Sparse Regression Learning by Aggregation and Langevin Monte-Carlo

Dalalyan, A. S.
CERTIS, Université Paris Est, Ecole des Ponts
dalalyan@certis.enpc.fr

Tsybakov, A. B.
CREST and LPMA, Université Paris 6
alexandre.tsybakov@ensae.fr

Abstract

We consider the problem of regression learning for deterministic design and independent random errors. We start by proving a sharp PAC-Bayesian type bound for the exponentially weighted aggregate (EWA) under the expected squared empirical loss. For a broad class of noise distributions the presented bound is valid whenever the temperature parameter β of the EWA is larger than or equal to $4\sigma^2$, where σ^2 is the noise variance. A remarkable feature of this result is that it is valid even for unbounded regression functions and the choice of the temperature parameter depends exclusively on the noise level.

Next, we apply this general bound to the problem of aggregating the elements of a finite-dimensional linear space spanned by a dictionary of functions $\phi_1, \dots, \phi_M$. We allow M to be much larger than the sample size n but we assume that the true regression function can be well approximated by a sparse linear combination of functions ϕ_j. Under this sparsity scenario, we propose an EWA with a heavy tailed prior and we show that it satisfies a sparsity oracle inequality with leading constant one.

Finally, we propose several Langevin Monte-Carlo algorithms to approximately compute such an EWA when the number M of aggregated functions can be large. We discuss in some detail the convergence of these algorithms and present numerical experiments that confirm our theoretical findings.

1 Introduction

In recent years a great deal of attention has been devoted to learning in high-dimensional models under the sparsity scenario. This typically assumes that, in addition to the sample, we have a finite dictionary of very large cardinality such that a small set of its elements provides a nearly complete description of the underlying model. Here, the words "large" and "small" are understood in comparison with the sample size. Sparse learning methods have been successfully applied in bioinformatics, financial engineering, image processing, etc. (see, e.g., the survey in [Yu07]).

A popular model in this context is linear regression. We observe n pairs $(\boldsymbol{X}_1, Y_1), \dots, (\boldsymbol{X}_n, Y_n)$, where each $\boldsymbol{X}_i$ – called the predictor – belongs to $\mathbb{R}^M$ and Y_i – called the response – is scalar and satisfies $Y_i = \boldsymbol{X}_i^\top \boldsymbol{\lambda}_0 + \xi_i$ with some zero-mean noise ξ_i. The goal is to develop inference on the unknown vector $\boldsymbol{\lambda}_0 \in \mathbb{R}^M$.

In many applications of linear regression the dimension of $\boldsymbol{X}_i$ is much larger than the sample size, i.e., $M \gg n$. It is well-known that in this case classical procedures, such as the least squares estimator, do not work. One of the most compelling ways for dealing with the situation where $M \gg n$ is to suppose that the sparsity assumption is fulfilled, i.e., that $\boldsymbol{\lambda}_0$ has only few coordinates different from 0. This assumption is helpful at least for two reasons: The model becomes easier to interpret and the consistent estimation of $\boldsymbol{\lambda}_0$ becomes possible if the number of non-zero coordinates is small enough.

During the last decade several learning methods exploiting the sparsity assumption have been discussed in the literature. The ℓ_1-penalized least squares (Lasso) is by far the most studied one and its statistical properties are now well understood (cf., e.g., [BRT08, BTW06, BTW07a, BTW07b, MB06, vdG08, ZH08] and the references cited therein). The Lasso is particularly attractive by its low computational cost. For instance, one can use the LARS algorithm [EHJT04], which is quite popular. Other procedures based on closely related ideas include the Elastic Net [ZH05], the Dantzig selector [CT07] and the least squares with entropy penalization [Kol08]. However, one important limitation of these procedures is that they are provably consistent under rather restrictive assumptions on the Gram matrix associated to the predictors, such as the mutual coherence assumption [DET06], the uniform uncertainty principle [CT06], the irrepresentable [ZY06] or the restricted eigenvalue [BRT08] conditions. This is somewhat unsatisfactory, since it is known that, at least in theory, there exist estimators attaining optimal accuracy of prediction under almost no assumption on the Gram matrix. This is, in particular, the case for the ℓ_0-penalized least squares estimator [BTW07a, Thm. 3.1]. However, the computation of this estimator is an NP-hard problem. We finally mention the paper [WR07], which brings to attention the fact that the empirical Bayes estimator in Gaussian regression with Gaussian prior can effectively recover the sparsity pattern. This method is realized in [WR07] via the EM alorithm. However, its theoretical properties are not explored, and it is

not clear what are the limits of application of the method beyond the considered set of numerical examples.

In [DT07, DT08] we proposed another approach to learning under the sparsity scenario, which consists in using an exponentially weighted aggregate (EWA) with a properly chosen sparsity-favoring prior. There exists an extensive literature on EWA. Some recent results focusing on the statistical properties can be found in [Alq08, Aud08, Cat07, JRT08, LB06, Yan04]. Procedures with exponential weighting received much attention in the literature on the on-line learning, see [CBCG04, HKW98, Vov90], the monograph [CBL06] and the references cited therein.

The main message of [DT07, DT08] is that the EWA with a properly chosen prior is able to deal with the sparsity issue. In particular, [DT07, DT08] prove that such an EWA satisfies a sparsity oracle inequality (SOI), which is more powerful than the best known SOI for other common procedures of sparse recovery. An important point is that almost no assumption on the Gram matrix is required. In the present work we extend this analysis in two directions. First, we prove a sharp PAC-Bayesian bound for a large class of noise distributions, which is valid for the temperature parameter depending only on the noise distribution. We impose no restriction on the values of the regression function. This result is presented in Section 2. The consequences in the context of linear regression under sparsity assumption are discussed in Section 3.

The second problem that we analyze here is the computation of EWA with the sparsity prior. Since we want to deal with large dimensions M, computation of integrals over $\mathbb{R}^M$ in the definition of this estimator can be a hard problem. Therefore, we suggest an approximation based on Langevin Monte-Carlo (LMC). This is described in detail in Section 4. Section 5 contains numerical experiments that confirm fast convergence properties of the LMC and demonstrate a nice performance of the resulting estimators.

2 PAC-Bayesian type oracle inequality

Throughout this section, as well as in Section 3, we assume that we are given the data (Z_i, Y_i), $i = 1, \dots, n$, generated by the non-parametric regression model

$$Y_i = f(Z_i) + \xi_i, \qquad i = 1, \dots, n, \tag{1}$$

with deterministic design $Z_1, \dots, Z_n$ and random errors ξ_i. We use the vector notation $\mathbf{Y} = \boldsymbol{f} + \boldsymbol{\xi}$. Thus, in what follows, the function $f(\cdot)$ is identified with the vector $\boldsymbol{f} = (f(Z_1), \dots, f(Z_n))^\top$. The space $\mathcal{Z}$ containing the design points Z_i can be arbitrary and f is a mapping from $\mathcal{Z}$ to $\mathbb{R}$. For each function $h : \mathcal{Z} \to \mathbb{R}$, we denote by $\|h\|_n$ the empirical norm $\left(\frac{1}{n}\sum_{i=1}^n h(Z_i)^2\right)^{1/2}$. Along with these notation, we will denote by $\|\boldsymbol{v}\|_p$ the ℓ_p-norm of a vector $\boldsymbol{v} = (v_1, \dots, v_n) \in \mathbb{R}^n$, that is $\|\boldsymbol{v}\|_p^p = \sum_{i=1}^n |v_i|^p$, $1 \leqslant p < \infty$, $\|\boldsymbol{v}\|_\infty = \max_i |v_i|$ and $\|\boldsymbol{v}\|_0$ is the number of nonzero entries of $\boldsymbol{v}$. With this notation, $\|\boldsymbol{f}\|_2^2 = n\|f\|_n^2$.

The noise vector $\boldsymbol{\xi} = (\xi_1, \dots, \xi_n)^\top$ is assumed to have zero mean and independent identically distributed (iid) coordinates. We introduce the following assumption on the distribution of noise.

Assumption N. For any $\gamma > 0$ small enough, there exist a probability space and two random variables ξ and ζ defined on this probability space such that

i) ξ has the same distribution as the regression errors ξ_i,

ii) $\xi + \zeta$ has the same distribution as $(1+\gamma)\xi$ and the conditional expectation $\mathbf{E}[\zeta|\xi] = 0$,

iii) there exist a real number $t_0 \in (0, \infty]$ and a bounded Borel function $v : \mathbb{R} \to \mathbb{R}_+$ such that,

$$\lim_{\gamma \to 0} \sup_{(t,a) \in [-t_0, t_0] \times \mathrm{supp}(\xi)} \frac{\log \mathbf{E}[e^{t\zeta}|\xi = a]}{t^2 \gamma v(a)} = 1,$$

where $\mathrm{supp}(\xi)$ is the support of the distribution of ξ.

Assume that we are given a collection $\{f_\lambda : \lambda \in \Lambda\}$ of functions $f_\lambda : \mathcal{Z} \to \mathbb{R}$ that will serve as building blocks for the learning procedure.The set Λ is assumed to be equipped with a σ-algebra and the mappings $\lambda \mapsto f_\lambda(z)$ are assumed to be measurable with respect to this σ-algebra for all $z \in \mathcal{Z}$. Let π be a probability measure on Λ, called the prior, and let β be a positive real number, called the temperature parameter. We define the EWA by

$$\widehat{f}_n(z) = \int_\Lambda f_\lambda(z)\, \widehat{\pi}_{n,\beta}(d\lambda),$$

where $\widehat{\pi}_{n,\beta}$ is the (posterior) probability distribution

$$\widehat{\pi}_{n,\beta}(d\lambda) \propto \exp\big\{ -\beta^{-1}\|\mathbf{Y} - \boldsymbol{f}_\lambda\|_2^2 \big\}\, \pi(d\lambda).$$

We assume that the set Λ satisfies

$$(\lambda, \lambda') \in \Lambda^2 \quad \Longrightarrow \quad \max_i |f_\lambda(Z_i) - f_{\lambda'}(Z_i)| \leqslant L \tag{2}$$

for some $L \in [0, \infty]$. In the sequel, we use the convention $\frac{+\infty}{+\infty} = 0$ and we denote by $\|v\|_\infty$ the $L_\infty(\mathbb{R})$-norm of function v.

Theorem 1 *Let Assumption N be satisfied with some function v and let (2) hold. Then for any prior π, any probability measure p on Λ and any $\beta \geqslant \max(4\|v\|_\infty, 2L/t_0)$ we have*

$$\mathbf{E}[\|\widehat{f}_n - f\|_n^2] \leqslant \int_\Lambda \|f - f_\lambda\|_n^2\, p(d\lambda) + \frac{\beta \mathcal{K}(p, \pi)}{n},$$

where $\mathcal{K}(\cdot\,,\cdot)$ stands for the Kullback-Leibler divergence.

Proof: It suffices to prove the theorem for p such that

$$\int_\Lambda \|f_\lambda - f\|_n^2\, p(d\lambda) < \infty$$

and $p \ll \pi$ (implying $\mathcal{K}(p, \pi) < \infty$), since otherwise the result is trivial.

We first assume that $\beta > 4\|v\|_\infty$ and that $L < \infty$. Let $\gamma > 0$ be a small number. Let now $(\xi_1, \zeta_1), \dots, (\xi_n, \zeta_n)$ be a sequence of iid pairs of random variables defined on a common probability space such that (ξ_i, ζ_i) satisfy conditions i)-iii) of Assumption N for any i. The existence of these random variables is ensured by Assumption N. We use here the same notation ξ_i as in model (1), since it causes no ambiguity.

Set $\boldsymbol{h}_\lambda = f_\lambda - f$, $\widehat{h} = \widehat{f}_n - f$, $\boldsymbol{\zeta} = (\zeta_1, \dots, \zeta_n)^\top$, $U(\boldsymbol{h}, \boldsymbol{h}') = \|\boldsymbol{h}\|_2^2 + 2\boldsymbol{h}^\top \boldsymbol{h}'$ and $\Delta U(\boldsymbol{h}, \boldsymbol{h}', \boldsymbol{h}'') = (\|\boldsymbol{h}\|_2^2 -$

$\|\boldsymbol{h}'\|_2^2) + 2(\boldsymbol{h}-\boldsymbol{h}')^\top \boldsymbol{h}''$ for any pair $\boldsymbol{h}, \boldsymbol{h}', \boldsymbol{h}'' \in \mathbb{R}^n$. With this notation we have

$$\mathbf{E}[\|\widehat{f}_n - f\|_n^2] = \mathbf{E}[\|\widehat{h}\|_n^2] = \mathbf{E}\Big[\|\widehat{h}\|_n^2 + \frac{2}{n\gamma}\widehat{\boldsymbol{h}}^\top \boldsymbol{\zeta}\Big].$$

Therefore, $\mathbf{E}[\|\widehat{f}_n - f\|_n^2] = S + S_1$, where

$$S = -\frac{\beta}{n\gamma}\mathbf{E}\Big[\log\int_\Lambda \exp\Big(-\frac{\gamma U(\boldsymbol{h}_\lambda, \gamma^{-1}\boldsymbol{\zeta})}{\beta}\Big)\,\widehat{\pi}_{n,\beta}(d\lambda)\Big],$$

$$S_1 = \frac{\beta}{n\gamma}\mathbf{E}\Big[\log\int_\Lambda \exp\Big(-\frac{\gamma \Delta U(\boldsymbol{h}_\lambda, \widehat{\boldsymbol{h}}, \gamma^{-1}\boldsymbol{\zeta})}{\beta}\Big)\,\widehat{\pi}_{n,\beta}(d\lambda)\Big].$$

We first bound the term S. To this end, note that

$$\widehat{\pi}_{n,\beta}(d\lambda) = \frac{\exp\{-\beta^{-1}U(\boldsymbol{h}_\lambda, \boldsymbol{\xi})\}}{\int_\Lambda \exp\{-\beta^{-1}U(\boldsymbol{h}_w, \boldsymbol{\xi})\}\pi(dw)}\,\pi(d\lambda)$$

and therefore

$$\begin{aligned} S \;=\; & \frac{\beta}{n\gamma}\mathbf{E}\Big[\log\int_\Lambda \exp\big\{-\tfrac{1}{\beta}U(\boldsymbol{h}_\lambda, \boldsymbol{\xi})\big\}\pi(d\lambda)\Big] \\ & -\frac{\beta}{n\gamma}\mathbf{E}\Big[\log\int_\Lambda \exp\big\{-\tfrac{1+\gamma}{\beta}U\big(\boldsymbol{h}_\lambda, \tfrac{\boldsymbol{\xi}+\boldsymbol{\zeta}}{1+\gamma}\big)\big\}\pi(d\lambda)\Big]. \end{aligned}$$

By part ii) of Assumption N and the independence of vectors (ξ_i, ζ_i) for different values of i, the probability distribution of the vector $(\boldsymbol{\xi}+\boldsymbol{\zeta})/(1+\gamma)$ coincides with that of $\boldsymbol{\xi}$. Therefore, $(\boldsymbol{\xi}+\boldsymbol{\zeta})/(1+\gamma)$ may be replaced by $\boldsymbol{\xi}$ inside the second expectation. Now, using the Hölder inequality, we get

$$S \leqslant -\frac{\beta}{n(1+\gamma)}\mathbf{E}\Big[\log\int_\Lambda e^{-(1+\gamma)\beta^{-1}U(\boldsymbol{h}_\lambda, \boldsymbol{\xi})}\pi(d\lambda)\Big].$$

Next, by a convex duality argument [Cat04, p. 160], we find

$$S \leqslant \int_\Lambda \|h_\lambda\|_n^2\, p(d\lambda) + \frac{\beta\mathcal{K}(p,\pi)}{n(1+\gamma)}\,.$$

Let us now bound the term S_1. According to part iii) of Assumption N, there exists $\gamma_0 > 0$ such that $\forall \gamma \leqslant \gamma_0$,

$$\sup_{|t|\leqslant t_0}\frac{\log \mathbf{E}[e^{t\zeta}|\xi = a]}{t^2\gamma} \leqslant v(a)(1+o_\gamma(1)), \quad \forall\, a \in \mathbb{R}.$$

In what follows we assume that $\gamma \leqslant \gamma_0$. Since for every i, $|2\beta^{-1}(h_\lambda(Z_i) - \widehat{h}(Z_i))| \leqslant 2\beta^{-1}L \leqslant t_0$, using Jensen's inequality we get

$$\begin{aligned} S_1 &\leqslant \frac{\beta}{n\gamma}\mathbf{E}\Big[\log\int_\Lambda \exp\Big\{-\frac{n\gamma}{\beta}(\|h_\lambda\|_n^2 - \|\widehat{h}\|_n^2)\Big\}\,\theta_\lambda \\ &\qquad \times\mathbf{E}\Big(\exp\Big\{\sum_{i=1}^n 2\beta^{-1}(h_\lambda(Z_i) - \widehat{h}(Z_i))\zeta_i\Big\}\,|\boldsymbol{\xi}\Big)\pi(d\lambda)\Big] \\ &\leqslant \frac{\beta}{n\gamma}\mathbf{E}\Big[\log\int_\Lambda \exp\Big\{-\frac{n\gamma}{\beta}(\|h_\lambda\|_n^2 - \|\widehat{h}\|_n^2)\Big\}\,\theta_\lambda \\ &\qquad \times\exp\Big\{\frac{4n\|v\|_\infty\gamma}{\beta^2}\|h_\lambda - \widehat{h}\|_n^2(1+o_\gamma(1))\Big\}\,\pi(d\lambda)\Big]. \end{aligned}$$

For γ small enough ($\gamma \leqslant \tilde{\gamma}_0$), this entails that up to a positive multiplicative constant, the term S_1 is bounded by the expression $\mathbf{E}\big[\log\int_\Lambda \exp\big(-\frac{n\gamma V(h_\lambda,\widehat{h})}{\beta^2}\big)\theta_\lambda\pi(d\lambda)\big]$, where

$$V(h_\lambda, \widehat{h}) = \beta(\|h_\lambda\|_n^2 - \|\widehat{h}\|_n^2) + \frac{(\beta + 4\|v\|_\infty)}{2}\|h_\lambda - \widehat{h}\|_n^2.$$

Using [DT07, Lemma 3] and Jensen's inequality we obtain $S_1 \leqslant 0$ for any $\gamma \leqslant (\beta - 4\|v\|_\infty)/4nL$. Thus, we proved that

$$\mathbf{E}[\|\widehat{h}\|_n^2] \leqslant \int_\Lambda \|h_\lambda\|_n^2\, p\,(d\lambda) + \frac{\beta\,\mathcal{K}(p,\pi)}{n(1+\gamma)}$$

for any $\gamma \leqslant \tilde{\gamma}_0 \wedge (\beta - 4\|v\|_\infty)/4nL$. Letting γ tend to zero, we obtain

$$\mathbf{E}[\|\widehat{h}\|_n^2] \leqslant \int_\Lambda \|h_\lambda\|_n^2 p(d\lambda) + \frac{\beta\,\mathcal{K}(p,\pi)}{n}$$

for any $\beta > \max(4\|v\|_\infty, 2L/t_0)$. Fatou's lemma allows us to extend this inequality to the case $\beta = \max(4\|v\|_\infty, 2L/t_0)$.

To cover the case $L = +\infty, t_0 = +\infty$, we fix some $L_0 \in (0,\infty)$ and apply the obtained inequality to the truncated prior $\pi^{L'}(d\lambda) \propto \mathbb{1}_{\Lambda_{L'}}(\lambda)\pi(d\lambda)$, where $L' \in (L_0,\infty)$ and $\Lambda_{L'} = \{\lambda \in \Lambda : \max_i |f_\lambda(Z_i)| \leqslant L'\}$. We obtain that for any measure $p \ll \pi$ supported by Λ_{L_0},

$$\begin{aligned} \mathbf{E}[\|\widehat{h}^{L'}\|_n^2] &\leqslant \int_\Lambda \|h_\lambda\|_n^2\, p(d\lambda) + \frac{\beta\mathcal{K}(p,\pi^{L'})}{n} \\ &\leqslant \int_\Lambda \|h_\lambda\|_n^2\, p(d\lambda) + \frac{\beta\mathcal{K}(p,\pi)}{n}. \end{aligned}$$

One easily checks that $\widehat{h}^{L'}$ tends a.s. to $\widehat{h}$ and that the random variable $\sup_{L'>L_0}\|\widehat{h}^{L'}\|_n^2\mathbb{1}(\max_i|\xi_i| \leqslant C)$ is integrable for any fixed C. Therefore, by Lebesgue's dominated convergence theorem we get

$$\mathbf{E}[\|\widehat{h}\|_n^2\mathbb{1}(\max_i|\xi_i| \leqslant C)] \leqslant \int_\Lambda \|h_\lambda\|_n^2\, p(d\lambda) + \frac{\beta\mathcal{K}(p,\pi)}{n}.$$

Letting C tend to infinity and using Lebesgue's monotone convergence theorem we obtain the desired inequality for any probability measure p which is absolutely continuous w.r.t. π and is supported by Λ_{L_0} for some $L_0 > 0$. If $p(\Lambda_{L_0}) < 1$ for any $L_0 > 0$, one can replace p by its truncated version $p^{L'}$ and use Lebesgue's monotone convergence theorem to get the desired result. ■

An important point is that many symmetric distributions encountered in applications satisfy Assumption N with $v(a)$ being identically equal to the variance of the noise. This follows from the next remarks and their combinations.

Remark 1 (Gaussian noise) *If ξ_1 is drawn according to the Gaussian distribution $\mathcal{N}(0,\sigma^2)$, then for any $\gamma > 0$ one can choose ζ independently of ξ according to the Gaussian distribution $\mathcal{N}(0,(2\gamma+\gamma^2)\sigma^2)$. This results in $v(a) \equiv \sigma^2$ and, as a consequence, Theorem 1 holds for any $\beta \geqslant 4\sigma^2$. Note that this reduces to the Leung and Barron's [LB06] result if the prior π is discrete.*

Remark 2 (Rademacher noise) *If ξ_1 is drawn according to the Rademacher distribution, i.e. $\mathbf{P}(\xi_1 = \pm\sigma) = 1/2$, then for any $\gamma > 0$ one can define ζ as follows:*

$$\zeta = (1+\gamma)\sigma\,\mathrm{sgn}[\sigma^{-1}\xi - (1+\gamma)U] - \xi,$$

where U is distributed uniformly in $[-1,1]$ and is independent of ξ. This results in $v(a) \equiv \sigma^2$ and, as a consequence, Theorem 1 holds for any $\beta \geqslant 4\sigma^2 = 4\mathbf{E}[\xi_1^2]$.

Remark 3 (Stability by convolution) *Assume that ξ_1 and ξ_1' are two independent random variables. If ξ_1 and ξ_1' satisfy Assumption N with $t_0 = \infty$ and with functions $v(a)$ and $v'(a)$, then any linear combination $\alpha\xi_1 + \alpha'\xi_1'$ satisfies Assumption N with $t_0 = \infty$ and the v-function $\alpha^2 v(a) + (\alpha')^2 v'(a)$.*

Remark 4 (Uniform distribution) *The claim of preceding remark can be generalized to linear combinations of a countable set of random variables, provided that the series converges in the mean squared sense. In particular, if ξ_1 is drawn according to the symmetric uniform distribution with variance σ^2, then Assumption N is fulfilled with $t_0 = \infty$ and $v(a) \equiv \sigma^2$. This can be proved using the fact that ξ_1 has the same distribution as $\sum_{i=1}^{\infty} 2^{-i}\eta_i$, where η_i are iid Rademacher random variables. Thus, in this case the inequality of Theorem 1 is true for any $\beta \geqslant 4\sigma^2$.*

Remark 5 (Laplace noise) *If ξ_1 is drawn according to the Laplace distribution with variance σ^2, then for any $\gamma > 0$ one can choose ζ independently of ξ according to the distribution associated to the characteristic function*

$$\varphi(t) = \frac{1}{(1+\gamma)^2}\Big(1 + \frac{2\gamma+\gamma^2}{1+(1+\gamma)^2(\sigma t)^2/2}\Big).$$

One can observe that the distribution of ζ is a mixture of the Dirac distribution at zero and the Laplace distribution with variance $(1+\gamma)^2\sigma^2$. This results in $v(a) \equiv 2\sigma^2/(2-\sigma^2 t_0^2)$ and, as a consequence, by taking $t_0 = 1/\sigma^2$, we get that Theorem 1 holds for any $\beta \geqslant \max(8\sigma^2, 2L\sigma)$.

Remark 6 (Bounded symmetric noise) *Assume that the errors ξ_i are symmetric and that $P(|\xi_i| \leqslant B) = 1$ for some $B \in (0,\infty)$. Let $U \sim \mathcal{U}([-1,1])$ be a random variable independent of ξ. Then, $\zeta = (1+\gamma)|\xi|\,\mathrm{sgn}[\mathrm{sgn}(\xi) - (1+\gamma)U] - \xi$ satisfies Assumption N with $v(a) = a^2$. Since $\|v\|_\infty \leqslant B^2$, we obtain that Theorem 1 is valid for any $\beta \geqslant 4B^2$.*

3 Sparsity prior and SOI

In this section we introduce the sparsity prior and present a sparsity oracle inequality (SOI) derived from Theorem 1.

In what follows we assume that $\Lambda \subset \mathbb{R}^M$ for some positive integer M. We will use boldface letters to denote vectors and, in particular, the elements of Λ. For any square matrix A, let $\mathrm{Tr}(\mathsf{A})$ denote the trace (sum of diagonal entries) of A. Furthermore, we focus on the particular case where $\mathcal{F}_\Lambda$ is a convex subset of the vector space spanned by a finite number of measurable functions $\{\phi_j\}_{j=1,\dots,M}$. More specifically, we assume that, for some $R \in (0,+\infty]$,

$$\mathcal{F}_\Lambda = \Big\{ f_{\boldsymbol{\lambda}} = \sum_{j=1}^M \lambda_j \phi_j \;\Big|\; \boldsymbol{\lambda} \in \mathbb{R}^M \text{ satisfies } \|\boldsymbol{\lambda}\|_1 \leqslant R \Big\},$$

where $\|\boldsymbol{\lambda}\|_1 = \sum_j |\lambda_j|$ stands for the ℓ_1-norm. If, in addition, $f \in \mathcal{F}_\Lambda$, then model (1) reduces to that of linear regression defined in the Introduction. Indeed, it suffices to take

$$\boldsymbol{X}_i = (\phi_1(Z_i), \dots, \phi_M(Z_i))^\top, \quad i = 1, \dots, n.$$

This notation will be used in the rest of the paper along with the assumption that $\boldsymbol{X}_i$ are normalized so that all the diagonal entries of matrix $\frac{1}{n}\sum_{i=1}^n \boldsymbol{X}_i \boldsymbol{X}_i^\top$ are equal to one.

We allow M to be large, possibly much larger than the sample size n. If $M \gg n$, we have in mind that the sparsity assumption holds, i.e., there exists $\boldsymbol{\lambda}^* \in \mathbb{R}^M$ such that f in (1) is close to $f_{\boldsymbol{\lambda}^*}$ for some $\boldsymbol{\lambda}^*$ having only a small number of non-zero entries. We handle this situation via a suitable choice of prior π. Namely, we use a modification of the sparsity prior proposed in [DT07]. It should be emphasized right away that we will take advantage of sparsity for the purpose of prediction and not for data compression. In fact, even if the underlying model is sparse, we do not claim that our estimator is sparse as well, but we claim that it is quite accurate under very mild assumptions. On the other hand, some simulations demonstrate the sparsity of our estimator and the fact that it recovers correctly the true sparsity pattern in examples where the (restrictive) assumptions mentioned in the Introduction are satisfied (cf. Section 5). However, our theoretical results do not deal with this property.

To specify the sparsity prior π we need the Huber function $\bar\omega : \mathbb{R} \to \mathbb{R}$ defined by

$$\bar\omega(t) = \begin{cases} t^2, & \text{if } |t| \leqslant 1 \\ 2|t| - 1, & \text{otherwise.} \end{cases}$$

This function behaves very much like the absolute value of t, but has the advantage of being differentiable at every point $t \in \mathbb{R}$. Let τ and α be positive numbers. We define the *sparsity prior*

$$\pi(d\boldsymbol{\lambda}) = \frac{\tau^{2M}}{C_{\alpha,\tau,R}} \Big\{ \prod_{j=1}^M \frac{e^{-\bar\omega(\alpha\lambda_j)}}{(\tau^2+\lambda_j^2)^2} \Big\} \mathbb{1}(\|\boldsymbol{\lambda}\|_1 \leqslant R)\, d\boldsymbol{\lambda}, \quad (3)$$

where $C_{\alpha,\tau,R}$ is the normalizing constant.

Since the sparsity prior (3) looks somewhat complicated, an heuristical explanation is in order. Let us assume that R is large and α is small so that the functions $e^{-\bar\omega(\alpha\lambda_j)}$ and $\mathbb{1}(\|\boldsymbol{\lambda}\|_1 \leqslant R)$ are approximately equal to one. With this in mind, we can notice that π is close to the distribution of $\sqrt{2}\tau\mathbf{Y}$, where $\mathbf{Y}$ is a random vector having iid coordinates drawn from Student's t-distribution with three degrees of freedom. In the examples below we choose a very small τ, smaller than $1/n$. Therefore, most of the coordinates of $\tau\mathbf{Y}$ are very close to zero. On the other hand, since Student's t-distribution has heavy tails, a few coordinates of $\tau\mathbf{Y}$ are quite far from zero.

These heuristics are illustrated by Figure 1 presenting the boxplots of one realization of a random vector in $\mathbb{R}^{10.000}$ with iid coordinates drawn from the scaled Gaussian, Laplace (double exponential) and Student $t(3)$ distributions. The scaling factor is such that the probability densities of the simulated distributions are equal to 100 at the origin. The boxplot which is most likely to represent a sparse vector corresponds to Student's $t(3)$ distribution.

The relevance of heavy tailed priors for dealing with sparsity has been emphasized by several authors (see [See08, Section 2.1] and references therein). However, most of this work focused on logarithmically concave priors, such as the multivariate Laplace distribution. Also in wavelet estimation on classes of "sparse" functions [JS05] and [Riv06] invoke quasi-Cauchy and Pareto priors. Bayes estimators with heavy-tailed priors in sparse Gaussian shift models are discussed in [AGP07].

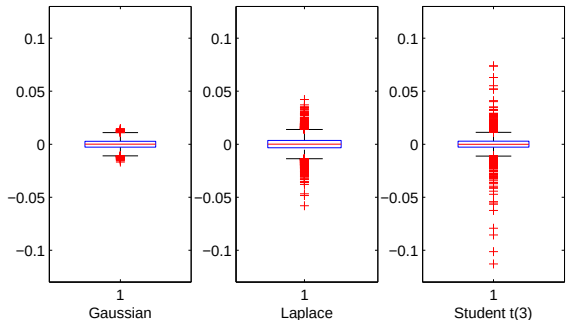

Figure 1: The boxplots of a sample of size 10^4 drawn from the scaled Gaussian (left panel), Laplace (central panel) and Student $t(3)$ (right panel) distributions. In all three cases the location parameter is 0 and the scale parameter is 10^{-2}.

The next theorem provides a SOI for the EWA with the sparsity prior (3).

Theorem 2 *Let Assumption N be satisfied with some function v and let (2) hold. Take the prior π defined in (3) and $\beta \geqslant \max(4\|v\|_\infty, 2L/t_0)$. Assume that $R > 2M\tau$ and $\alpha \leqslant 1/(4M\tau)$. Then for all $\boldsymbol{\lambda}^*$ such that $\|\boldsymbol{\lambda}^*\|_1 \leqslant R - 2M\tau$ we have*

$$\mathbf{E}[\|\widehat{f}_n - f\|_n^2] \leqslant \|f_{\boldsymbol{\lambda}^*} - f\|_n^2 + \frac{4\beta}{n}\sum_{j=1}^{M}\log\left(1+\tfrac{|\lambda_j^*|}{\tau}\right) + \frac{2\beta(\alpha\|\boldsymbol{\lambda}^*\|_1+1)}{n} + 4e\tau^2 M. \qquad (4)$$

Proof: Let us define the probability measure p_0 by

$$\frac{dp_0}{d\boldsymbol{\lambda}}(\boldsymbol{\lambda}) \propto \left(\frac{d\pi}{d\boldsymbol{\lambda}}(\boldsymbol{\lambda}-\boldsymbol{\lambda}^*)\right)\mathbb{1}_{B_1(2M\tau)}(\boldsymbol{\lambda}-\boldsymbol{\lambda}^*). \qquad (5)$$

Since $\|\boldsymbol{\lambda}^*\|_1 \leqslant R - 2M\tau$, the condition $\boldsymbol{\lambda} - \boldsymbol{\lambda}^* \in B_1(2M\tau)$ implies that $\boldsymbol{\lambda} \in B_1(R)$ and, therefore, p_0 is absolutely continuous w.r.t. the sparsity prior π. In view of Thm. 1, we have

$$\mathbf{E}[\|\widehat{f}_n - f\|_n^2] \leqslant \int_\Lambda \|f_{\boldsymbol{\lambda}} - f\|_n^2\, p_0(d\boldsymbol{\lambda}) + \frac{\beta\mathcal{K}(p_0,\pi)}{n}.$$

The facts that $f_{\boldsymbol{\lambda}}(Z_i) = \boldsymbol{X}_i^\top\boldsymbol{\lambda}$ and the diagonal entries of the matrix $\frac{1}{n}\sum_i \boldsymbol{X}_i\boldsymbol{X}_i^\top$ are equal to one, together with the symmetry of p_0 w.r.t. $\boldsymbol{\lambda}^*$ yield

$$\int_\Lambda \|f_{\boldsymbol{\lambda}} - f\|_n^2\, p_0(d\boldsymbol{\lambda}) = \|f_{\boldsymbol{\lambda}^*} - f\|_n^2 + \int_{\mathbb{R}^M}\|\boldsymbol{\lambda}-\boldsymbol{\lambda}^*\|_2^2\, p_0(d\boldsymbol{\lambda}).$$

To complete the proof, we use the following technical result.

Lemma 3 *For every integer M larger than 1, we have:*

$$\int_{\mathbb{R}^M}(\lambda_1 - \lambda_1^*)^2 p_0(d\boldsymbol{\lambda}) \leqslant 4\tau^2 e^{4M\alpha\tau},$$

$$\mathcal{K}(p_0,\pi) \leqslant 2(\alpha\|\boldsymbol{\lambda}^*\|_1 + 1) + 4\sum_{j=1}^{M}\log(1+|\lambda_j^*|/\tau).$$

The proof of this lemma is omitted.

Inequality (4) follows from Lemma 3, since $\int_{\mathbb{R}^M}\|\boldsymbol{\lambda} - \boldsymbol{\lambda}^*\|_2^2\, p_0(d\boldsymbol{\lambda}) = M\int_{\mathbb{R}^M}(\lambda_1 - \lambda_1^*)^2\, p_0(d\boldsymbol{\lambda})$ and, under the assumptions of the theorem, $e^{4M\alpha\tau} \leqslant e$. ■

Theorem 2 can be used to choose the tuning parameters τ, α, R when $M \gg n$. The idea is to choose them such that both terms in the second line of (4) were of the order $O(1/n)$. This can be achieved, for example, by taking $\tau^2 \sim (Mn)^{-1}$ and $R = O(M\tau)$. Then the term $\frac{4\beta}{n}\sum_{j=1}^{M}\log(1+|\lambda_j^*|/\tau)$ becomes dominating. It is important that the number of nonzero summands in this term, M^*, is equal to the number of nonzero coordinates of $\boldsymbol{\lambda}^*$. Therefore, for sparse vectors $\boldsymbol{\lambda}^*$, this term is rather small, namely of the order $M^*(\log M)/n$, which is the optimal rate for problems of sparse recovery, cf. [BTW06, CT07, BTW07a, BRT08]. An important difference compared with these and other papers on ℓ_1-based sparse recovery is that in Theorem 2, we have no assumption on the dictionary $\{\phi_1, \ldots, \phi_M\}$.

4 Computation of the EW-aggregate by the Langevin Monte-Carlo

In this section we suggest Langevin Monte-Carlo (LMC) procedures to approximately compute the EWA with the sparsity prior when $M \gg n$.

4.1 Langevin Diffusion in continuous time

We start by describing a continuous-time Markov process, called the Langevin diffusion, that will play the key role in this section. Let $V : \mathbb{R}^M \to \mathbb{R}$ be a smooth function, which in what follows will be referred to as potential. We will assume that the gradient of V is locally Lipschitz and is at most of linear growth. This ensures that the stochastic differential equation (SDE)

$$d\boldsymbol{L}_t = \nabla V(\boldsymbol{L}_t)\,dt + \sqrt{2}\,d\boldsymbol{W}_t,\ \boldsymbol{L}_0 = \boldsymbol{\lambda}_0,\ t \geqslant 0 \qquad (6)$$

has a unique strong solution, called the Langevin diffusion. In the last display, $\boldsymbol{W}$ stands for an M-dimensional Brownian motion and $\boldsymbol{\lambda}_0$ is an arbitrary deterministic vector from $\mathbb{R}^M$. It is well known that the process $\{\boldsymbol{L}_t\}_{t\geqslant 0}$ is a homogeneous Markov process and a semimartingale, cf. [RW87, Thm. 12.1].

As a Markov process, $\boldsymbol{L}$ may be transient, null recurrent or positively recurrent. The latter case, which is the most important for us, corresponds to the process satisfying the law of large numbers and implies the existence of a stationary distribution. In other terms, if $\boldsymbol{L}$ is positively recurrent, there exists a probability distribution P_V on $\mathbb{R}^M$ such that the process $\boldsymbol{L}$ is stationary provided that the initial condition $\boldsymbol{\lambda}_0$ is drawn at random according P_V. A remarkable property of the Langevin diffusion—making it very attractive for computing high-dimensional integrals—is that its stationary distribution, if exists, has the density

$$p_V(\boldsymbol{\lambda}) \propto e^{V(\boldsymbol{\lambda})}, \qquad \boldsymbol{\lambda} \in \mathbb{R}^M,$$

w.r.t. the Lebesgue measure [Ken78, Thm. 10.1]. Furthermore, there exist directly verifiable conditions on the potential V yielding the positive recurrence of $\boldsymbol{L}$. The following proposition contains an example of such a condition.

Proposition 1 ([RS02], Thm 2.1) *Assume that the function V is bounded from above. If there is a twice continuously differentiable function $D : \mathbb{R}^M \to [1, \infty)$ and three positive constants a, b and r such that*

$$\nabla V(\boldsymbol{\lambda})^\top \nabla D(\boldsymbol{\lambda}) + \Delta D(\boldsymbol{\lambda}) \leqslant -aD(\boldsymbol{\lambda}) + b\mathbb{1}(\|\boldsymbol{\lambda}\|_2 \leqslant r), \tag{7}$$

for every $\boldsymbol{\lambda} \in \mathbb{R}^M$, then the Langevin diffusion $\boldsymbol{L}$ defined by (6) is D-geometrically ergodic, that is

$$\Big|\mathbf{E}[h(\boldsymbol{L}_t)|\boldsymbol{L}_0 = \boldsymbol{\lambda}_0] - \int_{\mathbb{R}^M} h(\boldsymbol{\lambda})\, p_V(d\boldsymbol{\lambda})\Big| \leqslant R_V D(\boldsymbol{\lambda}_0)\rho_V^t$$

for every function h satisfying $\|h/D\|_\infty \leqslant 1$ and for some constants $R_V > 0$ and $\rho_V \in (0, 1)$.

The function D satisfying condition (7) is often referred to as Lyapunov function and condition (7) is called drift condition towards the set $\{\boldsymbol{\lambda} : \|\boldsymbol{\lambda}\|_2 \leqslant r\}$. If satisfied, the drift condition ensures geometrical mixing. Specifically, for every function h such that $\|h^2/D\|_\infty \leqslant 1$ and for every $t, s > 0$,

$$\big|\mathbf{Cov}_{\boldsymbol{\lambda}_0}[h(\boldsymbol{L}_t), h(\boldsymbol{L}_s)]\big| \leqslant R_V D(\boldsymbol{\lambda}_0)\rho_V^{|t-s|}.$$

Combining this with the result of Proposition 1 it is not hard to check that if $\|h^2/D\|_\infty \leqslant 1$, then

$$\mathbf{E}_{\boldsymbol{\lambda}_0}\Big[\Big(\frac{1}{T}\int_0^T h(\boldsymbol{L}_t)dt - \int_{\mathbb{R}^M} h(\boldsymbol{\lambda})p_V(d\boldsymbol{\lambda})\Big)^2\Big] \leqslant \frac{C}{T}, \tag{8}$$

where C is some positive constant depending only on V. Note also that, in view of Proposition 1, the squared bias term in the bias-variance decomposition of the left hand side of (8) is of order $O(T^{-2})$. Thus, the main error term comes from the stochastic part.

4.2 Langevin diffusion associated to EWA

We assume that we are given $(\boldsymbol{X}_i, Y_i)$, $i = 1, \dots, n$, with $\boldsymbol{X}_i \in \mathbb{R}^M$ and $Y_i \in \mathbb{R}$. We wish to compute the expression

$$\widehat{\boldsymbol{\lambda}} = \frac{\int_{\mathbb{R}^M} \boldsymbol{\lambda} \exp\big\{-\beta^{-1}\|\mathbf{Y} - \mathbb{X}\boldsymbol{\lambda}\|_2^2\big\}\pi(d\boldsymbol{\lambda})}{\int_{\mathbb{R}^M} \exp\big\{-\beta^{-1}\|\mathbf{Y} - \mathbb{X}\boldsymbol{\lambda}\|_2^2\big\}\pi(d\boldsymbol{\lambda})}, \tag{9}$$

where $\mathbb{X} = (\boldsymbol{X}_1, \dots, \boldsymbol{X}_n)^\top$. In what follows, we deal with the prior

$$\pi(d\boldsymbol{\lambda}) \propto \prod_{j=1}^M \frac{e^{-\bar{\omega}(\alpha\lambda_j)}}{(\tau^2 + \lambda_j^2)^2}$$

assuming that $R = +\infty$. As proved in Sections 2 and 3, this choice leads to a sharp oracle inequality for a number of noise distributions. An equivalent form for writing (9) is $\widehat{\boldsymbol{\lambda}} = \int_{\mathbb{R}^M} \boldsymbol{\lambda} p_V(\boldsymbol{\lambda})\, d\boldsymbol{\lambda}$, $p_V(\boldsymbol{\lambda}) \propto e^{V(\boldsymbol{\lambda})}$ with $V(\boldsymbol{\lambda})$ being equal to

$$-\frac{\|\mathbf{Y} - \mathbb{X}\boldsymbol{\lambda}\|_2^2}{\beta} - \sum_{j=1}^M \Big\{2\log(\tau^2 + \lambda_j^2) + \bar{\omega}(\alpha\lambda_j)\Big\}. \tag{10}$$

Simple computations show that $D(\boldsymbol{\lambda}) = e^{\alpha\|\boldsymbol{\lambda}\|_2}$ is a function for which the drift condition (7) is fulfilled. A nice property of this Lyapunov function is the inequality $\|\boldsymbol{\lambda}\|_\infty^2 \leqslant \alpha^{-1}D(\boldsymbol{\lambda})$. It guarantees that (8) is satisfied for the functions $h(\boldsymbol{\lambda}) = \lambda_i$. So, let us define the Langevin diffusion $\boldsymbol{L}_t$ as solution of (6) with the potential V given in (10) and the initial condition $\boldsymbol{L}_0 = 0$. In what follows we will consider only this particular diffusion process. We define the average value

$$\bar{\boldsymbol{L}}_T = \frac{1}{T}\int_0^T \boldsymbol{L}_t\, dt, \qquad T \geqslant 0.$$

According to (8) this average value converges as $T \to \infty$ to the vector $\widehat{\boldsymbol{\lambda}}$ that we want to compute. Clearly, it is much easier to compute $\bar{\boldsymbol{L}}_T$ than $\widehat{\boldsymbol{\lambda}}$. Indeed, $\widehat{\boldsymbol{\lambda}}$ involves integrals in M dimensions, whereas $\bar{\boldsymbol{L}}_T$ is a one-dimensional integral over a finite interval. Of course, to compute such an integral one needs to discretize the Langevin diffusion. This is done in the next subsection.

4.3 Discretization

Since the sample paths of a diffusion process are Hölder continuous, it is easy to show that the Riemann sum approximation

$$\bar{\boldsymbol{L}}_T^R = \frac{1}{T}\sum_{i=0}^{N-1} \boldsymbol{L}_{T_i}\,(T_{i+1} - T_i),$$

with $0 = T_0 < T_1 < \ldots < T_N = T$ converges to $\bar{\boldsymbol{L}}_T$ in mean square when the sampling is sufficiently dense, that is when $\max_i |T_{i+1} - T_i|$ is small. However, when simulating the diffusion sample path in practice, it is impossible to follow exactly the dynamics determined by (6). Usually one needs to discretize the SDE to make the computation of its solution possible.

A natural discretization for the SDE (6) is proposed by the Euler scheme with a constant step of discretization $h > 0$, defined as

$$\boldsymbol{L}_{k+1}^E = \boldsymbol{L}_k^E + h\nabla V(\boldsymbol{L}_k^E) + \sqrt{2h}\,\boldsymbol{W}_k,\ \boldsymbol{L}_0^E = 0, \tag{11}$$

for $k = 0, 1, \dots, [T/h] - 1$, where $\boldsymbol{W}_1, \boldsymbol{W}_2, \dots$ are i.i.d. standard Gaussian random vectors in $\mathbb{R}^M$ and $[x]$ stands for the integer part of $x \in \mathbb{R}$. Obviously, the sequence $\{\boldsymbol{L}_k; k \geqslant 0\}$ defines a discrete-time Markov process. Furthermore, one can show that this Markov process can be extrapolated to a continuous-time diffusion-type process which converges in distribution to the Langevin diffusion as $h \to 0$. Here extrapolation means the construction of a process $\{\tilde{\boldsymbol{L}}_t^E; t \in [0, T]\}$ satisfying $\tilde{\boldsymbol{L}}_{kh}^E = \boldsymbol{L}_k^E$ for every $k = 0, \dots, [T/h]$. Such a process $\tilde{\boldsymbol{L}}^E$ can be defined as a solution of the SDE

$$d\tilde{\boldsymbol{L}}^E = \sum_{k=0}^{[T/h]-1} \mathbb{1}_{[k,k+1)}(t/h)\nabla V(\boldsymbol{L}_k^E)\, dt + \sqrt{2}\, d\boldsymbol{W}_t,\ t \geqslant 0.$$

This amounts to connecting the successive values of the Markov chain by independent Brownian bridges. The Girsanov formula implies that the Kullback-Leibler divergence of the distribution of the process $\{\boldsymbol{L}_t; t \in [0, T]\}$ from the distribution of $\{\tilde{\boldsymbol{L}}_t^E; t \in [0, T]\}$ tends to zero as h tends to zero. Therefore, it makes sense to approximate $\bar{\boldsymbol{L}}_T$ by

$$\bar{\boldsymbol{L}}_{T,h}^E = \frac{1}{[T/h]}\sum_{k=0}^{[T/h]-1} \boldsymbol{L}_k^E.$$

This discretization algorithm is easily implementable and, for small values of h, $\bar{\boldsymbol{L}}_{T,h}^E$ is very close to the integral $\widehat{\boldsymbol{\lambda}} =$

$\int \boldsymbol{\lambda}\, p_V(\boldsymbol{\lambda})\, d\boldsymbol{\lambda}$ of interest. However, for some values of h, which may eventually be small but not enough, the Markov process $\{\boldsymbol{L}_k^E; k \geqslant 0\}$ is transient and therefore, the sum in the definition of $\bar{\boldsymbol{L}}_{T,h}^E$ explodes. To circumvent this problem, one can either modify the Markov chain $\boldsymbol{L}_k^E$ by incorporating a Metropolis-Hastings correction, or take a smaller h and restart the computations. The former approach has the advantage of guaranteeing the convergence to the desired distribution. However, it considerably slows down the algorithm because of a significant probability of rejection at each step of discretization. Of course, the second approach, where we just take a smaller h, also slows down the algorithm but at least we keep some control on its time of execution.

5 Implementation and experimental results

In this section we give more details on the implementation of the LMC for computing the EW-aggregate in the linear regression model.

5.1 Implementation

The input of the algorithm we are going to describe is the triplet $(\mathbf{Y}, \mathbb{X}, \sigma)$ and the tuning parameters $(\alpha, \beta, \tau, h, T)$, where

- $\mathbf{Y}$ is the n-vector of values of the response variable,
- $\mathbb{X}$ is the $n \times M$ matrix of predictor variables,
- σ is the noise level,
- β is the temperature parameter of the EW-aggregate,
- α and τ are the parameters of the sparsity prior,
- h and T are the parameters of the LMC algorithm.

The output of the proposed algorithm is a vector $\widehat{\boldsymbol{\lambda}} \in \mathbb{R}^M$ such that, for every $\boldsymbol{x} \in \mathbb{R}^M$, $\boldsymbol{x}^\top \widehat{\boldsymbol{\lambda}}$ provides a prediction for the unobservable value of the response variable corresponding to $\boldsymbol{x}$. The pseudo-code of the algorithm is given below.

Input: Observations $(\mathbf{Y}, \mathbb{X}, \sigma)$ and parameters $(\alpha, \beta, \tau, h, T)$
Output: The vector $\widehat{\boldsymbol{\lambda}}$
Set

```
[n,M]=size(X);
L=zeros(M,1);
lambda=zeros(M,1);
H=0;
```

Calculate

```
XX=X'*X;
Xy=X'*y;
```

while `H` *is less than* `T` **do**

```
nablaV=(2/β)*(Xy-XX*L)-α*ω̄'(αL);
nablaV=nablaV-4*L./(τ^2+L.^2);
L=L+h*nablaV+sqrt(2*h)*randn(M,1);
H=H+h;
lambda=lambda+h*L/T;
```

end
return `lambda`

5.1.1 Choice of T

Since the convergence rate of $\bar{\boldsymbol{L}}_T$ to $\widehat{\boldsymbol{\lambda}}$ is of the order $T^{-1/2}$ and the best rate of convergence an estimator can achieve is $n^{-1/2}$, it is natural to set $T = n$. This choice of T has the advantage of being simple for implementation, but it has the drawback of being not scale invariant. A better strategy for choosing T is to continue the procedure until the convergence is observed.

5.1.2 Choice of h

We choose the step of discretization in the form:

$$h = \beta/(Mn) = \beta/\mathrm{Tr}(\mathbb{X}^\top \mathbb{X}).$$

More details on the choice of h and T will be given in a future work.

5.1.3 Choice of β, τ and α

In our simulations we use the parameter values

$$\alpha = 0, \qquad \beta = 4\sigma^2, \qquad \tau = 4\sigma/(\mathrm{Tr}(\mathbb{X}^\top \mathbb{X}))^{1/2}\,.$$

These values of β and τ are derived from the theory developed above. However, we take here $\alpha = 0$ and not $\alpha > 0$ as suggested in Section 3. We introduced there $\alpha > 0$ for theoretical convenience, in order to guarantee the geometric mixing of the Langevin diffusion. Numerous simulations show that mixing properties of the Langevin diffusion are preserved with $\alpha = 0$ as well.

5.2 Numerical experiments

We present below two examples of application of the EWA with LMC for simulated data sets. In both examples we give also the results obtained by the Lasso procedure (rather as a benchmark, than for comparing the two procedures). The main goal of this section is to illustrate the predictive ability of the EWA and to show that it can be easily computed for relatively large dimensions of the problem. In all examples, the Lasso estimators are computed with the tuning parameter equal to $\sigma\sqrt{8 \log M/n}$ (cf. [BRT08]).

5.2.1 Example 1

Consider the model $\mathbf{Y} = \mathbb{X}\boldsymbol{\lambda}^* + \sigma\boldsymbol{\xi}$, where $\mathbb{X}$ is a $M \times n$ matrix with independent entries, such that each entry is a Rademacher random variable. Such matrices are particularly well suited for applications in compressed sensing. The noise $\boldsymbol{\xi} \in \mathbb{R}^n$ is a vector of independent standard Gaussian random variables. The vector $\boldsymbol{\lambda}^*$ is chosen to be S-sparse, where S is much smaller than M. W. l. o. g. we consider vectors $\boldsymbol{\lambda}^*$ such that only first S coordinates are different from 0; more precisely, $\lambda_j^* = \mathbb{1}(j \leqslant S)$. Following [CT07], we choose $\sigma^2 = S/9$. We run our procedure for several values of S and M. The results of 500 replications are summarized in Table 1. A typical scatterplot of estimated coefficients for $M = 500$, $n = 200$ and $S = 20$ is presented in Fig. 2.

An interesting observation is that the EWA selects the set of nonzero coordinates of $\boldsymbol{\lambda}^*$ even better than the Lasso does. In fact, the approximate sparsity of the EWA is not very surprising, since in the noise-free linear models with orthogonal matrix $\mathbb{X}$, the symmetry of the prior implies that the EWA estimates the zero coordinates without error.

5.2.2 Example 2

Consider model (1) where Z_i are independent random variables uniformly distributed in the unit square $[0, 1]^2$ and ξ_i are iid $\mathcal{N}(0, \sigma^2)$ random variables. For an integer $k > 0$,

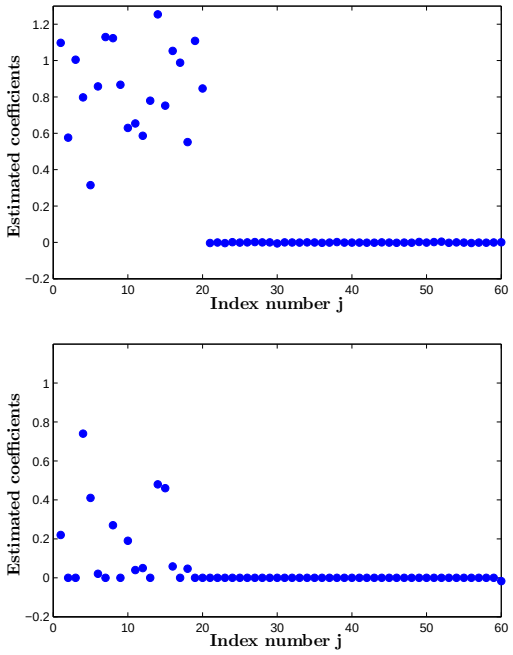

Figure 2: A typical result of the EWA (top panel) and the Lasso (bottom panel) in the setup of Example 1 with $n = 200$, $M = 500$ and $S = 20$. For better visibility, only 60 first coefficients are plotted.

	$M = 200$		$M = 500$	
	EWA	**Lasso**	**EWA**	**Lasso**
$n = 200$	0.022	0.657	0.019	0.746
$S = 5$	(0.013)	(0.155)	(0.012)	(0.157)
$n = 200$	0.109	2.790	0.117	3.299
$S = 10$	(0.052)	(0.617)	(0.051)	(0.730)
$n = 200$	1.790	11.70	3.045	12.990
$S = 20$	(1.467)	(1.672)	(2.015)	(1.607)

Table 1: Average loss $\|\widehat{\boldsymbol{\lambda}} - \boldsymbol{\lambda}^*\|^2$ of the estimators obtained by the EW-aggregate and the Lasso in Example 1. The standard deviation is given in parentheses.

we consider the indicator functions of rectangles with sides parallel to the axes and having as left-bottom vertex the origin and as right-top vertex a point of the form $(i/k, j/k)$, $(i, j) \in \mathbb{N}^2$. Formally, we define ϕ_j by

$$\phi_{(i-1)k+j}(x) = \mathbb{1}_{[0,i]\times[0,j]}(kx), \qquad \forall\, x \in [0,1]^2.$$

The underlying image f we are trying to recover is taken as a superposition of a small number of rectangles of this form, that is $f(x) = \sum_{\ell=1}^{k^2} \lambda^*_\ell \phi_\ell(x)$, for all $x \in [0,1]^2$ with some $\boldsymbol{\lambda}^*$ having a small ℓ_0-norm. We set $k = 15$, $\|\boldsymbol{\lambda}^*\|_0 = 3$, $\lambda^*_{10} = \lambda^*_{100} = \lambda^*_{200} = 1$. Thus, the cardinality of the dictionary is $M = k^2 = 225$.

In this example the functions ϕ_j are strongly correlated and therefore the assumptions like restricted isometry or low coherence are not fulfilled. Nevertheless, the Lasso succeeds in providing an accurate prediction (cf. Table 2). Furthermore, the Lasso with the theoretically justified choice of the

	EWA	**Lasso**	**Ideal LG**
$\sigma = 2, n = 100$	0.210	0.759	0.330
$T = 1$	(0.072)	(0.562)	(0.145)
$\sigma = 4, n = 100$	0.420	2.323	0.938
$T = 1$	(0.222)	(1.257)	(0.631)
$\sigma = 2, n = 200$	0.187	0.661	0.203
$T = 1$	(0.048)	(0.503)	(0.086)
$\sigma = 4, n = 200$	0.278	2.230	0.571
$T = 1$	(0.132)	(1.137)	(0.324)

Table 2: Average loss $\int_{[0,1]^2} \left(\sum_j (\widehat{\lambda}_j - \lambda^*_j)\phi_j(x)\right)^2 dx$ of the the EWA, the Lasso and the ideal LG procedures in Example 2. The standard deviation is given in parentheses.

smoothing parameter $\sigma\sqrt{8 \log M/n}$ is not much worse than the ideal Lasso-Gauss (LG) estimator. We call the LG estimator the ordinary least squares estimator in the reduced model where only the predictor variables selected at a preliminary Lasso step are kept. Of course, the performance of the LG procedure depends on the initial choice of the tuning parameter for the Lasso step. In our simulations, we use its ideal (oracle) value minimizing the prediction error and, therefore, we call the resulting procedure the ideal LG estimator.

As expected, the EWA has a smaller predictive risk than the Lasso estimator. A very good news is the supremacy of the EWA with respect to the ideal LG. Of course, the LG procedure is faster. However, even from this point of view the EWA is rather attractive, since it takes less than two seconds to compute it in the present example.

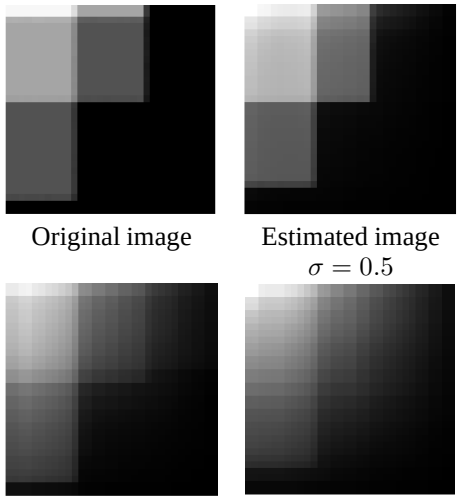

Original image | Estimated image $\sigma = 0.5$

Estimated image $\sigma = 1$ | Estimated image $\sigma = 2$

Figure 3: The original image (left panel) and the EWA estimated image from noisy observations with $\sigma = 0.5$, $\sigma = 1$ and $\sigma = 2$, respectively. In all cases $n = 200$ and $k = 15$.

6 Conclusion and outlook

This paper contains two contributions: New oracle inequalities for EWA, and the LMC method for approximate computation of the EWA. The first oracle inequality presented

in this work is in the line of the PAC-Bayesian bounds initiated by McAllester [McA03]. It is valid for any prior distribution and gives a bound on the risk of the EWA with an arbitrary family of functions. Next, we derive another inequality, which is adapted to the sparsity scenario and called the sparsity oracle inequality (SOI). In order to obtain it, we propose a prior distribution favoring sparse representations. The resulting EWA is shown to behave almost as well as the best possible linear combination within a residual term proportional to $M^*(\log M)/n$, where M is the true dimension, M^* is the number of atoms entering in the best linear combination and n is the sample size. A remarkable fact is that this inequality is obtained under no condition on the relationship between different atoms.

Sparsity oracle inequalities similar to that of Theorem 2 are valid for the penalized empirical risk minimizers (ERM) with a ℓ_0-penalty (proportional to the number of atoms involved in the representation). It is also well known that the problem of computing the ℓ_0-penalized ERM is NP-hard. In contrast with this, we have shown that the numerical evaluation of the suggested EWA is a computationally tractable problem. We demonstrated that it can be efficiently solved by the LMC algorithm. Numerous simulations we did (some of which are included in this work) confirm our theoretical findings and, furthermore, suggest that the EWA is able to efficiently select the sparsity pattern. Theoretical justification of this fact, as well as more thorough investigation of the choice of parameters involved in the LMC algorithm, are interesting topics for future research.

References

[AGP07] F. Abramovich, V. Grinshtein, and M. Pensky. On optimality of Bayesian testimation in the normal means problem. *Ann. Statist.*, 35:2261–2286, 2007.

[Alq08] P. Alquier. Pac-Bayesian bounds for randomized empirical risk minimizers. *Math. Methods Statist.*, 17(4):1–26, 2008.

[Aud08] J.-Y. Audibert. Fast learning rates in statistical inference through aggregation, 2008. Ann. Statist., to appear.

[BRT08] P. Bickel, Y. Ritov, and A. B. Tsybakov. Simultaneous analysis of Lasso and Dantzig selector, 2008. To appear in The Annals of Statistics.

[BTW06] F. Bunea, A. B. Tsybakov, and M. Wegkamp. Aggregation and sparsity via l_1 penalized least squares. In *Learning theory*, volume 4005 of *Lecture Notes in Comput. Sci.*, pages 379–391. Springer, Berlin, 2006.

[BTW07a] F. Bunea, A. B. Tsybakov, and M. Wegkamp. Aggregation for Gaussian regression. *Ann. Statist.*, 35(4):1674–1697, 2007.

[BTW07b] F. Bunea, A.B. Tsybakov, and M.H. Wegkamp. Sparsity oracle inequalities for the Lasso. *Electronic J. of Statist.*, 1:169–194, 2007.

[Cat04] O. Catoni. *Statistical learning theory and stochastic optimization*. Lecture Notes in Mathematics. Springer-Verlag, Berlin, 2004.

[Cat07] O. Catoni. *Pac-Bayesian Supervised Classification: The Thermodynamics of Statistical Learning*, volume 56. IMS Lecture Notes Monograph Series, December 03 2007.

[CBCG04] N. Cesa-Bianchi, A. Conconi, and C. Gentile. On the generalization ability of on-line learning algorithms. *IEEE Trans. Inform. Theory*, 50(9):2050–2057, 2004.

[CBL06] N. Cesa-Bianchi and G. Lugosi. *Prediction, learning, and games*. Cambridge University Press, Cambridge, 2006.

[CT06] E. Candès and T. Tao. Near-optimal signal recovery from random projections: universal encoding strategies? *IEEE Trans. Inform. Theory*, 52(12):5406–5425, 2006.

[CT07] E. Candès and T. Tao. The Dantzig selector: statistical estimation when p is much larger than n. *Ann. Statist.*, 35(6):2313–2351, 2007.

[DET06] D. Donoho, M. Elad, and V. Temlyakov. Stable recovery of sparse overcomplete representations in the presence of noise. *IEEE Trans. Inform. Theory*, 52(1):6–18, 2006.

[DT07] A. S. Dalalyan and A. B. Tsybakov. Aggregation by exponential weighting and sharp oracle inequalities. In *Learning theory*, volume 4539 of *Lecture Notes in Comput. Sci.*, pages 97–111. Springer, Berlin, 2007.

[DT08] A. S. Dalalyan and A. B. Tsybakov. Aggregation by exponential weighting, sharp PAC-Bayesian bounds and sparsity. *Mach. Learn.*, 72(1-2):39–61, 2008.

[EHJT04] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani. Least angle regression. *Ann. Statist.*, 32(2):407–499, 2004.

[HKW98] D. Haussler, J. Kivinen, and M. Warmuth. Sequential prediction of individual sequences under general loss functions. *IEEE Trans. Inform. Theory*, 44(5):1906–1925, 1998.

[JRT08] A. Juditsky, Ph. Rigollet, and A.B. Tsybakov. Learning by mirror averaging. *Ann. Statist.*, 36:2183–2206, 2008.

[JS05] I. Johnstone and B.W. Silverman. Empirical Bayes selection of wavelet thresholds. *Ann. Statist*, 33:1700–1752, 2005.

[Ken78] J. Kent. Time-reversible diffusions. *Adv. in Appl. Probab.*, 10(4):819–835, 1978.

[Kol08] V. Koltchinskii. Sparse recovery in convex hulls via entropy penalization, 2008. Ann. Statist., to appear.

[LB06] G. Leung and A. Barron. Information theory and mixing least-squares regressions. *IEEE Trans. Inform. Theory*, 52(8):3396–3410, 2006.

[MB06] N. Meinshausen and P. Bühlmann. High-dimensional graphs and variable selection with the Lasso. *Ann. Statist.*, 34(3):1436–1462, 2006.

[McA03] D. McAllester. Pac-Bayesian stochastic model selection. *Machine Learning*, 51(1):5–21, 2003.

[Riv06] V. Rivoirard. Non linear estimation over weak Besov spaces and minimax Bayes method. *Bernoulli*, 12(4):609–632, 2006.

[RS02] G. Roberts and O. Stramer. Langevin diffusions and Metropolis-Hastings algorithms. *Methodol.*

Comput. Appl. Probab., 4(4):337–357, 2002.

[RW87] L. Rogers and D. Williams. *Diffusions, Markov processes, and martingales. Vol. 2*. Probability and Mathematical Statistics. John Wiley & Sons Inc., New York, 1987.

[See08] M. W. Seeger. Bayesian inference and optimal design for the sparse linear model. *J. Mach. Learn. Res.*, 9:759–813, 2008.

[vdG08] S. van de Geer. High-dimensional generalized linear models and the Lasso. *Ann. Statist.*, 36(2):614–645, 2008.

[Vov90] V. Vovk. Aggregating strategies. In *COLT: Proceedings of the Workshop on Computational Learning Theory, Morgan Kaufmann Publishers*, pages 371–386, 1990.

[WR07] David P. Wipf and Bhaskar D. Rao. An empirical Bayesian strategy for solving the simultaneous sparse approximation problem. *IEEE Trans. Signal Process.*, 55(7, part 2):3704–3716, 2007.

[Yan04] Y. Yang. Aggregating regression procedures to improve performance. *Bernoulli*, 10(1):25–47, 2004.

[Yu07] B. Yu. Embracing statistical challenges in the information technology age. *Technometrics*, 49(3):237–248, 2007.

[ZH05] H. Zou and T. Hastie. Regularization and variable selection via the elastic net. *J. R. Stat. Soc. Ser. B*, 67(2):301–320, 2005.

[ZH08] C.H. Zhang and J. Huang. The sparsity and biais of the Lasso selection in high-dimensional linear regression. *Ann. Statist.*, 36:1567–1594, 2008.

[ZY06] P. Zhao and B. Yu. On model selection consistency of Lasso. *J. Mach. Learn. Res.*, 7:2541–2563, 2006.