show that while the standard Bernstein condition does not hold, one can relax it to develop fast convergence rates for $\mathcal{R}^\sigma(f_{z,\sigma}) - \mathcal{R}^\sigma(f^*)$. Let us start with recapping the Bernstein condition in learning theory.

4.1. Bernstein conditions in learning theory

Originally introduced in [2] in the context of empirical risk minimization, the standard Bernstein condition can be restated as follows: a set \mathcal{F} of random variables is said to satisfy the (β, B) -Bernstein condition with $0 < \beta \leq 1$ and $B > 0$ if for any $f \in \mathcal{F}$, it holds that $\mathbb{E}f^2 \leq B(\mathbb{E}f)^\beta$. In other words, the second moment of the random variable (and so the variance) can be upper bounded by its first moment. Later, the standard Bernstein condition was generalized and extended into various other Bernstein-like conditions for analyzing learning algorithms in different contexts; see e.g., [5, 24, 26]. It turns out that the Bernstein condition and its variants play an important role in establishing fast convergence rates for learning algorithms of interest because they provide tight upper bounds for the variance of the random variables induced by the resulting estimators.

In the context of Huber regression, a Bernstein-like condition is also desired in order to establish fast convergence rates for the excess generalization error $\mathcal{R}^\sigma(f_{z,\sigma}) - \mathcal{R}^\sigma(f^*)$. However, as shown in the preceding section, without further distributional restrictions to the noise variable, f^* may not be the optimal hypothesis that minimizes $\mathcal{R}^\sigma(f)$ over the measurable function space \mathcal{M} . Consequently, $\mathcal{R}^\sigma(f) - \mathcal{R}^\sigma(f^*)$ is not necessarily positive. As a result, the usual Bernstein condition, namely, $\mathbb{E}\xi^2 \leq B(\mathbb{E}\xi)^\beta$ for $\xi \in \mathcal{F}_H$ with constants $B > 0$ and $0 < \beta \leq 1$, may not hold. This brings barriers to the development of fast convergence rates

for $\mathcal{R}^\sigma(f) - \mathcal{R}^\sigma(f^*)$. To circumvent this problem, in our study, we shall establish a relaxed Bernstein condition, which takes the form

$$\mathbb{E}\xi^2 \leq B(\mathbb{E}\xi)^\beta + g(\sigma),$$

with $B > 0$, $0 < \beta \leq 1$, and g a nonnegative function of σ . A motivating observation for establishing such a relaxed Bernstein condition is that the gap between $\mathbb{E}(\xi)$ and $\|f - f^*\|_{2,\rho}^2$ can be upper bounded by $\mathcal{O}(\sigma^{-\epsilon})$ for any $f \in \mathcal{H}$, as stated in [Theorem 1](#).

4.2. A relaxed Bernstein condition

To establish our relaxed Bernstein condition, we prove the following variance bound for $\xi \in \mathcal{F}_{\mathcal{H}}$.

Theorem 2. *Let [Assumption 1](#) hold and $\sigma > \max\{2M, 1\}$. For any measurable function $f : \mathcal{X} \rightarrow \mathbb{R}$ with $\|f\|_\infty \leq M$, the random variable $\xi = \ell_\sigma(Y - f(X)) - \ell_\sigma(Y - f^*(X))$ satisfies*

$$\mathbb{E}\xi^2 \leq c_1 \|f - f^*\|_{2,\rho}^{\frac{2(\epsilon-1)_+}{\epsilon+1}} + c_2 \sigma^{1-\epsilon},$$

where c_1 and c_2 are positive constants independent of σ or f .

Recall the results in [Theorem 1](#) that states $\|f - f^*\|_{2,\rho}^2 \leq \mathbb{E}\xi + c_\epsilon \sigma^{-\epsilon}$. This in connection with [Theorem 2](#) immediately yields the following relaxed Bernstein condition:

$$\mathbb{E}\xi^2 \leq c_1(\mathbb{E}\xi)^{\frac{(\epsilon-1)_+}{\epsilon+1}} + c_1(c_\epsilon \sigma^{-\epsilon})^{\frac{(\epsilon-1)_+}{\epsilon+1}} + c_2 \sigma^{1-\epsilon}.$$

This relaxed Bernstein condition will be crucial in establishing error bounds and fast exponential-type convergence rates for the Huber regression estimator $f_{z,\sigma}$.

5. Assessing Huber regression under weak moment conditions

In this section, we present fast exponential-type convergence rates for the Huber regression estimator under the $(1 + \epsilon)$ -moment condition in [Assumption 1](#). Specifically, we are interested in bounding the $L_{\rho,\mathcal{X}}^2$ -distance between $f_{z,\sigma}$ and f^* .

To state the result, we introduce the following capacity assumption. For any $\eta > 0$, let $\mathcal{N}(\mathcal{H}, \eta)$ denote the covering number of \mathcal{H} by the balls of radius η in $C(\mathcal{X})$, that is,

$$\mathcal{N}(\mathcal{H}, \eta) = \min \left\{ k \in \mathbb{N} : \text{there exist } f_j \in \mathcal{H}, j = 1, \dots, k \text{ such that } \mathcal{H} \subset \bigcup_{j=1}^k B(f_j, \eta) \right\},$$

where $B(f_j, \eta) = \{f \in C(\mathcal{X}) : \|f - f_j\|_\infty < \eta\}$. Our capacity condition is stated as follows.

Assumption 2. There exist positive constants q and c such that $\log \mathcal{N}(\mathcal{H}, \eta) \leq c\eta^{-q}$, $\forall \eta > 0$.

Generalization error bounds in terms of the covering number argument under [Assumption 2](#) is typical in statistical learning theory; see e.g., [\[1, 7, 23\]](#) and references therein.

To state our results on the convergence rates, we introduce a new function $f_{\mathcal{H}}$, which is defined as

$$f_{\mathcal{H}} = \arg \min_{f \in \mathcal{H}} \|f - f^*\|_{2,\rho}^2.$$

The function $f_{\mathcal{H}}$ is the optimal function in \mathcal{H} that one may expect in approximating the truth function f^* . The distance $\|f_{\mathcal{H}} - f^*\|_{2,\rho}^2$ can be regarded as the approximation error when working with the hypothesis space \mathcal{H} and so corresponds to the bias caused by the choice of the hypothesis space \mathcal{H} .

Theorem 3. *Suppose that Assumptions 1 and 2 hold and let $\sigma > \max\{2M, 1\}$. Let $f_{z,\sigma}$ be produced by (1). For any $0 < \delta < 1$, with probability at least $1 - \delta$, it holds that*

$$\|f_{z,\sigma} - f^*\|_{2,\rho}^2 \lesssim \|f_{\mathcal{H}} - f^*\|_{2,\rho}^2 + \log(2/\delta)\Psi(n, \epsilon, \sigma),$$

where

$$\Psi(n, \epsilon, \sigma) := \begin{cases} \frac{1}{\sigma^\epsilon} + \frac{\sigma}{n^{1/(q+1)}}, & \text{if } 0 < \epsilon \leq 1, \\ \frac{1}{\sigma^\epsilon} + \left(\frac{\sigma^{q+2\epsilon}}{n^{1+\epsilon}} \right)^{1/(q+1)}, & \text{if } \epsilon > 1. \end{cases}$$

The proof of Theorem 3 is based on a ratio probability inequality and standard learning theory argument [1, 7, 8, 23, 27], where the results established in Theorems 1 and 2 play crucial roles.

The error bound in Theorem 3 involves three components: the approximation error due to the imperfect choice of the hypothesis space \mathcal{H} , the inherent bias caused by the integrated parameter σ , and the sample error. In practice, the hypothesis space could be chosen by structural risk minimization so that the approximation error decreases to a tolerably small level [10, 27]. The value of σ affects both the inherent bias and the sample error. The best choice depends on the sample size, the moment condition, and the capacity of the hypothesis space. To see this, consider a special case when $f^* \in \mathcal{H}$ so that the approximation error $\|f_{\mathcal{H}} - f^*\|_{2,\rho}^2$ disappears. With properly chosen σ values, we immediately obtain the following convergence rates.

Corollary 4. *Under the assumptions of Theorem 3, let $f^* \in \mathcal{H}$ and σ be chosen as $\sigma = n^{\Phi(\epsilon, q)}$ with*

$$\Phi(\epsilon, q) = \begin{cases} \frac{1}{(1+\epsilon)(1+q)}, & \text{if } 0 < \epsilon \leq 1, \\ \frac{1+\epsilon}{q(1+\epsilon)^2 + \epsilon(\epsilon+3)}, & \text{if } \epsilon > 1. \end{cases}$$

For any $0 < \delta < 1$, with probability at least $1 - \delta$, we have

$$\|f_{z,\sigma} - f^*\|_{2,\rho}^2 \lesssim \log(2/\delta)n^{-\epsilon\Phi(\epsilon, q)}.$$

According to Corollary 4, with properly chosen diverging σ values, we obtain exponential-type convergence rates for $f_{z,\sigma}$. In particular, if the noise variable is bounded or sub-Gaussian and hence the moment condition in Assumption 1 holds for any $\epsilon > 0$, one can select an arbitrarily large ϵ to obtain convergence rates of order arbitrarily close to $O(n^{-\frac{1}{1+q}})$. As a comparison, recall that for least square estimators convergence rates of order $O(n^{-\frac{1}{1+q/2}})$ can

usually be established; see e.g., [1,3,23] and references therein. Moreover, note that with weaker moment conditions, i.e., smaller ϵ values, one gets slower convergence rates for $f_{z,\sigma}$, indicating increased sacrifice for robustness. This coincides with our intuitive understanding of robust regression estimators. On the other hand, if f^* is smooth enough and one selects a smooth hypothesis space (such as the reproducing kernel Hilbert spaces induced by radial basis kernels or neural networks with smooth activate functions), then $q \rightarrow 0$ and the difference between Huber regression and least square method could be minimal, indicating less sacrifice necessary for learning smooth functions. Finally, we stress that we obtain exponential convergence rates even for the case when $0 < \epsilon < 1$ where the distribution of the conditional variable $Y|X$ does not possess finite variance and a least square based estimator cannot even be defined. This further explains the robustness of Huber regression estimators. We also remark that the parameter choice suggested by [Theorem 3](#) and [Corollary 4](#) requires the knowledge of the constant in [Assumption 1](#), which is usually not available. In practice, σ may be chosen through cross validation or the balance principle [4,17,19].

6. Proofs of theorems

6.1. Proof of [Theorem 1](#)

Let $f : \mathcal{X} \rightarrow \mathbb{R}$ with $\|f\|_\infty \leq M$ be a measurable function. For any $\sigma > \max\{2M, 1\}$, we denote the two events I_Y and II_Y as follows

$$I_Y := \left\{ y : |y| \geq \frac{\sigma}{2} \right\}, \quad \text{and} \quad II_Y := \left\{ y : |y| < \frac{\sigma}{2} \right\}.$$

Noticing that

$$\int_{\mathcal{X}} \int_{\mathcal{Y}} [(y - f(x))^2 - (y - f^*(x))^2] d\rho(y|x) d\rho_{\mathcal{X}}(x) = \|f - f^*\|_{2,\rho}^2,$$

we have

$$\begin{aligned} & \left| [\mathcal{R}^\sigma(f) - \mathcal{R}^\sigma(f^*)] - \|f - f^*\|_{2,\rho}^2 \right| \\ &= \left| \int_{\mathcal{X}} \int_{\mathcal{Y}} [\ell_\sigma(y - f(x)) - \ell_\sigma(y - f^*(x))] - [(y - f(x))^2 - (y - f^*(x))^2] d\rho(y|x) d\rho_{\mathcal{X}}(x) \right| \\ &\leq \left| \int_{\mathcal{X}} \int_{I_Y \cup II_Y} [\ell_\sigma(y - f(x)) - \ell_\sigma(y - f^*(x))] d\rho(y|x) d\rho_{\mathcal{X}}(x) \right| \\ &\quad + \left| \int_{\mathcal{X}} \int_{I_Y \cup II_Y} [(y - f(x))^2 - (y - f^*(x))^2] d\rho(y|x) d\rho_{\mathcal{X}}(x) \right|. \end{aligned}$$

For any $(x, y) \in \mathcal{X} \times II_Y$, since $\sigma > \max\{2M, 1\}$, we see that

$$|y - f(x)| \leq |y| + \|f\|_\infty < \sigma,$$

and

$$|y - f^*(x)| \leq |y| + \|f^*\|_\infty < \sigma.$$

Consequently, for any $(x, y) \in \mathcal{X} \times II_Y$, we have

$$[\ell_\sigma(y - f(x)) - \ell_\sigma(y - f^*(x))] - [(y - f(x))^2 - (y - f^*(x))^2] = 0,$$

and hence

$$\begin{aligned}
& \left| [\mathcal{R}^\sigma(f) - \mathcal{R}^\sigma(f^*)] - \|f - f^*\|_{2,\rho}^2 \right| \\
& \leq \left| \int_{\mathcal{X}} \int_{I_Y} [\ell_\sigma(y - f(x)) - \ell_\sigma(y - f^*(x))] d\rho(y|x) d\rho_{\mathcal{X}}(x) \right| \\
& \quad + \left| \int_{\mathcal{X}} \int_{I_Y} [(y - f(x))^2 - (y - f^*(x))^2] d\rho(y|x) d\rho_{\mathcal{X}}(x) \right|. \tag{6}
\end{aligned}$$

Recall that the Huber loss (2) is Lipschitz continuous with Lipschitz constant 2σ . The first term of the right-hand side of Eq. (6) can be upper bounded as follows

$$\begin{aligned}
& \left| \int_{\mathcal{X}} \int_{I_Y} [\ell_\sigma(y - f(x)) - \ell_\sigma(y - f^*(x))] d\rho(y|x) d\rho_{\mathcal{X}}(x) \right| \\
& \leq 2\sigma \left| \int_{\mathcal{X}} \int_{I_Y} |f(x) - f^*(x)| d\rho(y|x) d\rho_{\mathcal{X}}(x) \right| \\
& \leq 2\sigma \|f - f^*\|_\infty \Pr(I_Y).
\end{aligned}$$

The quantity $\Pr(I_Y)$ can be bounded by applying Markov's inequality which yields

$$\Pr(I_Y) \leq \frac{2^{1+\epsilon} \mathbb{E}(|Y|^{1+\epsilon})}{\sigma^{1+\epsilon}}. \tag{7}$$

Therefore, we have

$$\int_{\mathcal{X}} \int_{I_Y} [\ell_\sigma(y - f(x)) - \ell_\sigma(y - f^*(x))] d\rho(y|x) d\rho_{\mathcal{X}}(x) \leq \frac{2^{2+\epsilon} \|f - f^*\|_\infty \mathbb{E}(|Y|^{1+\epsilon})}{\sigma^\epsilon}. \tag{8}$$

The second term in the right-hand side of Eq. (6) can be upper bounded as follows:

$$\begin{aligned}
& \left| \int_{\mathcal{X}} \int_{I_Y} [(y - f(x))^2 - (y - f^*(x))^2] d\rho(y|x) d\rho_{\mathcal{X}}(x) \right| \\
& \leq \|f - f^*\|_\infty \int_{\mathcal{X}} \int_{I_Y} |2y - f(x) - f^*(x)| d\rho(y|x) d\rho_{\mathcal{X}}(x) \\
& \leq \|f - f^*\|_\infty \int_{\mathcal{X}} \int_{I_Y} (2|y| + \|f^*\|_\infty + \|f\|_\infty) d\rho(y|x) d\rho_{\mathcal{X}}(x) \\
& \leq \|f - f^*\|_\infty \left(2 \int_{I_Y} |y| d\rho(y) + (\|f^*\|_\infty + \|f\|_\infty) \Pr(I_Y) \right).
\end{aligned}$$

By applying Hölder inequality and recalling the estimate in (7), we have

$$\int_{I_Y} |y| d\rho(y) \leq (\Pr(I_Y))^{\frac{\epsilon}{1+\epsilon}} (\mathbb{E}(|Y|^{1+\epsilon}))^{\frac{1}{1+\epsilon}} \leq \frac{2^\epsilon \mathbb{E}(|Y|^{1+\epsilon})}{\sigma^\epsilon}.$$

As a result, we conclude that

$$\begin{aligned} & \left| \int_{\mathcal{X}} \int_{I_Y} [(y - f(x))^2 - (y - f^*(x))^2] d\rho(y|x) d\rho_{\mathcal{X}}(x) \right| \\ & \leq \frac{2^{1+\epsilon} \|f - f^*\|_{\infty} \mathbb{E}(|Y|^{1+\epsilon})}{\sigma^{\epsilon}} + \frac{2^{1+\epsilon} (\|f\|_{\infty} + \|f^*\|_{\infty})^2 \mathbb{E}(|Y|^{1+\epsilon})}{\sigma^{1+\epsilon}}. \end{aligned} \quad (9)$$

From (8) and (9), $|\mathcal{R}^{\sigma}(f) - \mathcal{R}^{\sigma}(f^*)| - \|f - f^*\|_{2,\rho}^2$ can be upper bounded by

$$\frac{(2^{2+\epsilon} + 2^{1+\epsilon}) \|f - f^*\|_{\infty} \mathbb{E}(|Y|^{1+\epsilon})}{\sigma^{\epsilon}} + \frac{2^{1+\epsilon} (\|f\|_{\infty} + \|f^*\|_{\infty})^2 \mathbb{E}(|Y|^{1+\epsilon})}{\sigma^{1+\epsilon}}.$$

Therefore, the desired estimate (5) holds with $c_{\epsilon} = 2^{4+\epsilon} M \mathbb{E}(|Y|^{1+\epsilon})$. This completes the proof of [Theorem 1](#).

6.2. Proof of [Theorem 2](#)

Let $f : \mathcal{X} \rightarrow \mathbb{R}$ with $\|f\|_{\infty} \leq M$ be a measurable function. For any $\sigma > \max\{2M, 1\}$, we again consider the following two events

$$I_Y := \left\{ y : |y| \geq \frac{\sigma}{2} \right\}, \quad \text{and} \quad II_Y := \left\{ y : |y| < \frac{\sigma}{2} \right\}.$$

Based on the above notation, it is obvious to see the following decomposition

$$\begin{aligned} \mathbb{E} \xi^2 &= \int_{\mathcal{X} \times \mathcal{Y}} (\ell_{\sigma}(y - f(x)) - \ell_{\sigma}(y - f^*(x)))^2 d\rho(x, y) \\ &= \int_{\mathcal{X} \times I_Y} (\ell_{\sigma}(y - f(x)) - \ell_{\sigma}(y - f^*(x)))^2 d\rho(x, y) \\ &\quad + \int_{\mathcal{X} \times II_Y} (\ell_{\sigma}(y - f(x)) - \ell_{\sigma}(y - f^*(x)))^2 d\rho(x, y) \\ &:= Q_1 + Q_2. \end{aligned}$$

The first term Q_1 can be easily bounded by applying the Lipschitz continuity property of the Huber loss (2) and Markov's inequality:

$$Q_1 \leq 4\sigma^2 \int_{\mathcal{X} \times I_Y} (f(x) - f^*(x))^2 d\rho(x, y) \leq 16M^2 \sigma^2 \Pr(I_Y) \leq 2^{5+\epsilon} M^2 \mathbb{E}|Y|^{1+\epsilon} \sigma^{1-\epsilon}.$$

To bound the second term Q_2 , noticing that for any $(x, y) \in \mathcal{X} \times II_Y$, since $\sigma > \max\{2M, 1\}$, we have

$$|y - f(x)| \leq |y| + \|f\|_{\infty} < \sigma,$$

and

$$|y - f^*(x)| \leq |y| + \|f^*\|_{\infty} < \sigma.$$

By the definition of the Huber loss ℓ_{σ} , for any $(x, y) \in \mathcal{X} \times II_Y$, we have

$$\ell_{\sigma}(y - f(x)) - \ell_{\sigma}(y - f^*(x)) = (y - f(x))^2 - (y - f^*(x))^2.$$

Therefore,

$$Q_2 = \int_{\mathcal{X} \times II_Y} ((y - f(x))^2 - (y - f^*(x))^2)^2 d\rho(x, y)$$

$$= \int_{\mathcal{X}} \int_{\Pi_Y} (f(x) - f^*(x))^2 (2y - f(x) - f^*(x))^2 d\rho(y|x) d\rho_{\mathcal{X}}(x).$$

If $\epsilon > 1$, applying Hölder's inequality, we obtain

$$\begin{aligned} Q_2 &\leq \int_{\mathcal{X}} \int_Y (f(x) - f^*(x))^2 (2y - f(x) - f^*(x))^2 d\rho(y|x) d\rho_{\mathcal{X}}(x) \\ &\leq \|f - f^*\|_{\infty}^{\frac{4}{1+\epsilon}} \mathbb{E} \left(|f(x) - f^*(x)|^{\frac{2\epsilon-2}{1+\epsilon}} (2|y| + 2M)^2 \right) \\ &\leq 64(M+1)^2 (\mathbb{E}|Y|^{1+\epsilon} + M^2 + 1) \|f - f^*\|_{2,\rho}^{\frac{2\epsilon-2}{1+\epsilon}}. \end{aligned}$$

If $0 < \epsilon \leq 1$, we have the following estimate

$$\begin{aligned} Q_2 &\leq 48M^2 \int_{\mathcal{X}} \int_{\Pi_Y} (|y|^{1+\epsilon} |y|^{1-\epsilon} + M^2) d\rho(y|x) d\rho_{\mathcal{X}}(x) \\ &\leq 48M^2 (\mathbb{E}|Y|^{1+\epsilon} + M^{1+\epsilon}) \sigma^{1-\epsilon}. \end{aligned}$$

Combining the above estimates for Q_1 and Q_2 , we come to the conclusion that

$$\mathbb{E}\xi^2 \leq c_1 \|f - f^*\|_{2,\rho}^{\frac{2(\epsilon-1)_+}{\epsilon+1}} + c_2 \sigma^{1-\epsilon},$$

with $c_1 = 64(M+1)^2 (\mathbb{E}|Y|^{1+\epsilon} + M^2 + 1)$ and $c_2 = 48M^2 (\mathbb{E}|Y|^{1+\epsilon} + M^{1+\epsilon}) + 16M^2 \mathbb{E}|Y|^{1+\epsilon}$. This completes the proof of [Theorem 2](#).

6.3. Proof of [Theorem 3](#)

We first prove a ratio inequality in [Section 6.3.1](#) which plays an important role in the proof of [Theorem 3](#). The detailed proof of [Theorem 3](#) will then be given in [Section 6.3.2](#). To proceed, for any measurable function $f : \mathcal{X} \rightarrow \mathbb{R}$, we denote

$$\mathcal{R}_{\mathbf{z}}^{\sigma}(f) = \frac{1}{n} \sum_{i=1}^n \ell_{\sigma}(y_i - f(x_i)),$$

and recall the notation

$$f_{\mathcal{H},\sigma} = \arg \min_{f \in \mathcal{H}} \mathcal{R}^{\sigma}(f).$$

6.3.1. A ratio inequality

Proposition 5. *Let $\sigma > \max\{2M, 1\}$. Under [Assumptions 1](#) and [2](#), for any $\gamma > \frac{c\epsilon}{\sigma\epsilon}$, we have*

$$\Pr \left\{ \sup_{f \in \mathcal{H}} \frac{|[\mathcal{R}^{\sigma}(f) - \mathcal{R}^{\sigma}(f^*)] - [\mathcal{R}_{\mathbf{z}}^{\sigma}(f) - \mathcal{R}_{\mathbf{z}}^{\sigma}(f^*)]|}{\sqrt{\mathcal{R}^{\sigma}(f) - \mathcal{R}^{\sigma}(f^*) + 2\gamma}} > 4\sqrt{\gamma} \right\} \leq \mathcal{N} \left(\mathcal{H}, \frac{\gamma}{2\sigma} \right) e^{-\Theta(n,\gamma,\sigma)},$$

where

$$\Theta(n, \gamma, \sigma) = \begin{cases} \frac{n\gamma}{c'_1 \sigma}, & \text{if } 0 < \epsilon \leq 1, \\ \frac{n\gamma}{c'_2 \sigma^{\frac{2\epsilon}{\epsilon+1}}}, & \text{if } \epsilon > 1, \end{cases}$$

with c'_1 and c'_2 being two positive constants independent of n , γ , or σ that will be explicitly specified in the proof.

To prove [Proposition 5](#), we need the following Bernstein concentration inequality that is frequently employed in the literature of learning theory.

Lemma 6. *Let ξ be a random variable on a probability space \mathcal{Z} with variance σ_\star^2 satisfying $|\xi - \mathbb{E}\xi| \leq M_\xi$ almost surely for some constant M_ξ and for all $z \in \mathcal{Z}$. Then for all $\lambda > 0$,*

$$\Pr \left\{ \frac{1}{n} \sum_{i=1}^n \xi(z_i) - \mathbb{E}\xi \geq \lambda \right\} \leq \exp \left\{ -\frac{n\lambda^2}{2(\sigma_\star^2 + \frac{1}{3}M_\xi\lambda)} \right\}.$$

Proof of Proposition 5. Recall that $\mathcal{F}_\mathcal{H}$ denotes the following set of random variables

$$\mathcal{F}_\mathcal{H} = \left\{ \xi \mid \xi(X, Y) = \ell_\sigma(Y - f(X)) - \ell_\sigma(Y - f^\star(X)), \ f \in \mathcal{H} \right\}.$$

For each $\xi \in \mathcal{F}_\mathcal{H}$, by the fact that the Huber loss (2) is Lipschitz continuous with Lipschitz constant 2σ , we have

$$\|\xi\|_\infty \leq 2\sigma \|f - f^\star\|_\infty \leq 4M\sigma \text{ and } \|\xi - \mathbb{E}\xi\|_\infty \leq 4\sigma \|f - f^\star\|_\infty \leq 8M\sigma.$$

According to [Theorem 2](#), we know that

$$\mathbb{E}\xi^2 \leq c_1 \|f - f^\star\|_{2,\rho}^{\frac{2(\epsilon-1)_+}{\epsilon+1}} + c_2 \sigma^{1-\epsilon}.$$

By [Assumption 2](#), we know that there exist a finite positive integer $J = \mathcal{N}(\mathcal{H}, \frac{\gamma}{2\sigma})$ and $\{f_j\}_{j=1}^J \subset \mathcal{H}$ such that $B(f_j, \eta)$, $j = 1, \dots, J$ form a $\frac{\gamma}{2\sigma}$ -cover of \mathcal{H} . We next show that for each $j = 1, \dots, J$, it holds that

$$\Pr \left\{ \frac{|[\mathcal{R}^\sigma(f_j) - \mathcal{R}^\sigma(f^\star)] - [\mathcal{R}_z^\sigma(f_j) - \mathcal{R}_z^\sigma(f^\star)]|}{\sqrt{\mathcal{R}^\sigma(f_j) - \mathcal{R}^\sigma(f^\star) + 2\gamma}} > \sqrt{\gamma} \right\} \leq e^{-\Theta(n, \gamma, \sigma)} \quad (10)$$

for $\gamma > \frac{c_\epsilon}{\sigma^\epsilon}$. To see this, we apply the Bernstein inequality in [Lemma 6](#) to the following random variables

$$\xi_j(X, Y) = \ell_\sigma(Y - f_j(X)) - \ell_\sigma(Y - f^\star(X)), \ j = 1, \dots, J,$$

and obtain

$$\begin{aligned} & \Pr \left\{ |[\mathcal{R}^\sigma(f_j) - \mathcal{R}^\sigma(f^\star)] - [\mathcal{R}_z^\sigma(f_j) - \mathcal{R}_z^\sigma(f^\star)]| > \sqrt{\gamma} \sqrt{\mathcal{R}^\sigma(f_j) - \mathcal{R}^\sigma(f^\star) + 2\gamma} \right\} \\ & \leq \Pr \left\{ |[\mathcal{R}^\sigma(f_j) - \mathcal{R}^\sigma(f^\star)] - [\mathcal{R}_z^\sigma(f_j) - \mathcal{R}_z^\sigma(f^\star)]| > \mu_j \sqrt{\gamma} \right\} \\ & \leq \exp \left\{ -\frac{n\gamma\mu_j^2}{(8M/3 + c_1 + c_2) \left(\sqrt{\gamma}\mu_j\sigma + \sigma^{1-\epsilon} + \|f_j - f^\star\|_{2,\rho}^{\frac{2(\epsilon-1)_+}{\epsilon+1}} \right)} \right\}, \end{aligned} \quad (11)$$

where $\mu_j^2 := \mathcal{R}^\sigma(f_j) - \mathcal{R}^\sigma(f^\star) + 2\gamma$. Since $\gamma > \frac{c_\epsilon}{\sigma^\epsilon}$, by [Theorem 1](#), we have for $j = 1, \dots, J$,

$$\begin{aligned} \mu_j^2 &= \mathcal{R}^\sigma(f_j) - \mathcal{R}^\sigma(f^\star) + 2\gamma > \mathcal{R}^\sigma(f_j) - \mathcal{R}^\sigma(f^\star) + c_\epsilon\sigma^{-\epsilon} + \gamma \\ &\geq \|f_j - f^\star\|_{2,\rho}^2 + \gamma \geq \gamma. \end{aligned} \quad (12)$$

We proceed with the proof by considering the two cases when $0 < \epsilon \leq 1$ and when $\epsilon > 1$. If $0 < \epsilon \leq 1$, by the assumption $\sigma > 1$ and (12), we have

$$\begin{aligned} & \frac{n\gamma\mu_j^2}{(8M/3 + c_1 + c_2) \left(\sqrt{\gamma}\mu_j\sigma + \sigma^{1-\epsilon} + \|f_j - f^*\|_{2,\rho}^{\frac{2(\epsilon-1)_+}{\epsilon+1}} \right)} \\ & > \frac{n\gamma\mu_j^2}{2(8M/3 + c_1 + c_2) \left(\sqrt{\gamma}\mu_j\sigma + \sigma^{1-\epsilon} \right)} > \frac{n\gamma}{c'_1\sigma}, \end{aligned}$$

where $c'_1 = 2(1 + c_\epsilon^{-1})(8M/3 + c_1 + c_2)$. If $\epsilon > 1$, note that (12) implies $\mu_j^2 > \|f_j - f^*\|_{2,\rho}^2$ and $\mu_j^2 > \gamma > c_\epsilon\sigma^{-\epsilon}$. We have

$$\begin{aligned} & \frac{n\gamma\mu_j^2}{(8M/3 + c_1 + c_2) \left(\sqrt{\gamma}\mu_j\sigma + \sigma^{1-\epsilon} + \|f_j - f^*\|_{2,\rho}^{\frac{2(\epsilon-1)_+}{\epsilon+1}} \right)} \\ & = \frac{n\gamma\mu_j^2}{(8M/3 + c_1 + c_2) \left(\sqrt{\gamma}\mu_j\sigma + \sigma^{1-\epsilon} + \|f_j - f^*\|_{2,\rho}^{\frac{2(\epsilon-1)_+}{\epsilon+1}} \right)} \\ & > \frac{n\gamma\mu_j^2}{(8M/3 + c_1 + c_2) \left(\sqrt{\gamma}\mu_j\sigma + \sigma^{1-\epsilon} + \mu_j^{\frac{2(\epsilon-1)}{\epsilon+1}} \right)} > \frac{n\gamma}{c'_2\sigma^{\frac{2\epsilon}{\epsilon+1}}}, \end{aligned}$$

where $c'_2 = (8M/3 + c_1 + c_2) \left(1 + c_\epsilon^{-1} + c_\epsilon^{-\frac{2}{\epsilon+1}} \right)$. Combining the above estimates for two cases and recalling (11), we thus have proved the result in (10).

Denote the two events A and B , respectively, as

$$A = \left\{ \sup_{f \in \mathcal{H}} \frac{|[\mathcal{R}^\sigma(f) - \mathcal{R}^\sigma(f^*)] - [\mathcal{R}_z^\sigma(f) - \mathcal{R}_z^\sigma(f^*)]|}{\sqrt{\mathcal{R}^\sigma(f) - \mathcal{R}^\sigma(f^*) + 2\gamma}} \leq 4\sqrt{\gamma} \right\}$$

and

$$B = \bigcap_{j=1}^J \left\{ \frac{|[\mathcal{R}^\sigma(f_j) - \mathcal{R}^\sigma(f^*)] - [\mathcal{R}_z^\sigma(f_j) - \mathcal{R}_z^\sigma(f^*)]|}{\sqrt{\mathcal{R}^\sigma(f_j) - \mathcal{R}^\sigma(f^*) + 2\gamma}} \leq \sqrt{\gamma} \right\}.$$

We next prove $B \subset A$. To this end, we assume that the event B occurs, that is, for all $j = 1, \dots, J$,

$$|[\mathcal{R}^\sigma(f_j) - \mathcal{R}^\sigma(f^*)] - [\mathcal{R}_z^\sigma(f_j) - \mathcal{R}_z^\sigma(f^*)]| \leq \sqrt{\gamma (\mathcal{R}^\sigma(f_j) - \mathcal{R}^\sigma(f^*) + 2\gamma)}. \quad (13)$$

Recall $B(f_j, \eta)$, $j = 1, \dots, J$, is a $\frac{\gamma}{2\sigma}$ -cover of \mathcal{H} . For every $f \in \mathcal{H}$, there exists $f_j \in \mathcal{H}$ such that $\|f - f_j\|_\infty \leq \frac{\gamma}{2\sigma}$. Since $\gamma > c_\epsilon\sigma^{-\epsilon}$, we have

$$\mathcal{R}^\sigma(f) - \mathcal{R}^\sigma(f^*) + 2\gamma \geq \|f - f^*\|_{2,\rho}^2 - c_\epsilon\sigma^{-\epsilon} + 2\gamma \geq \gamma. \quad (14)$$

Therefore, by the Lipschitz continuity of the Huber loss, for each $j = 1, \dots, J$, we have

$$|\mathcal{R}^\sigma(f) - \mathcal{R}^\sigma(f_j)| \leq 2\sigma\|f - f_j\|_\infty \leq \gamma \leq \sqrt{\gamma (\mathcal{R}^\sigma(f) - \mathcal{R}^\sigma(f^*) + 2\gamma)}, \quad (15)$$

and

$$|\mathcal{R}_z^\sigma(f) - \mathcal{R}_z^\sigma(f_j)| \leq 2\sigma\|f - f_j\|_\infty \leq \gamma \leq \sqrt{\gamma (\mathcal{R}^\sigma(f) - \mathcal{R}^\sigma(f^*) + 2\gamma)}. \quad (16)$$

By (15) and (14), we also have

$$\begin{aligned}
\mathcal{R}^\sigma(f_j) - \mathcal{R}^\sigma(f^*) + 2\gamma &= \left(\mathcal{R}^\sigma(f_j) - \mathcal{R}^\sigma(f) \right) + \mathcal{R}^\sigma(f) - \mathcal{R}^\sigma(f^*) + 2\gamma \\
&\leq \sqrt{\gamma(\mathcal{R}^\sigma(f) - \mathcal{R}^\sigma(f^*) + 2\gamma)} + \mathcal{R}^\sigma(f) - \mathcal{R}^\sigma(f^*) + 2\gamma \\
&\leq 2(\mathcal{R}^\sigma(f) - \mathcal{R}^\sigma(f^*) + 2\gamma).
\end{aligned} \tag{17}$$

Combining the estimates (15), (16), (17) with the assumption (13), we obtain

$$\begin{aligned}
&|[\mathcal{R}^\sigma(f) - \mathcal{R}^\sigma(f^*)] - [\mathcal{R}_z^\sigma(f) - \mathcal{R}_z^\sigma(f^*)]| \\
&\leq |\mathcal{R}^\sigma(f) - \mathcal{R}^\sigma(f_j)| + |\mathcal{R}^\sigma(f_j) - \mathcal{R}^\sigma(f^*) - [\mathcal{R}_z^\sigma(f_j) - \mathcal{R}_z^\sigma(f^*)]| + |\mathcal{R}_z^\sigma(f_j) - \mathcal{R}_z^\sigma(f)| \\
&\leq 2\sqrt{\gamma(\mathcal{R}^\sigma(f) - \mathcal{R}^\sigma(f^*) + 2\gamma)} + \sqrt{\gamma(\mathcal{R}^\sigma(f_j) - \mathcal{R}^\sigma(f^*) + 2\gamma)} \\
&\leq 4\sqrt{\gamma(\mathcal{R}^\sigma(f) - \mathcal{R}^\sigma(f^*) + 2\gamma)}.
\end{aligned} \tag{18}$$

Since (18) holds for every $f \in \mathcal{H}$, we have proved $B \subset A$ or equivalently $A^c \subset B^c$. This together with (10) leads to

$$\begin{aligned}
&\Pr \left\{ \sup_{f \in \mathcal{H}} \frac{|[\mathcal{R}^\sigma(f) - \mathcal{R}^\sigma(f^*)] - [\mathcal{R}_z^\sigma(f) - \mathcal{R}_z^\sigma(f^*)]|}{\sqrt{\mathcal{R}^\sigma(f) - \mathcal{R}^\sigma(f^*) + 2\gamma}} > 4\sqrt{\gamma} \right\} \\
&= \Pr(A^c) \leq \Pr(B^c) \\
&\leq \sum_{j=1}^J \Pr \left\{ \frac{|[\mathcal{R}^\sigma(f_j) - \mathcal{R}^\sigma(f^*)] - [\mathcal{R}_z^\sigma(f_j) - \mathcal{R}_z^\sigma(f^*)]|}{\sqrt{\mathcal{R}^\sigma(f_j) - \mathcal{R}^\sigma(f^*) + 2\gamma}} > \sqrt{\gamma} \right\} \\
&\leq \mathcal{N}(\mathcal{H}, \gamma\sigma^{-1}/2) e^{-\Theta(n, \gamma, \sigma)}.
\end{aligned}$$

This completes the proof of [Proposition 5](#). \square

6.3.2. Proof of [Theorem 3](#)

We first prove that for any $0 < \delta < 1$, with probability at least $1 - \delta/2$, we have

$$[\mathcal{R}^\sigma(f_{z,\sigma}) - \mathcal{R}^\sigma(f^*)] - [\mathcal{R}_z^\sigma(f_{z,\sigma}) - \mathcal{R}_z^\sigma(f^*)] - \frac{1}{2}[\mathcal{R}^\sigma(f_{z,\sigma}) - \mathcal{R}^\sigma(f^*)] \leq 9\gamma_0,$$

where γ_0 is given by

$$\gamma_0 := \begin{cases} \mathcal{O}\left(\frac{1}{\sigma^\epsilon} + \log\left(\frac{2}{\delta}\right)\frac{\sigma}{n^{1/(q+1)}}\right), & \text{if } 0 < \epsilon \leq 1, \\ \mathcal{O}\left(\frac{1}{\sigma^\epsilon} + \log\left(\frac{2}{\delta}\right)\left(\frac{\sigma^{q+2\epsilon}}{n^{1+\epsilon}}\right)^{1/(q+1)}\right), & \text{if } \epsilon > 1. \end{cases} \tag{19}$$

Note that [Proposition 5](#) implies that, for any $\gamma > c_\epsilon\sigma^{-\epsilon}$,

$$\sup_{f \in \mathcal{H}} \frac{|[\mathcal{R}^\sigma(f) - \mathcal{R}^\sigma(f^*)] - [\mathcal{R}_z^\sigma(f) - \mathcal{R}_z^\sigma(f^*)]|}{\sqrt{\mathcal{R}^\sigma(f) - \mathcal{R}^\sigma(f^*) + 2\gamma}} < 4\sqrt{\gamma} \tag{20}$$

holds with probability at least

$$1 - \mathcal{N}\left(\mathcal{H}, \frac{\gamma}{2\sigma}\right) e^{-\Theta(n, \gamma, \sigma)}.$$

We know from [Assumption 2](#) that

$$\mathcal{N}\left(\mathcal{H}, \frac{\gamma}{2\sigma}\right) \lesssim \exp\left\{2^q c\sigma^q / \gamma^q\right\}.$$

Consequently, the event (20) holds with probability at least $1 - \exp\{2^q c\sigma^q \gamma^{-q} - \Theta(n, \gamma, \sigma)\}$. For any $0 < \delta < 1$, let

$$\exp\{2^q c\sigma^q \gamma^{-q} - \Theta(n, \gamma, \sigma)\} = \delta/2,$$

or equivalently

$$2^q c\sigma^q \gamma^{-q} - \Theta(n, \gamma, \sigma) = \log(\delta/2).$$

The equation has a unique positive solution γ^* satisfying

$$\gamma^* \lesssim \begin{cases} \log\left(\frac{2}{\delta}\right) \left(\frac{\sigma}{n^{1/(q+1)}}\right), & \text{if } \epsilon \leq 1, \\ \log\left(\frac{2}{\delta}\right) \left(\frac{\sigma^{q+\frac{2\epsilon}{1+\epsilon}}}{n}\right)^{1/(q+1)}, & \text{if } \epsilon > 1. \end{cases}$$

Choose $\gamma_0 = c_\epsilon \sigma^{-\epsilon} + \gamma^*$. Then γ_0 satisfies the condition (19) and for any $0 < \delta < 1$, with probability at least $1 - \delta/2$, it holds that

$$\sup_{f \in \mathcal{H}} \frac{|[\mathcal{R}^\sigma(f) - \mathcal{R}^\sigma(f^*)] - [\mathcal{R}_z^\sigma(f) - \mathcal{R}_z^\sigma(f^*)]|}{\sqrt{\mathcal{R}^\sigma(f) - \mathcal{R}^\sigma(f^*) + 2\gamma_0}} \leq 4\sqrt{\gamma_0},$$

which immediately yields

$$\begin{aligned} [\mathcal{R}^\sigma(f_{z,\sigma}) - \mathcal{R}^\sigma(f^*)] - [\mathcal{R}_z^\sigma(f_{z,\sigma}) - \mathcal{R}_z^\sigma(f^*)] &\leq 4\sqrt{\gamma_0} \sqrt{\mathcal{R}^\sigma(f_{z,\sigma}) - \mathcal{R}^\sigma(f^*) + 2\gamma_0} \\ &\leq \frac{1}{2} (\mathcal{R}^\sigma(f_{z,\sigma}) - \mathcal{R}^\sigma(f^*)) + 9\gamma_0. \end{aligned} \quad (21)$$

By a similar procedure we can prove that for any $0 < \delta < 1$, with probability at least $1 - \delta/2$, it holds that

$$[\mathcal{R}_z^\sigma(f_{\mathcal{H},\sigma}) - \mathcal{R}_z^\sigma(f^*)] - [\mathcal{R}^\sigma(f_{\mathcal{H},\sigma}) - \mathcal{R}^\sigma(f^*)] \leq \frac{1}{2} (\mathcal{R}^\sigma(f_{\mathcal{H},\sigma}) - \mathcal{R}^\sigma(f^*)) + 9\gamma_0.$$

This in connection with the fact that $\mathcal{R}^\sigma(f_{\mathcal{H},\sigma}) \leq \mathcal{R}^\sigma(f_{\mathcal{H},\sigma})$ and Theorem 1 implies that for any $0 < \delta < 1$, with probability at least $1 - \delta/2$, we have

$$[\mathcal{R}_z^\sigma(f_{\mathcal{H},\sigma}) - \mathcal{R}_z^\sigma(f^*)] - [\mathcal{R}^\sigma(f_{\mathcal{H},\sigma}) - \mathcal{R}^\sigma(f^*)] \leq \frac{1}{2} \|f_{\mathcal{H}} - f^*\|_{2,\rho}^2 + 10\gamma_0. \quad (22)$$

Combining the two estimates in (21) and (22), we come to the conclusion that for any $0 < \delta < 1$, with probability at least $1 - \delta$, it holds that

$$\begin{aligned} &[\mathcal{R}^\sigma(f_{z,\sigma}) - \mathcal{R}^\sigma(f_{\mathcal{H},\sigma})] - [\mathcal{R}_z^\sigma(f_{z,\sigma}) - \mathcal{R}_z^\sigma(f_{\mathcal{H},\sigma})] \\ &\leq \frac{1}{2} [\mathcal{R}^\sigma(f_{z,\sigma}) - \mathcal{R}^\sigma(f^*)] + \frac{1}{2} \|f_{\mathcal{H}} - f^*\|_{2,\rho}^2 + 19\gamma_0. \end{aligned} \quad (23)$$

On the other hand, from the definitions of $f_{\mathcal{H},\sigma}$, $f_{\mathcal{H}}$, and $f_{z,\sigma}$, we have

$$\begin{aligned} &\mathcal{R}^\sigma(f_{z,\sigma}) - \mathcal{R}^\sigma(f^*) \\ &= [\mathcal{R}^\sigma(f_{z,\sigma}) - \mathcal{R}^\sigma(f_{\mathcal{H},\sigma})] + [\mathcal{R}^\sigma(f_{\mathcal{H},\sigma}) - \mathcal{R}^\sigma(f^*)] \\ &\leq [\mathcal{R}^\sigma(f_{z,\sigma}) - \mathcal{R}^\sigma(f_{\mathcal{H},\sigma})] - [\mathcal{R}_z^\sigma(f_{z,\sigma}) - \mathcal{R}_z^\sigma(f_{\mathcal{H},\sigma})] + \mathcal{R}^\sigma(f_{\mathcal{H}}) - \mathcal{R}^\sigma(f^*) \\ &\leq [\mathcal{R}^\sigma(f_{z,\sigma}) - \mathcal{R}^\sigma(f_{\mathcal{H},\sigma})] - [\mathcal{R}_z^\sigma(f_{z,\sigma}) - \mathcal{R}_z^\sigma(f_{\mathcal{H},\sigma})] + \|f_{\mathcal{H}} - f^*\|_{2,\rho}^2 + c_\epsilon \sigma^{-\epsilon}, \end{aligned}$$

where the first inequality is due to the following two facts

$$\mathcal{R}_z^\sigma(f_{z,\sigma}) \leq \mathcal{R}_z^\sigma(f_{\mathcal{H},\sigma}) \quad \text{and} \quad \mathcal{R}^\sigma(f_{\mathcal{H},\sigma}) \leq \mathcal{R}^\sigma(f_{\mathcal{H}}).$$

By (23), we know that for any $0 < \delta < 1$, with probability at least $1 - \delta$, it holds that

$$\mathcal{R}^\sigma(f_{\mathbf{z},\sigma}) - \mathcal{R}^\sigma(f^*) \leq \frac{1}{2} \left(\mathcal{R}^\sigma(f_{\mathbf{z},\sigma}) - \mathcal{R}^\sigma(f^*) \right) + \frac{3}{2} \|f_{\mathcal{H}} - f^*\|_{2,\rho}^2 + 20\gamma_0,$$

which implies that

$$\mathcal{R}^\sigma(f_{\mathbf{z},\sigma}) - \mathcal{R}^\sigma(f^*) \leq 3 \|f_{\mathcal{H}} - f^*\|_{2,\rho}^2 + 40\gamma_0$$

holds with probability at least $1 - \delta$. By [Theorem 1](#) again, we conclude that for any $0 < \delta < 1$, with probability at least $1 - \delta$, it holds that

$$\|f_{\mathbf{z},\sigma} - f^*\|_{2,\rho}^2 \lesssim \|f_{\mathcal{H}} - f^*\|_{2,\rho}^2 + \gamma_0.$$

Recalling the definition of Ψ and noticing that $\gamma_0 \lesssim \log(2/\delta)\Psi$, we complete the proof of [Theorem 3](#).

7. Concluding remarks

In this paper, we studied the Huber regression problem by investigating the empirical risk minimization scheme induced by the Huber loss. In a statistical learning setup, our study answered the four fundamental questions raised in the introduction: the \mathcal{R}^σ -risk consistency is insufficient in ensuring their convergence to the mean regression function; the scale parameter σ plays a trade-off role in bias and learnability; fast exponential-type convergence rates can be established under $(1 + \epsilon)$ -moment conditions ($\epsilon > 0$) by relaxing the standard Bernstein condition and allowing some additional small bias term; the merit of Huber regression in terms of the robustness can be reflected by its learnability under the $(1 + \epsilon)$ -moment conditions which are considered to be weak conditions in that heavy-tailed noise can be accommodated in regression problems. Moreover, it was shown that with higher moment conditions being imposed, one can obtain faster convergence rates. In the above senses, we conducted a complete and systematic statistical learning assessment of Huber regression estimators.

We remark that in the present study a general hypothesis space \mathcal{H} is considered. In practice, the implementation of learning with Huber regression requires to specify a particular hypothesis space. It can be a reproducing kernel Hilbert space, a neural network, or other families of functions. Functions in such a hypothesis space are generally not uniformly bounded. Regularization could be used to restrict the searching region of the Huber regression scheme and consequently controls the capacity of the working hypothesis space. The techniques developed in this study may still be applicable to assessing the regularized Huber regression schemes. Additionally, the development of these techniques for assessing Huber regression estimators may also shed light on the analysis of other robust regression schemes.

Acknowledgments

The work by Yunlong Feng is partially supported by the Simons Foundation #572064 and the National Science Foundation of the USA DMS-2111080. The work by Qiang Wu is partially supported by the Simons Foundation #712916, the National Science Foundation of the USA DMS-2110826, and the MTSU Non-Instructional Assignment Grant. The two authors made equal contributions to this paper and are listed alphabetically.

References

- [1] M. Anthony, P.L. Bartlett, *Neural Network Learning: Theoretical Foundations*, Cambridge University Press, 2009.
- [2] P.L. Bartlett, S. Mendelson, Empirical minimization, *Probab. Theory Related Fields* 135 (3) (2006) 311–334.
- [3] F. Bauer, S. Pereverzev, L. Rosasco, On regularization algorithms in learning theory, *J. Complexity* 23 (1) (2007) 52–72.
- [4] A. Caponnetto, Y. Yao, Cross-validation based adaptation for regularization operators in learning theory, *Anal. Appl.* 8 (02) (2010) 161–183.
- [5] G. Chinot, G. Lecué, M. Lerasle, Robust statistical learning with Lipschitz and convex loss functions, *Probab. Theory Related Fields* 176 (2019) 897–940.
- [6] A. Christmann, I. Steinwart, Consistency and robustness of kernel-based regression in convex risk minimization, *Bernoulli* 13 (3) (2007) 799–819.
- [7] F. Cucker, D.X. Zhou, *Learning Theory: An Approximation Theory Viewpoint*, Cambridge University Press, 2007.
- [8] L. Devroye, L. Györfi, G. Lugosi, *A Probabilistic Theory of Pattern Recognition*, Springer Science & Business Media, 2013.
- [9] J. Fan, Y. Guo, B. Jiang, Adaptive Huber regression on Markov-dependent data, *Stochastic Process. Appl.* (2019) in press.
- [10] J. Friedman, T. Hastie, R. Tibshirani, *The Elements of Statistical Learning*, Springer Series in Statistics, New York, 2001.
- [11] F.R. Hampel, E.M. Ronchetti, P.J. Rousseeuw, W.A. Stahel, *Robust Statistics: The Approach Based on Influence Functions*, John Wiley & Sons, 2011.
- [12] X. He, Q.-M. Shao, A general Bahadur representation of M-estimators and its application to linear regression with nonstochastic designs, *Ann. Statist.* 24 (6) (1996) 2608–2630.
- [13] P.J. Huber, Robust estimation of a location parameter, *Ann. Math. Stat.* 35 (1) (1964) 73–101.
- [14] P.J. Huber, Robust regression: asymptotics, conjectures and Monte Carlo, *Ann. Statist.* 1 (5) (1973) 799–821.
- [15] P.J. Huber, E. Ronchetti, *Robust Statistics*, Wiley, 2009.
- [16] P.-L. Loh, Statistical consistency and asymptotic normality for high-dimensional robust M-estimators, *Ann. Statist.* 45 (2) (2017) 866–896.
- [17] S. Lu, P. Mathé, S.V. Pereverzev, Balancing principle in supervised learning for a general regularization scheme, *Appl. Comput. Harmon. Anal.* 48 (1) (2020) 123–148.
- [18] R. Maronna, D. Martin, V. Yohai, *Robust Statistics: Theory and Methods*, John Wiley & Sons, Chichester, 2006.
- [19] G. Myleiko, S. Pereverzev Jr., S. Solodky, Regularized Nyström subsampling in regression and ranking problems under general smoothness assumptions, *Anal. Appl.* 17 (03) (2019) 453–475.
- [20] S. Portnoy, Asymptotic behavior of M-estimators of p regression parameters when p^2/n is large. I. Consistency, *Ann. Statist.* 12 (4) (1984) 1298–1309.
- [21] L. Rosasco, E.D. Vito, A. Caponnetto, M. Piana, A. Verri, Are loss functions all the same? *Neural Comput.* 16 (5) (2004) 1063–1076.
- [22] P.J. Rousseeuw, A.M. Leroy, *Robust Regression and Outlier Detection*, John Wiley & Sons, 2005.
- [23] I. Steinwart, A. Christmann, *Support Vector Machines*, Springer, New York, 2008.
- [24] I. Steinwart, C. Scovel, Fast rates for support vector machines using Gaussian kernels, *Ann. Statist.* 35 (2) (2007) 575–607.
- [25] Q. Sun, W.-X. Zhou, J. Fan, Adaptive Huber regression, *J. Amer. Statist. Assoc.* 115 (529) (2019) 254–265.
- [26] T. Van Erven, P.D. Grünwald, N.A. Mehta, M.D. Reid, R.C. Williamson, Fast rates in statistical and online learning, *J. Mach. Learn. Res.* 16 (1) (2015) 1793–1861.
- [27] V. Vapnik, *Statistical Learning Theory*, Wiley-Interscience, 1998.
- [28] V.J. Yohai, R.A. Maronna, Asymptotic behavior of M-estimators for the linear model, *Ann. Statist.* 7 (2) (1979) 258–268.