

Attacks on Robust Distributed Learning Schemes via Sensitivity Curve Maximization

Christian A. Schroth*, Stefan Vlaski[†] and Abdelhak M. Zoubir*

*Signal Processing Group, Technische Universität Darmstadt, Germany

[†]Department of Electrical and Electronic Engineering, Imperial College London, UK

*{schroth, zoubir}@spg.tu-darmstadt.de, [†]s.vlaski@imperial.ac.uk

Abstract—Distributed learning paradigms, such as federated or decentralized learning, allow a collection of agents to solve global learning and optimization problems through limited local interactions. Most such strategies rely on a mixture of local adaptation and aggregation steps, either among peers or at a central fusion center. Classically, aggregation in distributed learning is based on averaging, which is statistically efficient, but susceptible to attacks by even a small number of malicious agents. This observation has motivated a number of recent works, which develop robust aggregation schemes by employing robust variations of the mean. We present a new attack based on sensitivity curve maximization (SCM), and demonstrate that it is able to disrupt existing robust aggregation schemes by injecting small, but effective perturbations.

Keywords—Decentralized learning, federated learning, robust aggregation, byzantine robustness, sensitivity curve


I. INTRODUCTION

Distributed learning paradigms, such as federated or decentralized learning are growing in importance as data is increasingly available in dispersed locations, while central aggregation of data is infeasible due to concerns around privacy or communication efficiency¹. In cooperative settings, well-designed distributed algorithms can match the performance of a centralized benchmark, which has access to all data at a single location [1]–[4]. At the same time, however, their reliance on linear averaging for model aggregation renders most of these strategies susceptible to attacks by even a small number of deviating agents. This observation has motivated the development of robust aggregation schemes for distributed learning, primarily focusing on federated network structures. Popular methods are based on robust variations of the mean, such as the trimmed mean and median [5] or weighted geometric median [6], RFA [7], Krum [8] and Bulyan [9]. A recent survey can be found in [10].

Compared to federated learning, the literature of robust decentralized learning is more sparse. ByRDIE [11] and BRIDGE [12] use a coordinate-wise trimmed mean, where the latter has variants using a coordinate-wise median, Krum or a combination thereof. The Iterative Outlier Scissor (IOS) [13], iteratively discards a number of weights, which are furthest away from the weighted average. The number of discarded weights is equal to the number of estimated Byzantine outliers. In [14], a robust and efficient method for element-wise robust aggregation based on the Biweight Tukey M-estimator was presented. Self Centered Clipping (SCC) [15] is a variant of the

¹The work of Christian A. Schroth and Abdelhak M. Zoubir has been funded by the LOEWE initiative (Hesse, Germany) within the emergenCITY centre and by DFG Project under Grant 431431951 / ZO 215/19-1.

¹We use “distributed” for any structure where data remains local at individual agents, which includes federated and decentralized architectures. The term “decentralized” is used for a network without fusion center. In the literature, “distributed” and “decentralized” are sometimes used interchangeably.



(a) SCs of sample mean, median and Biweight Tukey (b) SCs of sample mean, Talwar and α -trimmed mean

Fig. 1: Overview of sensitivity curves (SCs) for different aggregation schemes. Tukey and the trimmed means are tuned, such that they achieve 95% efficiency.

trimmed-mean where the distance is measured from the local nodes’ own weight and weights beyond a specified threshold are clipped. All of these aggregation schemes have in common that they rely in one way or another on a distance measure to detect and reject outliers.

These schemes are quite effective against simple and/or non-intelligent attacks, but a byzantine attacker which is aware of the used aggregation scheme, will be able to exploit this information and will intelligently attack the aggregation scheme. As it was shown in “A little is enough” (ALIE) [16], most robust federated learning methods can be disturbed by injecting a carefully selected value, which is small enough to be not detected by the defense mechanism, but large enough to disturb the estimation. [17] proposed a method, where the malicious nodes try to generate weights, which will lead to the largest deviation in the inverse direction. In [18], it is shown that the inner product between the true gradient and the robust estimator has to be non-negative, otherwise a malicious agent could invert the direction of the gradient. More attack mechanisms are presented in [19], [20].

A novel defense against malicious attacks is MixTailor [21], which does not rely on one static aggregation scheme, but randomly draws an aggregation scheme from a set of predefined schemes. This prevents even omniscient attackers from crafting a perfect outlier, as they only could inject an optimal outlier in expectation.

Motivated by the observations of ALIE, we develop an attack scheme on robust aggregators by maximizing their sensitivity curve (SC). The sensitivity curve describes the influence of an outlier on the estimated value. The two extremes can be observed in Figure 1a, where for the sample mean the influence of an outlier increases linearly with its value, while for the median the value of an outlier has almost no influence. In what follows, we will develop an attack

scheme which maximizes the influence on the SC. Our contributions are:

- We propose a novel attack scheme for a broad range of robust aggregation schemes, by tailoring perturbations to the sensitivity curve of the aggregator.
- We demonstrate numerically that the proposed scheme is able to disrupt most popular aggregation schemes in the literature, even those based on robust variations of the mean, such as the median.

The remainder of the paper is organized as follows. Decentralized learning, different robust aggregation schemes and the sensitivity curve are introduced in Section II. Our proposed attack scheme is given in Section III. Section IV gives an overview of the experimental validations and finally, a conclusion and outlook are given in Section V.

II. PRELIMINARIES

A. Decentralized Learning

We will consider a collection of K agents, and associate with each agent a local objective function

$$J_k(w) = \mathbb{E}Q(w; \mathbf{x}_k) \quad (1)$$

where w denotes the parameters or weights of a model maintained by each agent, and \mathbf{x}_k is a random variable representing the data at agent k . Most federated and decentralized learning algorithms aim to pursue an optimal model according to the consensus optimization problem

$$J(w) = \frac{1}{K} \sum_{k=1}^K J_k(w). \quad (2)$$

In decentralized learning, agents only exchange intermediate estimates on a peer-to-peer basis, without communicating with a central fusion center. For example, the ATC-diffusion algorithm takes the form [1], [2]

$$\phi_{k,i} = \mathbf{w}_{k,i-1} - \mu \widehat{\nabla} J_k(\mathbf{w}_{k,i-1}) \quad (3)$$

$$\mathbf{w}_{k,i} = \sum_{\ell \in \mathcal{N}_k} a_{\ell k} \phi_{\ell,i} \quad (4)$$

where $\mu > 0$ denotes the step-size and \mathcal{N}_k denotes the neighborhood of node k , including node k . The local adaptation (3) is driven by a stochastic gradient approximation $\widehat{\nabla} J_k(\mathbf{w}_{k,i-1})$. It is common to choose as $\widehat{\nabla} J_k(\mathbf{w}_{k,i-1}) = \nabla Q(\mathbf{w}_{k,i-1}; \mathbf{x}_{k,i})$, where $\mathbf{x}_{k,i}$ denotes the sample available to agent k at time i . Commonly, the aggregation step in (4) is a non-robust sample mean. The choice of the sample mean is well justified in benign scenarios without malicious agents, as it usually exhibits the highest statistical efficiency, and hence the fastest rate of convergence. In scenarios with malicious agents, this is susceptible to significant degeneration of the performance. By replacing the non-robust average in Equation (4) with a general aggregation rule $\text{AGG}(\cdot)$ as in

$$\mathbf{w}_{k,i} = \text{AGG}(\phi_{\ell,i}), \quad \ell \in \mathcal{N}_k, \quad (5)$$

it is possible to introduce a variety of different aggregation schemes, including schemes which are robust against outliers.

B. Aggregation Schemes

For simplicity we will limit the investigation in this paper to element-wise aggregation schemes. For a set of samples

$$\mathcal{Y} = \{y_1, y_2, \dots, y_N\}:$$

1) *Sample mean*:

$$\hat{\mu} = \frac{1}{|\mathcal{Y}|} \sum_{y_n \in \mathcal{Y}} y_n \quad (6)$$

Achieves a breakdown point of 0, meaning that a single malicious sample can result in arbitrary deterioration of performance [22]. This choice of aggregator recovers (4).

2) *Median*: Select the median of the samples in \mathcal{Y} . Achieves a breakdown point of 0.5, meaning that 50% of samples can be corrupted before the median breaks down. This choice of aggregator recovers [5].

3) *α -trimmed mean*: The αN largest and smallest values are discarded, resulting in a total $2\alpha N$ values being removed. Achieves a breakdown point of α . This choice of aggregator is employed in [5], [11], [13].

4) *Talwar [23]*: Also called Huber type-skipped mean [24], falls into the class of M-estimators which solve for a location estimate

$$\sum_{y_n \in \mathcal{Y}} \psi \left(\frac{y_n - \hat{\mu}}{\hat{\sigma}} \right) = 0 \quad (7)$$

iteratively with a fixed point algorithm for $\hat{\mu}$ starting with a robust initial location and scale estimate and with

$$\psi(x) = \begin{cases} x & , |x| \leq c \\ 0 & , |x| > c. \end{cases} \quad (8)$$

As initial estimates commonly the median absolute deviation (mad) and median are chosen.

5) *Biweight Tukey*: The well known M-estimator [22, p. 11] with

$$\psi(x) = \begin{cases} x \left(1 - \frac{x^2}{c^2}\right)^2 & , |x| \leq c \\ 0 & , |x| > c. \end{cases} \quad (9)$$

achieves a breakdown point of close to 0.5, while being more efficient than the median. This type of aggregator is considered in [14].

C. Sensitivity Curve

Each aggregator can be associated with a sensitivity curve (SC), which measures the influence a single maliciously designed sample can exhibit on the estimator. We will use these sensitivity curves to determine worst case expressions for malicious perturbations. Formally, the SC or empirical influence function [25] for an estimator $\text{AGG}(\cdot)$ and samples $\mathcal{Y} = \{y_1, y_2, \dots, y_{N-1}\}$ of length $N - 1$ is defined as

$$\text{SC}(\mathcal{Y}, z) = N(\text{AGG}(\mathcal{Y} \cup z) - \text{AGG}(\mathcal{Y})). \quad (10)$$

It describes the bias of the estimator when an additional observation z is added to a sample of size $N - 1$. Examples of SCs of the considered aggregation schemes are shown in Figure 1.

For our further investigations, it is useful to extend the definition of the SC to account for P identical outliers $\mathcal{Z} = \{z \cdot \mathbb{1}_P\}$ where $\mathbb{1}_P$ is a repeater function and with $\mathcal{Y} = \{y_1, y_2, \dots, y_{N-P}\}$ as

$$\text{SC}(\mathcal{Y}, \mathcal{Z}) = N(\text{AGG}(\mathcal{Y} \cup \mathcal{Z}) - \text{AGG}(\mathcal{Y})). \quad (11)$$

We will use this definition to quantify the effect that P malicious agents can have when coordinating to induce maximum bias in the aggregation. We note that the SC is not necessarily fixed, but is a function on the underlying random samples, e.g. for the median,

the SC depends on the distance between the random samples and therefore can take larger or smaller values. But the fundamental shape is determined by the aggregator and is not influenced by the malicious samples.

III. PROPOSED SENSITIVITY CURVE MAXIMIZATION ATTACK

In [16] the ALIE attack model for federated learning is proposed, which illustrates the fact that most robust aggregation rules detect and reject outliers by a distance measure. Assuming that the samples, sent to node k from its neighborhood, are normally distributed, ALIE estimates the mean and variance of this distribution. Equipped with those estimates the cdf is used to estimate a value, which will lead to the rejection of benign nodes. The number of rejected benign nodes is chosen, such that the malicious nodes will gain the majority in the neighborhood.

Following the idea of injecting small values, which will ultimately lead to an estimation bias and hence, a breakdown, we propose an attack model which is based on sensitivity curve maximization (SCM). As introduced in the previous section, the SC describes the influence of an outlier on the estimate. Finding the minimal value, which maximizes the SC and injecting this value into the aggregation, will lead to the greatest possible distortion of the estimate. To estimate this optimal value z_m^{opt} an omniscient byzantine attacker is assumed, but it should be possible to only use the estimates of the malicious nodes, to gain an estimated optimal attack value.

The set of all benign values for agent k in dimension m is defined as $\mathcal{Y}_m = \{\phi_{l,i}(m)\}_{l \in \mathcal{N}_k^b}$ and the set of malicious values is defined as $\mathcal{Z}_m = \{z_m \cdot \mathbb{1}_{|\mathcal{N}_k^m|}\}$, where \mathcal{N}_k^b and \mathcal{N}_k^m denote the benign neighborhood and the malicious neighborhood, respectively. Finding the smallest value z_m^{opt} , which maximizes the SC, requires solving

$$z_m^{\text{opt}} = \min_{z_m \in \mathcal{C}_m} |z_m| \quad (12)$$

$$\mathcal{C}_m \triangleq \arg \max_{z_m} (\text{SC}(\mathcal{Y}_m, \mathcal{Z}_m))$$

for every dimension m . In the sequel, dedicated attack schemes for the α -trimmed mean, Talwar and Biweight Tukey are proposed.

α -trimmed mean: As αN_k of largest and smallest weights will be removed, the $(N_k - \alpha N_k - 1)$ -th weight of the ascending sorted weights is chosen. Then the $(N_k - \alpha N_k - 1)$ -th to the $(N_k - \alpha N_k - P - 1)$ -th weights are replaced with the $(N_k - \alpha N_k - 1)$ -th weight. Instead of using the $(N_k - \alpha N_k - 1)$ -th weight, it is also possible to use the $(N_k - \alpha N_k)$ -th weight minus some small value. This will have the effect that all malicious weights will be kept and will have a maximal influence on the aggregated weight.

Talwar and Biweight Tukey: Both attack schemes will lead to a similar solution, which only differs in a constant c_0 . By solving

$$c_0 = \arg \max_x \psi(x) \quad (13)$$

with $\psi(x)$ from Equations (8) and (9), we obtain for Talwar $c_0 = c$ and for Biweight Tukey $c_0 = \frac{c}{\sqrt{5}}$. As M-estimators subtract a robust location estimate and normalize with a robust scale estimate, we have to account for these shifts by calculating the initial value as

$$z_m^{\text{init}} = c_0 \cdot \text{mad}(\mathcal{Y}_m) + \text{median}(\mathcal{Y}_m) \quad (14)$$

where $\text{mad}(x)$ denotes the median absolute deviation. Injecting these values into the aggregation step will influence the robust location and




Fig. 2: SCs of the presented attack schemes. The black triangles denote the optimal outlier, which maximizes the respective SC.




Fig. 3: Training loss evolution for all presented aggregation schemes without malicious agents.

scale estimates. Hence, we have to account for this shift by calculating the optimal outlier with $\mathcal{Z}_m^{\text{init}} = \{z_m^{\text{init}} \cdot \mathbb{1}_{|\mathcal{N}_k^m|}\}$, resulting in

$$z_m^{\text{opt}} = c_0 \cdot \text{mad}(\mathcal{Y} \cup \mathcal{Z}_m^{\text{init}}) + \text{median}(\mathcal{Y}_m \cup \mathcal{Z}_m^{\text{init}}). \quad (15)$$

One might expect this to lead to an iterative calculation, but this is not the case, as the median is primarily influenced by the number of values added to the set and not by the actual values. The resulting values z_m^{opt} are shown in Figure 2 as black triangles. Clearly, it would be also possible to use the values mirrored through the origin. Finally, all malicious nodes in the neighborhood of node k will send these maximal influential values and, hence, will cause the largest bias possible in the aggregation. Applying this method over multiple rounds of learning, will lead to a small but steady distortion of the estimates, as we will show in the next section.

IV. SIMULATIONS

The simulations are performed on a network with $K = 32$ agents arranged in an Erdős–Rényi graph with an edge probability of 70%. Each agent observes data following a linear model of the form

$$\mathbf{d}_k = \mathbf{u}_k^\top \mathbf{w}^o + \mathbf{v}_k \quad (16)$$

with the regressors $\mathbf{u}_k \in \mathbb{R}^{10}$ being independently identical distributed as $\mathbf{u}_k \sim \mathcal{N}(0, \mathbf{I}_{10})$. The noise is distributed as $\mathbf{v}_k \sim \mathcal{N}(0, \sigma_v^2)$ with $\sigma_v^2 = 0.01$. In the simulations the following two assumptions are made, which are commonly found in the literature:

- The graph is generated such that for every benign node k , the majority of the neighborhood \mathcal{N}_k is benign.
- Each agent employs a Huber loss function which guarantees $\|\nabla J_k(w)\| < \infty$.

Different numbers of outliers and four different attack schemes were investigated. The deployed attack schemes include: large value (LV), α -SCM, Talwar-SCM and Tukey-SCM. Here, LV injects a large constant value, which can be viewed as approaching the sensitivity curve maximization (SCM) strategy for the sample mean.

A. Efficiency

The asymptotic relative efficiency (ARE) describes the performance loss to the optimal maximum likelihood estimator [22, p. 22]. Assuming an underlying normal distribution, the mean has an ARE of 1 and the median has an ARE of $\frac{2}{\pi} = 0.64$. For the other estimators the ARE can be tuned based on their tuning variables. Here, all aggregators are tuned such that they achieve an efficiency of 0.95. Hence, we choose $\alpha = 0.0688$ for the α -trimmed mean [26, p. 35], $c = 2.7955$ for Talwar [23], [27] and $c = 4.685$ for biweight




Fig. 4: Training loss over iterations. Columns show different amount of outliers. Rows show different attack schemes.

Tukey [22, p. 23]. This allows for a fair comparison of the different aggregation schemes.

B. Results

In Figure 3, the training loss without malicious agents is shown. It can be observed that the sample mean converges fastest and the median slowest, which is consistent with the fact that they exhibit the highest and lowest efficiency respectively. The effects of a different amount of outliers and different attack schemes are shown in Figure 4. The columns depict different contamination rates, whereas the rows depict different attack schemes. It can be observed that the aggregation scheme which is under attack usually has the worst performance or a significantly worse performance, compared to other attack schemes. For example, in the second column under the LV attack the worst performance is shown by the sample mean, under α -SCM the worst performance is shown by the α -trimmed mean, under Talwar-SCM the worst performance is shown by the Talwar aggregator and under Tukey-SCM the worst performance is also shown by the Talwar aggregator, but closely followed by the Tukey aggregator, which shows a worse performance compared to the other attack schemes. Even under moderate contamination rates the aggregation schemes fail, if a targeted SCM is applied. We

conclude that carefully targeted attacks can be effective even against aggregation schemes which are robust to simpler attack models.

V. CONCLUSION AND OUTLOOK

We presented a novel attack on decentralized learning algorithms by injecting an outlier which maximizes the sensitivity curve of the used aggregation scheme. The simulations show that all considered aggregation schemes diverge. These results suggest that many robust aggregation schemes, despite their short-term robustness, can be driven to divergence through small but targeted perturbations.

Efficient defenses against the presented SC maximization attack might be a stronger reliance on each nodes own estimated weights, e.g. self-centered clipping [15]. Another approach could be the obfuscation of the used aggregation scheme, e.g. by randomly choosing a scheme out of a predefined set [21]. This could be combined with randomly adapting the tuning variables, which would make it more difficult for malicious agents to design optimal attack patterns.

REFERENCES

- [1] J. Chen and A. H. Sayed, "On the Learning Behavior of Adaptive Networks—Part I: Transient Analysis," *IEEE Transactions on Information Theory*, vol. 61, no. 6, pp. 3487–3517, 2015.
- [2] A. Sayed, "Adaptation, Learning, and Optimization over Networks," *Foundations and Trends® in Machine Learning*, vol. 7, no. 4-5, pp. 311–801, 2014.
- [3] X. Lian, C. Zhang, H. Zhang, C.-J. Hsieh, W. Zhang, and J. Liu, "Can Decentralized Algorithms Outperform Centralized Algorithms? A Case Study for Decentralized Parallel Stochastic Gradient Descent," in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, ser. NIPS'17, Curran Associates Inc, 2017, pp. 5336–5346.
- [4] A. Nedic, A. Olshevsky, and M. G. Rabbat, "Network Topology and Communication-Computation Tradeoffs in Decentralized Optimization," *Proceedings of the IEEE*, vol. 106, no. 5, pp. 953–976, 2018.
- [5] D. Yin, Y. Chen, K. Ramchandran, and P. Bartlett, *Byzantine-Robust Distributed Learning: Towards Optimal Statistical Rates*, 2018.
- [6] X. Wang, H. Zhang, A. Bilal, H. Long, and X. Liu, "WGM-dSAGA: Federated Learning Strategies with Byzantine Robustness Based on Weighted Geometric Median," *Electronics*, vol. 12, no. 5, p. 1190, 2023.
- [7] K. Pillutla, S. M. Kakade, and Z. Harchaoui, "Robust Aggregation for Federated Learning," *IEEE Transactions on Signal Processing*, vol. 70, pp. 1142–1154, 2022.
- [8] P. Blanchard, E. M. El Mhamdi, R. Guerraoui, and J. Stainer, "Machine Learning with Adversaries: Byzantine Tolerant Gradient Descent," in *Advances in Neural Information Processing Systems*, I. Guyon, U. V. Luxburg, S. Bengio, et al., Eds., vol. 30, Curran Associates, Inc, 2017. [Online]. Available: <https://proceedings.neurips.cc/paper/2017/file/f4b9ec30ad9f68f89b29639786cb62ef-Paper.pdf>.
- [9] E. M. E. Mhamdi, R. Guerraoui, and S. Rouault, *The Hidden Vulnerability of Distributed Learning in Byzantium*, 2018.
- [10] N. Rodríguez-Barroso, D. Jiménez-López, M. V. Luzón, F. Herrera, and E. Martínez-Cámara, "Survey on federated learning threats: Concepts, taxonomy on attacks and defences, experimental study and challenges," *Information Fusion*, vol. 90, pp. 148–173, 2023.
- [11] Z. Yang and W. U. Bajwa, "ByRDIE: Byzantine-Resilient Distributed Coordinate Descent for Decentralized Learning," *IEEE Transactions on Signal and Information Processing over Networks*, vol. 5, no. 4, pp. 611–627, 2019.
- [12] C. Fang, Z. Yang, and W. U. Bajwa, *BRIDGE: Byzantine-resilient Decentralized Gradient Descent*, 2019.
- [13] Z. Wu, T. Chen, and Q. Ling, *Byzantine-Resilient Decentralized Stochastic Optimization with Robust Aggregation Rules*, 2022.
- [14] S. Vlaski, C. Schroth, M. Muma, and A. M. Zoubir, "Robust and Efficient Aggregation for Distributed Learning," in *2022 30th European Signal Processing Conference (EUSIPCO)*, IEEE, 2022, pp. 817–821.
- [15] L. He, S. P. Karimireddy, and M. Jaggi, *Byzantine-Robust Decentralized Learning via Self-Centered Clipping*, 2022.
- [16] G. Baruch, M. Baruch, and Y. Goldberg, "A Little Is Enough: Circumventing Defenses For Distributed Learning," in *Advances in Neural Information Processing Systems 32 (NeurIPS 2019)*, H. Wallach, H. Larochelle, A. Beygelzimer, F. Alché-Buc, E. Fox, and R. Garnett, Eds., Curran Associates, Inc, 2019. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2019/file/ec1c59141046cd1866bbcbdfb6ae31d4-Paper.pdf.
- [17] M. Fang, X. Cao, J. Jia, and N. Z. Gong, *Local Model Poisoning Attacks to Byzantine-Robust Federated Learning*, 2019.
- [18] C. Xie, O. Koyejo, and I. Gupta, "Fall of Empires: Breaking Byzantine-tolerant SGD by Inner Product Manipulation," in *Proceedings of The 35th Uncertainty in Artificial Intelligence Conference*, R. P. Adams and V. Gogate, Eds., ser. Proceedings of Machine Learning Research, vol. 115, PMLR, 2020, pp. 261–270. [Online]. Available: <https://proceedings.mlr.press/v115/xie20a.html>.
- [19] H. Wang, K. Sreenivasan, S. Rajput, et al., *Attack of the Tails: Yes, You Really Can Backdoor Federated Learning*, 2020.
- [20] L. Lyu, H. Yu, and Q. Yang, *Threats to Federated Learning: A Survey*, 2020.
- [21] A. Ramezani-Kebrya, I. Tabrizian, F. Faghri, and P. Popovski, "MixTailor: Mixed Gradient Aggregation for Robust Learning Against Tailored Attacks," *Transactions on Machine Learning Research*, 2022.
- [22] A. M. Zoubir, V. Koivunen, E. Ollila, and M. Muma, *Robust Statistics for Signal Processing*. Cambridge University Press, 2018, vol. 25.
- [23] D. de Menezes, D. M. Prata, A. R. Secchi, and J. C. Pinto, "A review on robust M-estimators for regression analysis," *Computers & Chemical Engineering*, vol. 147, p. 107254, 2021.
- [24] M. J. Hinich and P. P. Talwar, "A Simple Method for Robust Regression," *Journal of the American Statistical Association*, vol. 70, no. 349, pp. 113–119, 1975.
- [25] F. R. Hampel, E. M. Ronchetti, P. J. Rousseeuw, and W. A. Stahel, *Robust Statistics: The Approach Based on Influence Functions* (Wiley series in probability and mathematical statistics Probability and mathematical statistics). New York, NY: Wiley, 1986.
- [26] F. Mosteller and J. W. Tukey, *Data analysis and regression: A second course in statistics* (Addison-Wesley series in behavioral science Quantitative methods). Reading, Mass.: Addison-Wesley, 1977.
- [27] P. W. Holland and R. E. Welsch, "Robust regression using iteratively reweighted least-squares," *Communications in Statistics - Theory and Methods*, vol. 6, no. 9, pp. 813–827, 1977.