



HAL
open science

Sharp Oracle Inequalities for Low-complexity Priors

Tung Duy Luu, Jalal Fadili, Christophe Chesneau

► **To cite this version:**

Tung Duy Luu, Jalal Fadili, Christophe Chesneau. Sharp Oracle Inequalities for Low-complexity Priors. 2017. hal-01422476v4

HAL Id: hal-01422476

<https://hal.science/hal-01422476v4>

Preprint submitted on 1 Sep 2017 (v4), last revised 27 Sep 2018 (v5)

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.

Sharp Oracle Inequalities for Low-complexity Priors

Tung Duy Luu*

Jalal Fadili*

Christophe Chesneau†

Abstract

In this paper, we consider a high-dimensional statistical estimation problem in which the number of parameters is comparable or larger than the sample size. We present a unified analysis of the performance guarantees of exponential weighted aggregation and penalized estimators with a general class of data losses and priors which encourage objects which conform to some notion of simplicity/complexity. More precisely, we show that these two estimators satisfy sharp oracle inequalities for prediction ensuring their good theoretical performances. We also highlight the differences between them. When the noise is random, we provide oracle inequalities in probability using concentration inequalities. These results are then applied to several instances including the Lasso, the group Lasso, their analysis-type counterparts, the ℓ_∞ and the nuclear norm penalties. All our estimators can be efficiently implemented using proximal splitting algorithms.

Key words. High-dimensional estimation, exponential weighted aggregation, penalized estimation, oracle inequality, low-complexity models.

AMS subject classifications. 62G07 62G20

1 Introduction

1.1 Problem statement

Our statistical context is the following. Let $\mathbf{y} = (\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_n)$ be n identically distributed observations with common marginal distribution, and $\mathbf{X} \in \mathbb{R}^{n \times p}$ a deterministic design matrix. The goal is to estimate a parameter vector $\boldsymbol{\theta} \in \mathbb{R}^p$ of the observations marginal distribution based on the data \mathbf{y} and \mathbf{X} .

Let $F : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ be a loss function supposed to be smooth and convex that assigns to each $\boldsymbol{\theta} \in \mathbb{R}^p$ a cost $F(\mathbf{X}\boldsymbol{\theta}, \mathbf{y})$. Let $\boldsymbol{\theta}_0 \in \text{Argmin}_{\boldsymbol{\theta} \in \mathbb{R}^p} \mathbb{E} [F(\mathbf{X}\boldsymbol{\theta}, \mathbf{y})]$ be any minimizer of the population risk. We regard $\boldsymbol{\theta}_0$ as the true parameter. A usual instance of this statistical setting is the standard linear regression model based on n pairs $(\mathbf{y}_i, \mathbf{X}_i)$ of response-covariate that are linked linearly $\mathbf{y} = \mathbf{X}\boldsymbol{\theta}_0 + \boldsymbol{\xi}$, and $F(\mathbf{u}, \mathbf{y}) = \frac{1}{2} \|\mathbf{y} - \mathbf{u}\|_2^2$.

Our goal is to provide general oracle inequalities in prediction for two estimators of $\boldsymbol{\theta}_0$: the penalized estimator and exponential weighted aggregation. Whenever $p > n$, the estimation problem is ill-posed since the rectangular matrix \mathbf{X} has a kernel of dimension at least $p - n$. To circumvent this difficulty, we will exploit the prior that $\boldsymbol{\theta}_0$ has some low-complexity structure (among which sparsity and low-rank are the most popular). That is, even if the ambient dimension p of $\boldsymbol{\theta}_0$ is very large, its intrinsic dimension is much smaller than the sample size n . This makes it possible to build estimates $\mathbf{X}\hat{\boldsymbol{\theta}}$ with good provable performance guarantees under appropriate conditions. There has been a flurry of research on the use of low-complexity regularization in ill-posed recovery problems in various areas including statistics and machine learning.

*Normandie Univ, ENSICAEN, CNRS, GREYC, France, Email: {duy-tung.luu, Jalal.Fadili}@ensicaen.fr.

†Normandie Univ, UNICAEN, CNRS, LMNO, France, Email: christophe.chesneau@unicaen.fr.

1.2 Variational/Penalized Estimators

Regularization is now a central theme in many fields including statistics, machine learning and inverse problems. It allows one to impose on the set of candidate solutions some prior structure on the object to be estimated. This regularization ranges from squared Euclidean or Hilbertian norms to non-Hilbertian norms (e.g. ℓ_1 norm for sparse objects, or nuclear norm for low-rank matrices) that have sparked considerable interest in the recent years. In this paper, we consider the class of estimators obtained by solving the convex optimization problem¹

$$\hat{\boldsymbol{\theta}}_n^{\text{PEN}} \in \underset{\boldsymbol{\theta} \in \mathbb{R}^p}{\text{Argmin}} \{V_n(\boldsymbol{\theta}) \stackrel{\text{def}}{=} \frac{1}{n}F(\mathbf{X}\boldsymbol{\theta}, \mathbf{y}) + \lambda_n J(\boldsymbol{\theta})\}, \quad (1.1)$$

where the regularizing penalty J is a proper closed convex function that promotes some specific notion of simplicity/low-complexity, and $\lambda_n > 0$ is the regularization parameter. A prominent member covered by (1.1) is the Lasso [13, 52, 39, 22, 8, 5, 7, 31] and its variants such the analysis/fused Lasso [47, 53], SLOPE [6, 49] or group Lasso [2, 67, 1, 65]. Another example is the nuclear norm minimization for low rank matrix recovery motivated by various applications including robust PCA, phase retrieval, control and computer vision [41, 10, 26, 11]. See [37, 7, 61, 58] for generalizations and comprehensive reviews.

1.3 Exponential Weighted Aggregation (EWA)

An alternative to the the variational estimator (1.1) is the aggregation by exponential weighting, which consists in substituting averaging for minimization. The aggregators are defined via the probability density function

$$\mu_n(\boldsymbol{\theta}) = \frac{\exp(-V_n(\boldsymbol{\theta})/\beta)}{\int_{\Theta} \exp(-V_n(\boldsymbol{\omega})/\beta)d\boldsymbol{\omega}}, \quad (1.2)$$

where $\beta > 0$ is called temperature parameter. If all $\boldsymbol{\theta}$ are candidates to estimate the true vector $\boldsymbol{\theta}_0$, then $\Theta = \mathbb{R}^p$. The aggregate is thus defined by

$$\hat{\boldsymbol{\theta}}_n^{\text{EWA}} = \int_{\mathbb{R}^p} \boldsymbol{\theta} \mu_n(\boldsymbol{\theta}) d\boldsymbol{\theta}. \quad (1.3)$$

Aggregation by exponential weighting has been widely considered in the statistical and machine learning literatures, see e.g. [19, 16, 15, 20, 38, 66, 42, 33, 27, 24] to name a few. $\hat{\boldsymbol{\theta}}_n^{\text{EWA}}$ can also be interpreted as the posterior conditional mean in the Bayesian sense if $F/(n\beta)$ is the negative-loglikelihood associated to the noise $\boldsymbol{\xi}$ with the prior density $\pi(\boldsymbol{\theta}) \propto \exp(-\lambda_n J(\boldsymbol{\theta})/\beta)$.

1.4 Oracle inequalities

Oracle inequalities, which are at the heart of our work, quantify the quality of an estimator compared to the best possible one among a family of estimators. These inequalities are well adapted in the scenario where the prior penalty promotes some notion of low-complexity (e.g. sparsity, low rank, etc.). Given two vectors $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$, let $R_n(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2)$ be a nonnegative error measure between their predictions, respectively $\mathbf{X}\boldsymbol{\theta}_1$ and $\mathbf{X}\boldsymbol{\theta}_2$. A popular example is the averaged prediction squared error $\frac{1}{n} \|\mathbf{X}\boldsymbol{\theta}_1 - \mathbf{X}\boldsymbol{\theta}_2\|_2^2$, where $\|\cdot\|_2$ is the ℓ_2 norm. R_n will serve as a measure of the performance of the estimators $\hat{\boldsymbol{\theta}}_n^{\text{EWA}}$ and $\hat{\boldsymbol{\theta}}_n^{\text{PEN}}$. More precisely,

¹To avoid trivialities, the set of minimizers is assumed non-empty, which holds for instance if J is also coercive.

we aim to prove that $\hat{\boldsymbol{\theta}}_n^{\text{EWA}}$ and $\hat{\boldsymbol{\theta}}_n^{\text{PEN}}$ mimic as much as possible the best possible model. This idea is materialized in the following type of inequalities (stated here for EWA)

$$R_n(\hat{\boldsymbol{\theta}}_n^{\text{EWA}}, \boldsymbol{\theta}_0) \leq C \inf_{\boldsymbol{\theta} \in \mathbb{R}^p} (R_n(\boldsymbol{\theta}, \boldsymbol{\theta}_0) + \Delta_{n,p,\lambda_n,\beta}(\boldsymbol{\theta})), \quad (1.4)$$

where $C \geq 1$ is the leading constant of the oracle inequality and the remainder term $\Delta_{n,p,\lambda_n,\beta}(\boldsymbol{\theta})$ depends on the performance of the estimator, the complexity of $\boldsymbol{\theta}$, the sample size n , the dimension p , and the regularization and temperature parameters (λ_n, β) . An estimator with good oracle properties would correspond to C close to 1 (ideally, $C = 1$, in which case the inequality is said “sharp”), and $\Delta_{n,p,\lambda_n,\beta}(\boldsymbol{\theta})$ is small and decreases rapidly to 0 as $n \rightarrow +\infty$.

1.5 Contributions

We provide a unified analysis where we capture the essential ingredients behind the low-complexity priors promoted by J , relying on sophisticated arguments from convex analysis and our previous work [25, 57, 59, 56, 58]. Our main contributions are summarized as follows:

- We show that the EWA estimator $\hat{\boldsymbol{\theta}}_n^{\text{EWA}}$ in (1.2) and the variational/penalized estimator $\hat{\boldsymbol{\theta}}_n^{\text{PEN}}$ in (1.1) satisfy (deterministic) sharp oracle inequalities for prediction with optimal remainder term, for general data losses F beyond the usual quadratic one, and J is a proper finite-valued sublinear function (i.e. J is finite-valued convex and positively homogeneous). We also highlight the differences between the two estimators in terms of the corresponding bounds.
- When the observations are random, we prove oracle inequalities in probability. The theory is non-asymptotic in nature, as it yields explicit bounds that hold with high probability for finite sample sizes, and reveals the dependence on dimension and other structural parameters of the model.
- For the standard linear model with Gaussian or sub-Gaussian noise, and a quadratic loss, we deliver refined versions of these oracle inequalities in probability. We underscore the role of the Gaussian width, a concept that captures important geometric characteristics of sets in \mathbb{R}^n .
- These results yield naturally a large number of corollaries when specialized to penalties routinely used in the literature, among which the Lasso, the group Lasso, their analysis-type counterparts (fused (group) Lasso), the ℓ_∞ and the nuclear norms. Soem of these corollaries are known and others novel.

The estimators $\hat{\boldsymbol{\theta}}_n^{\text{EWA}}$ and $\hat{\boldsymbol{\theta}}_n^{\text{PEN}}$ can be easily implemented thanks to the framework of proximal splitting methods, and more precisely forward-backward type splitting. While the latter is well-known to solve (1.1) [58], its application within a proximal Langevin Monte-Carlo algorithm to compute $\hat{\boldsymbol{\theta}}_n^{\text{EWA}}$ with provable guarantees has been recently developed by the authors in [24] to sample from log-semiconcave densities², see also [23] for log-concave densities.

1.6 Relation to previous work

Our oracle inequality for $\hat{\boldsymbol{\theta}}_n^{\text{EWA}}$ extends the work of [17] with an unprecedented level of generality, far beyond the Lasso and the nuclear norm. Our prediction sharp oracle inequality for $\hat{\boldsymbol{\theta}}_n^{\text{PEN}}$ specializes to that of [50]

²In a forthcoming paper, this framework was extended to cover the even more general class of prox-regular functions.

in the case of the Lasso (see also the discussion in [18] and references therein) and that of [32] for the case of the nuclear norm. Our work also goes much beyond that in [61] on weakly decomposable priors, where we show in particular that there is no need to impose decomposability on the regularizer, since it is rather an intrinsic property of it.

1.7 Paper organization

Section 2 states our main assumptions on the data loss and the prior penalty. All the concepts and notions are exemplified on some penalties some of which are popular in the literature. In Section 3, we prove our main oracle inequalities, and their versions in probability. We then tackle the case of linear regression with quadratic data loss in Section 4. Concepts from convex analysis that are essential to this work are gathered in Section A. A key intermediate result in the proof of our main results is established in Section B with an elegant argument relying on Moreau-Yosida regularization.

1.8 Notations

Vectors and matrices For a d -dimensional Euclidean space \mathbb{R}^d , we endow it with its usual inner product $\langle \cdot, \cdot \rangle$ and associated norm $\|\cdot\|_2$. \mathbf{Id}_d is the identity matrix on \mathbb{R}^d . For $p \geq 1$, $\|\cdot\|_p$ will denote the ℓ_p norm of a vector with the usual adaptation for $p = +\infty$.

In the following, if T is a vector space, P_T denotes the orthogonal projector on T , and

$$\boldsymbol{\theta}_T = P_T \boldsymbol{\theta} \quad \text{and} \quad \mathbf{X}_T = \mathbf{X} P_T.$$

For a finite set \mathcal{C} we denote $|\mathcal{C}|$ its cardinality. For $I \subset \{1, \dots, p\}$, we denote by I^c its complement. $\boldsymbol{\theta}_I$ is the subvector whose entries are those of $\boldsymbol{\theta}$ restricted to the indices in I , and \mathbf{X}_I the submatrix whose columns are those of \mathbf{X} indexed by I . For any matrix \mathbf{X} , \mathbf{X}^\top denotes its transpose and \mathbf{X}^+ its Moore-Penrose pseudo-inverse. For a linear operator \mathbf{A} , \mathbf{A}^* is its adjoint.

Sets For a nonempty set $\mathcal{C} \in \mathbb{R}^p$, we denote $\overline{\text{conv}}(\mathcal{C})$ the closure of its convex hull, and $\iota_{\mathcal{C}}$ its indicator function, i.e. $\iota_{\mathcal{C}}(\boldsymbol{\theta}) = 0$ if $\boldsymbol{\theta} \in \mathcal{C}$ and $+\infty$ otherwise. For a nonempty convex set \mathcal{C} , its *affine hull* $\text{aff}(\mathcal{C})$ is the smallest affine manifold containing it. It is a translate of its *parallel subspace* $\text{par}(\mathcal{C})$, i.e. $\text{par}(\mathcal{C}) = \text{aff}(\mathcal{C}) - \boldsymbol{\theta} = \mathbb{R}(\mathcal{C} - \boldsymbol{\theta})$; for any $\boldsymbol{\theta} \in \mathcal{C}$. The *relative interior* $\text{ri}(\mathcal{C})$ of a convex set \mathcal{C} is the interior of \mathcal{C} for the topology relative to its affine hull.

Functions A function $f : \mathbb{R}^p \rightarrow \mathbb{R} \cup \{+\infty\}$ is closed (or lower semicontinuous) if so is its epigraph. It is coercive if $\lim_{\|\boldsymbol{\theta}\|_2 \rightarrow +\infty} f(\boldsymbol{\theta}) = +\infty$, and strongly coercive if $\lim_{\|\boldsymbol{\theta}\|_2 \rightarrow +\infty} f(\boldsymbol{\theta}) / \|\boldsymbol{\theta}\|_2 = +\infty$. The effective domain of f is $\text{dom}(f) = \{\boldsymbol{\theta} \in \mathbb{R}^p : f(\boldsymbol{\theta}) < +\infty\}$ and f is proper if $\text{dom}(f) \neq \emptyset$ as is the case when it is finite-valued. A function is said sublinear if it is convex and positively homogeneous. The Legendre-Fenchel conjugate of f is $f^*(\mathbf{z}) = \sup_{\boldsymbol{\theta} \in \mathbb{R}^p} \langle \mathbf{z}, \boldsymbol{\theta} \rangle - f(\boldsymbol{\theta})$. For f proper, the functions (f, f^*) obey the Fenchel-Young inequality

$$f(\boldsymbol{\theta}) + f^*(\mathbf{z}) \geq \langle \mathbf{z}, \boldsymbol{\theta} \rangle, \quad \forall (\boldsymbol{\theta}, \mathbf{z}) \in \mathbb{R}^p \times \mathbb{R}^p. \quad (1.5)$$

When f is a proper lower semicontinuous and convex function, (f, f^*) is actually the best pair for which this inequality cannot be tightened. For a function g on \mathbb{R}_+ , the function $g^+ : a \in \mathbb{R}_+ \mapsto g^+(a) = \sup_{t \geq 0} at - g(t)$ is called the monotone conjugate of g . The pair (g, g^+) obviously obeys (1.5) on $\mathbb{R}_+ \times \mathbb{R}_+$.

For a C^1 -smooth function f , $\nabla f(\boldsymbol{\theta})$ is its (Euclidean) gradient. For a bivariate function $g : (\boldsymbol{\eta}, \mathbf{y}) \in \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ that is C^2 with respect to the first variable $\boldsymbol{\eta}$, for any \mathbf{y} , we will denote $\nabla g(\boldsymbol{\eta}, \mathbf{y})$ the gradient of g at $\boldsymbol{\eta}$ with respect to the first variable.

The *subdifferential* $\partial f(\boldsymbol{\theta})$ of a convex function f at $\boldsymbol{\theta}$ is the set

$$\partial f(\boldsymbol{\theta}) = \{ \boldsymbol{\eta} \in \mathbb{R}^p : f(\boldsymbol{\theta}') \geq f(\boldsymbol{\theta}) + \langle \boldsymbol{\eta}, \boldsymbol{\theta}' - \boldsymbol{\theta} \rangle, \quad \forall \boldsymbol{\theta}' \in \text{dom}(f) \}.$$

An element of $\partial f(\boldsymbol{\theta})$ is a subgradient. If the convex function f is differentiable at $\boldsymbol{\theta}$, then its only subgradient is its gradient, i.e. $\partial f(\boldsymbol{\theta}) = \{ \nabla f(\boldsymbol{\theta}) \}$.

The *Bregman divergence* associated to a convex function f at $\boldsymbol{\theta}$ with respect to $\boldsymbol{\eta} \in \partial f(\boldsymbol{\theta}) \neq \emptyset$ is

$$D_f^{\boldsymbol{\eta}}(\bar{\boldsymbol{\theta}}, \boldsymbol{\theta}) = f(\bar{\boldsymbol{\theta}}) - f(\boldsymbol{\theta}) - \langle \boldsymbol{\eta}, \bar{\boldsymbol{\theta}} - \boldsymbol{\theta} \rangle.$$

The Bregman divergence is in general nonsymmetric. It is also nonnegative by convexity. When f is differentiable at $\bar{\boldsymbol{\theta}}$, we simply write $D_f(\boldsymbol{\theta}, \bar{\boldsymbol{\theta}})$ (which is, in this case, also known as the Taylor distance).

2 Estimation with low-complexity penalties

The estimators $\hat{\boldsymbol{\theta}}_n^{\text{PEN}}$ and $\hat{\boldsymbol{\theta}}_n^{\text{EWA}}$ in (1.1) and (1.3) require two essential ingredients: the data loss term F and the prior penalty J . We here specify the class of such functions covered in our work, and provide illustrating examples.

2.1 Data loss

The class of loss functions F that we consider obey the following assumptions:

(H.1) $F(\cdot, \mathbf{y}) : \mathbb{R}^n \rightarrow \mathbb{R}$ is $C^1(\mathbb{R}^n)$ and uniformly convex for all \mathbf{y} of modulus φ , i.e.

$$F(\mathbf{v}, \mathbf{y}) \geq F(\mathbf{u}, \mathbf{y}) + \langle \nabla F(\mathbf{u}, \mathbf{y}), \mathbf{v} - \mathbf{u} \rangle + \varphi(\|\mathbf{v} - \mathbf{u}\|_2),$$

where $\varphi : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is a convex non-decreasing function that vanishes only at 0.

(H.2) For any $\bar{\boldsymbol{\theta}} \in \mathbb{R}^p$ and $\mathbf{y} \in \mathbb{R}^n$, $\int_{\mathbb{R}^p} \exp(-F(\mathbf{X}\boldsymbol{\theta}, \mathbf{y})/(n\beta)) |\langle \nabla F(\mathbf{X}\boldsymbol{\theta}, \mathbf{y}), \mathbf{X}(\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}) \rangle| d\boldsymbol{\theta} < +\infty$.

Recall that by Lemma A.1, the monotone conjugate φ^+ of φ is a proper, closed, convex, strongly coercive and non-decreasing function on \mathbb{R}_+ that vanishes at 0. Moreover, $\varphi^{++} = \varphi$. φ^+ is finite-valued on \mathbb{R}_+ if φ is strongly coercive, and it vanishes only at 0 under e.g. Lemma A.1(iii).

The class of data loss functions in (H.1) is fairly general. It is reminiscent of the negative log-likelihood in the regular exponential family. For the moment assumption (H.2) to be satisfied, it is sufficient that

$$\int_{\mathbb{R}^p} \exp(-\varphi(\|\mathbf{X}\boldsymbol{\theta}\|_2)/(n\beta)) \|\nabla F(\mathbf{X}\boldsymbol{\theta} + \mathbf{u}^*, \mathbf{y})\|_2 \|\mathbf{X}\boldsymbol{\theta} + (\mathbf{u}^* - \mathbf{X}\bar{\boldsymbol{\theta}})\|_2 d\boldsymbol{\theta} < +\infty,$$

where \mathbf{u}^* be a minimizer of $F(\cdot, \mathbf{y})$, which is unique by uniform convexity. We here provide an example.

Example 2.1. Consider the case where³ $\varphi(t) = t^q/q$, $q \in]1, +\infty[$, or equivalently $\varphi^+(t) = t^{q^*}/q_*$ where $1/q + 1/q_* = 1$. For $q = q_* = 2$, (H.1) amounts to saying that $F(\cdot, \mathbf{y})$ is strongly convex for all \mathbf{y} . In

³We consider a scaled version of φ for simplicity, but the same conclusions remain valid if we take $\varphi(t) = Ct^q/q$, with $C > 0$.

particular, [3, Proposition 10.13] shows that $F(\mathbf{u}, \mathbf{y}) = \|\mathbf{u} - \mathbf{y}\|_2^q/q$ is uniformly convex for $q \in [2, +\infty[$ with modulus $\varphi(t) = C_q t^q/q$, where $C_q > 0$ is a constant that depends solely on q .

For **(H.2)** to be verified, it is sufficient that

$$\int_{\mathbb{R}^p} \exp(-\|\mathbf{X}\boldsymbol{\theta}\|_2^q/(qn\beta)) \|\nabla F(\mathbf{X}\boldsymbol{\theta} + \mathbf{u}^*), \mathbf{y}\|_2 \|(\mathbf{X}\boldsymbol{\theta} + \mathbf{u}^*) - \mathbf{X}\bar{\boldsymbol{\theta}}\|_2 d\boldsymbol{\theta} < +\infty.$$

In particular, taking $F(\mathbf{u}, \mathbf{y}) = \|\mathbf{u} - \mathbf{y}\|_2^q/q$, $q \in [2, +\infty[$, we have $\|\nabla F(\mathbf{u}, \mathbf{y})\|_2 = \|\mathbf{u} - \mathbf{y}\|_2^{q-1}$, and thus **(H.2)** holds since

$$\int_{\mathbb{R}^p} \exp(-\|\mathbf{X}\boldsymbol{\theta}\|_2^q/(qn\beta)) \|\mathbf{y} - (\mathbf{X}\boldsymbol{\theta} + \mathbf{u}^*)\|_2^{q-1} \|\mathbf{X}\bar{\boldsymbol{\theta}} - (\mathbf{X}\boldsymbol{\theta} + \mathbf{u}^*)\|_2 d\boldsymbol{\theta} < +\infty.$$

2.2 Prior penalty

Recall the main definitions and results from convex analysis that are collected in Section A. Our main assumption on J is the following.

(H.3) $J : \mathbb{R}^p \rightarrow \mathbb{R}$ is the gauge of a non-empty convex compact set containing the origin as an interior point.

By Lemma A.3, this assumption is equivalent to saying that $J \stackrel{\text{def}}{=} \gamma_C$ is proper, convex, positively homogeneous, finite-valued and coercive. In turn, J is locally Lipschitz continuous on \mathbb{R}^p . Observe also that by virtue of Lemma A.4 and Lemma A.2, the polar gauge $J^\circ \stackrel{\text{def}}{=} \gamma_{C^\circ}$ enjoys the same properties as J in **(H.3)**.

2.3 Decomposability of the prior penalty

We are now in position to provide an important characterization of the subdifferential mapping of a function J satisfying **(H.3)**. This characterization will play a pivotal role in our proof of the oracle inequality.

We start by defining some essential geometrical objects that were introduced in [57].

Definition 2.1 (Model Subspace). *Let $\boldsymbol{\theta} \in \mathbb{R}^p$. We denote by $e_\boldsymbol{\theta}$ as*

$$e_\boldsymbol{\theta} = \text{P}_{\text{aff}(\partial J(\boldsymbol{\theta}))}(0).$$

We denote

$$S_\boldsymbol{\theta} = \text{par}(\partial J(\boldsymbol{\theta})) \quad \text{and} \quad T_\boldsymbol{\theta} = S_\boldsymbol{\theta}^\perp.$$

$T_\boldsymbol{\theta}$ is coined the model subspace of $\boldsymbol{\theta}$ associated to J .

It can be shown, see [57, Proposition 5], that $\boldsymbol{\theta} \in T_\boldsymbol{\theta}$, hence the name model subspace. When J is differentiable at $\boldsymbol{\theta}$, we have $e_\boldsymbol{\theta} = \nabla J(\boldsymbol{\theta})$ and $T_\boldsymbol{\theta} = \mathbb{R}^p$. When J is the ℓ_1 -norm (Lasso), the vector $e_\boldsymbol{\theta}$ is nothing but the sign of $\boldsymbol{\theta}$. Thus, $e_\boldsymbol{\theta}$ can be viewed as a generalization of the sign vector. Observe also that $e_\boldsymbol{\theta} = \text{P}_{T_\boldsymbol{\theta}}(\partial J(\boldsymbol{\theta}))$, and thus $e_\boldsymbol{\theta} \in T_\boldsymbol{\theta} \cap \text{aff}(\partial J(\boldsymbol{\theta}))$. However, in general, $e_\boldsymbol{\theta} \notin \partial J(\boldsymbol{\theta})$.

We now provide a fundamental equivalent description of the subdifferential of J at $\boldsymbol{\theta}$ in terms of $e_\boldsymbol{\theta}$, $T_\boldsymbol{\theta}$, $S_\boldsymbol{\theta}$ and the polar gauge J° .

Theorem 2.1. *Let J satisfy **(H.3)**. Let $\boldsymbol{\theta} \in \mathbb{R}^p$ and $f_\boldsymbol{\theta} \in \text{ri}(\partial J(\boldsymbol{\theta}))$.*

(i) The subdifferential of J at $\boldsymbol{\theta}$ reads

$$\begin{aligned}\partial J(\boldsymbol{\theta}) &= \text{aff}(\partial J(\boldsymbol{\theta})) \cap \mathcal{C}^\circ \\ &= \{\boldsymbol{\eta} \in \mathbb{R}^n : \boldsymbol{\eta}_{T_\theta} = e_\theta \quad \text{and} \quad \inf_{\tau \geq 0} \max(J^\circ(\tau e_\theta + \boldsymbol{\eta}_{S_\theta} + (\tau - 1) P_{S_\theta} f_\theta), \tau) \leq 1\}.\end{aligned}$$

(ii) For any $\boldsymbol{\omega} \in \mathbb{R}^p$, $\exists \boldsymbol{\eta} \in \partial J(\boldsymbol{\theta})$ such that

$$J(\boldsymbol{\omega}_{S_\theta}) = \langle \boldsymbol{\eta}_{S_\theta}, \boldsymbol{\omega}_{S_\theta} \rangle.$$

Proof. (i) This follows by piecing together [57, Theorem 1, Proposition 4 and Proposition 5(iii)].

(ii) From [57, Proposition 5(iv)], we have

$$\sigma_{\partial J(\boldsymbol{\theta}) - f_\theta}(\boldsymbol{\omega}) = J(\boldsymbol{\omega}_{S_\theta}) - \langle P_{S_\theta} f_\theta, \boldsymbol{\omega}_{S_\theta} \rangle.$$

Thus there exists a supporting point $\boldsymbol{v} \in \partial J(\boldsymbol{\theta}) - f_\theta \subset S_\theta$ with normal vector $\boldsymbol{\omega}$ [3, Corollary 7.6(iii)], i.e.

$$\sigma_{\partial J(\boldsymbol{\theta}) - f_\theta}(\boldsymbol{\omega}) = \langle \boldsymbol{v}, \boldsymbol{\omega}_{S_\theta} \rangle.$$

Taking $\boldsymbol{\eta} = \boldsymbol{v} + f_\theta$ concludes the proof. □

Remark 2.1. The coercivity assumption in (H.3) is not needed for Theorem 2.1 to hold.

The decomposability of described in Theorem 2.1(i) depends on the particular choice of the mapping $\boldsymbol{\theta} \mapsto f_\theta \in \text{ri}(\partial J(\boldsymbol{\theta}))$. An interesting situation is encountered when $e_\theta \in \text{ri}(J(\boldsymbol{\theta}))$, so that one can choose $f_\theta = e_\theta$. Strong gauges, see [57, Definition 6], are precisely a class of gauges for which this situation occurs, and in this case, Theorem 2.1(i) has the simpler form

$$\partial J(\boldsymbol{\theta}) = \text{aff}(\partial J(\boldsymbol{\theta})) \cap \mathcal{C}^\circ = \{\boldsymbol{\eta} \in \mathbb{R}^n : \boldsymbol{\eta}_{T_\theta} = e_\theta \quad \text{and} \quad J^\circ(\boldsymbol{\eta}_{S_\theta}) \leq 1\}. \quad (2.1)$$

The Lasso, group Lasso and nuclear norms are typical examples of (symmetric) strong gauges. However, analysis sparsity penalties (e.g. the fused Lasso) or the ℓ_∞ -penalty are not strong gauges, though they obviously satisfy (H.3). See the next section for a detailed discussion.

2.4 Calculus with the prior family

The family of penalties complying with (H.3) form a robust class enjoying important calculus rules. In particular it is closed under the sum and composition with an injective linear operator as we now prove.

Lemma 2.1. *The set of functions satisfying (H.3) is closed under addition⁴ and pre-composition by an injective linear operator. More precisely, the following holds:*

(i) Let J and G be two gauges satisfying (H.3). Then $H \stackrel{\text{def}}{=} J + G$ also obeys (H.3). Moreover,

(a) $T_\theta^H = T_\theta^J \cap T_\theta^G$ and $e_\theta^H = P_{T_\theta^H}(e_\theta^J + e_\theta^G)$, where T_θ^J and e_θ^J (resp. T_θ^G and e_θ^G) are the model subspace and vector at $\boldsymbol{\theta}$ associated to J (resp. G);

⁴It is obvious that the same holds with any positive linear combination.

$$(b) H^\circ(\boldsymbol{\omega}) = \max_{\rho \in [0,1]} \overline{\text{conv}}(\inf(\rho J^\circ(\boldsymbol{\omega}), (1-\rho)G^\circ(\boldsymbol{\omega}))).$$

(ii) Let J be a gauge satisfying **(H.3)**, and $\mathbf{D} : \mathbb{R}^q \rightarrow \mathbb{R}^p$ be surjective. Then $H \stackrel{\text{def}}{=} J \circ \mathbf{D}^\top$ also fulfills **(H.3)**. Moreover,

- (a) $T_\theta^H = \text{Ker}(\mathbf{D}_{S_u}^\top)$ and $e_\theta^H = P_{T_\theta^H} \mathbf{D} e_u^J$, where T_u^J and e_u^J are the model subspace and vector at $\mathbf{u} \stackrel{\text{def}}{=} \mathbf{D}^\top \theta$ associated to J ;
- (b) $H^\circ(\boldsymbol{\omega}) = J^\circ(\mathbf{D}^+ \boldsymbol{\omega})$, where $\mathbf{D}^+ = \mathbf{D}^\top (\mathbf{D} \mathbf{D}^\top)^{-1}$.

The outcome of Lemma 2.1 is naturally expected. For instance, assertion (i) states that combining several penalties/priors will promote objects living on the intersection of the respective low-complexity models. Similarly, for (ii), one promotes low-complexity in the image of the analysis operator \mathbf{D}^\top . It then follows that one has not to deploy an ad hoc analysis when linearly pre-composing or combining (or both) several penalties (e.g. ℓ_1 +nuclear norms for recovering sparse and low-rank matrices) since our unified analysis in Section 3 will apply to them just as well.

Proof. (i) Convexity, positive homogeneity, coercivity and finite-valuedness are straightforward.

(a) This is [57, Proposition 8(i)-(ii)].

(b) We have from Lemma A.4 and calculus rules on support functions,

$$\begin{aligned} H^\circ(\boldsymbol{\omega}) &= \sigma_{J(\theta)+G(\theta) \leq 1}(\boldsymbol{\omega}) = \sup_{J(\theta)+G(\theta) \leq 1} \langle \boldsymbol{\omega}, \theta \rangle = \max_{\rho \in [0,1]} \sup_{J(\theta) \leq \rho, G(\theta) \leq 1-\rho} \langle \boldsymbol{\omega}, \theta \rangle \\ &\stackrel{([28, \text{Theorem V.3.3.3}])}{=} \max_{\rho \in [0,1]} \overline{\text{conv}}(\inf(\sigma_{J(\theta) \leq \rho}(\boldsymbol{\omega}), \sigma_{G(\theta) \leq 1-\rho}(\boldsymbol{\omega}))) \\ &\stackrel{(\text{Positive homogeneity})}{=} \max_{\rho \in [0,1]} \overline{\text{conv}}(\inf(\rho \sigma_{J(\theta) \leq 1}(\boldsymbol{\omega}), (1-\rho) \sigma_{G(\theta) \leq 1}(\boldsymbol{\omega}))) \\ &\stackrel{(\text{Lemma A.4})}{=} \max_{\rho \in [0,1]} \overline{\text{conv}}(\inf(\rho J^\circ(\boldsymbol{\omega}), (1-\rho)G^\circ(\boldsymbol{\omega}))). \end{aligned}$$

(ii) Again, Convexity, positive homogeneity and finite-valuedness are immediate. Coercivity holds by injectivity of \mathbf{D}^\top .

(a) This is [57, Proposition 10(i)-(ii)].

(b) Denote $J = \gamma_{\mathcal{C}}$. We have

$$\begin{aligned} H^\circ(\boldsymbol{\omega}) &= \sup_{\mathbf{D}^\top \theta \in \mathcal{C}} \langle \boldsymbol{\omega}, \theta \rangle \\ &\stackrel{(\mathbf{D}^\top \text{ is injective})}{=} \sup_{\mathbf{D}^\top \theta \in \mathcal{C}} \langle \mathbf{D}^+ \boldsymbol{\omega}, \mathbf{D}^\top \theta \rangle \\ &= \sup_{\mathbf{u} \in \mathcal{C} \cap \text{Span}(\mathbf{D}^\top)} \langle \mathbf{D}^+ \boldsymbol{\omega}, \mathbf{u} \rangle \\ &\stackrel{([28, \text{Theorem V.3.3.3}] \text{ and Lemma A.4})}{=} \overline{\text{conv}}(\inf(J^\circ(\mathbf{D}^+ \boldsymbol{\omega}), \iota_{\text{Ker}(\mathbf{D})}(\mathbf{D}^+ \boldsymbol{\omega}))) \\ &= J^\circ(\mathbf{D}^+ \boldsymbol{\omega}). \end{aligned}$$

where in the last equality, we used the fact that $\mathbf{D}^+ \boldsymbol{\omega} \in \text{Span}(\mathbf{D}^\top) = \text{Ker}(\mathbf{D})^\perp$, and thus $\iota_{\text{Ker}(\mathbf{D})}(\mathbf{D}^+ \boldsymbol{\omega}) = +\infty$ unless $\boldsymbol{\omega} = 0$, and J° is continuous and convex by **(H.3)** and Lemma A.4. \square

2.5 Examples

2.5.1 Lasso

The Lasso regularization is used to promote the sparsity of the minimizers, see [7] for a comprehensive review. It corresponds to choosing J as the ℓ_1 -norm

$$J(\boldsymbol{\theta}) = \|\boldsymbol{\theta}\|_1 = \sum_{i=1}^p |\boldsymbol{\theta}_i|. \quad (2.2)$$

It is also referred to as ℓ_1 -synthesis in the signal processing community, in contrast to the more general ℓ_1 -analysis sparsity penalty detailed below.

We denote $(\mathbf{a}_i)_{1 \leq i \leq p}$ the canonical basis of \mathbb{R}^p and $\text{supp}(\boldsymbol{\theta}) \stackrel{\text{def}}{=} \{i \in \{1, \dots, p\} : \boldsymbol{\theta}_i \neq 0\}$. Then,

$$T_{\boldsymbol{\theta}} = \text{Span}\{(\mathbf{a}_i)_{i \in \text{supp}(\boldsymbol{\theta})}\}, \quad (e_{\boldsymbol{\theta}})_i = \begin{cases} \text{sign}(\boldsymbol{\theta}_i) & \text{if } i \in \text{supp}(\boldsymbol{\theta}) \\ 0 & \text{otherwise} \end{cases}, \quad \text{and } J^\circ = \|\cdot\|_\infty. \quad (2.3)$$

2.5.2 Group Lasso

The group Lasso has been advocated to promote sparsity by groups, i.e. it drives all the coefficients in one group to zero together hence leading to group selection, see [2, 67, 1, 65] to cite a few. The group Lasso penalty with L groups reads

$$J(\boldsymbol{\theta}) = \|\boldsymbol{\theta}\|_{1,2} \stackrel{\text{def}}{=} \sum_{i=1}^L \|\boldsymbol{\theta}_{b_i}\|_2. \quad (2.4)$$

where $\bigcup_{i=1}^L b_i = \{1, \dots, p\}$, $b_i, b_j \subset \{1, \dots, p\}$, and $b_i \cap b_j = \emptyset$ whenever $i \neq j$. Define the group support as $\text{supp}_{\mathcal{B}}(\boldsymbol{\theta}) \stackrel{\text{def}}{=} \{i \in \{1, \dots, L\} : \boldsymbol{\theta}_{b_i} \neq 0\}$. Thus, one has

$$T_{\boldsymbol{\theta}} = \text{Span}\{(\mathbf{a}_j)_{\{j : \exists i \in \text{supp}_{\mathcal{B}}(\boldsymbol{\theta}), j \in b_i\}}\}, \quad (e_{\boldsymbol{\theta}})_{b_i} = \begin{cases} \frac{\boldsymbol{\theta}_{b_i}}{\|\boldsymbol{\theta}_{b_i}\|_2} & \text{if } i \in \text{supp}_{\mathcal{B}}(\boldsymbol{\theta}) \\ 0 & \text{otherwise} \end{cases}, \quad \text{and } J^\circ(\boldsymbol{\omega}) = \max_{i \in \{1, \dots, L\}} \|\boldsymbol{\omega}_{b_i}\|_2. \quad (2.5)$$

2.5.3 Analysis (group) Lasso

One can push the structured sparsity idea one step further by promoting group/block sparsity through a linear operator, i.e. analysis-type sparsity. Given a linear operator $\mathbf{D} : \mathbb{R}^q \rightarrow \mathbb{R}^p$ (seen as a matrix), the analysis group sparsity penalty is

$$J(\boldsymbol{\theta}) = \|\mathbf{D}^\top \boldsymbol{\theta}\|_{1,2}. \quad (2.6)$$

This encompasses the 2-D isotropic total variation [47]. For when all groups of cardinality one, we have the analysis- ℓ_1 penalty (a.k.a. general Lasso), which encapsulates several important penalties including that of the 1-D total variation [47], and the fused Lasso [53]. The overlapping group Lasso [29] is also a special case of (2.4) by taking \mathbf{D}^\top to be an operator that extract the blocks [40, 14] (in which case \mathbf{D} has even orthogonal rows).

Let $\Lambda_\theta = \bigcup_{i \in \text{supp}_B(D^\top \theta)} b_i$ and Λ_θ^c its complement. From Lemma 2.1(ii) and (2.5), we get

$$T_\theta = \text{Ker}(D_{\Lambda_\theta^c}^\top), \quad e_\theta = P_{T_\theta} D e_{D^\top \theta}^{\|\cdot\|_{1,2}} \quad \text{where} \quad \left(e_{D^\top \theta}^{\|\cdot\|_{1,2}} \right)_{b_i} = \begin{cases} \frac{(D^\top \theta)_{b_i}}{\|(D^\top \theta)_{b_i}\|_2} & \text{if } i \in \text{supp}_B(D^\top \theta) \\ 0 & \text{otherwise.} \end{cases} \quad (2.7)$$

If, in addition, D is surjective, then by virtue of Lemma 2.1(ii) we also have

$$J^\circ(\omega) = \|D^+ \omega\|_{\infty,2} \stackrel{\text{def}}{=} \max_{i \in \{1, \dots, L\}} \|(D^+ \omega)_{b_i}\|_2 \quad (2.8)$$

2.5.4 Anti-sparsity

If the vector to be estimated is expected to be flat (anti-sparse), this can be captured using the ℓ_∞ norm (a.k.a. Tchebychev norm) as prior

$$J(\theta) = \|\theta\|_\infty = \max_{i \in \{1, \dots, p\}} |\theta_i|. \quad (2.9)$$

The ℓ_∞ regularization has found applications in several fields [30, 36, 48]. Suppose that $\theta \neq 0$, and define the saturation support of θ as $I_\theta^{\text{sat}} \stackrel{\text{def}}{=} \{i \in \{1, \dots, p\} : |\theta_i| = \|\theta\|_\infty\} \neq \emptyset$. From [57, Proposition 14], we have

$$T_\theta = \{\bar{\theta} \in \mathbb{R}^p : \bar{\theta}_{I_\theta^{\text{sat}}} \in \mathbb{R} \text{sign}(\theta_{I_\theta^{\text{sat}}})\}, \quad (e_\theta)_i = \begin{cases} \text{sign}(\theta_i)/|I_\theta^{\text{sat}}| & \text{if } i \in I_\theta^{\text{sat}} \\ 0 & \text{otherwise} \end{cases}, \quad \text{and} \quad J^\circ = \|\cdot\|_1. \quad (2.10)$$

2.5.5 Nuclear norm

The natural extension of low-complexity priors to matrices $\theta \in \mathbb{R}^{p_1 \times p_2}$ is to penalize the singular values of the matrix. Let $\text{rank}(\theta) = r$, and $\theta = U \text{diag}(\lambda(\theta)) V^\top$ be a reduced rank- r SVD decomposition, where $U \in \mathbb{R}^{p_1 \times r}$ and $V \in \mathbb{R}^{p_2 \times r}$ have orthonormal columns, and $\lambda(\theta) \in (\mathbb{R}_+ \setminus \{0\})^r$ is the vector of singular values $(\lambda_1(\theta), \dots, \lambda_r(\theta))$ in non-increasing order. The nuclear norm of θ is

$$J(\theta) = \|\theta\|_* = \|\lambda(\theta)\|_1. \quad (2.11)$$

This penalty is the best convex surrogate to enforce a low-rank prior. It has been widely used for various applications [41, 10, 9, 26, 11].

Following e.g. [56, Example 21], we have

$$T_\theta = \{UA^\top + BV^\top : A \in \mathbb{R}^{p_2 \times r}, B \in \mathbb{R}^{p_1 \times r}\}, \quad e_\theta = UV^\top \quad \text{and} \quad J^\circ(\omega) = \|\omega\|_{2 \rightarrow 2} = \|\lambda(\omega)\|_\infty. \quad (2.12)$$

3 Oracle inequalities for a general loss

Before delving into the details, in the sequel, we will need a bit of notations.

We recall T_θ and e_θ the model subspace and vector associated to θ (see Definition 2.1). Denote $S_\theta = T_\theta^\perp$. Given two coercive finite-valued gauges $J_1 = \gamma_{C_1}$ and $J_2 = \gamma_{C_2}$, and a linear operator A , we define $\|A\|_{J_1 \rightarrow J_2}$ the operator bound as

$$\|A\|_{J_1 \rightarrow J_2} = \sup_{\theta \in C_1} J_2(A\theta).$$

Note that $\|\mathbf{A}\|_{J_1 \rightarrow J_2}$ is bounded (this follows from Lemma A.3(v)). Furthermore, we have from Lemma A.4 that

$$\|\mathbf{A}\|_{J_1 \rightarrow J_2} = \sup_{\boldsymbol{\theta} \in \mathcal{C}_1} \sup_{\boldsymbol{\omega} \in \mathcal{C}_2^\circ} \langle \mathbf{A}^\top \boldsymbol{\omega}, \boldsymbol{\theta} \rangle = \sup_{\boldsymbol{\omega} \in \mathcal{C}_2^\circ} \sup_{\boldsymbol{\theta} \in \mathcal{C}_1} \langle \mathbf{A}^\top \boldsymbol{\omega}, \boldsymbol{\theta} \rangle = \sup_{\boldsymbol{\omega} \in \mathcal{C}_2^\circ} J_1^\circ(\mathbf{A}^\top \boldsymbol{\omega}) = \|\mathbf{A}^\top\|_{J_2^\circ \rightarrow J_1^\circ}.$$

In the following, whenever it is clear from the context, to lighten notation when J_i is a norm, we write the subscript of the norm instead of J_i (e.g. p for the ℓ_p norm, $*$ for the nuclear norm, etc.).

Our main result will involve a measure of well-conditionedness of the design matrix \mathbf{X} when restricted to some subspace T . More precisely, for $c > 0$, we introduce the coefficient

$$\Upsilon(T, c) = \inf_{\{\boldsymbol{\omega} \in \mathbb{R}^p : J(\boldsymbol{\omega}_S) < cJ(\boldsymbol{\omega}_T)\}} \frac{\|\mathbf{P}_T\|_{2 \rightarrow J} \|\mathbf{X}\boldsymbol{\omega}\|_2}{n^{1/2}(J(\boldsymbol{\omega}_T) - J(\boldsymbol{\omega}_S)/c)}. \quad (3.1)$$

This generalizes the compatibility factor introduced in [62] for the Lasso (and used in [17]). The experienced reader may have recognized that this factor is reminiscent of the null space property and restricted injectivity that play a central role in the analysis of the performance guarantees of variational/penalized estimators (1.1); see [25, 57, 59, 56, 58]. One can see in particular that $\Upsilon(T, c)$ is larger than the smallest singular value of \mathbf{X}_T .

The oracle inequalities will be provided in terms of the loss

$$R_n(\boldsymbol{\theta}, \boldsymbol{\theta}_0) = \frac{1}{n} D_F(\mathbf{X}\boldsymbol{\theta}, \mathbf{X}\boldsymbol{\theta}_0).$$

3.1 Oracle inequality for $\hat{\boldsymbol{\theta}}_n^{\text{EWA}}$

We are now ready to establish our first main result: an oracle inequality for the EWA estimator (1.3).

Theorem 3.1. *Consider the EWA estimator $\hat{\boldsymbol{\theta}}_n^{\text{EWA}}$ in (1.3) with the density (1.2), where F and J satisfy Assumptions (H.1)-(H.2) and (H.3). Then, for any $\tau > 1$ such that $\lambda_n \geq \tau J^\circ(-\mathbf{X}^\top \nabla F(\mathbf{X}\boldsymbol{\theta}_0, \mathbf{y}))/n$, the following holds,*

$$R_n(\hat{\boldsymbol{\theta}}_n^{\text{EWA}}, \boldsymbol{\theta}_0) \leq \inf_{\boldsymbol{\theta} \in \mathbb{R}^p} \left(R_n(\boldsymbol{\theta}, \boldsymbol{\theta}_0) + \frac{1}{n} \varphi^+ \left(\frac{\lambda_n \sqrt{n} (\tau J^\circ(\mathbf{e}_\theta) + 1) \|\mathbf{P}_{T_\theta}\|_{2 \rightarrow J}}{\tau \Upsilon\left(T_\theta, \frac{\tau J^\circ(\mathbf{e}_\theta) + 1}{\tau - 1}\right)} \right) \right) + p\beta. \quad (3.2)$$

Remark 3.1.

1. It should be emphasized that Theorem 3.1 is actually a deterministic statement for a fixed choice of λ_n . Probabilistic analysis will be required when the result is applied to particular statistical models as we will see later. For this, we will use concentration inequalities in order to provide bounds that hold with high probability over the data.
2. The oracle inequality is sharp. The remainder in it has two terms. The first one encodes the complexity of the model promoted by J . The second one, $p\beta$, captures the influence of the temperature parameter. In particular, taking β sufficiently small of the order $O((pn)^{-1})$, this term becomes $O(n^{-1})$.

3. When $\varphi(t) = \nu t^2/2$, i.e. $F(\cdot, \mathbf{y})$ is ν -strongly convex, then $\varphi^+(t) = t^2/(2\nu)$, and the reminder term becomes

$$\frac{\lambda_n^2 (\tau J^\circ(e_\theta) + 1)^2 \|\mathbf{P}_{T_\theta}\|_{2 \rightarrow J}^2}{2\tau^2 \nu \Upsilon \left(T_\theta, \frac{\tau J^\circ(e_\theta) + 1}{\tau - 1} \right)^2}. \quad (3.3)$$

If, moreover, ∇F is also κ -Lipschitz continuous, then it can be shown that $R_n(\boldsymbol{\theta}, \boldsymbol{\theta}_0)$ is equivalent to a quadratic loss. This means that the oracle inequality in Theorem 3.1 can be stated in terms of the quadratic prediction error. However, the inequality is not anymore sharp in this case as a constant factor equal to the condition number $\kappa/\nu \geq 1$ naturally multiplies the right-hand side.

4. If J is such that $e_\theta \in \partial J(\boldsymbol{\theta}) \subset \mathcal{C}^\circ$ (typically for a strong gauge by (2.1)), then $J^\circ(e_\theta) \leq 1$ (in fact an equality if $\boldsymbol{\theta} \neq 0$). Thus the term $J^\circ(e_\theta)$ can be omitted in (3.2).
5. A close inspection of the proof of Theorem 3.1 reveals that the term $p\beta$ can be improved to the smaller bound

$$p\beta + \left(V_n(\widehat{\boldsymbol{\theta}}_n^{\text{EWA}}) - \mathbb{E}_{\mu_n} [V_n(\boldsymbol{\theta})] \right),$$

where the upper-bound is a consequence of Jensen inequality.

Proof. By convexity of J and assumption (H.1), we have for any $\boldsymbol{\eta} \in \partial V_n(\boldsymbol{\theta})$ and any $\bar{\boldsymbol{\theta}} \in \mathbb{R}^p$,

$$D_{V_n}^{\boldsymbol{\eta}}(\bar{\boldsymbol{\theta}}, \boldsymbol{\theta}) \geq \frac{1}{n} \varphi(\|\mathbf{X}\bar{\boldsymbol{\theta}} - \mathbf{X}\boldsymbol{\theta}\|_2).$$

Since φ is non-decreasing and convex, $\varphi \circ \|\cdot\|_2$ is a convex function. Thus, taking the expectation w.r.t. to μ_n on both sides and using Jensen inequality, we get

$$\begin{aligned} V_n(\bar{\boldsymbol{\theta}}) &\geq \mathbb{E}_{\mu_n} [V_n(\boldsymbol{\theta})] + \mathbb{E}_{\mu_n} [\langle \boldsymbol{\eta}, \bar{\boldsymbol{\theta}} - \boldsymbol{\theta} \rangle] + \frac{1}{n} \mathbb{E}_{\mu_n} [\varphi(\|\mathbf{X}\bar{\boldsymbol{\theta}} - \mathbf{X}\boldsymbol{\theta}\|_2)] \\ &\geq V_n(\widehat{\boldsymbol{\theta}}_n^{\text{EWA}}) + \mathbb{E}_{\mu_n} [\langle \boldsymbol{\eta}, \bar{\boldsymbol{\theta}} - \boldsymbol{\theta} \rangle] + \frac{1}{n} \varphi(\|\mathbf{X}\bar{\boldsymbol{\theta}} - \mathbf{X}\widehat{\boldsymbol{\theta}}_n^{\text{EWA}}\|_2). \end{aligned}$$

This holds for any $\boldsymbol{\eta} \in \partial V_n(\boldsymbol{\theta})$, and in particular at the minimal selection $(\partial V_n(\boldsymbol{\theta}))^0$ (see Section B for details). It then follows from the pillar result in Proposition B.1⁵ that

$$\mathbb{E}_{\mu_n} \left[\langle (\partial V_n(\boldsymbol{\theta}))^0, \bar{\boldsymbol{\theta}} - \boldsymbol{\theta} \rangle \right] = -p\beta.$$

We thus deduce the inequality

$$V_n(\widehat{\boldsymbol{\theta}}_n^{\text{EWA}}) - V_n(\bar{\boldsymbol{\theta}}) \leq p\beta - \frac{1}{n} \varphi(\|\mathbf{X}\widehat{\boldsymbol{\theta}}_n^{\text{EWA}} - \mathbf{X}\bar{\boldsymbol{\theta}}\|_2), \quad \forall \bar{\boldsymbol{\theta}} \in \mathbb{R}^p. \quad (3.4)$$

By definition of the Bregman divergence, we have

$$\begin{aligned} R_n(\widehat{\boldsymbol{\theta}}_n^{\text{EWA}}, \boldsymbol{\theta}_0) - R_n(\boldsymbol{\theta}, \boldsymbol{\theta}_0) &= \frac{1}{n} \left(F(\mathbf{X}\widehat{\boldsymbol{\theta}}_n^{\text{EWA}}, \mathbf{y}) - F(\mathbf{X}\boldsymbol{\theta}, \mathbf{y}) + \langle -\mathbf{X}^\top \nabla F(\mathbf{X}\boldsymbol{\theta}_0, \mathbf{y}), \widehat{\boldsymbol{\theta}}_n^{\text{EWA}} - \boldsymbol{\theta} \rangle \right) \\ &= \left(V_n(\widehat{\boldsymbol{\theta}}_n^{\text{EWA}}) - V_n(\boldsymbol{\theta}) \right) + \frac{1}{n} \langle -\mathbf{X}^\top \nabla F(\mathbf{X}\boldsymbol{\theta}_0, \mathbf{y}), \widehat{\boldsymbol{\theta}}_n^{\text{EWA}} - \boldsymbol{\theta} \rangle \\ &\quad - \lambda_n (J(\widehat{\boldsymbol{\theta}}_n^{\text{EWA}}) - J(\boldsymbol{\theta})). \end{aligned}$$

⁵In the appendix, we provide a self-contained proof based on a novel Moreau-Yosida regularization argument. In [17, Corollary 1 and 2], an alternative proof is given using an absolute continuity argument since μ_n is locally Lipschitz, hence a Sobolev function.

By virtue of the duality inequality (A.1), we have

$$\begin{aligned} R_n(\widehat{\boldsymbol{\theta}}_n^{\text{EWA}}, \boldsymbol{\theta}_0) - R_n(\boldsymbol{\theta}, \boldsymbol{\theta}_0) &\leq \left(V_n(\widehat{\boldsymbol{\theta}}_n^{\text{EWA}}) - V_n(\boldsymbol{\theta}) \right) + \frac{1}{n} J^\circ \left(-\mathbf{X}^\top \nabla F(\mathbf{X}\boldsymbol{\theta}_0, \mathbf{y}) \right) J(\widehat{\boldsymbol{\theta}}_n^{\text{EWA}} - \boldsymbol{\theta}) \\ &\quad - \lambda_n (J(\widehat{\boldsymbol{\theta}}_n^{\text{EWA}}) - J(\boldsymbol{\theta})) \\ &\leq \left(V_n(\widehat{\boldsymbol{\theta}}_n^{\text{EWA}}) - V_n(\boldsymbol{\theta}) \right) + \frac{\lambda_n}{\tau} \left(J(\widehat{\boldsymbol{\theta}}_n^{\text{EWA}} - \boldsymbol{\theta}) - \tau (J(\widehat{\boldsymbol{\theta}}_n^{\text{EWA}}) - J(\boldsymbol{\theta})) \right). \end{aligned}$$

Denote $\boldsymbol{\omega} = \widehat{\boldsymbol{\theta}}_n^{\text{EWA}} - \boldsymbol{\theta}$. By virtue of (H.3), Theorem 2.1 and (A.1), we obtain

$$\begin{aligned} J(\boldsymbol{\omega}) - \tau (J(\widehat{\boldsymbol{\theta}}_n^{\text{EWA}}) - J(\boldsymbol{\theta})) &\leq J(\boldsymbol{\omega}_{T_\theta}) + J(\boldsymbol{\omega}_{S_\theta}) - \tau \langle e_\theta, \boldsymbol{\omega}_{T_\theta} \rangle - \tau J(\boldsymbol{\omega}_{S_\theta}) \\ &\leq J(\boldsymbol{\omega}_{T_\theta}) + J(\boldsymbol{\omega}_{S_\theta}) + \tau J^\circ(e_\theta) J(\boldsymbol{\omega}_{T_\theta}) - \tau J(\boldsymbol{\omega}_{S_\theta}) \\ &= (\tau J^\circ(e_\theta) + 1) J(\boldsymbol{\omega}_{T_\theta}) - (\tau - 1) J(\boldsymbol{\omega}_{S_\theta}) \\ &\leq (\tau J^\circ(e_\theta) + 1) \left(J(\boldsymbol{\omega}_{T_\theta}) - \frac{\tau-1}{\tau J^\circ(e_\theta)+1} J(\boldsymbol{\omega}_{S_\theta}) \right). \end{aligned}$$

This inequality together with (3.4) (applied with $\bar{\boldsymbol{\theta}} = \boldsymbol{\theta}$) and (3.1) yield

$$\begin{aligned} R_n(\widehat{\boldsymbol{\theta}}_n^{\text{EWA}}, \boldsymbol{\theta}_0) - R_n(\boldsymbol{\theta}, \boldsymbol{\theta}_0) &\leq p\beta - \frac{1}{n} \varphi(\|\mathbf{X}\boldsymbol{\omega}\|_2) + \frac{\lambda_n (\tau J^\circ(e_\theta) + 1) \|\mathbb{P}_{T_\theta}\|_{2 \rightarrow J} \|\mathbf{X}\boldsymbol{\omega}\|_2}{n^{1/2} \tau \Upsilon \left(T_\theta, \frac{\tau J^\circ(e_\theta)+1}{\tau-1} \right)} \\ &\leq p\beta + \frac{1}{n} \varphi^+ \left(\frac{\lambda_n \sqrt{n} (\tau J^\circ(e_\theta) + 1) \|\mathbb{P}_{T_\theta}\|_{2 \rightarrow J}}{\tau \Upsilon \left(T_\theta, \frac{\tau J^\circ(e_\theta)+1}{\tau-1} \right)} \right), \end{aligned}$$

where we applied Fenchel-Young inequality (1.5) to get the last bound. Taking the infimum over $\boldsymbol{\theta} \in \mathbb{R}^p$ yields the desired result. \square

Partly smooth functions Theorem 3.1 has a nice instantiation for the case of partly smooth functions convex functions. These functions have been thoroughly studied recently in [57, 59, 56, 58] for various statistical and inverse problems. In particular, a finite-valued convex function J is partly smooth at a point $\boldsymbol{\theta}$ relative to an active set $\mathcal{M} \ni \boldsymbol{\theta}$, if \mathcal{M} is a C^2 -smooth submanifold of \mathbb{R}^p , and J behaves smoothly along \mathcal{M} and sharply transverse to it. In addition, the partial smoothness active submanifold is always unique. There are two consequences of partial smoothness. First, the sharpness property is equivalently characterized by $\mathcal{T}_\theta(\mathcal{M}) = T_\theta$, where $\mathcal{T}_\theta(\mathcal{M})$ is the tangent space of \mathcal{M} at $\boldsymbol{\theta}$. Second, e_θ coincides with the Riemannian gradient of J along \mathcal{M} . Moreover, both these two properties locally persist at all nearby points of $\boldsymbol{\theta}$ on \mathcal{M} ; see [56, 58] for details. As an example, a C^2 -smooth function at $\boldsymbol{\theta}$ is partly smooth relative to the whole space \mathbb{R}^p . All popular penalty functions discussed in Sections 2.5 are also partly smooth (see [56, 58]).

Let's denote \mathcal{M} the set of all possible partial smoothness active submanifolds associated to J . An interesting fact is that there are many situations where \mathcal{M} contains finitely many active submanifolds (that form a stratification of \mathbb{R}^p), as is the case for the examples of Section 2.5. With this notation at hand, the oracle inequality (3.2) now reads

$$R_n(\widehat{\boldsymbol{\theta}}_n^{\text{EWA}}, \boldsymbol{\theta}_0) \leq \inf_{\substack{\mathcal{M} \in \mathcal{M} \\ \boldsymbol{\theta} \in \mathcal{M}}} \left(R_n(\boldsymbol{\theta}, \boldsymbol{\theta}_0) + \frac{1}{n} \varphi^+ \left(\frac{\lambda_n \sqrt{n} (\tau J^\circ(e_\theta) + 1) \|\mathbb{P}_{\mathcal{T}_\theta(\mathcal{M})}\|_{2 \rightarrow J}}{\tau \Upsilon \left(\mathcal{T}_\theta(\mathcal{M}), \frac{\tau J^\circ(e_\theta)+1}{\tau-1} \right)} \right) \right) + p\beta.$$

3.2 Oracle inequality for $\hat{\boldsymbol{\theta}}_n^{\text{PEN}}$

The next result establishes that $\hat{\boldsymbol{\theta}}_n^{\text{PEN}}$ satisfies a sharp prediction oracle inequality that we will compare to (3.2).

Theorem 3.2. *Consider the penalized estimator $\hat{\boldsymbol{\theta}}_n^{\text{PEN}}$ in (1.1), where F and J satisfy Assumptions (H.1) and (H.3). Then, for any $\tau > 1$ such that $\lambda_n \geq \tau J^\circ(-\mathbf{X}^\top \nabla F(\mathbf{X}\boldsymbol{\theta}_0, \mathbf{y}))/n$, the following holds,*

$$R_n(\hat{\boldsymbol{\theta}}_n^{\text{PEN}}, \boldsymbol{\theta}_0) \leq \inf_{\boldsymbol{\theta} \in \mathbb{R}^p} \left(R_n(\boldsymbol{\theta}, \boldsymbol{\theta}_0) + \frac{1}{n} \varphi^+ \left(\frac{\lambda_n \sqrt{n} (\tau J^\circ(e_{\boldsymbol{\theta}}) + 1) \|\mathbb{P}_{\mathcal{T}_{\boldsymbol{\theta}}(\mathcal{M})}\|_{2 \rightarrow J}}{\tau \Upsilon(\mathcal{T}_{\boldsymbol{\theta}}(\mathcal{M}), \frac{\tau J^\circ(e_{\boldsymbol{\theta}}) + 1}{\tau - 1})} \right) \right). \quad (3.5)$$

Proof. The proof follows the same lines as that of Theorem 3.1 except that we use the fact that $\hat{\boldsymbol{\theta}}_n^{\text{PEN}}$ is a global minimizer of V_n , i.e. $0 \in \partial V_n(\hat{\boldsymbol{\theta}}_n^{\text{PEN}})$. Indeed, we have for any $\boldsymbol{\theta} \in \mathbb{R}^p$

$$V_n(\boldsymbol{\theta}) \geq V_n(\hat{\boldsymbol{\theta}}_n^{\text{PEN}}) + \frac{1}{n} \varphi(\|\mathbf{X}\boldsymbol{\theta} - \mathbf{X}\hat{\boldsymbol{\theta}}_n^{\text{PEN}}\|_2). \quad (3.6)$$

Continuing exactly as just after (3.4), replacing $\hat{\boldsymbol{\theta}}_n^{\text{EWA}}$ with $\hat{\boldsymbol{\theta}}_n^{\text{PEN}}$ and invoking (3.6) instead of (3.4), we arrive at the claimed result. \square

Remark 3.2.

1. Observe that the penalized estimator $\hat{\boldsymbol{\theta}}_n^{\text{PEN}}$ does not require the moment assumption (H.2) for (3.5) to hold. The convexity assumption on φ in (H.1), which was important to apply Jensen's inequality in the proof of (3.2), is not needed either to get (3.5).
2. As we remarked for Theorem 3.1, Theorem 3.2 is also a deterministic statement for a fixed choice of λ_n that holds for any minimizer $\hat{\boldsymbol{\theta}}_n^{\text{PEN}}$, which is not unique in general. The condition on λ_n is similar to the one in [37] where authors established different guarantees for $\hat{\boldsymbol{\theta}}_n^{\text{PEN}}$.

One clearly sees that the difference between the prediction performance of $\hat{\boldsymbol{\theta}}_n^{\text{EWA}}$ and $\hat{\boldsymbol{\theta}}_n^{\text{PEN}}$ lies in the term $p\beta$ (or rather its lower-bound in Remark 3.1-5). Thus letting $\beta \rightarrow 0$ in (3.2), one recovers the oracle inequality (3.5) of penalized estimators. In particular, for $\beta = O((pn)^{-1})$, this is on the order $O(n^{-1})$.

3.3 Oracle inequalities in probability

It remains to check when the event $\mathcal{E} = \{\lambda_n \geq \tau J^\circ(-\mathbf{X}^\top \nabla F(\mathbf{X}\boldsymbol{\theta}_0, \mathbf{y}))/n\}$ holds with high probability when \mathbf{y} is random. We will use concentration inequalities in order to provide bounds that hold with high probability over the data. Toward this goal, we will need the following assumption.

(H.4) $\mathbf{y} = (\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_n)$ are independent and identically distributed observations, and $F(\mathbf{u}, \mathbf{y}) = \sum_{i=1}^n f_i(\mathbf{u}_i, \mathbf{y}_i)$, $f_i : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$. Moreover,

- (i) $\mathbb{E}[|f_i((\mathbf{X}\boldsymbol{\theta}_0)_i, \mathbf{y}_i)|] < +\infty, \forall 1 \leq i \leq n$;
- (ii) $|f'_i((\mathbf{X}\boldsymbol{\theta}_0)_i, t)| \leq g(t)$, where $\mathbb{E}[g(\mathbf{y}_i)] < +\infty, \forall 1 \leq i \leq n$;
- (iii) Bernstein moment condition: $\forall 1 \leq i \leq n$ and all integers $m \geq 2$, $\mathbb{E}[|f'_i((\mathbf{X}\boldsymbol{\theta}_0)_i, \mathbf{y}_i)|^m] \leq m! \kappa^{m-2} \sigma_i^2 / 2$ for some constants $\kappa > 0, \sigma_i > 0$ independent of n .

Observe that under **(H.4)**, and by virtue of Lemma A.4(iv) and [28, Proposition V.3.3.4], we have

$$J^\circ(-\mathbf{X}^\top \nabla F(\mathbf{X}\boldsymbol{\theta}_0, \mathbf{y})) = \sigma_{\mathcal{C}}(-\mathbf{X}^\top \nabla F(\mathbf{X}\boldsymbol{\theta}_0, \mathbf{y})) = \sup_{\mathbf{z} \in \mathcal{C}} - \sum_{i=1}^n f'_i((\mathbf{X}\boldsymbol{\theta}_0)_i, \mathbf{y}_i) \mathbf{z}_i. \quad (3.7)$$

Thus, checking the event \mathcal{E} amounts to establishing a deviation inequality for the supremum of an empirical process⁶ above its mean under the weak Bernstein moment condition **(H.4)**(iii), which essentially requires that the $f'_i((\mathbf{X}\boldsymbol{\theta}_0)_i, \mathbf{y}_i)$ have sub-exponential tails, We will first tackle the case where \mathcal{C} is the convex hull of a finite set (i.e. \mathcal{C} is a polytope).

3.3.1 Polyhedral penalty

We here suppose that is a finite gauge of $\mathcal{C} = \overline{\text{conv}}(\mathcal{V})$, where \mathcal{V} is finite, i.e. \mathcal{C} is a polytope with vertices [44, Corollary 19.1.1]. Our first oracle inequality in probability is the following.

Proposition 3.1. *Consider the estimators $\hat{\boldsymbol{\theta}}_n^{\text{EWA}}$ and $\hat{\boldsymbol{\theta}}_n^{\text{PEN}}$, where F and $J \stackrel{\text{def}}{=} \gamma_{\mathcal{C}}$ satisfy Assumptions **(H.1)**, **(H.2)**, **(H.3)** and **(H.4)**, and \mathcal{C} is a polytope with vertices \mathcal{V} . Suppose that $\text{rank}(\mathbf{X}) = n$ and $\max_{\mathbf{v} \in \mathcal{V}} \|\mathbf{X}\mathbf{v}\|_\infty \leq 1$, and take*

$$\lambda_n \geq \tau \sigma \sqrt{\frac{2\delta \log(|\mathcal{V}|)}{n}} \left(1 + \sqrt{2}\kappa/\sigma \sqrt{\frac{\delta \log(|\mathcal{V}|)}{n}} \right),$$

for some $\tau > 1$ and $\delta > 1$. Then (3.2) and (3.5) hold with probability at least $1 - 2|\mathcal{V}|^{1-\delta}$.

Proof. In view of Assumptions **(H.1)** and **(H.4)**, one can differentiate under the expectation sign (Leibniz rule) to conclude that $\mathbb{E}[F(\mathbf{X}\cdot, \mathbf{y})]$ is C^1 at $\boldsymbol{\theta}_0$ and $\nabla \mathbb{E}[F(\mathbf{X}\boldsymbol{\theta}_0, \mathbf{y})] = \mathbf{X}^\top \mathbb{E}[\nabla F(\mathbf{X}\boldsymbol{\theta}_0, \mathbf{y})]$. As $\boldsymbol{\theta}_0$ minimizes the population risk, one has $\nabla \mathbb{E}[F(\mathbf{X}\boldsymbol{\theta}_0, \mathbf{y})] = 0$. Using the rank assumption on \mathbf{X} , we deduce that

$$\mathbb{E}[f'_i((\mathbf{X}\boldsymbol{\theta}_0)_i, \mathbf{y}_i)] = 0, \quad \forall 1 \leq i \leq n.$$

Moreover, (3.7) specializes to

$$J^\circ(-\mathbf{X}^\top \nabla F(\mathbf{X}\boldsymbol{\theta}_0, \mathbf{y})) = \sup_{\mathbf{z} \in \mathcal{X}(\mathcal{V})} - \sum_{i=1}^n f'_i((\mathbf{X}\boldsymbol{\theta}_0)_i, \mathbf{y}_i) \mathbf{z}_i.$$

Let $t = \lambda_n n / \tau$. By the union bound and (3.7), we have

$$\begin{aligned} \mathbb{P}\left(J^\circ(-\mathbf{X}^\top \nabla F(\mathbf{X}\boldsymbol{\theta}_0, \mathbf{y})) \geq t\right) &\leq \mathbb{P}\left(\max_{\mathbf{z} \in \mathcal{X}(\mathcal{V})} - \sum_{i=1}^n f'_i((\mathbf{X}\boldsymbol{\theta}_0)_i, \mathbf{y}_i) \mathbf{z}_i \geq t\right) \\ &\leq |\mathcal{V}| \max_{\mathbf{z} \in \mathcal{X}(\mathcal{V})} \mathbb{P}\left(\left|\sum_{i=1}^n f'_i((\mathbf{X}\boldsymbol{\theta}_0)_i, \mathbf{y}_i) \mathbf{z}_i\right| \geq t\right). \end{aligned}$$

The random variables $(f'_i((\mathbf{X}\boldsymbol{\theta}_0)_i, \mathbf{y}_i) \mathbf{z}_i)_i$ are zero-mean independent, and $\forall i$ and $m \geq 2$

$$\mathbb{E}\left[|f'_i((\mathbf{X}\boldsymbol{\theta}_0)_i, \mathbf{y}_i) \mathbf{z}_i|^m\right] \leq |\mathbf{z}_i|^m m! \kappa^{m-2} \sigma_i^2 / 2 \leq \max_{\mathbf{v} \in \mathcal{V}} \|\mathbf{X}\mathbf{v}\|_\infty^m m! \kappa^{m-2} \sigma_i^2 / 2 \leq m! \kappa^{m-2} \sigma_i^2 / 2.$$

⁶As $\mathcal{X}(\mathcal{C})$ is compact, it has a dense countable subset.

We are then in position to apply the Bernstein inequality to get

$$\mathbb{P}\left(J^\circ(-\mathbf{X}^\top \nabla F(\mathbf{X}\boldsymbol{\theta}_0, \mathbf{y})) \geq t\right) \leq 2|\mathcal{V}| \exp\left(-\frac{t^2}{2(\kappa t + n\sigma^2)}\right),$$

where $\sigma^2 = \max_{1 \leq i \leq n} \sigma_i^2$. Every t such that

$$t \geq \sqrt{\delta \log(|\mathcal{V}|)} \left(\kappa \sqrt{\delta \log(|\mathcal{V}|)} + \sqrt{\kappa^2 \delta \log(|\mathcal{V}|) + 2n\sigma^2} \right),$$

satisfies $t^2 \geq 2\delta \log(|\mathcal{V}|)(\kappa t + n\sigma^2)$. Applying the trivial inequality $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ to the bound on t , we conclude. \square

Lasso To lighten the notation, let $I_\boldsymbol{\theta} = \text{supp}(\boldsymbol{\theta})$. From (2.3), it is easy to see that

$$\|\mathbb{P}_{T_\boldsymbol{\theta}}\|_{2 \rightarrow 1} = \sqrt{|I_\boldsymbol{\theta}|} \quad \text{and} \quad J^\circ(e_\boldsymbol{\theta}) = \|\text{sign}(\boldsymbol{\theta}_{I_\boldsymbol{\theta}})\|_\infty \leq 1,$$

where last bound holds as an equality whenever $\boldsymbol{\theta} \neq 0$. Further the ℓ_1 norm is the gauge of the cross-polytope (i.e. the unit ℓ_1 ball). Its vertex set \mathcal{V} is the set of unit-norm one-sparse vectors $(\pm \mathbf{a}_i)_{1 \leq i \leq p}$, where we recall $(\mathbf{a}_i)_{1 \leq i \leq p}$ the canonical basis. Thus

$$|\mathcal{V}| = 2p \quad \text{and} \quad \max_{\mathbf{v} \in \mathcal{V}} \|\mathbf{X}\mathbf{v}\|_2 = \max_{1 \leq i \leq p} \|\mathbf{X}_i\|_2.$$

Inserting this into Proposition 3.1, we obtain the following corollary.

Corollary 3.1. *Consider the estimators $\hat{\boldsymbol{\theta}}_n^{\text{EWA}}$ and $\hat{\boldsymbol{\theta}}_n^{\text{PEN}}$, where where J is the Lasso penalty and F satisfies Assumptions (H.1), (H.2) and (H.4). Suppose that $\text{rank}(\mathbf{X}) = n$ and $\max_i \|\mathbf{X}_i\|_\infty \leq 1$, and take*

$$\lambda_n \geq \tau \sigma \sqrt{\frac{2\delta \log(2p)}{n}} \left(1 + \sqrt{2}\kappa/\sigma \sqrt{\frac{\delta \log(2p)}{n}} \right),$$

for some $\tau > 1$ and $\delta > 1$. Then, with probability at least $1 - 2(2p)^{1-\delta}$, the following holds

$$R_n(\hat{\boldsymbol{\theta}}_n^{\text{EWA}}, \boldsymbol{\theta}_0) \leq \inf_{\substack{I \subset \{1, \dots, p\} \\ \boldsymbol{\theta}: \text{supp}(\boldsymbol{\theta})=I}} \left(R_n(\boldsymbol{\theta}, \boldsymbol{\theta}_0) + \frac{1}{n} \varphi^+ \left(\frac{\lambda_n \sqrt{n}(\tau+1)\sqrt{|I|}}{\tau \Upsilon(\text{Span}\{\mathbf{a}_i\}_{i \in I}, \frac{\tau+1}{\tau-1})} \right) \right) + p\beta, \quad (3.8)$$

and

$$R_n(\hat{\boldsymbol{\theta}}_n^{\text{PEN}}, \boldsymbol{\theta}_0) \leq \inf_{\substack{I \subset \{1, \dots, p\} \\ \boldsymbol{\theta}: \text{supp}(\boldsymbol{\theta})=I}} \left(R_n(\boldsymbol{\theta}, \boldsymbol{\theta}_0) + \frac{1}{n} \varphi^+ \left(\frac{\lambda_n \sqrt{n}(\tau+1)\sqrt{|I|}}{\tau \Upsilon(\text{Span}\{\mathbf{a}_i\}_{i \in I}, \frac{\tau+1}{\tau-1})} \right) \right). \quad (3.9)$$

For $\hat{\boldsymbol{\theta}}_n^{\text{PEN}}$, we recover a similar scaling for λ_n and the oracle inequality as in [60], though in the latter the oracle inequality is not sharp unlike ours. Note that the above oracle inequality extends readily to the case of analysis/fused Lasso $\|\mathbf{D}^\top \cdot\|_1$ where \mathbf{D} is surjective. We leave the details to the interested reader (see also the analysis group Lasso example in Section 4).

Anti-sparsity From Section 2.5.4, recall the saturation support $I_{\hat{\theta}}^{\text{sat}}$ of θ . From (2.10), we get

$$\|P_{T_{\theta}}\|_{2 \rightarrow \infty} = 1 \quad \text{and} \quad J^{\circ}(e_{\theta}) = \|\text{sign}(\theta_{I_{\hat{\theta}}^{\text{sat}}})\|_1 / |I_{\hat{\theta}}^{\text{sat}}| \leq 1,$$

with equality whenever $\theta \neq 0$. In addition, the ℓ_{∞} norm is the gauge of the hypercube whose vertex set is $\mathcal{V} = \{\pm 1\}^p$. Thus

$$|\mathcal{V}| = 2^p.$$

We have the following oracle inequalities.

Corollary 3.2. Consider the estimators $\hat{\theta}_n^{\text{EWA}}$ and $\hat{\theta}_n^{\text{PEN}}$, where where J is anti-sparsity penalty (2.9), and F satisfies Assumptions (H.1), (H.2) and (H.4). Suppose that $\text{rank}(\mathbf{X}) = n$ and $\max_{i,j} |\mathbf{X}_{i,j}| \leq 1/p$, and take

$$\lambda_n \geq \tau \sigma \sqrt{2\delta \log(2)} \sqrt{\frac{p}{n}} \left(1 + 2\kappa/\sigma \sqrt{\delta \log(2)} \sqrt{\frac{p}{n}} \right),$$

for some $\tau > 1$ and $\delta > 1$. Then, with probability at least $1 - 2(2p)^{1-\delta}$, the following holds

$$R_n(\hat{\theta}_n^{\text{EWA}}, \theta_0) \leq \inf_{\substack{I \subset \{1, \dots, p\} \\ \theta: I_{\hat{\theta}}^{\text{sat}} = I}} \left(R_n(\theta, \theta_0) + \frac{1}{n} \varphi^+ \left(\frac{\lambda_n \sqrt{n}(\tau+1)}{\tau \Upsilon(\{\bar{\theta} : \bar{\theta}_I \in \mathbb{R} \text{sign}(\theta_I)\}, \frac{\tau+1}{\tau-1})} \right) \right) + p\beta, \quad (3.10)$$

and

$$R_n(\hat{\theta}_n^{\text{PEN}}, \theta_0) \leq \inf_{\substack{I \subset \{1, \dots, p\} \\ \theta: I_{\hat{\theta}}^{\text{sat}} = I}} \left(R_n(\theta, \theta_0) + \frac{1}{n} \varphi^+ \left(\frac{\lambda_n \sqrt{n}(\tau+1)}{\tau \Upsilon(\{\bar{\theta} : \bar{\theta}_I \in \mathbb{R} \text{sign}(\theta_I)\}, \frac{\tau+1}{\tau-1})} \right) \right). \quad (3.11)$$

We are not aware of any result of this kind in the literature. The bound imposed on \mathbf{X} is similar to what is generally assumed in the vector quantization literature [36, 48].

3.3.2 General penalty

Extending the above reasoning to a general penalty requires a deviation inequality for the supremum of an empirical process in (3.7) under the Bernstein moment condition (H.4)(iii), but without the need of uniform boundedness. This can be achieved via generic chaining along a tree using entropy with bracketing; see [63, Theorem 8]. The resulting deviation bound will thus depend on the entropies with bracketing. These quantities capture the complexity of the set $\mathbf{X}(\mathcal{C})$ but are intricate to compute in general. This subject deserves further investigation that we leave to a future work.

Remark 3.3 (Group Lasso). Using the union bound, we have

$$\mathbb{P} \left(\max_{i \in \{1, \dots, L\}} \|\mathbf{X}_{b_i}^{\top} \boldsymbol{\xi}\|_2 \geq \lambda_n n / \tau \right) \leq \sum_{i=1}^L \mathbb{P} \left(\|\mathbf{X}_{b_i}^{\top} \boldsymbol{\xi}\|_2 \geq \lambda_n n / \tau \right).$$

This requires a concentration inequality for quadratic forms of independent random variables satisfying the Bernstein moment assumption above. We are not aware of any such a result. But if our moment condition is strengthened to

$$\mathbb{E} \left[|f'_i((\mathbf{X}\theta_0)_i, \mathbf{y}_i)|^{2m} \right] \leq m! \kappa^{2(m-1)} \sigma_i^2 / 2, \quad \forall 1 \leq i \leq n, \forall m \geq 1,$$

then one can use [4, Theorem 3]. Indeed, assume the normalization $\max_i \|\mathbf{X}_{b_i}^\top \mathbf{X}_{b_i}\|_{2 \rightarrow 2} \leq n$, which entails

$$\mathbb{E} \left[\|\mathbf{X}_{b_i}^\top \nabla F(\mathbf{X}\boldsymbol{\theta}_0, \mathbf{y})\|_2 \right] \leq \mathbb{E} \left[\|\mathbf{X}_{b_i}^\top \nabla F(\mathbf{X}\boldsymbol{\theta}_0, \mathbf{y})\|_2^2 \right]^{1/2} \leq \sigma \sqrt{Kn/2}.$$

It then follows that taking

$$\lambda_n \geq \tau \frac{\sigma \sqrt{K} + 16\kappa \sqrt{\delta \log(L)}}{n}, \quad \delta > 1,$$

the oracle inequalities (4.5) and (4.6) hold for the group Lasso with probability at least $1 - L^{1-\delta}$. A similar result can be proved for the analysis group Lasso just as well with a proper normalization assumption on \mathbf{X} (see Section 4.3.3).

4 Oracle inequalities for low-complexity linear regression

In this section, we consider the classical linear regression problem where the n response-covariate pairs $(\mathbf{y}_i, \mathbf{X}_i)$ are linked as

$$\mathbf{y} = \mathbf{X}\boldsymbol{\theta}_0 + \boldsymbol{\xi}, \quad (4.1)$$

where $\boldsymbol{\xi}$ is a noise vector. The data loss will be set to $F(\mathbf{u}, \mathbf{y}) = \frac{1}{2} \|\mathbf{y} - \mathbf{u}\|_2^2$. This in turn entails that $\varphi = \varphi^+ = \frac{1}{2}(\cdot)^2$ on \mathbb{R}_+ and $R_n(\boldsymbol{\theta}, \boldsymbol{\theta}_0) = \frac{1}{2n} \|\mathbf{X}\boldsymbol{\theta} - \mathbf{X}\boldsymbol{\theta}_0\|_2^2$.

In this section, we assume that the noise $\boldsymbol{\xi}$ is a zero-mean sub-Gaussian vector in \mathbb{R}^n with parameter σ . That is, its one-dimensional marginals $\langle \boldsymbol{\xi}, \mathbf{z} \rangle$ are sub-Gaussian random variables $\forall \mathbf{z} \in \mathbb{R}^n$, i.e. they satisfy

$$\mathbb{P}(|\langle \boldsymbol{\xi}, \mathbf{z} \rangle| \geq t) \leq 2e^{-t^2/(2\|\mathbf{z}\|_2^2\sigma^2)}, \quad \forall \mathbf{z} \in \mathbb{R}^n. \quad (4.2)$$

In this case, the bounds of Section 3.3 can be improved.

4.1 General penalty

As we will shortly show, the event \mathcal{E} will depend on the Gaussian width, a summary geometric quantity which, informally speaking, measures the size of the bulk of a set in \mathbb{R}^n .

Definition 4.1. *The Gaussian width of a subset $\mathcal{S} \subset \mathbb{R}^n$ is defined as*

$$w(\mathcal{S}) \stackrel{\text{def}}{=} \mathbb{E} [\sigma_{\mathcal{S}}(\mathbf{g})], \quad \text{where } \mathbf{g} \sim \mathcal{N}(0, \mathbf{Id}_n).$$

The concept of Gaussian width has appeared in the literature in different contexts. In particular, it has been used to establish sample complexity bounds to ensure exact recovery (noiseless case) and mean-square estimation stability (noisy case) for low-complexity penalized estimators from Gaussian measurements; see e.g. [46, 12, 54, 64, 58].

The Gaussian width has deep connections to convex geometry and it enjoys many useful properties. It is well-known that it is positively homogeneous, monotonic w.r.t. inclusion, and invariant under orthogonal transformations. Moreover, $w(\overline{\text{conv}}(\mathcal{S})) = w(\mathcal{S})$. From Lemma A.2(ii)-(iii), $w(\mathcal{S})$ is a non-negative finite quantity whenever the set \mathcal{S} is bounded and contains the origin.

We are now ready to state our oracle inequality in probability with sub-Gaussian noise.

Proposition 4.1. *Let the data generated by (4.1) where ξ is a zero-mean sub-Gaussian random vector with parameter σ . Consider the estimators $\widehat{\theta}_n^{\text{EWA}}$ and $\widehat{\theta}_n^{\text{PEN}}$, where F and $J \stackrel{\text{def}}{=} \gamma_{\mathcal{C}}$ satisfy Assumptions (H.1)-(H.2) and (H.3). Suppose that $\lambda_n \geq \frac{\tau \sigma c_1 \sqrt{2 \log(c_2/\delta) w(\mathbf{X}(\mathcal{C}))}}{n} <$, for some $\tau > 1$ and $0 < \delta < \min(c_2, 1)$, where c_1 and c_2 are positive absolute constants. Then with probability at least $1 - \delta$, (3.2) and (3.5) hold with the remainder term given by (3.3) with $\nu = 1$.*

The proof requires sophisticated ideas from the theory of generic chaining [51], but we only apply these results. The constants c_1 and c_2 can be traced back to the proof of these results as detailed in [51].

Proof. First, from (4.2), we have the bound

$$\mathbb{P}(|\langle \xi, z - z' \rangle| \geq t) \leq 2e^{-t^2/(2\|z - z'\|_2^2 \sigma^2)}, \quad \forall z, z' \in \mathbb{R}^n,$$

i.e. the increment condition [51, (0.4)] is verified. Thus combining (3.7) with the probability bound in [51, page 11], the generic chaining theorem [51, Theorem 1.2.6] and the majorizing measure theorem [51, Theorem 2.1.1], we have

$$\begin{aligned} \mathbb{P}\left(J^\circ(\mathbf{X}^\top \xi) \geq \lambda_n n / \tau\right) &\leq \mathbb{P}\left(\sup_{z \in \mathbf{X}(\mathcal{C})} \langle \xi, z \rangle \geq \sigma c_1 \sqrt{2 \log(c_2/\delta) w(\mathbf{X}(\mathcal{C}))}\right) \\ &\leq c_2 \exp\left(-\frac{\sigma^2 2 \log(c_2/\delta)}{2\sigma^2}\right) = \delta. \end{aligned}$$

□

If the noise is Gaussian, an enhanced version can be proved by invoking Gaussian concentration of Lipschitz functions [34].

Proposition 4.2. *Let the data generated by (4.1) with noise $\xi \sim \mathcal{N}(0, \sigma^2 \mathbf{Id}_n)$. Consider the estimators $\widehat{\theta}_n^{\text{EWA}}$ and $\widehat{\theta}_n^{\text{PEN}}$, where F and $J \stackrel{\text{def}}{=} \gamma_{\mathcal{C}}$ satisfy Assumptions (H.1)-(H.2) and (H.3). Suppose that $\lambda_n \geq \frac{(1+\delta)\tau \sigma w(\mathbf{X}(\mathcal{C}))}{n}$, for some $\tau > 1$ and $\delta > 0$. Then with probability at least $1 - \exp\left(-\frac{\delta^2 w(\mathbf{X}(\mathcal{C}))^2}{2\|\mathbf{X}\|_{J \rightarrow 2}^2}\right)$, (3.2) and (3.5) hold with the remainder term given by (3.3) with $\nu = 1$.*

Proof. Thanks to sublinearity (see Lemma A.3(i) and Lemma A.4), the function $\xi \mapsto J^\circ(\mathbf{X}^\top \xi)$ is Lipschitz continuous with Lipschitz constant $\|\mathbf{X}^\top\|_{2 \rightarrow J^\circ} = \|\mathbf{X}\|_{J \rightarrow 2}$. From (3.7), we also have

$$\mathbb{E}\left[J^\circ(\mathbf{X}^\top \xi)\right] = \sigma w(\mathbf{X}(\mathcal{C})).$$

Observe that $\mathbf{X}(\mathcal{C})$ is a convex compact set containing the origin. Setting $\epsilon = \lambda_n n / \tau - \sigma w(\mathbf{X}(\mathcal{C})) \geq \delta \sigma w(\mathbf{X}(\mathcal{C}))$, it follows from (3.7) and the Gaussian concentration of Lipschitz functions [34] that

$$\begin{aligned} \mathbb{P}\left(J^\circ(\mathbf{X}^\top \xi) \geq \lambda_n n / \tau\right) &\leq \mathbb{P}\left(J^\circ(\mathbf{X}^\top \xi) - \mathbb{E}\left[J^\circ(\mathbf{X}^\top \xi)\right] \geq \epsilon\right) \\ &\leq \mathbb{P}\left(J^\circ(\mathbf{X}^\top \xi / \sigma) - w(\mathbf{X}(\mathcal{C})) \geq \delta w(\mathbf{X}(\mathcal{C}))\right) \\ &\leq \exp\left(-\frac{\delta^2 w(\mathbf{X}(\mathcal{C}))^2}{2\|\mathbf{X}\|_{J \rightarrow 2}^2}\right). \end{aligned}$$

□

Estimating theoretically the Gaussian width of a set⁷ is a non-trivial problem that has been extensively studied in the areas of probability in Banach spaces and stochastic processes. There are classical bounds on the Gaussian width (Sudakov's and Dudley's inequalities), but they are difficult to estimate in most cases and neither of these bounds is tight for all sets. When the set is a convex cone (intersected with a sphere), tractable estimates based on polarity arguments were proposed in, e.g., [12].

4.2 Polyhedral penalty

When \mathcal{C} is a polytope, enhanced oracle inequalities can be obtained by invoking a simple union bound argument.

Proposition 4.3. *Let the data generated by (4.1) where $\boldsymbol{\xi}$ is a zero-mean sub-Gaussian random vector with parameter σ . Consider the estimators $\hat{\boldsymbol{\theta}}_n^{\text{EWA}}$ and $\hat{\boldsymbol{\theta}}_n^{\text{PEN}}$, where F and $J \stackrel{\text{def}}{=} \gamma_{\mathcal{C}}$ satisfy Assumptions (H.1)-(H.2) and (H.3), and moreover \mathcal{C} is a polytope with vertices \mathcal{V} . Suppose that $\lambda_n \geq \frac{\tau\sigma(\max_{\mathbf{v} \in \mathcal{V}} \|\mathbf{X}\mathbf{v}\|_2)\sqrt{2\delta \log(|\mathcal{V}|)}}{n}$, for some $\tau > 1$ and $\delta > 1$. Then with probability at least $1 - 2|\mathcal{V}|^{1-\delta}$, (3.2) and (3.5) hold with the remainder term given by (3.3) with $\nu = 1$.*

In particular, if $\max_{\mathbf{v} \in \mathcal{V}} \|\mathbf{X}\mathbf{v}\|_2 \leq \sqrt{n}$, the first term in the choice of λ_n scales as $\tau\sigma\sqrt{\frac{2\delta \log(|\mathcal{V}|)}{n}}$.

Proof. From (3.7) we have

$$J^\circ(\mathbf{X}^\top \boldsymbol{\xi}) = \max_{\mathbf{v} \in \mathcal{C}} \langle \mathbf{X}\mathbf{v}, \boldsymbol{\xi} \rangle = \max_{\mathbf{v} \in \mathcal{V}} \langle \mathbf{X}\mathbf{v}, \boldsymbol{\xi} \rangle,$$

where in the last inequality, we used the fact that a convex function attains its maximum on \mathcal{C} at an extreme point \mathcal{V} . Let $\epsilon = \sigma(\max_{\mathbf{v} \in \mathcal{V}} \|\mathbf{X}\mathbf{v}\|_2)\sqrt{2\delta \log(|\mathcal{V}|)}$. By the union bound, (4.2) and (3.7), we have

$$\begin{aligned} \mathbb{P}\left(J^\circ(\mathbf{X}^\top \boldsymbol{\xi}) \geq \lambda_n n / \tau\right) &\leq \mathbb{P}\left(\max_{\mathbf{v} \in \mathcal{V}} \langle \mathbf{X}\mathbf{v}, \boldsymbol{\xi} \rangle \geq \epsilon\right) \\ &\leq |\mathcal{V}| \max_{\mathbf{v} \in \mathcal{V}} \mathbb{P}(\langle \mathbf{X}\mathbf{v}, \boldsymbol{\xi} \rangle \geq \epsilon) \\ &\leq |\mathcal{V}| \max_{\mathbf{v} \in \mathcal{V}} \mathbb{P}(|\langle \mathbf{X}\mathbf{v}, \boldsymbol{\xi} \rangle| \geq \epsilon) \\ &\leq 2|\mathcal{V}| \exp\left(-\epsilon^2 / (2\sigma^2 \max_{\mathbf{v} \in \mathcal{V}} \|\mathbf{X}\mathbf{v}\|_2^2)\right) \\ &\leq 2|\mathcal{V}|^{1-\delta}. \end{aligned}$$

□

4.3 Applications

In this section, we exemplify our oracle inequalities for the penalties described in Section 2.5.

4.3.1 Lasso

Recall the derivations for the Lasso in Section 3.3.1. We obtain the following corollary of Proposition 4.3.

⁷Not to mention its image with a linear operator as for $\mathbf{X}(\mathcal{C})$.

Corollary 4.1. *Let the data generated by (4.1) where ξ is a zero-mean sub-Gaussian random vector with parameter σ . Assume that \mathbf{X} is such that $\max_i \|\mathbf{X}_i\|_2 \leq \sqrt{n}$. Consider the estimators $\hat{\boldsymbol{\theta}}_n^{\text{EWA}}$ and $\hat{\boldsymbol{\theta}}_n^{\text{PEN}}$, where J is the Lasso penalty (2.2) and F satisfies Assumptions (H.1)-(H.2). Suppose that $\lambda_n \geq \tau\sigma\sqrt{\frac{2\delta\log(2p)}{n}}$, for some $\tau > 1$ and $\delta > 1$. Then, with probability at least $1 - 2(2p)^{1-\delta}$, the following holds*

$$R_n(\hat{\boldsymbol{\theta}}_n^{\text{EWA}}, \boldsymbol{\theta}_0) \leq \inf_{\substack{I \subset \{1, \dots, p\} \\ \boldsymbol{\theta}: \text{supp}(\boldsymbol{\theta})=I}} \left(R_n(\boldsymbol{\theta}, \boldsymbol{\theta}_0) + \frac{\lambda_n^2(\tau+1)^2|I|}{2\tau^2\nu\Upsilon\left(\text{Span}\{\mathbf{a}_i\}_{i \in I, \frac{\tau+1}{\tau-1}}\right)^2} \right) + p\beta, \quad (4.3)$$

and

$$R_n(\hat{\boldsymbol{\theta}}_n^{\text{PEN}}, \boldsymbol{\theta}_0) \leq \inf_{\substack{I \subset \{1, \dots, p\} \\ \boldsymbol{\theta}: \text{supp}(\boldsymbol{\theta})=I}} \left(R_n(\boldsymbol{\theta}, \boldsymbol{\theta}_0) + \frac{\lambda_n^2(\tau+1)^2|I|}{2\tau^2\nu\Upsilon\left(\text{Span}\{\mathbf{a}_i\}_{i \in I, \frac{\tau+1}{\tau-1}}\right)^2} \right). \quad (4.4)$$

The remainder term grows as $\frac{|I|\log(p)}{n}$. The oracle inequality (4.4) recovers [17, Theorem 1] in the exactly sparse case, and (4.4) the one in [50, Theorem 4] (see also [18, Theorem 2]). It is worth mentioning, however, that [17, Theorem 1] handles the inexactly sparse case while we do not.

4.3.2 Group Lasso

Recall the notations in Section 2.5.2, and denote $I_\theta = \text{supp}_B(\boldsymbol{\theta})$ the number of active blocks in $\boldsymbol{\theta}$. From (2.5), we have

$$\|\mathbb{P}_{T_\theta}\|_{2 \rightarrow J} = \sqrt{|I_\theta|} \quad \text{and} \quad J^\circ(e_\theta) = \|e_\theta\|_{\infty, 2} \leq 1,$$

where the last bound holds as an equality whenever $\boldsymbol{\theta} \neq 0$.

We have the following oracle inequalities as corollaries of Proposition 4.1 and Proposition 4.2.

Corollary 4.2. *Let the data generated by (4.1). Consider the estimators $\hat{\boldsymbol{\theta}}_n^{\text{EWA}}$ and $\hat{\boldsymbol{\theta}}_n^{\text{PEN}}$, where F satisfies Assumptions (H.1)-(H.2), and J is the group Lasso (2.4) with L non-overlapping blocks of equal size K . Assume that \mathbf{X} is such that $\max_i \|\mathbf{X}_{b_i}^\top \mathbf{X}_{b_i}\|_{2 \rightarrow 2} \leq n$.*

(i) ξ is a zero-mean sub-Gaussian random vector with parameter σ : suppose that $\lambda_n \geq 3/2\tau\sigma c_1 \frac{\sqrt{\log(c_2/\delta)}(\sqrt{K} + \sqrt{2\log(L)})}{\sqrt{n}}$, for some $\tau > 1$ and $0 < \delta < \min(c_2, 1)$, where c_1 and c_2 are the positive absolute constants in Proposition 4.1. Then, with probability at least $1 - \delta$, the following holds

$$R_n(\hat{\boldsymbol{\theta}}_n^{\text{EWA}}, \boldsymbol{\theta}_0) \leq \inf_{\substack{I \subset \{1, \dots, L\} \\ \boldsymbol{\theta}: \text{supp}_B(\boldsymbol{\theta})=I}} \left(R_n(\boldsymbol{\theta}, \boldsymbol{\theta}_0) + \frac{\lambda_n^2(\tau+1)^2|I|}{2\tau^2\nu\Upsilon\left(\text{Span}\{\mathbf{a}_j\}_{j \in b_i, i \in I, \frac{\tau+1}{\tau-1}}\right)^2} \right) + p\beta, \quad (4.5)$$

and

$$R_n(\hat{\boldsymbol{\theta}}_n^{\text{PEN}}, \boldsymbol{\theta}_0) \leq \inf_{\substack{I \subset \{1, \dots, L\} \\ \boldsymbol{\theta}: \text{supp}_B(\boldsymbol{\theta})=I}} \left(R_n(\boldsymbol{\theta}, \boldsymbol{\theta}_0) + \frac{\lambda_n^2(\tau+1)^2|I|}{2\tau^2\nu\Upsilon\left(\text{Span}\{\mathbf{a}_j\}_{j \in b_i, i \in I, \frac{\tau+1}{\tau-1}}\right)^2} \right). \quad (4.6)$$

(ii) $\xi \sim \mathcal{N}(0, \sigma^2 \mathbf{Id}_n)$: suppose that $\lambda_n \geq \tau\sigma \frac{\sqrt{K} + \sqrt{2\delta\log(L)}}{\sqrt{n}}$, for some $\tau > 1$ and $\delta > 1$. Then, with probability at least $1 - L^{1-\delta}$, (4.5) and (4.6) hold.

The first remainder term is on the order $\frac{|I|(\sqrt{K} + \sqrt{2\log(L)})^2}{n}$. This is similar to the scaling that has been provided in the literature for EWA with other group sparsity priors and noises [43, 24]. Similar rates were given for $\hat{\boldsymbol{\theta}}_n^{\text{PEN}}$ with the group Lasso in [37, 35, 61].

Proof.

(i) This is a consequence of Proposition 4.1, for which we need to bound

$$w(\mathbf{X}(\mathcal{C})) = \mathbb{E} \left[\max_{i \in \{1, \dots, L\}} \|\mathbf{X}_{b_i}^\top \mathbf{g}\|_2 \right].$$

We first have, for any block b_i

$$\mathbb{E} \left[\|\mathbf{X}_{b_i}^\top \mathbf{g}\|_2 \right] \leq \mathbb{E} \left[\|\mathbf{X}_{b_i}^\top \mathbf{g}\|_2^2 \right]^{1/2} \leq \sqrt{Kn}.$$

Furthermore, $\|\mathbf{X}_{b_i}^\top \cdot\|_2$ is Lipschitz continuous with Lipschitz constant $\|\mathbf{X}_{b_i}\|_{2 \rightarrow 2} \leq \sqrt{n}$. Thus the union bound and Gaussian concentration of Lipschitz functions [34] yield, for any $t > 0$,

$$\mathbb{P} \left(\max_{i \in \{1, \dots, L\}} \|\mathbf{X}_{b_i}^\top \mathbf{g}\|_2 \geq \sqrt{Kn} + t \right) \leq \sum_{i=1}^L \mathbb{P} \left(\|\mathbf{X}_{b_i}^\top \boldsymbol{\xi}\|_2 - \mathbb{E} \left[\|\mathbf{X}_{b_i}^\top \boldsymbol{\xi}\|_2 \right] \geq t \right) \leq L \exp \left(-\frac{t^2}{2n} \right).$$

Let $\kappa = \sqrt{Kn} + \sqrt{2n \log(L)}$. $w(\mathbf{X}(\mathcal{C}))$ can be expressed as

$$\begin{aligned} w(\mathbf{X}(\mathcal{C})) &= \int_0^\infty \mathbb{P} \left(\max_{i \in \{1, \dots, L\}} \|\mathbf{X}_{b_i}^\top \mathbf{g}\|_2 \geq s \right) ds \leq \int_0^\kappa ds + \int_\kappa^\infty e^{-\frac{(s-\sqrt{Kn})^2 - 2n \log(L)}{2n}} ds \\ &= \kappa + \sqrt{n} \int_{\kappa/\sqrt{n}}^\infty e^{-\frac{(s-\sqrt{K})^2 - 2 \log(L)}{2}} ds \\ &\leq \kappa + \sqrt{n} \int_{\kappa/\sqrt{n}}^\infty e^{-\frac{s-\kappa/\sqrt{n}}{2}} ds = \kappa + \sqrt{n}/2 \leq 3\kappa/2. \end{aligned}$$

(ii) The proof follows the lines of Proposition 4.2 where we additionally use the union bound. Indeed,

$$\begin{aligned} \mathbb{P} \left(\max_{i \in \{1, \dots, L\}} \|\mathbf{X}_{b_i}^\top \boldsymbol{\xi}\|_2 \geq \lambda_n n / \tau \right) &\leq \sum_{i=1}^L \mathbb{P} \left(\|\mathbf{X}_{b_i}^\top \boldsymbol{\xi}\|_2 - \mathbb{E} \left[\|\mathbf{X}_{b_i}^\top \boldsymbol{\xi}\|_2 \right] \geq \lambda_n n / \tau - \mathbb{E} \left[\|\mathbf{X}_{b_i}^\top \boldsymbol{\xi}\|_2 \right] \right) \\ &\leq \sum_{i=1}^L \mathbb{P} \left(\|\mathbf{X}_{b_i}^\top \boldsymbol{\xi}\|_2 - \mathbb{E} \left[\|\mathbf{X}_{b_i}^\top \boldsymbol{\xi}\|_2 \right] \geq \lambda_n n / \tau - \sigma \sqrt{Kn} \right) \\ &\leq \sum_{i=1}^L \mathbb{P} \left(\|\mathbf{X}_{b_i}^\top \boldsymbol{\xi}\|_2 - \mathbb{E} \left[\|\mathbf{X}_{b_i}^\top \boldsymbol{\xi}\|_2 \right] \geq \sigma \sqrt{2\delta n \log(L)} \right) \\ &\leq L \exp(-\delta \log(L)) = L^{1-\delta}, \end{aligned}$$

where used the Gaussian concentration of Lipschitz functions [34] in the last inequality. □

We observe in passing that another way to prove the oracle inequalities in the sub-Gaussian is to use Dudley's inequality on the sphere in \mathbb{R}^K after applying a union bound on the L blocks.

4.3.3 Analysis group Lasso

We now turn to the prior penalty (2.6). Recall the notations in Section 2.5.3, and remind $\Lambda_\theta = \bigcup_{i \in \text{supp}_B(\mathbf{D}^\top \theta)} b_i$. We assume that \mathbf{D} is a frame of \mathbb{R}^p , hence surjective, meaning that there exist $c, d > 0$ such that for any $\omega \in \mathbb{R}^p$

$$d \|\omega\|_2^2 \leq \|\mathbf{D}^\top \omega\|_2^2 \leq c \|\omega\|_2^2.$$

This together with (2.7)-(2.8) and Cauchy-Schwarz inequality entail

$$\begin{aligned} \|\mathbb{P}_{T_\theta}\|_{2 \rightarrow J} &= \sup_{\|\omega_{T_\theta}\|_2 \leq 1} \|\mathbf{D}^\top \omega_{T_\theta}\|_{1,2} \leq \sqrt{c} \sup_{\|\mathbf{D}^\top \omega_{T_\theta}\|_2 \leq 1} \|\mathbf{D}^\top \omega_{T_\theta}\|_{1,2} \\ &= \sqrt{c} \sup_{\|\mathbf{D}_{\Lambda_\theta}^\top \omega_{T_\theta}\|_2 \leq 1} \|\mathbf{D}_{\Lambda_\theta}^\top \omega_{T_\theta}\|_{1,2} \\ &= \sqrt{c} \sqrt{|\text{supp}_B(\mathbf{D}^\top \theta)|}. \end{aligned}$$

Note, however, that from (2.7), we do not have in general $\left\| \mathbf{D}^+ \mathbb{P}_{\text{Ker}(\mathbf{D}_{\Lambda_\theta}^\top)} \mathbf{D} e_{\mathbf{D}^\top \theta}^{\|\cdot\|_{1,2}} \right\|_{\infty,2} \leq 1$.

With exactly the same arguments to those for proving Corollary 4.2, replacing \mathbf{X} by $\mathbf{X}\mathbf{D}$, we arrive at the following oracle inequalities.

Corollary 4.3. *Let the data generated by (4.1). Consider the estimators $\hat{\theta}_n^{\text{EWA}}$ and $\hat{\theta}_n^{\text{PEN}}$, where F satisfies Assumptions (H.1)-(H.2), and J is the analysis group Lasso (2.6) with L blocks of equal size K . Assume that \mathbf{D} is a frame, and \mathbf{X} is such that $\max_i \|\mathbf{D}_{b_i}^\top \mathbf{X}^\top \mathbf{X} \mathbf{D}_{b_i}\|_{2 \rightarrow 2} \leq n$.*

(i) ξ is a zero-mean sub-Gaussian random vector with parameter σ : suppose that $\lambda_n \geq 3\tau\sigma c_1 \frac{\sqrt{\log(c_2/\delta)}(\sqrt{K} + \sqrt{2\log(L)})}{\sqrt{n}}$, for some $\tau > 1$ and $0 < \delta < \min(c_2, 1)$, where c_1 and c_2 are the positive absolute constants in Proposition 4.1. Then, with probability at least $1 - \delta$, the following holds

$$R_n(\hat{\theta}_n^{\text{EWA}}, \theta_0) \leq \inf_{\substack{I \subset \{1, \dots, L\} \\ \theta: \text{supp}_B(\mathbf{D}^\top \theta) = I}} \left(R_n(\theta, \theta_0) + \frac{c\lambda_n^2 (\tau \|\mathbf{D}^+ \mathbb{P}_{\text{Ker}(\mathbf{D}_{\Lambda_\theta}^\top)} \mathbf{D} e_{\mathbf{D}^\top \theta}^{\|\cdot\|_{1,2}}\|_{\infty,2} + 1)^2 |I|}{2\tau^2 \nu \Upsilon \left(\text{Ker}(\mathbf{D}_{\Lambda_\theta}^\top), \frac{\tau \|\mathbf{D}^+ \mathbb{P}_{\text{Ker}(\mathbf{D}_{\Lambda_\theta}^\top)} \mathbf{D} e_{\mathbf{D}^\top \theta}^{\|\cdot\|_{1,2}}\|_{\infty,2} + 1}{\tau - 1} \right)^2} \right) + p\beta, \quad (4.7)$$

and

$$R_n(\hat{\theta}_n^{\text{PEN}}, \theta_0) \leq \inf_{\substack{I \subset \{1, \dots, L\} \\ \theta: \text{supp}_B(\mathbf{D}^\top \theta) = I}} \left(R_n(\theta, \theta_0) + \frac{c\lambda_n^2 (\tau \|\mathbf{D}^+ \mathbb{P}_{\text{Ker}(\mathbf{D}_{\Lambda_\theta}^\top)} \mathbf{D} e_{\mathbf{D}^\top \theta}^{\|\cdot\|_{1,2}}\|_{\infty,2} + 1)^2 |I|}{2\tau^2 \nu \Upsilon \left(\text{Ker}(\mathbf{D}_{\Lambda_\theta}^\top), \frac{\tau \|\mathbf{D}^+ \mathbb{P}_{\text{Ker}(\mathbf{D}_{\Lambda_\theta}^\top)} \mathbf{D} e_{\mathbf{D}^\top \theta}^{\|\cdot\|_{1,2}}\|_{\infty,2} + 1}{\tau - 1} \right)^2} \right). \quad (4.8)$$

(ii) $\boldsymbol{\xi} \sim \mathcal{N}(0, \sigma^2 \mathbf{Id}_n)$: suppose that $\lambda_n \geq \tau \sigma \frac{\sqrt{K} + \sqrt{2\delta \log(L)}}{\sqrt{n}}$, for some $\tau > 1$ and $\delta > 1$. Then, with probability at least $1 - L^{1-\delta}$, (4.7) and (4.8) hold.

To the best of our knowledge, this result is new to the literature. The scaling of the remainder term is the same as in [24] and [43] with analysis sparsity priors different from ours (the authors in the latter also assume that \mathbf{D} is invertible).

4.3.4 Anti-sparsity

Recall the derivations for the ℓ_∞ norm example in Section 3.3.1. We have the following oracle inequalities from Proposition 4.3.

Corollary 4.4. *Let the data generated by (4.1) where $\boldsymbol{\xi}$ is a zero-mean sub-Gaussian random vector with parameter σ . Assume that \mathbf{X} is such that $\max_{i,j} |\mathbf{X}_{i,j}| \leq 1/p$. Consider the estimators $\widehat{\boldsymbol{\theta}}_n^{\text{EWA}}$ and $\widehat{\boldsymbol{\theta}}_n^{\text{PEN}}$, where F satisfies Assumptions (H.1)-(H.2), and J is the anti-sparsity penalty (2.9). Suppose that $\lambda_n \geq \tau \sigma \sqrt{2\delta \log(2)} \sqrt{\frac{p}{n}}$, for some $\tau > 1$ and $\delta > 1$. Then, with probability at least $1 - 2e^{(1-\delta) \log(2)p}$, the following holds*

$$R_n(\widehat{\boldsymbol{\theta}}_n^{\text{EWA}}, \boldsymbol{\theta}_0) \leq \inf_{\substack{I \subset \{1, \dots, p\} \\ \boldsymbol{\theta}: I_{\boldsymbol{\theta}}^{\text{sat}} = I}} \left(R_n(\boldsymbol{\theta}, \boldsymbol{\theta}_0) + \frac{\lambda_n^2 (\tau+1)^2}{2\tau^2 \nu \Upsilon \left(\left\{ \bar{\boldsymbol{\theta}} : \bar{\boldsymbol{\theta}}_I \in \mathbb{R} \text{ sign}(\boldsymbol{\theta}_I) \right\}, \frac{\tau+1}{\tau-1} \right)^2} \right) + p\beta, \quad (4.9)$$

and

$$R_n(\widehat{\boldsymbol{\theta}}_n^{\text{PEN}}, \boldsymbol{\theta}_0) \leq \inf_{\substack{I \subset \{1, \dots, p\} \\ \boldsymbol{\theta}: I_{\boldsymbol{\theta}}^{\text{sat}} = I}} \left(R_n(\boldsymbol{\theta}, \boldsymbol{\theta}_0) + \frac{\lambda_n^2 (\tau+1)^2}{2\tau^2 \nu \Upsilon \left(\left\{ \bar{\boldsymbol{\theta}} : \bar{\boldsymbol{\theta}}_I \in \mathbb{R} \text{ sign}(\boldsymbol{\theta}_I) \right\}, \frac{\tau+1}{\tau-1} \right)^2} \right). \quad (4.10)$$

The first remainder term scales as $\frac{p}{n}$ which reflects that

4.3.5 Nuclear norm

We now turn to the nuclear norm case. Recall the notations of Section 2.5.5. For matrices $\boldsymbol{\theta} \in \mathbb{R}^{p_1 \times p_2}$, a measurement map \mathbf{X} takes the form of a linear operator whose i th component is given by the Frobenius scalar product

$$\mathbf{X}(\boldsymbol{\theta})_i = \text{tr}((\mathbf{X}^i)^\top \boldsymbol{\theta}) = \langle \mathbf{X}^i, \boldsymbol{\theta} \rangle_{\text{F}},$$

where \mathbf{X}^i is a matrix in $\mathbb{R}^{p_1 \times p_2}$. We denote $\|\cdot\|_{\text{F}}$ the associated norm. From (2.12), it is immediate to see that whenever $\boldsymbol{\theta} \neq 0$,

$$J^\circ(e_{\boldsymbol{\theta}}) = \|\mathbf{U}\mathbf{V}^\top\|_{2 \rightarrow 2} = 1.$$

Moreover, $\mathbf{P}_{T_{\boldsymbol{\theta}}}(\mathbf{A}) = \mathbf{U}\mathbf{U}^\top \mathbf{A} + (\mathbf{Id}_{p_1} - \mathbf{U}\mathbf{U}^\top) \mathbf{A} \mathbf{V}\mathbf{V}^\top$. Thus since

$$\|\boldsymbol{\theta}\|_* = \|\lambda(\boldsymbol{\theta})\|_1 \leq \sqrt{r} \|\lambda(\boldsymbol{\theta})\|_2 = \sqrt{r} \|\boldsymbol{\theta}\|_{\text{F}},$$

we deduce that

$$\|\mathbf{P}_{T_{\boldsymbol{\theta}}}\|_{\text{F} \rightarrow *} \leq 2\sqrt{r}.$$

To apply Proposition 4.1 and Proposition 4.2, we need to bound $w(\mathbf{X}(\mathcal{C}))$ (\mathcal{C} is the nuclear ball), or equivalently, to bound

$$\mathbb{E} [\|\mathbf{X}^*(\mathbf{g})\|_{2 \rightarrow 2}] = \mathbb{E} \left[\left\| \sum_{i=1}^n \mathbf{X}^i \mathbf{g}_i \right\|_{2 \rightarrow 2} \right], \quad \mathbf{g} \sim \mathcal{N}(0, \sigma^2 \mathbf{Id}_n),$$

which is the expectation of the operator norm of a random series with matrix coefficients. Thus using [55, Theorem 4.1.1(4.1.5)] to get this bound, and inserting it into Proposition 4.1 and Proposition 4.2, we get the following oracle inequalities for the nuclear norm. Define

$$v(\mathbf{X}) = \max \left(\left\| \sum_{i=1}^n \mathbf{X}^i (\mathbf{X}^i)^\top \right\|_{2 \rightarrow 2}, \left\| \sum_{i=1}^n (\mathbf{X}^i)^\top \mathbf{X}^i \right\|_{2 \rightarrow 2} \right).$$

Corollary 4.5. *Let the data generated by (4.1) with a linear operator $\mathbf{X} : \mathbb{R}^{p_1 \times p_2} \rightarrow \mathbb{R}^n$. Assume that $v(\mathbf{X}) \leq n$. Consider the estimators $\hat{\boldsymbol{\theta}}_n^{\text{EWA}}$ and $\hat{\boldsymbol{\theta}}_n^{\text{PEN}}$, where F satisfies Assumptions (H.1)-(H.2), and J is the nuclear norm (2.11).*

(i) $\boldsymbol{\xi}$ is a zero-mean sub-Gaussian random vector with parameter σ : suppose that $\lambda_n \geq 2\tau\sigma c_1 \sqrt{\frac{\log(c_2/\delta) \log(p_1+p_2)}{n}}$, for some $\tau > 1$ and $0 < \delta < \min(c_2, 1)$, where c_1 and c_2 are the positive absolute constants in Proposition 4.1. Then, with probability at least $1 - \delta$, the following holds

$$R_n(\hat{\boldsymbol{\theta}}_n^{\text{EWA}}, \boldsymbol{\theta}_0) \leq \inf_{\substack{r \in \{1, \dots, \min(p_1, p_2)\} \\ \boldsymbol{\theta}: \text{rank } \boldsymbol{\theta} = r}} \left(R_n(\boldsymbol{\theta}, \boldsymbol{\theta}_0) + \frac{2\lambda_n^2(\tau+1)^2 r}{\tau^2 \nu \Upsilon \left(T_{\boldsymbol{\theta}}, \frac{\tau+1}{\tau-1} \right)^2} \right) + p_1 p_2 \beta, \quad (4.11)$$

and

$$R_n(\hat{\boldsymbol{\theta}}_n^{\text{PEN}}, \boldsymbol{\theta}_0) \leq \inf_{\substack{r \in \{1, \dots, \min(p_1, p_2)\} \\ \boldsymbol{\theta}: \text{rank } \boldsymbol{\theta} = r}} \left(R_n(\boldsymbol{\theta}, \boldsymbol{\theta}_0) + \frac{2\lambda_n^2(\tau+1)^2 r}{\tau^2 \nu \Upsilon \left(T_{\boldsymbol{\theta}}, \frac{\tau+1}{\tau-1} \right)^2} \right). \quad (4.12)$$

(ii) $\boldsymbol{\xi} \sim \mathcal{N}(0, \sigma^2 \mathbf{Id}_n)$: suppose that $\lambda_n \geq (1 + \delta)\tau\sigma \sqrt{\frac{2 \log(p_1+p_2)}{n}}$, for some $\tau > 1$ and $\delta > 0$. Then, with probability at least $1 - (p_1 + p_2)^{-\delta^2}$, (4.11) and (4.12) hold.

The set over which the infimum is taken just reminds us that the nuclear norm is partly smooth (see above) relative to the constant rank manifold (which is a Riemannian submanifold of $\mathbb{R}^{p_1 \times p_2}$) [21, Theorem 3.19]. The first remainder term now scales as $\frac{r \log(p_1+p_2)}{n}$. In the iid Gaussian case, we recover the same rate as in [17, Theorem 3] for $\hat{\boldsymbol{\theta}}_n^{\text{EWA}}$ and in [32, Theorem 2] for $\hat{\boldsymbol{\theta}}_n^{\text{PEN}}$.

A Pre-requisites from convex analysis

We here collect some ingredients from convex analysis that are essential to our exposition.

Monotone conjugate

Lemma A.1. *Let g be a non-decreasing function on \mathbb{R}_+ that vanishes at 0. Then the following hold:*

- (i) g^+ is a proper closed convex and non-decreasing function on \mathbb{R}_+ that vanishes at 0.
- (ii) If g is also closed and convex, then $g^{++} = g$.
- (iii) Let $f : t \in \mathbb{R} \mapsto g(|t|)$ such that f is differentiable on \mathbb{R} , where g is finite-valued, strictly convex and strongly coercive. Then g^+ is likewise finite-valued, strictly convex, strongly coercive, and $f^* = g^+ \circ |\cdot|$ is differentiable on \mathbb{R} . In particular, both g and g^+ are strictly increasing on \mathbb{R}_+ .

Proof. (i) By [3, Proposition 13.11], g^+ is a closed convex function. We have $\inf_{t \geq 0} g(t) = -\sup_{t \geq 0} t \cdot 0 - g(t) = -g^+(0)$. Since g is non-decreasing and $g(0) = 0$, then $g^+(0) = -\inf_{t \geq 0} g(t) = -g(0) = 0$. In addition, by (1.5), we have $g^+(a) \geq a \cdot 0 - g(0) = 0, \forall a \in \mathbb{R}_+$. This shows that g^+ is non-negative and $\text{dom}(g^+) \neq \emptyset$, and in turn, it is also proper.

Let a, b in \mathbb{R}_+ such that $a < b$. Then

$$g^+(a) - g^+(b) = \left(\sup_{t \geq 0} ta - g(t) \right) - \left(\sup_{t' \geq 0} t'b - g(t') \right) \leq \sup_{t \geq 0} (ta - g(t) - tb + g(t)) = \sup_{t \geq 0} t(a - b) = 0.$$

That is, g^+ is non-decreasing on \mathbb{R}_+ .

- (ii) This follows from [44, Theorem 12.4].
- (iii) By definition of f , f is a finite-valued function on \mathbb{R} , strictly convex, differentiable and strongly coercive. It then follows from [28, Corollary X.4.1.4] that f^* enjoys the same properties. In turn, using the fact that both f and f^* are even, we have g^+ is strongly coercive, and strict convexity of f (resp. f^*) is equivalent to that of g (resp. g^+). Altogether, this shows the first claim. We now prove that g vanishes only at 0 (and similarly for g^+). As g is non-decreasing and strictly convex, we have, for any $\rho \in]0, 1[$ and a, b in \mathbb{R}_+ such that $a < b$,

$$g(a) \leq g(\rho a + (1 - \rho)b) < \rho g(a) + (1 - \rho)g(b) \leq \rho g(b) + (1 - \rho)g(b) = g(b).$$

□

Support function The *support function* of $C \subset \mathbb{R}^p$ is

$$\sigma_C(\omega) = \sup_{\theta \in C} \langle \omega, \theta \rangle.$$

We recall the following properties whose proofs can be found in e.g. [44, 28].

Lemma A.2. *Let C be a non-empty set.*

- (i) σ_C is proper lsc and sublinear.
- (ii) σ_C is finite-valued if and only if C is bounded.
- (iii) If $0 \in C$, then σ_C is non-negative.
- (iv) If C is convex and compact with $0 \in \text{int}(C)$, then σ_C is finite-valued and coercive.

Gauges and polars

Definition A.1 (Polar set). *Let \mathcal{C} be a nonempty convex set. The set \mathcal{C}° given by*

$$\mathcal{C}^\circ = \{\boldsymbol{\eta} \in \mathbb{R}^p : \langle \boldsymbol{\eta}, \boldsymbol{\theta} \rangle \leq 1 \quad \forall \boldsymbol{\theta} \in \mathcal{C}\}$$

is called the polar of \mathcal{C} .

The set \mathcal{C}° is closed convex and contains the origin. When \mathcal{C} is also closed and contains the origin, then it coincides with its bipolar, i.e. $\mathcal{C}^{\circ\circ} = \mathcal{C}$.

Let $\mathcal{C} \subseteq \mathbb{R}^p$ be a non-empty closed convex set containing the origin. The *gauge* of \mathcal{C} is the function $\gamma_{\mathcal{C}}$ defined on \mathbb{R}^p by

$$\gamma_{\mathcal{C}}(\boldsymbol{\theta}) = \inf \{\lambda > 0 : \boldsymbol{\theta} \in \lambda\mathcal{C}\}.$$

As usual, $\gamma_{\mathcal{C}}(\boldsymbol{\theta}) = +\infty$ if the infimum is not attained.

Lemma A.3 hereafter recaps the main properties of a gauge that we need. In particular, (ii) is a fundamental result of convex analysis that states that there is a one-to-one correspondence between gauge functions and closed convex sets containing the origin. This allows to identify sets from their gauges, and vice versa.

Lemma A.3.

(i) $\gamma_{\mathcal{C}}$ is a non-negative, lsc and sublinear function.

(ii) \mathcal{C} is the unique closed convex set containing the origin such that

$$\mathcal{C} = \{\boldsymbol{\theta} \in \mathbb{R}^p : \gamma_{\mathcal{C}}(\boldsymbol{\theta}) \leq 1\}.$$

(iii) $\gamma_{\mathcal{C}}$ is finite-valued if, and only if, $0 \in \text{int}(\mathcal{C})$, in which case $\gamma_{\mathcal{C}}$ is Lipschitz continuous.

(iv) $\gamma_{\mathcal{C}}$ is finite-valued and coercive if, and only if, \mathcal{C} is compact and $0 \in \text{int}(\mathcal{C})$.

See [57] for the proof.

Observe that thanks to sublinearity, local Lipschitz continuity valid for any finite-valued convex function is strengthened to global Lipschitz continuity. Moreover, $\gamma_{\mathcal{C}}$ is a norm, having \mathcal{C} as its unit ball, if and only if \mathcal{C} is bounded with nonempty interior and symmetric.

We now define the polar gauge.

Definition A.2 (Polar Gauge). *The polar of a gauge $\gamma_{\mathcal{C}}$ is the function $\gamma_{\mathcal{C}}^\circ$ defined by*

$$\gamma_{\mathcal{C}}^\circ(\boldsymbol{\omega}) = \inf \{\mu \geq 0 : \langle \boldsymbol{\theta}, \boldsymbol{\omega} \rangle \leq \mu \gamma_{\mathcal{C}}(\boldsymbol{\theta}), \forall \boldsymbol{\theta}\}.$$

An immediate consequence is that gauges polar to each other have the property

$$\langle \boldsymbol{\theta}, \boldsymbol{u} \rangle \leq \gamma_{\mathcal{C}}(\boldsymbol{\theta}) \gamma_{\mathcal{C}}^\circ(\boldsymbol{u}) \quad \forall (\boldsymbol{\theta}, \boldsymbol{u}) \in \text{dom}(\gamma_{\mathcal{C}}) \times \text{dom}(\gamma_{\mathcal{C}}^\circ), \quad (\text{A.1})$$

just as dual norms satisfy a duality inequality. In fact, polar pairs of gauges correspond to the best inequalities of this type.

Lemma A.4. *Let $\mathcal{C} \subseteq \mathbb{R}^p$ be a closed convex set containing the origin. Then,*

(ii) $\gamma_{\mathcal{C}}^\circ$ is a gauge function and $\gamma_{\mathcal{C}}^{\circ\circ} = \gamma_{\mathcal{C}}$.

(iii) $\gamma_{\mathcal{C}}^{\circ} = \gamma_{\mathcal{C}^{\circ}}$, or equivalently

$$\mathcal{C}^{\circ} = \{\boldsymbol{\theta} \in \mathbb{R}^p : \gamma_{\mathcal{C}}^{\circ}(\boldsymbol{\theta}) \leq 1\}.$$

(iv) The gauge of \mathcal{C} and the support function of \mathcal{C} are mutually polar, i.e.

$$\gamma_{\mathcal{C}} = \sigma_{\mathcal{C}^{\circ}} \quad \text{and} \quad \gamma_{\mathcal{C}^{\circ}} = \sigma_{\mathcal{C}}.$$

See [44, 28, 57] for the proof.

B Expectation of the inner product

We start with some definitions and notations that will be used in the proof. For a non-empty closed convex set $\mathcal{C} \in \mathbb{R}^p$, we denote $(\mathcal{C})^0$ its minimal selection, i.e. the element of minimal norm in \mathcal{C} . This element is of course unique. For a proper lsc and convex function f and $\gamma > 0$, its Moreau envelope (or Moreau-Yosida regularization) is defined by

$$\gamma f(\boldsymbol{\theta}) \stackrel{\text{def}}{=} \min_{\bar{\boldsymbol{\theta}} \in \mathbb{R}^p} \frac{1}{2\gamma} \|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}\|_2^2 + f(\bar{\boldsymbol{\theta}}).$$

The Moreau envelope enjoys several important properties that we collect in the following lemma.

Lemma B.1. *Let f be a finite-valued and convex function. Then*

- (i) $(\gamma f(\boldsymbol{\theta}))_{\gamma > 0}$ is a decreasing net, and $\forall \boldsymbol{\theta} \in \mathbb{R}^p$, $\gamma f(\boldsymbol{\theta}) \nearrow f(\boldsymbol{\theta})$ as $\gamma \searrow 0$.
- (ii) $\gamma f \in C^1(\mathbb{R}^p)$ with γ^{-1} -Lipschitz continuous gradient.
- (iii) $\forall \boldsymbol{\theta} \in \mathbb{R}^p$, $\nabla \gamma f(\boldsymbol{\theta}) \rightarrow (\partial f(\boldsymbol{\theta}))^0$ and $\|\nabla \gamma f(\boldsymbol{\theta})\|_2 \nearrow \|(\partial f(\boldsymbol{\theta}))^0\|_2$ as $\gamma \searrow 0$.

Proof. (i) [3, Proposition 12.32]. (ii) [3, Proposition 12.29]. (iii) By assumption, f is subdifferentiable everywhere and its subdifferential is a maximal monotone operator with domain \mathbb{R}^p , and the result follows from [3, Corollary 23.46(i)]. \square

We are now equipped to prove the following important result⁸.

Proposition B.1. *Let the density μ_n in (1.2), where*

- (a) F satisfies Assumptions (H.1)-(H.2);
- (b) J is a finite-valued lower-bounded convex function, and $\exists R > 0$ and $\rho \geq 0$, such that $\forall \boldsymbol{\theta} \in \mathbb{R}^p$, $\|(\partial J(\boldsymbol{\theta}))^0\|_2 \leq R \|\boldsymbol{\theta}\|_2^{\rho}$;
- (c) and V_n is coercive.

Then, $\forall \bar{\boldsymbol{\theta}} \in \mathbb{R}^p$,

$$\mathbb{E}_{\mu_n} \left[\langle (\partial V_n(\boldsymbol{\theta}))^0, \bar{\boldsymbol{\theta}} - \boldsymbol{\theta} \rangle \right] = -p\beta.$$

⁸The result will be proved using Moreau-Yosida regularization. Yet another alternative proof could be based on mollifiers for approximating subdifferentials.

This result covers of course the situation where J fulfills **(H.3)**. In this case, since $\partial J(\boldsymbol{\theta}) \subset \mathcal{C}^\circ$ by Theorem 2.1(i), we have $\rho = 0$ and $R = \text{diam}(\mathcal{C}^\circ)$, the diameter of the convex compact set \mathcal{C}° containing the origin. It can be shown that, when $F(\cdot, \mathbf{y})$ is strongly coercive, the coercivity assumption (c) can be equivalently stated as $J_\infty(\boldsymbol{\theta}) > 0, \forall \boldsymbol{\theta} \in \ker(\mathbf{X}) \setminus \{0\}$, where J_∞ is the recession/asymptotic function of J ; see e.g. [45].

Proof. Let $V_n^\gamma(\boldsymbol{\theta}) \stackrel{\text{def}}{=} \frac{1}{n}F(\mathbf{X}\boldsymbol{\theta}, \mathbf{y}) + \lambda_n^\gamma J(\boldsymbol{\theta})$ and define $\mu_n^\gamma(\boldsymbol{\theta}) \stackrel{\text{def}}{=} \exp(-V_n^\gamma(\boldsymbol{\theta})/\beta)/Z$, where $0 < Z < +\infty$ is the normalizing constant of the density μ_n . Assumption **(H.1)** and Lemma B.1(ii)-(iii) tell us that $V_n^\gamma \in C^1(\mathbb{R}^p)$ and $\nabla V_n^\gamma(\boldsymbol{\theta}) \rightarrow (\partial V_n(\boldsymbol{\theta}))^0$ as $\gamma \rightarrow 0$. Thus

$$\mathbb{E}_{\mu_n} \left[\langle (\partial V_n(\boldsymbol{\theta}))^0, \bar{\boldsymbol{\theta}} - \boldsymbol{\theta} \rangle \right] = \int_{\mathbb{R}^p} \lim_{\gamma \rightarrow 0} \langle \mu_n^\gamma(\boldsymbol{\theta}) \nabla V_n^\gamma(\boldsymbol{\theta}), \bar{\boldsymbol{\theta}} - \boldsymbol{\theta} \rangle d\boldsymbol{\theta}.$$

We now check that $\langle \mu_n^\gamma(\boldsymbol{\theta}) \nabla V_n^\gamma(\boldsymbol{\theta}), \bar{\boldsymbol{\theta}} - \boldsymbol{\theta} \rangle$ is dominated by an integrable function. From the definition of the Moreau envelope, we have

$$V_n^\gamma(\boldsymbol{\theta}) = \min_{\bar{\boldsymbol{\theta}} \in \mathbb{R}^p} \frac{1}{n}F(\mathbf{X}\bar{\boldsymbol{\theta}}, \mathbf{y}) + \lambda_n (J(\boldsymbol{\theta} - \bar{\boldsymbol{\theta}}) + \frac{1}{2\gamma} \|\bar{\boldsymbol{\theta}}\|_2^2).$$

From coercivity of V_n , the objective in the min is also coercive in $(\boldsymbol{\theta}, \bar{\boldsymbol{\theta}})$ by [45, Exercise 3.29(b)]. It then follows from [45, Theorem 3.31] that V_n^γ is also coercive. In turn, [45, Theorem 11.8(c) and 3.26(a)] allow to assert that for some $a \in]0, +\infty[$, $\exists b \in]-\infty, +\infty[$ such that for all $\gamma > 0$ and $\boldsymbol{\theta} \in \mathbb{R}^p$

$$\mu_n^\gamma(\boldsymbol{\theta}) \leq \exp(-a \|\boldsymbol{\theta}\|_2 - b)/Z. \quad (\text{B.1})$$

Lemma B.1-(iii) and assumption (b) on J entail that for any $\boldsymbol{\theta} \in \mathbb{R}^p$,

$$\|\nabla^\gamma J(\boldsymbol{\theta})\|_2 \leq \|(\partial J(\boldsymbol{\theta}))^0\|_2 \leq R \|\boldsymbol{\theta}\|_2^\rho.$$

Altogether, we have

$$\begin{aligned} |\langle \mu_n^\gamma(\boldsymbol{\theta}) \nabla V_n^\gamma(\boldsymbol{\theta}), \bar{\boldsymbol{\theta}} - \boldsymbol{\theta} \rangle| &\leq \mu_n^\gamma(\boldsymbol{\theta}) \left(|\langle \mathbf{X}^\top \frac{1}{n} \nabla F(\mathbf{X}\boldsymbol{\theta}, \mathbf{y}), \bar{\boldsymbol{\theta}} - \boldsymbol{\theta} \rangle| + \lambda_n \|\nabla^\gamma J(\boldsymbol{\theta})\|_2 \|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}\|_2 \right) \\ &\leq CZ^{-1} \exp(-F(\mathbf{X}\boldsymbol{\theta}, \mathbf{y})/(n\beta)) |\langle \frac{1}{n} \nabla F(\mathbf{X}\boldsymbol{\theta}, \mathbf{y}), \mathbf{X}(\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}) \rangle| \\ &\quad + (Z \exp b)^{-1} \lambda_n R \exp(-a \|\boldsymbol{\theta}\|_2) \|\boldsymbol{\theta}\|_2^\rho \|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}\|_2, \end{aligned}$$

where the constant $C > 0$ reflects the lower-boundedness of J . It is easy to see that the function in this upper-bound is integrable, where we also use **(H.2)**. Hence, we can apply the dominated convergence theorem to get

$$\mathbb{E}_{\mu_n} \left[\langle (\partial V_n(\boldsymbol{\theta}))^0, \bar{\boldsymbol{\theta}} - \boldsymbol{\theta} \rangle \right] = \lim_{\gamma \rightarrow 0} \int_{\mathbb{R}^p} \langle \mu_n^\gamma(\boldsymbol{\theta}) \nabla V_n^\gamma(\boldsymbol{\theta}), \bar{\boldsymbol{\theta}} - \boldsymbol{\theta} \rangle d\boldsymbol{\theta}.$$

Now, by simple differential calculus (chain and product rules), we have

$$\begin{aligned} \langle \mu_n^\gamma(\boldsymbol{\theta}) \nabla V_n^\gamma(\boldsymbol{\theta}), \bar{\boldsymbol{\theta}} - \boldsymbol{\theta} \rangle &= -\beta \langle \nabla \mu_n^\gamma(\boldsymbol{\theta}), \bar{\boldsymbol{\theta}} - \boldsymbol{\theta} \rangle \\ &= -\beta \sum_{i=1}^p \frac{\partial}{\partial \theta_i} (\mu_n^\gamma(\boldsymbol{\theta}) (\bar{\theta}_i - \theta_i)) - p\beta \mu_n^\gamma(\boldsymbol{\theta}). \end{aligned}$$

Integrating the first term, we get by Fubini theorem and the Newton-Leibniz formula

$$\begin{aligned} \int_{\mathbb{R}^{p-1}} \left(\int_{\mathbb{R}} \frac{\partial}{\partial \boldsymbol{\theta}_i} (\mu_n^\gamma(\boldsymbol{\theta})(\bar{\boldsymbol{\theta}}_i - \boldsymbol{\theta}_i)) d\boldsymbol{\theta}_i \right) d\boldsymbol{\theta}_1 \cdots d\boldsymbol{\theta}_{i-1} d\boldsymbol{\theta}_{i+1} \cdots d\boldsymbol{\theta}_p \\ = \int_{\mathbb{R}^{p-1}} [\mu_n^\gamma(\boldsymbol{\theta})(\bar{\boldsymbol{\theta}}_i - \boldsymbol{\theta}_i)]_{\mathbb{R}} d\boldsymbol{\theta}_1 \cdots d\boldsymbol{\theta}_{i-1} d\boldsymbol{\theta}_{i+1} \cdots d\boldsymbol{\theta}_p = 0, \end{aligned}$$

where we used coercivity of V_n^γ (see (B.1)) to conclude that $\lim_{|\boldsymbol{\theta}_i| \rightarrow +\infty} \mu_n^\gamma(\boldsymbol{\theta})(\bar{\boldsymbol{\theta}}_i - \boldsymbol{\theta}_i) = 0$. For the second term, we have from Lemma B.1(i) that $\mu_n^\gamma \rightarrow \mu_n$ as $\gamma \rightarrow 0$. Thus, arguing again as in (B.1), we can apply the dominated convergence theorem to conclude that

$$\lim_{\gamma \rightarrow 0} \int_{\mathbb{R}^p} \mu_n^\gamma(\boldsymbol{\theta}) d\boldsymbol{\theta} = \int_{\mathbb{R}^p} \mu_n(\boldsymbol{\theta}) d\boldsymbol{\theta} = 1.$$

This concludes the proof. □

Acknowledgement. This work was supported by Conseil Régional de Basse-Normandie and partly by Institut Universitaire de France.

References

- [1] F. Bach. Consistency of the group lasso and multiple kernel learning. *Journal of Machine Learning Research*, 9:1179–1225, 2008.
- [2] S. Bakin. Adaptive regression and model selection in data mining problems, 1999. Thesis (Ph.D.)–Australian National University, 1999.
- [3] H. H. Bauschke and P. L. Combettes. *Convex analysis and monotone operator theory in Hilbert spaces*. Springer, 2011.
- [4] P. Bellec. Concentration of quadratic forms under a bernstein moment assumption. Technical report, Ecole Polytechnique, 2014.
- [5] P. J. Bickel, Y. Ritov, and A. Tsybakov. Simultaneous analysis of lasso and Dantzig selector. *Annals of Statistics*, 37(4):1705–1732, 2009.
- [6] M. Bogdan, E. van den Berg, C. Sabatti, W. Su, and E. J. Candès. Slope – adaptive variable selection via convex optimization. *Annals of Applied Statistics*, 9(3):1103–1140, 2014.
- [7] P. Bühlmann and S. van de Geer. *Statistics for High-Dimensional Data: Methods, Theory and Applications*. Springer Series in Statistics. Springer-Verlag Berlin Heidelberg, 2011.
- [8] E. Candès and Y. Plan. Near-ideal model selection by ℓ_1 minimization. *Annals of Statistics*, 37(5A):2145–2177, 2009.
- [9] E. J. Candès, X. Li, Y. Ma, and J. Wright. Robust principal component analysis? *Journal of the ACM*, 58(3):11:1–11:37, June 2011.
- [10] E. J. Candès and B. Recht. Exact matrix completion via convex optimization. *Foundations of Computational mathematics*, 9(6):717–772, 2009.

- [11] E. J. Candès, T. Strohmer, and V. Voroninski. Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming. *Communications on Pure and Applied Mathematics*, 66(8):1241–1274, 2013.
- [12] V. Chandrasekaran, B. Recht, P. A. Parrilo, and A. Willsky. The convex geometry of linear inverse problems. *Foundations of Computational Mathematics*, 12(6):805–849, 2012.
- [13] S. Chen, D. Donoho, and M. Saunders. Atomic decomposition by basis pursuit. *SIAM journal on scientific computing*, 20(1):33–61, 1999.
- [14] X. Chen, Q. Lin, S. Kim, J. G. Carbonell, and E. P. Xing. An efficient proximal-gradient method for general structured sparse learning. *Preprint arXiv:1005.4717*, 2010.
- [15] A. Dalalyan and A. Tsybakov. Pac-bayesian bounds for the expected error of aggregation by exponential weights. Technical report, Université Paris 6, CREST and CERTIS, Ecole des Ponts ParisTech, 2009. personal communication.
- [16] A. Dalalyan and A. B. Tsybakov. Aggregation by exponential weighting, sharp pac-bayesian bounds and sparsity. *Mach. Learn.*, 72(1-2):39–61, Aug. 2008.
- [17] A. S. Dalalyan, E. Grappin, and Q. Paris. On the Exponentially Weighted Aggregate with the Laplace Prior. Technical report, arXiv:1611.08483, Nov. 2016.
- [18] A. S. Dalalyan, M. Hebiri, and J. Lederer. On the prediction performance of the lasso. *Bernoulli*, 23(1):552–581, 02 2017.
- [19] A. S. Dalalyan and A. B. Tsybakov. Aggregation by exponential weighting and sharp oracle inequalities. In *Proceedings of the 20th Annual Conference on Learning Theory, COLT’07*, pages 97–111, Berlin, Heidelberg, 2007. Springer-Verlag.
- [20] A. S. Dalalyan and A. B. Tsybakov. Sparse regression learning by aggregation and langevin monte-carlo. *J. Comput. Syst. Sci.*, 78(5):1423–1443, Sept. 2012.
- [21] A. Daniilidis, D. Drusvyatskiy, and A. S. Lewis. Orthogonal invariance and identifiability. Technical report, arXiv 1304.1198, 2013.
- [22] D. Donoho. For most large underdetermined systems of linear equations the minimal ℓ^1 -norm solution is also the sparsest solution. *Communications on Pure and Applied Mathematics*, 59(6):797–829, 2006.
- [23] A. Durmus, E. Moulines, and M. Pereyra. Sampling from convex non continuously differentiable functions, when Moreau meets Langevin. Preprint hal-01267115, Feb. 2016.
- [24] T. Duy Luu, J. M. Fadili, and C. Chesneau. PAC-Bayesian risk bounds for group-analysis sparse regression by exponential weighting. Technical report, hal-01367742, Sept. 2016.
- [25] M. J. Fadili, G. Peyré, S. Vaïter, C. Deledalle, and J. Salmon. Stable recovery with analysis decomposable priors. In *Proc. Sampta’13*, pages 113–116, 2013.
- [26] M. Fazel, H. Hindi, and S. P. Boyd. A rank minimization heuristic with application to minimum order system approximation. In *Proceedings of the American Control Conference*, volume 6, pages 4734–4739. IEEE, 2001.

- [27] B. Guedj and P. Alquier. Pac-bayesian estimation and prediction in sparse additive models. *Electron. J. Statist.*, 7:264–291, 2013.
- [28] J.-B. Hiriart-Urruty and C. Lemaréchal. *Convex Analysis And Minimization Algorithms*, volume I and II. Springer, 2001.
- [29] L. Jacob, G. Obozinski, and J.-P. Vert. Group lasso with overlap and graph lasso. In A. P. Danyluk, L. Bottou, and M. L. Littman, editors, *ICML'09*, volume 382, page 55, 2009.
- [30] H. Jégou, T. Furon, and J.-J. Fuchs. Anti-sparse coding for approximate nearest neighbor search. In *IEEE ICASSP*, pages 2029–2032, 2012.
- [31] V. Koltchinskii. Oracle inequalities in empirical risk minimization and sparse recovery problems. In *Lectures from the 38th Probability Summer School held in Saint-Flour*, volume 2033 of *Lecture Notes in Mathematics*. Springer, 2008.
- [32] V. Koltchinskii, K. Lounici, and A. B. Tsybakov. Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion. *Ann. Statist.*, 39(5):2302–2329, 10 2011.
- [33] G. Lecué. Simultaneous adaptation to the margin and to complexity in classification. *Ann. Statist.*, 35(4):1698–1721, 08 2007.
- [34] M. Ledoux. *The concentration of measure phenomenon*. Mathematical surveys and monographs. American Mathematical Society, Providence (R.I.), 2001. L'ISSN figurant sur le substitut de la page de titre 0076-5376 correspond à la revue *Mathematical surveys*. Le titre a changé en 1981 en *Mathematical surveys and monographs* et porte le numéro 0885-4653.
- [35] K. Lounici, M. Pontil, S. van de Geer, and A. B. Tsybakov. Oracle inequalities and optimal inference under group sparsity. *Ann. Statist.*, 39(4):2164–2204, 08 2011.
- [36] Y. Lyubarskii and R. Vershynin. Uncertainty principles and vector quantization. *IEEE Transactions on Information Theory*, 56(7):3491–3501, 2010.
- [37] S. Negahban, P. Ravikumar, M. J. Wainwright, and B. Yu. A unified framework for high-dimensional analysis of M-estimators with decomposable regularizers. *Statistical Science*, 27(4):538–557, December 2012.
- [38] A. Nemirovski. *Topics in non-parametric statistics*, 2000.
- [39] M. Osborne, B. Presnell, and B. Turlach. A new approach to variable selection in least squares problems. *IMA journal of numerical analysis*, 20(3):389–403, 2000.
- [40] G. Peyré, J. Fadili, and C. Chesneau. Adaptive Structured Block Sparsity Via Dyadic Partitioning. In *EUSIPCO*, Barcelona, Spain, Aug. 2011.
- [41] B. Recht, M. Fazel, and P. A. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM review*, 52(3):471–501, 2010.
- [42] P. Rigollet and A. Tsybakov. Linear and convex aggregation of density estimators. *Mathematical Methods of Statistics*, 16(3):260–280, 2007.

- [43] P. Rigollet and A. B. Tsybakov. Sparse estimation by exponential weighting. *Statist. Sci.*, 27(4):558–575, 11 2012.
- [44] R. Rockafellar. *Convex analysis*, volume 28. Princeton University Press, 1996.
- [45] R. T. Rockafellar and R. Wets. *Variational analysis*, volume 317. Springer Verlag, 1998.
- [46] M. Rudelson and R. Vershynin. On sparse reconstruction from Fourier and Gaussian measurements. *Communications on Pure and Applied Mathematics*, 61(8):1025–1045, 2008.
- [47] L. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(1-4):259–268, 1992.
- [48] C. Studer, W. Yin, and R. G. Baraniuk. Signal representations with minimum ℓ_∞ -norm. In *50th Annual Allerton Conference on Communication, Control, and Computing*, 2012.
- [49] W. Su and E. J. Candès. Slope is adaptive to unknown sparsity and asymptotically minimax. *Annals of Statistics*, 44(3):1038–1068, 2015.
- [50] T. Sun and C.-H. Zhang. Scaled sparse linear regression. *Biometrika*, 99(4):879, 2012.
- [51] M. Talagrand. *The generic chaining. Upper and lower bounds of stochastic processes*. Springer-Verlag, Berlin, 2005.
- [52] R. Tibshirani. Regression shrinkage and selection via the Lasso. *Journal of the Royal Statistical Society. Series B. Methodological*, 58(1):267–288, 1996.
- [53] R. Tibshirani, M. Saunders, S. Rosset, J. Zhu, and K. Knight. Sparsity and smoothness via the fused Lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(1):91–108, 2005.
- [54] J. Tropp. Convex recovery of a structured signal from independent random linear measurements. In *Sampling Theory, a Renaissance*. Birkhäuser, 2014.
- [55] J. A. Tropp. An introduction to matrix concentration inequalities. *Foundations and Trends in Machine Learning*, 8(1-2):1–230, 2015.
- [56] S. Vaiter, C. Deledalle, M. J. Fadili, G. Peyré, and C. Dossal. The degrees of freedom of partly smooth regularizers. *Annals of the Institute of Statistical Mathematics*, 2016. in press.
- [57] S. Vaiter, M. Golbabaee, M. J. Fadili, and G. Peyré. Model selection with low complexity priors. *Information and Inference: A Journal of the IMA (IMAIAI)*, 2015.
- [58] S. Vaiter, G. Peyré, and M. J. Fadili. Low complexity regularization of linear inverse problems. In G. Pfander, editor, *Sampling Theory, a Renaissance, Applied and Numerical Harmonic Analysis (ANHA)*. Birkhäuser/Springer, 2015.
- [59] S. Vaiter, G. Peyré, and M. J. Fadili. Model consistency of partly smooth regularizers. Technical report, Preprint Hal-00987293, 2015. submitted.
- [60] S. van de Geer. High-dimensional generalized linear models and the Lasso. *Annals of Statistics*, 36:614–645, 2008.

- [61] S. van de Geer. Weakly decomposable regularization penalties and structured sparsity. *Scandinavian Journal of Statistics*, 41(1):72–86, 2014.
- [62] S. van de Geer and P. Bühlmann. On the conditions used to prove oracle results for the lasso. *Electron. J. Statist.*, 3:1360–1392, 2009.
- [63] S. van de Geer and J. Lederer. The bernstein–orlicz norm and deviation inequalities. *Probab. Theory Relat. Fields*, 157:225–250, 2013.
- [64] R. Vershynin. Estimation in high dimensions: A geometric perspective. In *Sampling Theory, a Renaissance*. Birkhäuser, 2014.
- [65] F. Wei and J. Huang. Consistent group selection in high-dimensional linear regression. *Bernoulli*, 16(4):1369–1384, 2010.
- [66] Y. Yang. Aggregating regression procedures to improve performance. *Bernoulli*, 10(1):25–47, 2004.
- [67] M. Yuan and Y. Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1):49–67, 2006.