



Stein COnsistent Risk Estimator (SCORE) for hard thresholding

Charles-Alban Deledalle, Gabriel Peyré, Jalal M. Fadili

► To cite this version:

Charles-Alban Deledalle, Gabriel Peyré, Jalal M. Fadili. Stein COnsistent Risk Estimator (SCORE) for hard thresholding. Signal Processing with Adaptive Sparse Structured Representations., Jul 2013, Lausanne, Switzerland. hal-00776303v1

HAL Id: hal-00776303

<https://hal.science/hal-00776303v1>

Submitted on 24 Jan 2013 (v1), last revised 28 Apr 2015 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Stein CONSistent Risk Estimator (SCORE) for hard thresholding

Charles-Alban Deledalle
Institut de Mathématiques de Bordeaux
CNRS-Université Bordeaux I
Bordeaux, France
Email: cdeledal@math.u-bordeaux1.fr

Gabriel Peyré
CEREMADE
CNRS-Université Paris Dauphine
Paris, France
Email: gabriel.peyre@ceremade.dauphine.fr

Jalal Fadili
GREYC
CNRS-ENSICAEN
Caen, France
Email: jalal.fadili@greyc.ensicaen.fr

Abstract

In this work, we construct a risk estimator for hard thresholding which can be used as a basis to solve the difficult task of automatically selecting the threshold. As hard thresholding is not even continuous, Stein's lemma cannot be used to get an unbiased estimator of degrees of freedom, hence of the risk. We prove that under a mild condition, our estimator of the degrees of freedom, although biased, is consistent. Numerical evidence shows that our estimator outperforms another biased risk estimator proposed in [1].

I. INTRODUCTION

We observe a realisation $y \in \mathbb{R}^P$ of the normal random vector $Y = x_0 + W$, $W \sim \mathcal{N}(x_0, \sigma^2 \text{Id}_P)$. Given an estimator $y \mapsto x(y, \lambda)$ of x_0 evaluated at y and parameterized by λ , the associated Degree Of Freedom (DOF) is defined as [2]

$$df\{x\}(x_0, \lambda) \triangleq \sum_{i=1}^P \frac{\text{cov}(Y_i, x(Y_i, \lambda))}{\sigma^2}. \quad (1)$$

The DOF plays an important role in model/parameter selection. For instance, define the criterion

$$\|Y - x(Y, \lambda)\|^2 - P\sigma^2 + 2\sigma^2 \widehat{df}\{x\}(Y, \lambda). \quad (2)$$

In the rest, we denote div the divergence operator. If $x(\cdot, \lambda)$ is weakly differentiable w.r.t. its first argument with an essentially bounded gradient, Stein's lemma [3] implies that $\widehat{df}\{x\}(Y, \lambda) = \text{div}(x(Y, \lambda))$ and (2) (the SURE in this case) are respectively unbiased estimates of $df\{x\}(x_0, \lambda)$ and of the risk $\mathbb{E}_W \|x(Y, \lambda) - x_0\|^2$. In practice, (2) relies solely on the realisation y which is useful for selecting λ minimizing (2).

In this paper, we focus on Hard Thresholding (HT)

$$y \mapsto \text{HT}(y, \lambda)_i = \begin{cases} 0 & \text{if } |y_i| < \lambda, \\ y_i & \text{otherwise} \end{cases}. \quad (3)$$

HT is not even continuous, and the Stein's lemma does not apply, so that $df\{x\}(x_0, \lambda)$ and the risk cannot be unbiasedly estimated [1]. To overcome this difficulty, we build an estimator that, although biased, turns out to enjoy good asymptotic properties. In turn, this allows efficient selection of the threshold λ .

II. STEIN CONSISTENT RISK ESTIMATOR (SCORE)

Remark that the HT can be written as

$$\begin{aligned} \text{HT}(y, \lambda) &= \text{ST}(y, \lambda) + D(y, \lambda) \\ \text{where } \text{ST}(y, \lambda)_i &= \begin{cases} y_i + \lambda & \text{if } y_i < -\lambda \\ 0 & \text{if } -\lambda \leq y_i < +\lambda \\ y_i - \lambda & \text{otherwise} \end{cases} \\ \text{and } D(y, \lambda)_i &= \begin{cases} -\lambda & \text{if } y_i < -\lambda \\ 0 & \text{if } -\lambda \leq y_i < +\lambda \\ +\lambda & \text{otherwise} \end{cases}, \end{aligned}$$

where $y \mapsto \text{ST}(y, \lambda)$ is the soft thresholding operator. Soft thresholding is a Lipschitz continuous function of y with an essentially bounded gradient, and therefore, appealing to Stein's lemma, an unbiased estimator of its DOF is given by $\widehat{df}\{\text{ST}\}(Y, \lambda) = \text{div} \text{ST}(Y, \lambda)$. This DOF estimate at a realization y is known to be equal to $\#\{|y| > \lambda\}$, i.e., the number of entries of $|y|$ greater than λ (see [4], [5]). The mapping $y \mapsto D(y, \lambda)$ is piece-wise constant with discontinuities at $\pm\lambda$ so that Stein's lemma does not apply to estimate the DOF of hard thresholding. To circumvent this difficulty, we instead propose an estimator of the DOF of a smoothed version replacing $D(\cdot, \lambda)$ by $\mathcal{G}_h \star D(\cdot, \lambda)$ where \mathcal{G}_h is a Gaussian kernel of bandwidth $h > 0$ and \star is the convolution operator. In this case $\mathcal{G}_h \star D(\cdot, \lambda)$ is obviously C^∞ whose DOF can be unbiasedly estimated as $\text{div}(\mathcal{G}_h \star D(\cdot, \lambda))(Y)$. To reduce bias (this will be made clear from the proof), we have furthermore introduced a multiplicative constant, $\sqrt{\sigma^2 + h^2}/\sigma$, leading to the following DOF formula

$$y \mapsto \widehat{df}\{\text{HT}\}(y, \lambda, h) = \#\{|y| > \lambda\} + \frac{\lambda\sqrt{\sigma^2 + h^2}}{\sqrt{2\pi}\sigma h} \sum_{i=1}^P \left[\exp\left(-\frac{(y_i + \lambda)^2}{2h^2}\right) + \exp\left(-\frac{(y_i - \lambda)^2}{2h^2}\right) \right]. \quad (4)$$

We now give our two main results proved in Section IV.

Algorithm Risk estimation for Hard Thresholding

Inputs: observation $y \in \mathbb{R}^P$, threshold $\lambda > 0$

Parameters: noise variance $\sigma^2 > 0$

Output: solution x^*

Initialize $h \leftarrow \hat{h}(P)$

for all λ in the tested range **do**

 Compute $x \leftarrow \text{HT}(y, \lambda)$ using (3)

 Compute $\hat{df}\{\text{HT}\}(y, \lambda, h)$ using (4)

 Compute SCORE at y using (2)

end for

return $x^* \leftarrow x$ that provides the smallest SCORE

Fig. 1. Pseudo-algorithm for HT with SCORE-based threshold optimization.




Fig. 2. Risk and its SCORE estimate with respect to the threshold λ .

Theorem 1: Let $Y = x_0 + W$ for $W \sim \mathcal{N}(x_0, \sigma^2 \text{Id}_P)$. Take $\hat{h}(P)$ such that $\lim_{P \rightarrow \infty} \hat{h}(P) = 0$ and $\lim_{P \rightarrow \infty} P^{-1} \hat{h}(P)^{-1} = 0$. Then $\text{plim}_{P \rightarrow \infty} \frac{1}{P} \left(\hat{df}\{\text{HT}\}(Y, \lambda, \hat{h}(P)) - df\{\text{HT}\}(x_0, \lambda) \right) = 0$. In particular

1. $\lim_{P \rightarrow \infty} \mathbb{E}_W \left[\frac{1}{P} \hat{df}\{\text{HT}\}(Y, \lambda, \hat{h}(P)) \right] = \lim_{P \rightarrow \infty} \frac{1}{P} df\{\text{HT}\}(x_0, \lambda)$, and
2. $\lim_{P \rightarrow \infty} \mathbb{V}_W \left[\frac{1}{P} \hat{df}\{\text{HT}\}(Y, \lambda, \hat{h}(P)) \right] = 0$,

where \mathbb{V}_W is the variance w.r.t. W .

We now turn to a straightforward corollary of this theorem.

Corollary 1: Let $Y = x_0 + W$ for $W \sim \mathcal{N}(x_0, \sigma^2 \text{Id}_P)$, and assume that $\|x_0\|_4 = o(P^{1/2})$. Take $\hat{h}(P)$ such that $\lim_{P \rightarrow \infty} \hat{h}(P) = 0$ and $\lim_{P \rightarrow \infty} P^{-1} \hat{h}(P)^{-1} = 0$. Then, the Stein CONSistent Risk Estimator (SCORE) evaluated at a realization y of Y

$$\text{SCORE}\{x\}(y, \lambda, \hat{h}(P)) = \sum_{i=1}^P \left((y_i^2 - \sigma^2) + I(|y_i| > \lambda)(2\sigma^2 - y_i^2) + 2\sigma \frac{\lambda \sqrt{\sigma^2 + \hat{h}(P)^2}}{\sqrt{2\pi} \hat{h}(P)} \left[\exp\left(-\frac{(y_i + \lambda)^2}{2\hat{h}(P)^2}\right) + \exp\left(-\frac{(y_i - \lambda)^2}{2\hat{h}(P)^2}\right) \right] \right)$$

is such that $\text{plim}_{P \rightarrow \infty} \frac{1}{P} \left(\text{SCORE}\{x\}(Y, \lambda, \hat{h}(P)) - \mathbb{E}_W \|\text{HT}(Y, \lambda) - x_0\|^2 \right) = 0$.

Fig. 1 summarizes the pseudo-code when applying SCORE to automatically find the optimal threshold λ that minimizes SCORE in a predefined (non-empty) range.

III. EXPERIMENTS AND CONCLUSIONS

Fig. 2 shows the evolution of the true risk, the SCORE and the risk estimator of [1] as a function of λ where x_0 is a compressible vector of length $P = 2\text{E}5$ whose sorted values in magnitude decay as $|x_0|_{(i)} = 1/i^\gamma$ for $\gamma > 0$, and we have chosen σ such that the SNR of y is of about 5.65dB and $\hat{h}(P) = 6\sigma/P^{1/3} \approx \sigma/10$. The optimal λ is found around the minimum of the true risk.

Future work will concern a deeper investigation of the choice of $\hat{h}(P)$, comparison with other biased risk estimators, and extensions to other non-continuous estimators and inverse problems.

IV. PROOF

We first derive a closed-form expression for the DOF of HT.

Lemma 1: Let $Y = x_0 + W$ where $W \sim \mathcal{N}(x_0, \sigma^2 \text{Id}_P)$. The DOF of HT is given by

$$df\{\text{HT}\}(x_0, \lambda) = \left(P - \frac{1}{2} \sum_{i=1}^P \left[\text{erf}\left(\frac{(x_0)_i + \lambda}{\sqrt{2}\sigma}\right) - \text{erf}\left(\frac{(x_0)_i - \lambda}{\sqrt{2}\sigma}\right) \right] \right) + \frac{\lambda}{\sqrt{2\pi}\sigma} \sum_{i=1}^P \left[\exp\left(-\frac{((x_0)_i + \lambda)^2}{2\sigma^2}\right) + \exp\left(-\frac{((x_0)_i - \lambda)^2}{2\sigma^2}\right) \right]. \quad (5)$$

Proof: According to [1], we have

$$df\{\text{HT}\}(x_0, \lambda) = \mathbb{E}_W [\#\{|Y| > \lambda\}] + \lambda/\sigma^2 \mathbb{E}_W \left[\sum_{i=1}^P \text{sign}(Y_i) W_i I(|Y_i| > \lambda) \right]$$

where $\text{sign}(\cdot)$ is the sign function and $I(\omega)$ is the indicator for an event ω . Integrating w.r.t. to the zero-mean Gaussian density of variance σ^2 yields the closed form of the expectation terms. ■

We now turn to the proof of our theorem.

Proof: The first part of (5) corresponds to $\mathbb{E}_W[\#\{|Y| > \lambda\}]$, and can then be obviously unbiasedly estimated from an observation y by $\#\{|y| > \lambda\}$. Let A be the function defined, for $(t, a) \in \mathbb{R}^2$, by

$$A(t, a) = \frac{\sqrt{\sigma^2 + h^2}}{h} \exp\left(-\frac{(t - a)^2}{2h^2}\right).$$

By classical convolution properties of Gaussians, we have

$$\begin{aligned} \mathbb{E}_{W_i}[A(Y_i, a)] &= \exp\left(-\frac{((x_0)_i - a)^2}{2(\sigma^2 + h^2)}\right), \quad \text{and} \\ \mathbb{V}_{W_i}[A(Y_i, a)] &= \frac{\sigma^2 + h^2}{h\sqrt{2\sigma^2 + h^2}} \exp\left(-\frac{((x_0)_i - a)^2}{2\sigma^2 + h^2}\right) - \exp\left(-\frac{((x_0)_i - a)^2}{\sigma^2 + h^2}\right). \end{aligned}$$

Taking $h = \hat{h}(P)$ and assuming $\lim_{P \rightarrow \infty} \hat{h}(P) = 0$ shows that

$$\lim_{P \rightarrow \infty} \mathbb{E}_W \left[\frac{1}{P} \sum_{i=1}^P A(Y_i, a) \right] = \lim_{P \rightarrow \infty} \frac{1}{P} \sum_{i=1}^P \mathbb{E}_W [A(Y_i, a)] = \lim_{P \rightarrow \infty} \frac{1}{P} \sum_{i=1}^P \exp\left(-\frac{((x_0)_i - a)^2}{2\sigma^2}\right).$$

Since from (4), we have

$$\hat{df}\{\text{HT}\}(Y, \lambda, h) = \#\{|Y| > \lambda\} + \frac{\lambda}{\sqrt{2\pi}\sigma} \sum_{i=1}^P [A(Y_i, \lambda) + A(Y_i, -\lambda)]$$

and using Lemma 1, statement 1. follows.

For statement 2., the Cauchy-Schwartz inequality implies that

$$\mathbb{V}_W \left[\frac{1}{P} \hat{df}\{\text{HT}\}(Y, \lambda, h) \right]^{1/2} \leq \frac{\mathbb{V}_W [\#\{|Y| > \lambda\}]^{1/2}}{P} + \frac{\lambda^2}{2\pi P} \sum_{i=1}^P \left[\mathbb{V}_W [A(Y_i, \lambda)]^{1/2} + \mathbb{V}_W [A(Y_i, -\lambda)]^{1/2} \right].$$

$\#\{|Y_i| > \lambda\} \sim_{\text{iid}} \text{Bin}(P, 1 - p)$ whose variance is $Pp(1 - p)$, where $p = \frac{1}{2} \left(\text{erf}\left(\frac{(x_0)_i + \lambda}{\sqrt{2}\sigma}\right) - \text{erf}\left(\frac{(x_0)_i - \lambda}{\sqrt{2}\sigma}\right) \right)$. It follows that

$$\lim_{P \rightarrow \infty} \mathbb{V}_W \left[\frac{1}{P} \#\{|Y| > \lambda\} \right] = 0.$$

Taking again $h = \hat{h}(P)$ with $\lim_{P \rightarrow \infty} \hat{h}(P) = 0$ and $\lim_{P \rightarrow \infty} P^{-1} \hat{h}(P)^{-1} = 0$, yields

$$\lim_{P \rightarrow \infty} \mathbb{V}_W \left[\frac{1}{P} \sum_{i=1}^P A(Y_i, a) \right] = \lim_{P \rightarrow \infty} \frac{1}{P^2} \sum_{i=1}^P \mathbb{V}_W [A(Y_i, a)] = 0,$$

where we used the fact that the random variables Y_i are uncorrelated. This establishes 2.. Consistency (i.e. convergence in probability) follows from traditional arguments by invoking Chebyshev inequality and using asymptotic unbiasedness and vanishing variance established in 1. and 2.. ■

Let us now prove the corollary.

Proof: By assumption, $\lim_{P \rightarrow \infty} \hat{h}(P) = 0$. Thus by virtue of statement 1. of Theorem 1 and specializing (2) to the case of HT gives

$$\begin{aligned} \lim_{P \rightarrow \infty} \mathbb{E}_W \left[\frac{1}{P} \text{SCORE}\{\text{HT}\}(y, \lambda, \hat{h}(P)) \right] &= \lim_{P \rightarrow \infty} \frac{1}{P} \mathbb{E}_W \left[\|Y - \text{HT}(Y, \lambda)\|^2 - P\sigma^2 + 2\sigma^2 \hat{df}\{\text{HT}\}(Y, \lambda, \hat{h}(P)) \right] \\ &= \lim_{P \rightarrow \infty} \frac{1}{P} \mathbb{E}_W \|Y - \text{HT}(Y, \lambda)\|^2 - \sigma^2 + 2\sigma^2 \lim_{P \rightarrow \infty} \frac{1}{P} \mathbb{E}_W \hat{df}\{\text{HT}\}(Y, \lambda, \hat{h}(P)) \\ &= \lim_{P \rightarrow \infty} \frac{1}{P} \mathbb{E}_W \|Y - \text{HT}(Y, \lambda)\|^2 - \sigma^2 + 2\sigma^2 \lim_{P \rightarrow \infty} \frac{1}{P} \mathbb{E}_W \hat{df}\{\text{HT}\}(x_0, \lambda) \\ &= \lim_{P \rightarrow \infty} \frac{1}{P} \mathbb{E}_W \|\text{HT}(y, \lambda) - x_0\|^2 \end{aligned}$$

where we used the fact that all the limits of the expectations are finite. The Cauchy-Schwartz inequality again yields

$$\begin{aligned}
\mathbb{V}_W \left[\frac{1}{P} \text{SCORE}\{\text{HT}\}(Y, \lambda, \hat{h}(P)) \right]^{1/2} &\leq \mathbb{V}_W \left[\frac{1}{P} (\|Y - \text{HT}(Y, \lambda)\|^2) \right]^{1/2} + 2\sigma^2 \mathbb{V}_W \left[\frac{1}{P} \hat{df}\{\text{HT}\}(Y, \lambda, \hat{h}(P)) \right]^{1/2} \\
&= \frac{1}{P} \left(\sum_{i=1}^P \mathbb{V}_{W_i} [|Y_i|^2 I(|Y_i| < \lambda)] \right)^{1/2} + 2\sigma^2 \mathbb{V}_W \left[\frac{1}{P} \hat{df}\{\text{HT}\}(Y, \lambda, \hat{h}(P)) \right]^{1/2} \\
&\leq \frac{1}{P} \left(\sum_{i=1}^P \mathbb{E}_{W_i} |Y_i|^4 \right)^{1/2} + 2\sigma^2 \mathbb{V}_W \left[\frac{1}{P} \hat{df}\{\text{HT}\}(Y, \lambda, \hat{h}(P)) \right]^{1/2} \\
&= \left(\frac{\|x_0\|_4^4}{P^2} + 6\sigma^2 \frac{\|x_0\|^2}{P^2} + \frac{3\sigma^4}{P} \right)^{1/2} + 2\sigma^2 \mathbb{V}_W \left[\frac{1}{P} \hat{df}\{\text{HT}\}(Y, \lambda, \hat{h}(P)) \right]^{1/2} \\
&\leq \left(\left(\frac{\|x_0\|_4^2}{P} \right)^2 + 6\sigma^2 \frac{\|x_0\|_4^2}{P} + \frac{3\sigma^4}{P} \right)^{1/2} + 2\sigma^2 \mathbb{V}_W \left[\frac{1}{P} \hat{df}\{\text{HT}\}(Y, \lambda, \hat{h}(P)) \right]^{1/2}.
\end{aligned}$$

As by assumption, $\lim_{P \rightarrow \infty} P^{-1} \hat{h}(P)^{-1} = 0$ and $\|x_0\|_4 = o(P^{1/2})$, the variance of SCORE vanishes as $P \rightarrow \infty$. We conclude using the same convergence in probability arguments used at the end of the proof of Theorem 1. \blacksquare

REFERENCES

- [1] M. Jansen, “Information criteria for variable selection under sparsity,” Technical report, ULB, Tech. Rep., 2011.
- [2] B. Efron, “How biased is the apparent error rate of a prediction rule?” *Journal of the American Statistical Association*, vol. 81, no. 394, pp. 461–470, 1986.
- [3] C. Stein, “Estimation of the mean of a multivariate normal distribution,” *The Annals of Statistics*, vol. 9, no. 6, pp. 1135–1151, 1981.
- [4] D. Donoho and I. Johnstone, “Adapting to Unknown Smoothness Via Wavelet Shrinkage,” *Journal of the American Statistical Association*, vol. 90, no. 432, pp. 1200–1224, 1995.
- [5] H. Zou, T. Hastie, and R. Tibshirani, “On the “degrees of freedom” of the lasso,” *The Annals of Statistics*, vol. 35, no. 5, pp. 2173–2192, 2007.