This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Ollila, Esa; Mian, Ammar

Block-wise minimization-majorization algorithm for huber's criterion: Sparse learning and applications

Published in:

Proceedings of the 2020 IEEE 30th International Workshop on Machine Learning for Signal Processing, MLSP 2020

DOI:

[10.1109/MLSP49062.2020.9231538](https://doi.org/10.1109/MLSP49062.2020.9231538)

Published: 01/09/2020

Document Version

Peer reviewed version

Please cite the original version:

Ollila, E., & Mian, A. (2020). Block-wise minimization-majorization algorithm for huber's criterion: Sparse learning and applications. In *Proceedings of the 2020 IEEE 30th International Workshop on Machine Learning for Signal Processing, MLSP 2020* [9231538] (IEEE International Workshop on Machine Learning for Signal Processing). IEEE. <https://doi.org/10.1109/MLSP49062.2020.9231538>

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

BLOCK-WISE MINIMIZATION-MAJORIZATION ALGORITHM FOR HUBER'S CRITERION: SPARSE LEARNING AND APPLICATIONS

Esa Ollila and Ammar Mian

Department of Signal Processing and Acoustics, Aalto University, Finland

ABSTRACT

Huber's criterion can be used for robust joint estimation of regression and scale parameters in the linear model. Huber's [1] motivation for introducing the criterion stemmed from non-convexity of the joint maximum likelihood objective function as well as non-robustness (unbounded influence function) of the associated ML-estimate of scale. In this paper, we illustrate how the original algorithm proposed by Huber can be set within the block-wise minimization majorization framework. In addition, we propose novel data-adaptive step sizes for both the location and scale, which are further improving the convergence. We then illustrate how Huber's criterion can be used for sparse learning of underdetermined linear model using the iterative hard thresholding approach. We illustrate the usefulness of the algorithms in an image denoising application and simulation studies.

Index Terms— Huber's criterion, robust regression, sparse learning, minimization-majorization algorithm.

1. INTRODUCTION

Consider having N measurements or *outputs* (responses) $y_i \in \mathbb{R}$ and each output is associated with a p -dimensional vector of *inputs* (predictors) $x_i^\top = (x_{i1}, \dots, x_{ip}) \in \mathbb{R}^p$. We assume a linear regression model, where input-output relationship is described by

$$y_i = x_i^\top \beta + e_i, \quad i = 1, \dots, N, \quad (1)$$

where the random error terms e_i , $i = 1, \dots, N$, are independent and identically distributed (i.i.d.) and account for both the modelling and measurement errors. The distribution of the errors is assumed to be symmetric with probability density function (p.d.f.) $f(e) = (1/\sigma)f_0(e/\sigma)$, where σ denotes the scale parameter of the distribution and $f_0(e)$ denote the standardized (unit scale) distribution. The goal is to estimate the vector $\beta = (\beta_1, \dots, \beta_p)^\top \in \mathbb{R}^p$ of *regression coefficients* and the *scale* parameter σ given the data $\{(y_i, x_i); i = 1, \dots, N\}$. In many applications, the scale is a nuisance parameter, and the primary interest is on estimation of β .

The linear model can be conveniently expressed using matrix-vector notations. We use the convention that matrices are represented by bold uppercase letters while vectors will not be bold, except when they have N components. We may then express the set of N input p -vectors compactly via the $N \times p$ design matrix, $\mathbf{X} = (x_1 \ x_2 \ \dots \ x_N)^\top = (\mathbf{x}_1 \ \mathbf{x}_2 \ \dots \ \mathbf{x}_p)$. The convention distinguishes a p -vector of inputs x_i for the i^{th} observation from the N -vectors $\mathbf{x}_i$ consisting of all the observations on i^{th} variable. We collect the outputs into a vector $\mathbf{y} = (y_1, \dots, y_N)^\top$ and error terms into vector $\mathbf{e}$. Thus the linear model (1) now rewrites as $\mathbf{y} = \mathbf{X}\beta + \mathbf{e}$.

We return to the problem of joint robust estimation of regression and scale parameters in the linear model using Huber's criterion [1]. In this paper, we elaborate on recent derivations in [2] showing how the original algorithm proposed by Huber can be derived as a block-wise minimization majorization algorithm [3]. In addition, we propose new novel data-adaptive step sizes for both the location and scale parameter updates, which are illustrated to improve the convergence. We then describe the use of Huber's criterion in sparse signal recovery (SSR) [4, 5] problem using normalized iterative hard thresholding (NIHT) approach [6, 7]. Robust performance of Huber's sparse and non-sparse estimates of regression and scale are illustrated with a simulation study and in robust image denoising application and Matlab and Python toolboxes are made available in github.com/AmmarMian/huber_mm_framework. We hope that the present paper and the toolboxes are able to bring more attention to Huber's criterion which is scalable robust approach for many practical large-scale problems. Finally, we note that sparse/non-sparse linear estimation in the linear regression model using Huber's criterion has been considered in [8, 9, 2]. The derived MM framework for Huber's criterion can potentially lead to new research directions using theory developed in [10, 3].

2. MAXIMUM LIKELIHOOD APPROACH

In his seminal work, Huber [11] derived a family of univariate heavy-tailed distributions, which he called the "least favorable distributions" (LFDs). The LFD corresponds to a sym-

This work was supported by the Academy of Finland under Grant 298118.

metric unimodal distribution which follows a Gaussian distribution in the middle, and a double exponential distribution in the tails. The standardized (unit scale) p.d.f. of LFD is $f_0(x) = e^{-\rho_c(x)}$, where $\rho_c(x)$ is the *Huber's loss function*,

$$\rho_c(x) = \frac{1}{2} \times \begin{cases} |x|^2, & \text{for } |x| \leq c \\ 2c|x| - c^2, & \text{for } |x| > c, \end{cases} \quad x \in \mathbb{R}, \quad (2)$$

where c is a user-defined *threshold* that influences the degree of robustness. Huber's loss (2) is a hybrid of the (Least-squares) LS- and the least absolute deviation (LAD) loss functions using the LS-loss function for relatively small errors and LAD loss function for relatively large errors. Furthermore, in the limit $c \rightarrow \infty$, the loss function reduces to LS-loss $\rho_c(x) = \frac{1}{2}x^2$. Huber's loss function is convex and differentiable, and the derivative of the loss function,

$$\psi_c(x) = \begin{cases} x, & \text{for } |x| \leq c \\ c \operatorname{sign}(x), & \text{for } |x| > c, \end{cases}$$

will be referred to as the *score function* in the sequel. Note that $\psi_c(\cdot)$ is a winsorizing function.

The threshold c is usually chosen so that the minimizer in the regression-only problem (σ being fixed) attains a user-defined asymptotic relative efficiency (ARE) w.r.t. the LS estimate (LSE) under Gaussian errors. In order to obtain 95% (or 85%) ARE for the Gaussian noise case, the threshold are chosen according to $c_{.95} = 1.345$ and $c_{.85} = 0.7317$.

Assume now that the error terms e_i in the linear model (1) follow the LFD, so with the standardized distribution given previously. Then consider finding the ML estimates (MLEs) of the unknown parameters β and σ via minimizing the negative log-likelihood function of the data:

$$L_{\text{ML}}(\beta, \sigma) = N \ln \sigma + \sum_{i=1}^N \rho_c \left(\frac{y_i - x_i^\top \beta}{\sigma} \right). \quad (3)$$

The problem is that the negative log-likelihood function is not convex in (β, σ) . This is easy to see by simply noting that $L_{\text{ML}}(\beta, \sigma)$ is not convex in σ for a fixed β . Another problem is that the associated scale estimate will not be robust, e.g., possessing a bounded influence function. The problem of non-robustness of the associated MLE and the non-convexity of the ML criterion (3), lead Huber to consider an alternative criterion function discussed in detail in the following sections.

3. HUBER'S CRITERION

In order to avoid the problems associated with (3) described earlier, Huber [1] proposed an alternative criterion function

$$L(\beta, \sigma) = N(\alpha\sigma) + \sum_{i=1}^N \rho_c \left(\frac{y_i - x_i^\top \beta}{\sigma} \right) \sigma \quad (4)$$

where $\alpha > 0$ is a fixed scaling factor. We refer to (4) as *Huber's criterion* and the minimizer $(\hat{\beta}, \hat{\sigma})$ of $L(\beta, \sigma)$ as *Huber's joint estimates of regression and scale*. Unlike (3), Huber's criterion function (4) is jointly convex in (β, σ) . In addition to this, the minimizer $\hat{\beta}$ preserves the same theoretical robustness properties (such as bounded influence function) as the estimator in the regression-only problem where the scale parameter is known.

An important concept when optimizing the Huber's criterion is the pseudo-residual which is in essence a winsorized version of the conventional residual $\mathbf{r} = \mathbf{y} - \mathbf{X}\beta$. The *pseudo-residual* is associated with the score function, and defined as

$$\mathbf{r}_\psi \equiv \mathbf{r}_\psi(\beta, \sigma) = \psi_c \left(\frac{\mathbf{y} - \mathbf{X}\beta}{\sigma} \right) \sigma, \quad (5)$$

where $\psi_c(\mathbf{r}/\sigma) = (\psi_c(r_1/\sigma), \dots, \psi_c(r_N/\sigma))^\top$. In the case of LS-loss (the case $c \rightarrow \infty$), $\psi_c(x) = x$, and $\mathbf{r}_\psi$ coincides with the conventional residual vector, so $\mathbf{r}_\psi = \mathbf{r}$. Scaling by σ in (5) maps the residuals back to the original scale of the data.

Since the optimization problem (4) is convex, the global minimizer $(\hat{\beta}, \hat{\sigma})$ is a stationary point of (4). Thus, $(\hat{\beta}, \hat{\sigma})$ can be found by solving the *M-estimating equations*, obtained by setting the gradient of $L(\beta, \sigma)$ to zero

$$\begin{aligned} \nabla_\beta L(\hat{\beta}, \hat{\sigma}) &= 0 \Leftrightarrow \mathbf{x}_j^\top \hat{\mathbf{r}}_\psi = 0, \quad j = 1, \dots, p \\ \nabla_\sigma L(\hat{\beta}, \hat{\sigma}) &= 0 \Leftrightarrow \frac{1}{N} \sum_{i=1}^N \chi_c \left(\frac{y_i - x_i^\top \hat{\beta}}{\hat{\sigma}} \right) = \alpha, \end{aligned} \quad (6)$$

where $\hat{\mathbf{r}}_\psi = \mathbf{r}_\psi(\hat{\beta}, \hat{\sigma})$ and $\chi_c : \mathbb{R} \rightarrow \mathbb{R}_0^+$ is defined as

$$\chi_c(x) = \psi_c(x)x - \rho_c(x) = \frac{1}{2}\psi_c(x)^2. \quad (7)$$

It is instructive to consider the estimating equations (6) in the case of LS loss $\rho_c(x) = \frac{1}{2}x^2$. In this case, the score function is $\psi_c(x) = x$ and $\hat{\mathbf{r}}_\psi = \hat{\mathbf{r}}$, and the first equation in (6) becomes the conventional normal equations $\mathbf{x}_j^\top \hat{\mathbf{r}} = 0$, for $j = 1, \dots, p \Leftrightarrow \mathbf{X}^\top (\mathbf{y} - \mathbf{X}\hat{\beta}) = 0$. Hence, the minimizer $\hat{\beta}$ of Huber's criterion is simply the LS estimate (LSE) of regression $\hat{\beta} = (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{y}$. Furthermore, the χ_c function in (7) is simply $\chi_c(x) = \frac{1}{2}x^2$ and the 2nd estimating equation in (6) reduces to $\hat{\sigma}^2 = \frac{1}{N(2\alpha)} \sum_{i=1}^N (y_i - x_i^\top \hat{\beta})^2$. Thus, If we choose $\alpha = \frac{1}{2}$, then the solution $\hat{\sigma}$ coincides with the classical MLE of scale for Gaussian errors, the standard deviation of the residuals. Interestingly, the two quite different criterion functions, the Huber's criterion (4) and the ML criterion function (3), share the same unique global (joint) minimum when the LS-loss function is used. Therefore, Huber's criterion can be seen as a method to convexify the Gaussian ML criterion function.

The scaling factor α in (4) is used to ensure that $\hat{\sigma}$ is Fisher-consistent for the unknown scale σ when $\{e_i\}_{i=1}^N \stackrel{i.i.d.}{\sim}$

$\mathcal{N}(0, \sigma^2)$. Due to (6), this is achieved if

$$\alpha = \mathbb{E}[\chi_c(e)] = \frac{c^2}{2}(1 - F_{\chi_1^2}(c^2)) + \frac{1}{2}F_{\chi_3^2}(c^2), \quad (8)$$

where $F_{\chi_k^2}$ denotes the c.d.f. of chi-squared distribution with k degrees of freedom and $e \sim \mathcal{N}(0, 1)$. Hence, when using the LS loss, $\chi_c(e) = \frac{1}{2}e^2$, and Fisher consistency is obtained if $\alpha = \frac{1}{2}\mathbb{E}[e^2] = \frac{1}{2}$

4. MINIMIZATION-MAJORIZATION ALGORITHM

A block-wise Minimization-Majorization algorithm works similarly to MM algorithms [3, 12, 13]. Let θ be partitioned into $\theta = (\theta_1, \theta_2)$, where $\theta_1 \in \Theta_1, \theta_2 \in \Theta_2$ and $\Theta = \Theta_1 \times \Theta_2$ and we wish to find a minimizer of a real-valued function $L(\theta) = L(\theta_1, \theta_2)$. At $(n+1)$ th iteration, the blocks are updated in a cyclic manner as follows:

$$\theta_2^{(n+1)} = T_2(\theta_1^{(n)}, \theta_2^{(n)}), \quad (9)$$

$$T_2(\theta_1^{(n)}, \theta_2^{(n)}) = \arg \min_{\theta_2 \in \Theta_2} g_2(\theta_2 | \theta_1^{(n)}, \theta_2^{(n)})$$

$$\theta_1^{(n+1)} = T_1(\theta_1^{(n)}, \theta_2^{(n+1)}), \quad (10)$$

$$T_1(\theta_1^{(n)}, \theta_2^{(n+1)}) = \arg \min_{\theta_1 \in \Theta_1} g_1(\theta_1 | \theta_1^{(n)}, \theta_2^{(n+1)})$$

where the majorization surrogate functions $g_i(\theta_i | \theta_1', \theta_2') = g_i(\theta_i | \theta')$, $i \in \{1, 2\}$ satisfy

$$g_i(\theta_i' | \theta_1', \theta_2') = L(\theta_1', \theta_2') \quad \forall (\theta_1', \theta_2') \in \Theta, \quad (11)$$

$$g_1(\theta_1 | \theta_1', \theta_2') \geq L(\theta_1, \theta_2') \quad \forall \theta_1 \in \Theta_1, \forall (\theta_1', \theta_2') \in \Theta, \quad (12)$$

$$g_2(\theta_2 | \theta_1', \theta_2') \geq L(\theta_1', \theta_2) \quad \forall \theta_2 \in \Theta_2, \forall (\theta_1', \theta_2') \in \Theta. \quad (13)$$

Furthermore, when L and g_i are differentiable functions, it is possible to impose the constraint

$$\nabla_{\theta_i} g_i(\theta_i | \theta_1', \theta_2')|_{\theta_i=\theta_i'} = \nabla_{\theta_i} L(\theta)|_{\theta=\theta'}, \quad i \in \{1, 2\}. \quad (14)$$

Under some regularity conditions, the sequence obtained by iterating the steps above is a stationary point of $L(\theta)$ if it lies in the interior of Θ .

The challenging part in designing a block-wise MM algorithm is, naturally, in finding appropriate surrogate functions $g_i(\cdot | \theta')$, $i \in \{1, 2\}$. A common choice is a quadratic function of the form $b_0 + b_1\theta + b_2\theta^2$ as it results in a simple update formula in (10), (9).

We now construct an MM algorithm for obtaining the stationary solution $(\hat{\beta}, \hat{\sigma})$ in (6) of Huber's criterion $L(\beta, \sigma) = \sum_i \rho_c(r_i/\sigma)\sigma + N(\alpha\sigma)$. For this purpose, we will let β' and σ' denote values of previous iterates, and write $\mathbf{r}' = \mathbf{y} - \mathbf{X}\beta'$ and $\mathbf{r}'_\psi = \psi(\mathbf{r}'/\sigma')\sigma'$ for the corresponding residual and pseudo-residual.

For the scale term, we will use the following majorization surrogate function:

$$g_2(\sigma | \beta', \sigma') = a' + b' \frac{1}{\sigma} + N\alpha\sigma. \quad (15)$$

In (15), $a' + b'\sigma^{-1}$ is used to majorize $\sum_i \rho_c(r_i/\sigma)\sigma$, where a' and b' are constants that depend on the previous iterates β' and σ' . These terms can be found by solving the pair of equations (11) and (14) which yields (after simplifying) the following surrogate function

$$g_2(\sigma | \beta', \sigma') = L(\beta', \sigma') + N\alpha(\sigma - \sigma') + \sum_{i=1}^N \chi_c\left(\frac{r'_i}{\sigma'}\right) \left(\frac{(\sigma')^2}{\sigma} - \sigma'\right). \quad (16)$$

Next we turn into constructing a majorization surrogate function $g_1(\beta | \beta', \sigma')$. For this purpose, consider a surrogate function for $\rho_c(r_i/\sigma')$ by using

$$\rho_M\left(\frac{r_i}{\sigma'}\right) = a'_i + b'_i \frac{r_i}{\sigma'} + \frac{1}{2} \frac{r_i^2}{(\sigma')^2},$$

where the constants a'_i and b'_i depend on the previous iterates β' and σ' and are found by solving the pair of equations (11) and (14) w.r.t. a'_i and b'_i . After finding these solutions, we obtain a surrogate function of the form

$$\begin{aligned} g_1(\beta | \beta', \sigma') &= N(\alpha\sigma') + \sum_{i=1}^N \left(a'_i + b'_i \frac{r_i}{\sigma'} + \frac{1}{2} \frac{r_i^2}{(\sigma')^2} \right) \sigma' \\ &= \text{const} + \sum_{i=1}^N \left([r'_{\psi,i} - r'_i] r_i + \frac{1}{2} r_i^2 \right) \frac{1}{\sigma'} \end{aligned}$$

where the constant term does not depend on β .

4.1. MM Algorithm for Huber's Criterion

Next we prove that g_1 and g_2 are valid surrogate functions, so verify (12) and (13). First note that the difference function is

$$h_2(\sigma) = g_2(\sigma | \beta', \sigma') - L(\beta', \sigma) = a_0 + \frac{b_0}{\sigma} - \sum_{i=1}^N \rho_c\left(\frac{r'_i}{\sigma}\right) \sigma$$

for some constants a_0 and b_0 . Then note that $h_2(\sigma)$ is a convex function in $1/\sigma$ since the first term is linear and $\rho_c(x)/x$ is a concave function in $x \geq 0$. Furthermore, since $h_2(\sigma') = 0$, it follows that $h_2(\sigma) \geq 0$, i.e., $g_2(\sigma | \beta', \sigma') \geq L(\beta', \sigma)$, for all $\sigma > 0$. This proves that g_2 verifies (13). Next consider the difference function

$$\begin{aligned} h_1(\beta) &= g_1(\beta | \beta', \sigma') - L(\beta, \sigma') \\ &= \text{cst} + \sum_{i=1}^N \left([r'_{\psi,i} - r'_i] r_i + \frac{1}{2} r_i^2 \right) \frac{1}{\sigma'} - \sum_{i=1}^N \rho_c\left(\frac{r_i}{\sigma'}\right) \sigma'. \end{aligned}$$

The Hessian matrix of the difference function h_1 is

$$\mathbf{H}_1 = \frac{\partial^2 h_1}{\partial \beta \partial \beta^\top} = \frac{1}{\sigma'} \sum_{i=1}^N \left\{ 1 - \psi'_c\left(\frac{r_i}{\sigma'}\right) \right\} x_i x_i^\top.$$

Algorithm 1: hubreg, solves (4) via MM.

input : $(\mathbf{y}, \mathbf{X})$, threshold $c > 0$, initial guess $(\beta^{(0)}, \sigma^{(0)})$, $\epsilon > 0$, $\mu^{(0)} = 0$ and $\lambda^{(0)} = 1$
initialize: $N_{\text{iter}} \in \mathbb{N}$, compute $\alpha = \alpha(c)$ in (8) and $\mathbf{X}^+$
for $n = 0, 1, \dots, N_{\text{iter}}$ **do**
1 Update the residual $\mathbf{r}^{(n)} = \mathbf{y} - \mathbf{X}\beta^{(n)}$ and τ in (18)
2 Update the step size for scale:

$$\lambda^{(n+1)} = \lambda^{(n)} + \frac{\log \left\| \psi_c \left(\frac{\mathbf{r}^{(n)}}{\sigma^{(n)} \tau^{\lambda^{(n)}}} \right) \frac{1}{\sqrt{2\alpha N}} \right\|}{\log \tau}$$

3 Update scale $\sigma^{(n+1)} = \sigma^{(n)} \tau^{\lambda^{(n+1)}}$ and δ in (19)
4 Update the step size for regression:

$$\mathbf{z} = \mathbf{X}\delta, \quad \mathbf{w} = w \left(\frac{\mathbf{r}^{(n)} - \mu^{(n)} \mathbf{z}}{\sigma^{(n+1)}} \right)$$

$$\mu^{(n+1)} = \|\mathbf{z}\|_{\mathbf{w}}^{-2} \langle \mathbf{r}^{(n)}, \mathbf{z} \rangle_{\mathbf{w}}$$

5 Update the regression vector:

$$\beta^{(n+1)} = \beta^{(n)} + \mu^{(n+1)} \delta$$
if $\|\mu^{(n+1)} \delta\| / \|\beta^{(n)}\| < \epsilon$ & $|\tau^{\lambda^{(n+1)}} - 1| < \epsilon$ **then**
 return $(\hat{\beta}, \hat{\sigma}) \leftarrow (\beta^{(n+1)}, \sigma^{(n+1)})$
output : $(\hat{\beta}, \hat{\sigma})$, the minimizer of $L(\beta, \sigma)$.

Note that $0 \leq \psi'_c(x) \leq 1$. Thus the matrix $\mathbf{H}_1$ is a positive semi-definite matrix, and thus the difference function h_1 is a convex function with a minimum at β' . These results and the fact that $h_1(\beta') = 0$ imply that $g_1(\beta|\beta', \sigma') \geq L(\beta, \sigma')$, $\forall \beta$. In the next theorem we obtain the minimizers. The proof is omitted due to lack of space.

Theorem 1. *The MM update of scale is*

$$\sigma^{(n+1)} = \arg \min_{\sigma > 0} g_2(\sigma|\beta^{(n)}, \sigma^{(n)}) = \sigma^{(n)} \tau, \quad (17)$$

where

$$\tau = \frac{1}{\sqrt{2N\alpha}} \left\| \psi_c \left(\frac{\mathbf{r}^{(n)}}{\sigma^{(n)}} \right) \right\| \quad (18)$$

and $\mathbf{r}^{(n)} = \mathbf{y} - \mathbf{X}\beta^{(n)}$. The MM update for regression is

$$\beta^{(n+1)} = \arg \min_{\beta \in \mathbb{R}^{p+1}} g_1(\beta|\beta^{(n)}, \sigma^{(n+1)}) = \beta^{(n)} + \delta,$$

where

$$\delta = \mathbf{X}^+ \psi_c \left(\frac{\mathbf{r}^{(n)}}{\sigma^{(n+1)}} \right) \sigma^{(n+1)} \quad (19)$$

and $\mathbf{X}^+ = (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top$.

The updates of Theorem 1 thus form the basic frame for the MM algorithm that is described in algorithm 1.

4.2. Step size selection

An additional change to the MM algorithm is the introduction of the step sizes in line 2 and line 4.

Since the update for β can be viewed as a gradient descent move towards the direction δ , one may try to identify an optimal step size that maximally reduces the objective function at each iteration. In other words, we minimize Huber's criterion $L(\beta, \sigma)$ with $\sigma = \sigma^{(n+1)}$ and β fixed at $\beta = \beta^{(n)} + \mu\delta$:

$$\begin{aligned} \mu^{(n+1)} &= \arg \min_{\mu} \sum_{i=1}^N \rho_c \left(\frac{y_i - x_i^\top (\beta^{(n)} + \mu\delta)}{\sigma^{(n+1)}} \right) \\ &= \arg \min_{\mu} \sum_{i=1}^N \rho_c \left(\frac{r_i^{(n)} - \mu x_i^\top \delta}{\sigma^{(n+1)}} \right). \end{aligned} \quad (20)$$

Solving (20) is equivalent to computing Huber's M -estimator of regression with auxiliary scale $\sigma^{(n+1)}$ in the simple (one predictor) linear regression model with response vector $\mathbf{r}^{(n)}$ and regressor $\mathbf{z} = \mathbf{X}\delta$. A standard approach for finding the (unique) minimizer of the convex optimization problem in (20) is the *iteratively reweighted LS (IRWLS) algorithm* [14], which iterates the steps

$$\mathbf{w} \leftarrow w \left(\frac{\mathbf{r}^{(n)} - \mu \mathbf{z}}{\sigma^{(n+1)}} \right), \quad \mu \leftarrow \|\mathbf{z}\|_{\mathbf{w}}^{-2} \langle \mathbf{r}^{(n)}, \mathbf{z} \rangle_{\mathbf{w}} \quad (21)$$

until convergence, given an initial value of μ to start the iterations. Here $\langle \mathbf{a}, \mathbf{b} \rangle_{\mathbf{w}} = \sum_{i=1}^N a_i b_i w_i$ and $\|\mathbf{a}\|_{\mathbf{w}}^2 = \langle \mathbf{a}, \mathbf{a} \rangle_{\mathbf{w}}$. Instead of iterating until convergence, we use a 1-step estimator, which correspond to a single iteration of (21) with initial value given by previous value of step size $\mu^{(n)}$. This results to the update shown in line 4 of algorithm 1.

Next we turn our attention to the step size for scale. First we note that the update $\sigma^{(n+1)} = \sigma^{(n)} \tau$ can be viewed as a gradient descent move in the log-space (after change of variables $a = \log \sigma$, i.e., $\sigma = e^a$), $a^{(n+1)} = a^{(n)} + \lambda \log(\tau)$ with stepsize $\lambda = 1$. Thus employing a data adaptive step size $\lambda = \lambda^{(n+1)}$ translates to update $\sigma^{(n+1)} = \sigma^{(n)} \tau^{\lambda^{(n+1)}}$. We then identify an optimal step size that maximally reduces the objective function in each iteration. In other words, we minimize Huber's criterion $L(\beta, \sigma)$ with β fixed at $\beta = \beta^{(n)}$ and $\sigma = \sigma^{(n)} \tau^\lambda$:

$$\lambda^{(n+1)} = \arg \min_{\lambda} N \alpha \tau^\lambda + \sum_{i=1}^N \rho_c \left(\frac{\tilde{r}_i^{(n)}}{\tau^\lambda} \right) \tau^\lambda \quad (22)$$

where $\tilde{r}_i^{(n)} = r_i^{(n)} / \sigma^{(n)} = (y_i - x_i^\top \beta^{(n)}) / \sigma^{(n)}$. As was shown in Theorem 1, an MM algorithm that finds the optimum of (22) would iterate

$$\tau^\lambda \leftarrow \tau^\lambda \left\| \psi_c \left(\frac{\tilde{\mathbf{r}}^{(n)}}{\tau^\lambda} \right) \frac{1}{\sqrt{2N\alpha}} \right\|$$

starting from some initial value for λ until convergence. The iteration update above can be written equivalently as

$$\lambda \leftarrow \lambda + \log \left\| \psi_c(\tilde{\mathbf{r}}^{(n)} / \tau^\lambda) \frac{1}{\sqrt{2N\alpha}} \right\| / \log \tau. \quad (23)$$

Again, instead of iterating (23) until convergence, we use a 1-step estimator, which correspond to a single iteration of (23) with initial value given by previous value of step size $\lambda^{(n)}$. This results to the update shown in [line 2 of algorithm 1](#).

5. SPARSE LEARNING VIA HUBER'S CRITERION

Next consider a sparse signal recovery (SSR) problem [4, 5], where β in (1) is K -sparse, i.e., the *support* $\Gamma = \{i \in \{1, \dots, p\} : \beta_i \neq 0\}$ has at most K -nonzero elements, so $\|\beta\|_0 = |\Gamma| \leq K$. Furthermore, the dimensionality p can be greater than the number of measurements N . This means that the design matrix $\mathbf{X}$ can be underdetermined (and not of full rank). In the context of SSR, $\mathbf{X}$ is often referred to as *measurement matrix* or *dictionary* and β as the *signal vector*.

Traditional non-robust approaches aim at minimizing $\|\mathbf{y} - \mathbf{X}\beta\|^2$ w.r.t. β subject to $\|\beta\|_0 \leq K$. This is a combinatorial NP-hard problem. However, when the measurement matrix $\mathbf{X}$ satisfies certain coherence conditions, bounds on the recovery error are known for several reconstruction algorithms when the measurements are corrupted by noise with bounded norm. One such algorithm is the *normalized iterative hard-thresholding (NIHT)* algorithm [15].

The greedy NIHT approach for minimizing $L(\beta, \sigma)$ subject to $\|\beta\|_0 \leq K$ iterates the following steps [6]:

1. update $\mathbf{r}^{(n)} = \mathbf{y} - \mathbf{X}\beta^{(n)}$ and τ in (18).
2. update the step size $\lambda^{(n+1)}$ for scale.
3. update the scale $\sigma^{(n+1)} = \sigma^{(n)} \tau^{\lambda^{(n+1)}}$.
4. update the step size $\mu^{(n+1)}$ for signal vector.
5. update $\mathbf{r}_\psi^{(n+1)} = \psi_c \left(\frac{\mathbf{r}^{(n)}}{\sigma^{(n+1)}} \right) \sigma^{n+1}$.
6. $\beta^{(n+1)} = H_K(\beta^{(n)} + \mu^{(n+1)} \mathbf{X}^\top \mathbf{r}_\psi^{(n+1)})$.

until convergence, where $H_K(\beta)$ denotes the *hard thresholding operator*: it retains the K elements of vector β that are largest in absolute value and set the other elements to zero. The step size for scale and regression can be found as described in [algorithm 1](#). This robust SSR method described above is the same as in [6] except we have incorporated the proposed step size for scale. We refer to this method as HUBNIHT in the sequel. Extension of [6] to *simultaneous SSR* [5] problem has been covered in [7] and [16] for real- and complex-valued signals, respectively (and referred to as HUBSNIHT).

6. APPLICATIONS

6.1. Regression example

First we consider a simple regression experiment to validate the proposed MM framework for Huber's criterion. We consider a linear model as in (1) where $N = 500$, $p = 250$, the e_i are generated through standardized Gaussian distribution and σ is chosen so that the signal to noise ratio is $\text{SNR} = 20\text{dB}$. Additionally, some errors are introduced into the model by changing the sign of random samples y_i chosen according to a Bernoulli distribution with probability $\varepsilon \in [0, 1]$. We then compare the performance of the hubreg estimator ($c = 1.345$) of [algorithm 1](#) compared to the LSE for regression and standard deviation (SD) for scale by performing 2000 Monte-Carlo trials. The results displayed in [Figure 1](#) show that while the estimation is correct for both methods when there is no outliers, the hubreg estimator is more robust to the presence of errors.

Fig. 1. Evolution of estimation error versus sign-change error probability ε of outputs y_i ; $N = 500$, $p = 250$, $\text{SNR} = 20\text{dB}$, $c = 1.345$, 2000 trials.

6.2. Image denoising

To illustrate potential applications of the sparse framework presented in [section 5](#), we consider an image denoising problem of a grayscale image. Following a sliding windows approach, we consider the pixel values $\mathbf{y}$ on the local patch as the outputs of a linear model where the inputs $\mathbf{X}$ constitutes an overcomplete dictionary. Such dictionaries contain more atoms than their dimensions and are often redundant so a valid representation of the image must be sparse. Given this dictionary, an encoding $\hat{\beta}$ can be obtained by minimizing Huber's criterion with K -sparsity constraint using the HUBNIHT algorithm. Thanks to the robustness of this criterion to outliers, as shown previously, it is possible to reconstruct pixels values $\hat{\mathbf{y}} = \mathbf{X}\hat{\beta}$ that are less affected by the noise.

An example of denoising via sparse reconstruction by using Huber's criterion is given in [Figure 2](#). The dictionary, of size $p = 192$, has been constructed using standard Daubechies and Coiflets dictionaries often used in image processing applications and a window of size 8×8 has been

used. The difference between the two obtained results highlight that there is a compromise between sparsity to reduce noise and loss of details in the original image.

Fig. 2. i) Original Image. ii) Noisy image. iii) Reconstructed image ($K = 6$). iv) Reconstructed image ($K = 10$).

6.3. Dictionary learning

Finally, let us consider the problem of dictionary learning which aims at finding a sparse representation as in the SSR problem but also aims to learn the dictionary directly from the data. This can be formulated as the following optimization problem:

$$\min_{\beta, \mathbf{X} \in \mathcal{D}} \|\mathbf{y} - \mathbf{X}\beta\|^2 \quad \text{subject to} \quad \|\beta\|_0 \leq K, \quad (24)$$

where $\mathcal{D}$ is a set imposing constraints to the dictionary elements in order help resolve the scale invariance problem between $(\beta, \mathbf{X})$. Among those constraints, unit ℓ_2 -norm of each column of $\mathbf{X}$ is often used in the literature. Traditionally algorithms that solve (24) iterate between a step of sparse coding, where $\mathbf{X}$ being fixed, and a step of estimation of dictionary $\mathbf{X}$, where β is fixed (obtained in the sparse coding step). Popular algorithms include the Method of Optimal Directions (MOD) [17] or the K-SVD [18] algorithms. Both of the methods perform sparse coding using the matching pursuit algorithm and differ in the way the dictionary is learned from the data. Since the MM-framework presented in this paper has been successfully applied to sparse coding, it can replace the matching pursuit used in both approaches. This plug-in methodology is expected to inherit robustness properties from Huber's estimates. This research direction will be studied in the subsequent extended journal version of this work.

7. REFERENCES

- [1] P. J. Huber, *Robust Statistics*, Wiley, New York, 1981.
- [2] A. M. Zoubir, V. Koivunen, E. Ollila, and M. Muma, *Robust Statistics for Signal Processing*, Cambridge University Press, Cambridge, UK, Nov. 2018.
- [3] M. Razaviyayn, M. Hong, and Z.-Q. Luo, "A unified convergence analysis of block successive minimization methods for nonsmooth optimization," *SIAM Journal on Optimization*, vol. 23, no. 2, pp. 1126–1153, 2013.
- [4] D. L. Donoho, M. Elad, and V. N. Temlyakov, "Stable recovery of sparse overcomplete representations in the presence of noise," *IEEE Transactions on information theory*, vol. 52, no. 1, pp. 6–18, 2006.
- [5] M. F. Duarte and Y. C. Eldar, "Structured compressed sensing: From theory to applications," *IEEE Transactions on Signal Processing*, vol. 59, no. 9, pp. 4053–4085, 2011.
- [6] E. Ollila, H.-J. Kim, and V. Koivunen, "Robust iterative hard thresholding for compressed sensing," in *Proc. 2014 6th International Symposium on Communications, Control and Signal Processing (ISCCSP)*, Athens, Greece, May 21 – 23, 2014, pp. 226–229.
- [7] E. Ollila, *Modern Nonparametric, Robust and Multivariate Methods: A Festschrift in Honor of Professor Hannu Oja*, chapter Robust simultaneous sparse approximation, Springer, 2015.
- [8] A. B. Owen, "A robust hybrid of lasso and ridge regression," *Contemporary Mathematics*, vol. 443, pp. 59–72, 2007.
- [9] S. Lambert-Lacroix and L. Zwald, "Robust regression through the Huber's criterion and adaptive lasso penalty," *Electronic Journal of Statistics*, vol. 5, pp. 1015–1053, 2011.
- [10] E. D. Schifano, R. L. Strawderman, and M. T. Wells, "Majorization-minimization algorithms for nonsmoothly penalized objective functions," *Electronic Journal of Statistics*, vol. 4, pp. 1258–1299, 2010.
- [11] P. J. Huber, "Robust estimation of a location parameter," *The Annals of Mathematical Statistics*, vol. 35, no. 1, pp. 73–101, 1964.
- [12] D. R. Hunter and K. Lange, "A tutorial on MM algorithms," *The American Statistician*, vol. 58, no. 1, pp. 30–37, 2004.
- [13] Y. Sun, P. Babu, and D. P. Palomar, "Majorization-minimization algorithms in signal processing, communications, and machine learning," *IEEE Trans. Signal Process.*, vol. 65, no. 3, pp. 794–816, 2017.
- [14] R. A. Maronna, R. D. Martin, and V. J. Yohai, *Robust Statistics: Theory and Methods*, Wiley, New York, 2006.
- [15] T. Blumensath and M. E. Davies, "Normalized iterative hard thresholding: Guaranteed stability and performance," *IEEE Journal of selected topics in signal processing*, vol. 4, no. 2, pp. 298–309, 2010.
- [16] E. Ollila, "Multichannel sparse recovery of complex-valued signals using Huber's criterion," in *Proc. Compressed Sens. Theory Appl. Radar, Sonar, and Remote Sens. (CoSeRa)*, Pisa, Italy, 2015, pp. 32–36.
- [17] K. Egan, S. O. Aase, and J. Hakon Husoy, "Method of optimal directions for frame design," in *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 1999, vol. 5, pp. 2443–2446.
- [18] M. Aharon, M. Elad, and A. Bruckstein, "K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation," *IEEE Transactions on Signal Processing*, vol. 54, no. 11, pp. 4311–4322, 2006.