
Combinatorial Penalties: Which structures are preserved by convex relaxations?

Marwa El Halabi
LIONS, EPFL

Francis Bach
INRIA - ENS - PSL Research University

Volkan Cevher
LIONS, EPFL

Abstract

We consider the *homogeneous* and the *non-homogeneous* convex relaxations for combinatorial penalty functions defined on support sets. Our study identifies key differences in the tightness of the resulting relaxations through the notion of the *lower combinatorial envelope* of a set-function along with new necessary conditions for support identification. We then propose a general adaptive estimator for convex monotone regularizers, and derive new sufficient conditions for support recovery in the asymptotic setting.

1 Introduction

Over the last years, *sparsity* has been a key model in machine learning, signal processing, and statistics. While sparsity modelling is powerful, *structured sparsity* models further exploit domain knowledge by characterizing the interdependency between the non-zero coefficients of an unknown parameter vector w . For example, in certain applications domain knowledge may dictate that we should favor non-zero patterns corresponding to: unions of groups [32] in cancer prognosis from gene expression data; or complements of union of groups [21] in neuroimaging and background subtraction, or rooted connected trees [22, 37] in natural image processing. Incorporating such key prior information beyond just sparsity leads to significant improvements in estimation performance, noise robustness, interpretability and sample complexity [4].

Structured sparsity models are naturally encoded by combinatorial functions. However, direct combinatorial treatments often lead to intractable learning problems. Hence, we often use either non-convex greedy methods or continuous convex relaxations, where the combinatorial penalty is replaced by a tractable convex surrogate; *cf.*, [4, 20, 2].

In this paper, we adopt the convex approach because it benefits from a mature set of efficient numerical algorithms as well as strong analysis tools that rely on convex geometry in order to establish statistical efficiency. Convex formulations are also robust to model mis-specifications. Moreover, there is a rich set of convex penalties with structure-inducing properties already studied in the literature [36, 21, 23, 22, 37, 32]. For an overview, we refer the reader to [2] and references therein.

For choosing a convex relaxation, a systematic approach, already adopted in [1, 6, 31, 10], considers the *tightest* convex relaxation of combinatorial penalties expressing the desired structure. For instance, [1] shows that computing the tightest convex relaxation over the unit ℓ_∞ -ball is tractable for the ensemble of *monotone submodular functions*. Similarly, the authors in [10] demonstrates the tractability of such relaxation for combinatorial penalties that can be described via *totally unimodular* constraints.

A different principled approach in convex relaxations is proposed by [31], where the authors considered the tightest *homogeneous* convex relaxation of general set functions regularized by an ℓ_p -norm. The authors show, for instance, the resulting norm takes the form of a generalized latent group Lasso norm [32]. The homogeneity imposed in [31] naturally ensures the invariance of the regularizer to rescaling of the data. However, such requirement may cost a loss of structure as was observed in an example in [10]. This observation begs the question:

When do the *homogeneous* and *non-homogeneous* convex relaxations differ and which structures can be encoded by each?

In order to answer this question, we rigorously identify which combinatorial structures are preserved by the non-homogeneous relaxation in a manner similar to [31] for the homogeneous one. We further study the statistical properties of both relaxations. In particular, we consider the problem of support recovery in the context of regularized learning problems by these relaxed convex penalties, which was only investigated so far in special cases, e.g., for norms associated with submodular functions [1], or for the latent group Lasso norm [32].

To this end, this paper makes the following contributions:

- We derive formulations of the non-homogeneous tightest convex relaxation of general ℓ_p -regularized combinatorial penalties (Section 2.1). We show that any *monotone* set function is preserved by such relaxation, while the homogeneous relaxation only preserves a smaller subset of set-functions (Section 2.2).
- We identify necessary conditions for support recovery in learning problems regularized by general convex penalties (Section 3.1).
- We propose an adaptive weight estimation scheme and provide sufficient conditions for support recovery under the asymptotic regime (Section 3.2). This scheme does not require any irrepresentability condition and is applicable to general monotone convex regularizers.
- We identify sufficient conditions with respect to combinatorial penalties which ensure that the sufficient support recovery conditions hold with respect to the associated convex relaxations (Section 4).
- We illustrate numerically the effect on support recovery of the choice of the relaxation as well as the adaptive weights scheme (Section 5).

In the sequel, we defer all proofs to the Appendix.

Notation. Let $V = \{1, \dots, d\}$ be the ground set and $2^V = \{A | A \subseteq V\}$ be its power-set. Given $w \in \mathbb{R}^d$ and a set $J \subseteq V$, w_J denotes the vector in \mathbb{R}^d s.t., $[w_J]_i = w_i, i \in J$ and $[w_J]_i = 0, i \notin J$. Q_{JJ} is defined similarly for a matrix $Q \in \mathbb{R}^{d \times d}$. Accordingly, we let $\mathbb{1}_J$ be the indicator vector of the set J . We drop the subscript for $J = V$, so that $\mathbb{1}_V = \mathbb{1}$ denotes the vector of all ones. The notation J^c denotes the set complement of J with respect to V .

The operations $|w|, w \geq w'$ and $w \circ v$ are applied element-wise. For $p > 0$, the ℓ_p - (quasi) norm is given by $\|w\|_p = (\sum_{i=1}^d |w_i|^p)^{1/p}$, and $\|w\|_\infty = \max_i |w_i|$. For $p \in [1, \infty]$, we define the conjugate $q \in [1, \infty]$ via $\frac{1}{p} + \frac{1}{q} = 1$.

We call the set of non-zero elements of a vector w the support, denoted by $\text{supp}(w) = \{i : w_i \neq 0\}$. We use the notation from submodular analysis, $w(A) = \sum_{i \in A} w_i$. We write $\overline{\mathbb{R}}_+$ for $\mathbb{R}_+ \cup \{+\infty\}$. For a function $f : \mathbb{R}^d \rightarrow \overline{\mathbb{R}} = \mathbb{R} \cup \{+\infty\}$, we will denote by f^* its Fenchel-Legendre conjugate. We will denote by $\iota_S(w)$ the indicator function of the set S , taking value 0 on the set S and $+\infty$ outside it.

2 Combinatorial penalties and convex relaxations

We consider positive-valued set functions of the form $F : 2^V \rightarrow \overline{\mathbb{R}}_+$ such that $F(\emptyset) = 0, F(A) > 0, \forall A \subseteq V$ to encode structured sparsity models. For generality, we

do not assume a priori that F is *monotone* (i.e., $F(A) \leq F(B), \forall A \subseteq B$). However, as we argue in the sequel, convex relaxations of non-monotone set functions is hopeless.

The domain of F is defined as $\mathcal{D} := \{A : F(A) < +\infty\}$. We assume that it covers V , i.e., $\cup_{A \in \mathcal{D}} A = V$, which is equivalent to assuming that F is finite at singletons if F is monotone. A finite-valued set function F is *submodular* if and only if for all $A \subseteq B$ and $i \in B^c, F(B \cup \{i\}) - F(B) \leq F(A \cup \{i\}) - F(A)$ (see, e.g., [15, 2]). Unless explicitly stated, we do not assume that F is submodular.

We consider the same model in [31], parametrized by $w \in \mathbb{R}^d$, with general ℓ_p -regularized combinatorial penalties:

$$F_p(w) = \frac{1}{q} F(\text{supp}(w)) + \frac{1}{p} \|w\|_p^p$$

for $p \geq 1$, where the set function F controls the structure of the model in terms of allowed/favored non-zero patterns and the ℓ_p -norm serves to control the magnitude of the coefficients. Allowing F to take infinite values let us enforce hard constraints. For $p = \infty, F_p$ reduces to $F_\infty(w) = F(\text{supp}(w)) + \iota_{\|w\|_\infty \leq 1}(w)$. Considering the case $p \neq \infty$ is appealing to avoid the clustering artifacts of the values of the learned vector induced by the ℓ_∞ -norm.

Since such combinatorial regularizers lead to computationally intractable problems, we seek convex surrogate penalties that capture the encoded structure as much as possible. A natural candidate for a convex surrogate of F_p is then its *convex envelope* (largest convex lower bound), given by the biconjugate (the Fenchel conjugate of the Fenchel conjugate) F_p^{**} . Two general approaches are proposed in the literature to achieve just this; one requires the surrogate to also be positively homogeneous [31] and thus considers the convex envelope of the positively homogeneous envelope of F_p , given by $F(\text{supp}(w))^{1/q} \|w\|_p$, which we denote by Ω_p , the other computes instead the convex envelope of F_p directly [10, 1], which we denote by Θ_p . Note that from the definition of convex envelope, it holds that $\Theta_p \geq \Omega_p$.

2.1 Homogeneous and non-homogeneous convex envelopes

In [31], the homogeneous convex envelope Ω_p of F_p was shown to correspond to the *latent group Lasso norm* [32] with groups set to all elements of the power set 2^V . We recall this form of Ω_∞ in Lemma 1 as well as a variational form of Ω_p which highlights the relation between the two. Other variational forms can be found in the Appendix.

Lemma 1 ([31]). *The homogeneous convex envelope Ω_p of F_p is given by*

$$\Omega_p(w) = \inf_{\eta \in \mathbb{R}_+^d} \frac{1}{p} \sum_{j=1}^d \frac{|w_j|^p}{\eta_j^{p-1}} + \frac{1}{q} \Omega_\infty(\eta), \quad (1)$$

$$\Omega_\infty(w) = \min_{\alpha \geq 0} \left\{ \sum_{S \subseteq V} \alpha_S F(S) : \sum_{S \subseteq V} \alpha_S \mathbb{1}_S \geq |w| \right\}. \quad (2)$$

The non-homogeneous convex envelope Θ_p of F_p is only considered thus far in the case where $p = \infty$. [10] shows that $\Theta_\infty(w) = \inf_{\eta \in [0,1]^d} \{f(\eta) : \eta \geq |w|\}$ where f is any proper ($\text{dom}(f) \neq \emptyset$), lower semi-continuous (l.s.c.) convex extension of F , i.e., $f(\mathbf{1}_A) = F(A), \forall A \subseteq V$ (cf., Lemma 1 in [10]). A natural choice for f is the convex closure of F , which corresponds to the tightest convex extension of F on $[0, 1]^d$ (cf., Appendix for a more rigorous treatment).

Lemma 2 below presents this choice, deriving a new form of Θ_∞ that parallels (2). We also derive the non-homogeneous convex envelope Θ_p of F_p for any $p \geq 1$ and present the variational form relating it to Θ_∞ in Lemma 2. For simplicity, the variational form (3) presented below holds only for monotone functions F ; the general form and other variational forms that parallel the ones known for the homogeneous envelope are presented in the Appendix.

Lemma 2. *The non-homogeneous convex envelope Θ_p of F_p , for monotone functions F , is given by*

$$\Theta_p(w) = \inf_{\eta \in [0,1]^d} \frac{1}{p} \sum_{j=1}^d \frac{|w_j|^p}{\eta_j^{p-1}} + \frac{1}{q} \Theta_\infty(\eta), \quad (3)$$

$$\Theta_\infty(w) = \min_{\alpha \geq 0} \left\{ \sum_{S \subseteq V} \alpha_S F(S) : \sum_{S \subseteq V} \alpha_S \mathbf{1}_S \geq |w|, \sum_{S \subseteq V} \alpha_S = 1 \right\}. \quad (4)$$

The infima in (1) and (3), for $w \in \text{dom}(\Theta_p)$, can be replaced by a minimization, if we extend $b \rightarrow \frac{a}{b}$ by continuity in zero with $\frac{a}{0} = \infty$ if $a \neq 0$ and 0 otherwise, as suggested in [24] and [3]. Note that, for $p = 1$, both relaxations reduce to $\Omega_1 = \Theta_1 = \|\cdot\|_1$. Hence, the ℓ_1 -relaxations essentially lose the combinatorial structure encoded in F . We thus follow up on the case $p > 1$.

In order to decide when to employ Ω_p or Θ_p , it is of interest to study the respective properties of these two relaxations and to identify when they coincide. Remark 1 shows that the homogeneous and non-homogeneous envelopes are identical, for $p = \infty$, for monotone submodular functions.

Remark 1. *If F is a monotone submodular function then $\Theta_\infty(w) = \Omega_\infty(w) = f_L(|w|), \forall w \in [-1, 1]^d$, where f_L denotes the Lovász extension of F [28].*

The two relaxations do not coincide in general: Note the added constraints $\eta \in [0, 1]^d$ in (3) and the sum constraint on α in (4). Another clear difference to note is that Ω_p are norms that belong to the broad family of H-norms [29, 3], as shown in [31]. On the other hand, by virtue of being non-homogeneous, Θ_p are not norms in general. We illustrate below two interesting examples where Ω_p and Θ_p differ.

Example 1 (Berhu penalty). *Since the cardinality function $F(S) = |S|$ is a monotone submodular function, $\Theta_\infty(w) = \Omega_\infty(w) = \|w\|_1$. However, this is not the case for $p \neq \infty$. In particular, we consider the ℓ_2 -regularized*

cardinality function $F_2^{\text{card}}(w) = \frac{1}{2}\|w\|_0 + \frac{1}{2}\|w\|_2^2$. Figure 2.1 shows that the non-homogeneous envelope is tighter than the homogeneous one in this case. Indeed, Ω_2^{card} is simply the ℓ_1 -norm, while Θ_2^{card} is given by $[\Theta_2^{\text{card}}(w)]_i = |w_i|$ if $|w_i| \leq 1$ and $[\Theta_2^{\text{card}}(w)]_i = \frac{1}{2}|w_i|^2 + \frac{1}{2}$ otherwise. This penalty, called ‘‘Berhu,’’ is introduced in [33] to produce a robust ridge regression estimator and is shown to be the convex envelope of F_2^{card} in [25].

This kind of behavior, where the non-homogeneous relaxation Θ_p acts as an ℓ_1 -norm on the small coefficients and as $\frac{1}{p}\|w\|_p^p$ for large ones, is not limited to the Berhu penalty, but holds for general set functions. However the point where the penalty moves from one mode to the other depends on the structure of F and is different along each coordinate. This is easier to see via the second variational form of Θ_p presented in the Appendix. We further illustrate in the following example.

Example 2 (Range penalty). *Consider the range function defined as $\text{range}(A) = \max(A) - \min(A) + 1$ where $\max(A)$ ($\min(A)$) denotes maximal (minimal) element in A . This penalty allow us to favor the selection of interval non-zero patterns on a chain or rectangular patterns on grids. It was shown in [31] that $\Omega_p(w) = \|w\|_1$ for any $p \geq 1$. On the other hand, Θ_p has no closed form solution, but is different from ℓ_1 -norm. Figure 2.1 illustrates the balls of different radii of Θ_∞ and Θ_2 . We can see how the penalty morphs from ℓ_1 -norm to ℓ_∞ and squared ℓ_2 -norm respectively, with different ‘‘speed’’ along each coordinate. Looking carefully for example on the ball $\Theta_2(w) \leq 2$, we can see that the penalty acts as an ℓ_1 -norm along the (x, z) -plane and as a squared ℓ_2 -norm along the (y, z) -plane.*

We highlight other ways in which the two relaxations differ and their implications in the sequel.

In terms of computational efficiency, note that even though the formulations (1) and (3) are jointly convex in (w, η) , Ω_p and Θ_p can still be intractable to compute and to optimize.

However, for certain classes of functions, they are tractable. For example, since for monotone submodular functions, $\Omega_\infty = \Theta_\infty$ is the Lovász extension of F , as stated in Remark 1, then they can be efficiently computed by the greedy algorithm [2]. Moreover, efficient algorithms to compute Ω_p , the associated proximal operator and to solve learning problems regularized with Ω_p is proposed in [31]. Similarly, if F can be expressed by integer programs over totally unimodular constraints as in [10], then $\Omega_\infty, \Theta_\infty$ and their associated Fenchel-type operators can be computed efficiently by linear programs. Hence, we can use conditional gradient algorithms for numerical solutions.

2.2 Lower combinatorial envelopes

In this section, we are interested in analyzing which combinatorial structures are preserved by each relaxation. To




Figure 1: ℓ_2 -regularized cardinality example in one dimension (left) and two dimensions (middle: homogeneous, right: non-homogeneous).




Figure 2: Balls of different radii of the non-homogeneous ℓ_∞ -convex envelope of the range function (top): $\Theta_\infty(w) \leq 1$ (left), $\Theta_\infty(w) \leq 2$ (middle), $\Theta_\infty(w) \leq 3$ (right) and of its ℓ_2 -convex envelope (bottom): $\Theta_2(w) \leq 1$ (left), $\Theta_2(w) \leq 2$ (middle), $\Theta_2(w) \leq 4$ (right).

that end, we generalize the notion of *lower combinatorial envelope* (LCE) [31]. The homogeneous LCE F_- of F is defined as the set function which agrees with the ℓ_∞ -homogeneous convex relaxation of F at the vertices of the unit hypercube, i.e., $F_-(A) = \Omega_\infty(\mathbb{1}_A)$, $\forall A \subseteq V$.

For the non-homogeneous relaxation, we define the non-homogeneous LCE similarly as $\tilde{F}_-(A) = \Theta_\infty(\mathbb{1}_A)$. The ℓ_∞ -relaxation reflects most directly the combinatorial structure of the function F . Indeed, ℓ_p -relaxations only depend on F through the ℓ_∞ -relaxation as expressed in the variational forms (1) and (3).

We say Ω_∞ is a tight relaxation of F if $F = F_-$. Similarly, Θ_∞ is a tight relaxation of F if $\tilde{F}_- = F$. Ω_∞ and Θ_∞ are then *extensions* of F from $\{0, 1\}^d$ to \mathbb{R}^d ; in this sense, the relaxation is tight for all w of the form $w = \mathbb{1}_A$. Moreover, following the definition of convex envelope, the relaxation Ω_∞ (resp. Θ_∞) is the same for F and F_- (resp. F and \tilde{F}_-), and hence, the LCE can be interpreted as the combinatorial structure preserved by each convex relaxation.

The homogeneous relaxation can capture any monotone submodular function [31]. Since Ω_∞ is the Lovász extension [1] in this case, and hence, $F_-(A) = \Omega_\infty(\mathbb{1}_A) = f_L(\mathbb{1}_A) = F(A)$. Also, since the two ℓ_∞ -relaxations are identical for this class of functions, their LCEs are also equal, i.e., $\tilde{F}_-(A) = \Theta_\infty(\mathbb{1}_A) = \Omega_\infty(\mathbb{1}_A) = F(A)$.

The LCEs, however, are not equal in general. In fact, the non-homogeneous relaxation is tight for a larger class of functions. In particular, the following proposition shows that \tilde{F}_- is equal to the *monotonization* of F , that is $\tilde{F}_-(A) = \inf_{S \subseteq V} \{F(S) : A \subseteq S\}$, for all set functions F , and is thus equal to the function itself if F is monotone.

Proposition 1. *The non-homogeneous lower combinatorial envelope can be written as*

$$\begin{aligned} \tilde{F}_-(A) &= \Theta_\infty(\mathbb{1}_A) \\ &= \inf_{\alpha_S \in \{0,1\}} \left\{ \sum_{S \subseteq V} \alpha_S F(S) : \sum_{S \subseteq V} \alpha_S \mathbb{1}_S \geq \mathbb{1}_A, \sum_{S \subseteq V} \alpha_S = 1 \right\} \\ &= \inf_{S \subseteq V} \{F(S) : A \subseteq S\}. \end{aligned}$$

Proof. To see why we can restrict α_S to be integral, let $\mathcal{E} = \{S : \alpha_S > 0\}$, then $\forall T \subseteq V$ such that $\exists e \in A, e \notin T$, then $\sum_{\alpha_S > 0, S \neq T} \alpha_S = 1$ and hence $\alpha_T = 0$. Hence $\forall S \in \mathcal{E}$ we have $A \subseteq S$ and $\sum_{\alpha_S > 0} \alpha_S F(S) = \min_{\alpha_S > 0} F(S)$. \square

Proposition 1 argues that the non-homogeneous convex envelope is tight if and only if F is monotone. Two important practical implications follow from this result.

Given a target model that cannot be expressed by a monotone function, it is impossible to obtain a tight convex relaxation. Non-convex methods can be potentially better.

On the other hand, if the model can be expressed by a monotone non-submodular set function, the homogeneous function may not be tight, and hence, a non-homogeneous relaxation can be more useful. For instance, [31] shows that for any set function where $F(\{e\}) = 1$ for all singletons $e \in V$ and $F(A) \geq |A|$, $\forall A \subseteq V$, the homogeneous LCE $F_-(A) = |A|$ and accordingly Ω_p is the ℓ_1 -norm, thus losing completely the structure encoded in F .

We discuss three examples that fall in this class of functions, where the non-homogeneous relaxation is tight while the homogeneous one is not.

Example 3 (Range penalty). *Consider $\text{range}(A) = \max(A) - \min(A) + 1$. For $F(A) = \text{range}(A)$, we have $F_-(A) = |A|$, while $\tilde{F}_- = F$ by Prop. 1.*

Example 4 (Dispersive ℓ_0 -penalty). *Given a set of predefined groups $\{G_1, \dots, G_M\}$, consider the dispersive ℓ_0 -*

penalty, introduced by [10]: $F(A) = |A| + \iota_{B^T \mathbf{1}_{A \leq 1}}(A)$ where the columns of B correspond to the indicator vectors of the groups, i.e., $B_{V,i} = \mathbf{1}_{G_i}$. The dispersive penalty enforces the selection of sparse supports where no two non-zeros are selected from the same group. Neural sparsity models induce such structures [18]. In this case, we have $F_-(A) = |A|$, while $\tilde{F}_- = F$ by Prop. 1.

Example 5 (Weighted graph model). Given a graph $\mathcal{G} = (V, E)$, consider a relaxed version of the weighted graph model of [19]: $F(A) = |A| + \iota_{\gamma(F_A) \leq g, \omega(F_A) \leq B}(A)$, where $\gamma(F_A)$ is the number of connected components formed by the forest F_A corresponding to A and $\omega(F_A)$ is the total weight of edges in the forest F_A . This model describes a wide range of structures, including 1D-clustering, tree hierarchies, and the Earth Mover Distance model. We have $F_-(A) = |A|$, while $\tilde{F}_- = F$ by Prop. 1.

The last two examples belong to a natural class of structured sparsity penalties of the form $F(A) = |A| + \iota_{A \in \mathcal{M}}(A)$, which favors sparse non-zero patterns among a set \mathcal{M} of allowed patterns. If \mathcal{M} is down-monotone, i.e., $\forall A \in \mathcal{M}, \forall B \subseteq A, B \in \mathcal{M}$, then the non-homogeneous relaxation preserves its structure, i.e., $\tilde{F}_- = F$, while its homogeneous relaxation is oblivious to the hard constraints, with $F_-(A) = |A|$.

3 Sparsity-inducing properties of convex relaxations

The notion of LCE captures the combinatorial structure preserved by convex relaxations in a geometric sense. In this section, we characterize the preserved structure from a statistical perspective.

To this end, we consider the linear regression model $y = Xw^* + \epsilon$, where $X \in \mathbb{R}^{n \times d}$ is a fixed design matrix, $y \in \mathbb{R}^n$ is the response vector, and ϵ is a vector of iid random variables with mean 0 and variance σ^2 . Given $\lambda_n > 0$, we define \hat{w} as a minimizer of the regularized least-squares:

$$\min_{w \in \mathbb{R}^d} \frac{1}{2} \|y - Xw\|_2^2 + \lambda_n \Phi(w), \quad (5)$$

We are interested in the sparsity-inducing properties of Ω_p and Θ_p on the solutions of (5). In this section, we consider though the more general setting where Φ is any proper normalized ($\Phi(0) = 0$) convex function which is absolute, i.e., $\Phi(w) = \Phi(|w|)$ and monotonic in the absolute values of w , that is $|w| \geq |w'| \Rightarrow \Phi(w) \geq \Phi(w')$. In what follows, monotone functions refer to this notion of monotonicity.

We determine in Section 3.1 necessary conditions for support recovery in (5) and in Section 3.2 we provide sufficient conditions for support recovery and consistency of a variant of (5). As both Ω_p and Θ_p are normalized absolute monotone convex functions, the results presented in this section apply directly to them as a corollary.

For simplicity, we assume $Q = X^T X/n \succ 0$, thus \hat{w} is unique. This forbids the high-dimensional setting. We expect though the insights developed towards the presented results to contribute to understanding the high-dimensional learning setting, which we defer to a later work.

3.1 Continuous stable supports

Existing results on the consistency of special cases of the estimator (5) typically rely heavily on decomposition properties of Φ [30, 1, 32, 31]. The notions of decomposability assumed in these prior works are either too strong or too specific to be applicable to the general convex penalties Ω_p and Θ_p we are considering. Instead, we introduce a general weak notion of decomposability applicable to any absolute monotone convex regularizer.

Definition 1 (Decomposability). Given $J \subseteq V$ and $w \in \mathbb{R}^d$, $\text{supp}(w) \subseteq J$, we say that Φ is decomposable at w w.r.t J if $\exists M_J > 0$ such that $\forall \Delta \in \mathbb{R}^d, \text{supp}(\Delta) \subseteq J^c$,

$$\Phi(w + \Delta) \geq \Phi(w) + M_J \|\Delta\|_\infty.$$

For example, for the ℓ_1 -norm, this decomposability property holds for any $J \subseteq V$ and $w \in \mathbb{R}^d$, with $M_J = 1$.

It is reasonable to expect this property to hold at the solution \hat{w} of (5) and its support $\hat{J} = \text{supp}(\hat{w})$. Theorem 1 shows that this is indeed the case. In Section 3.2, we devise an estimation scheme able to recover supports J that satisfy this property at any $w \in \mathbb{R}^d$. This leads then to following notion of *continuous* stable supports, which characterizes supports with respect to the continuous penalty Φ . In Section 4, we relate this to the notion of *discrete* stable supports, which characterizes supports with respect to the combinatorial penalty F .

Definition 2 (Continuous stability). We say that $J \subseteq V$ is weakly stable w.r.t Φ if there exists $w \in \mathbb{R}^d, \text{supp}(w) = J$ such that Φ is decomposable at w wrt J . Furthermore, we say that $J \subseteq V$ is strongly stable w.r.t Φ if for all $w \in \mathbb{R}^d$ s.t. $\text{supp}(w) \subseteq J$, Φ is decomposable at w wrt J .

Theorem 1 considers slightly more general estimators than (5) and shows that weak stability is a necessary condition for a non-zero pattern to be allowed as a solution.

Theorem 1. The minimizer \hat{w} of $\min_{w \in \mathbb{R}^d} L(w) - z^\top w + \lambda \Phi(w)$, where L is a strongly-convex and smooth loss function and $z \in \mathbb{R}^d$ has a continuous density w.r.t to the Lebesgue measure, has a weakly stable support w.r.t. Φ , with probability one.

This new result extends and simplifies the result in [1] which considers the special case of quadratic loss functions and Φ being the ℓ_∞ -convex relaxation of a submodu-

lar function. The proof we present, in the Appendix, is also shorter and simpler.

Corollary 1. *Assume $y \in \mathbb{R}^d$ has a continuous density w.r.t to the Lebesgue measure, then the support of the minimizer \hat{w} of Eq. (5) is weakly stable wrt Φ with probability one.*

3.2 Adaptive estimation

Restricting the choice of regularizers in (5) to convex relaxations as surrogates to combinatorial penalties is motivated by computational tractability concerns. However, other non-convex regularizers such as ℓ_α -quasi-norms [26, 14] or more generally penalties of the form $\Phi(w) = \sum_{i=1}^d \phi(|w_i|)$, where ϕ is a monotone concave penalty [12, 8, 16] can be more advantageous than the convex ℓ_1 -norm. Such penalties are closer to the ℓ_0 -quasi-norm and penalize more aggressively small coefficients, thus they have a stronger sparsity-inducing effect than ℓ_1 -norm.

The authors in [24] extended such concave penalties to the ℓ_α/ℓ_2 -quasi-norm $\Phi(w) = \sum_{i=1}^M \|w_{G_i}\|_\alpha$ for some $\alpha \in (0, 1)$, which enforces sparsity at the group level more aggressively. We generalize this to $\Phi(|w|^\alpha)$ where Φ is any structured sparsity-inducing monotone convex regularizer.

These non-convex penalties lead to intractable estimation problems, but approximate solutions can be obtained by majorization-minimization algorithms, as suggested for e.g., in [13, 39, 5].

Lemma 3. *Let Φ be a monotone convex function, $\Phi(|w|^\alpha)$ admits the following majorizer, $\forall w^0 \in \mathbb{R}^d$, $\Phi(|w|^\alpha) \leq (1 - \alpha)\Phi(|w^0|^\alpha) + \alpha\Phi(|w^0|^{\alpha-1} \circ |w|)$, which is tight at w^0 .*

We consider the adaptive weight estimator (6) resulting from applying a 1-step majorization-minimization to (5),

$$\min_{w \in \mathbb{R}^d} \frac{1}{2} \|y - Xw\|_2^2 + \lambda_n \Phi(|w^0|^{\alpha-1} \circ |w|), \quad (6)$$

where w^0 is a \sqrt{n} -consistent estimator to w^* , that is converging to w^* at rate $1/\sqrt{n}$ (typically obtained from $w^0 = \mathbb{1}$ or ordinary least-squares).

We study sufficient support recovery and estimation consistency conditions for (6) for general convex monotone regularizers Φ . Such consistency results have been established for (6), in the classical asymptotic setting, only in the special case of ℓ_1 -norm in [38] and for the (non-adaptive) estimator (5) for homogeneous convex envelopes of monotone submodular functions, for $p = \infty$ in [1] and for general p in [31], in the high dimensional setting, and for latent group Lasso norm in [32], in the asymptotic setting.

Compared to prior works, the discussion of support recovery is complicated here by the fact that Φ is not necessarily a norm (e.g., if $\Phi = \Theta_p$) and only satisfies a weak notion of decomposability.

As in [38], we consider the classical asymptotic regime in which the model generating the data is of fixed finite dimension d while $n \rightarrow \infty$. As before, we assume $Q \succ 0$ and thus the minimizer of (6) is unique, we denote it by \hat{w} .

The following Theorem extends the results of [38] for the ℓ_1 -norm to any normalized absolute monotone convex regularizer if the true support satisfy the sufficient condition of strong stability in Definition 2. As we previously remarked this condition is trivially satisfied for the ℓ_1 -norm.

Theorem 2. *[Consistency and Support Recovery] Let $\Phi : \mathbb{R}^d \rightarrow \overline{\mathbb{R}}_+$ be a proper normalized absolute monotone convex function and denote by J the true support $J = \text{supp}(w^*)$. If $|w^*|^\alpha \in \text{int dom } \Phi$, J is strongly stable with respect to Φ and λ_n satisfies $\frac{\lambda_n}{\sqrt{n}} \rightarrow 0$, $\frac{\lambda_n}{n^{\alpha/2}} \rightarrow \infty$, then the estimator (6) is consistent and asymptotically normal, i.e., it satisfies*

$$\sqrt{n}(\hat{w}_J - w_J^*) \xrightarrow{d} \mathcal{N}(0, \sigma^2 Q_{JJ}^{-1}), \quad (7)$$

and

$$P(\text{supp}(\hat{w}) = J) \rightarrow 1. \quad (8)$$

Consistency results in most existing works are established under various necessary conditions on X , some of which are difficult to verify in practice, such as the *irrepresentability condition* (c.f., [38, 1, 32, 31]). Adding data-dependent weights does not require such conditions and allows recovery even in the correlated measurement matrix setup as illustrated in our numerical results (c.f., Sect. 5).

4 Sparsity-inducing properties of combinatorial penalties

In section 3, we derived necessary and sufficient conditions for support recovery defined with respect to the continuous convex penalties Ω_p and Θ_p . In this Section, we translate these to conditions with respect to the combinatorial penalties F_p themselves. Hence, the results of this section allows one to check which supports to expect to recover, without the need to compute the corresponding convex relaxation. To that end, we introduce in Section 4.1 discrete counterparts of weak and strong stability, and show in Section 4.2 that discrete strong stability is a sufficient, and in some cases necessary, condition for support recovery.

4.1 Discrete stable supports

We recall the concept of discrete stable sets [1], also referred to as *flat* or *closed* sets [27]. We refer to such sets as discrete weakly stable sets and introduce a stronger notion of discrete stability.

Definition 3 (Discrete stability). *Given a monotone set function $F : 2^V \rightarrow \overline{\mathbb{R}}_+$, a set $J \subseteq V$ is said to be weakly stable w.r.t F if $\forall i \in J^c$, $F(J \cup \{i\}) > F(J)$.*

A set $J \subseteq V$ is said to be strongly stable w.r.t F if $\forall A \subseteq J, \forall i \in J^c, F(A \cup \{i\}) > F(A)$.

Note that discrete stability imply in particular feasibility, i.e., $F(J) < +\infty$. Also, if F is a strictly monotone function, such as the cardinality function, then all supports are stable w.r.t F . It is interesting to note that for monotone submodular functions, weak and strong stability are equivalent. In fact, this equivalence holds for a more general class of functions, we call ρ -submodular.

Definition 4. A function $F : 2^V \rightarrow \mathbb{R}$ is ρ -submodular iff $\exists \rho \in (0, 1]$ s.t., $\forall B \subseteq V, A \subseteq B, i \in B^c$

$$\rho[F(B \cup \{i\}) - F(B)] \leq F(A \cup \{i\}) - F(A)$$

The notion of ρ -submodularity is a special case of the weakly DR-submodular property defined for continuous functions [17]. It is also related to the notion of weak submodularity (c.f., [7, 11]). We show in the appendix that ρ -submodularity is a stronger condition.

Proposition 2. If F is a finite-valued monotone function, F is ρ -submodular iff discrete weak stability is equivalent to strong stability.

Example 6. The range function $\text{range}(A) = \max(A) - \min(A) + 1$ is ρ -submodular with $\rho = \frac{1}{d-1}$.

4.2 Relation between discrete and continuous stability

This section provides several technical results relating the discrete and continuous notions of stability. It thus provides us with the necessary tools to characterize which supports can be correctly estimated w.r.t the combinatorial penalty itself, without going through its relaxations.

Proposition 3. Given any monotone set function F , all sets $J \subseteq V$ strongly stable w.r.t to F are also strongly stable w.r.t Ω_p and Θ_p .

It follows then by Theorem 2 that discrete strong stability is a sufficient condition for correct estimation.

Corollary 2. If Φ is equal to Ω_p or Θ_p for $p \in (1, \infty)$ and $\text{supp}(w^*) = J$ is strongly stable w.r.t F , then Theorem 2 holds, i.e., the adaptive estimator (6) is consistent and correctly recovers the support. This also holds for $p = \infty$ if we further assume that $\|w^*\|_\infty < 1$.

Furthermore, if F is ρ -submodular, then by Proposition 2, it is enough for $\text{supp}(w^*) = J$ to be weakly stable w.r.t F for Corollary 2 to hold. Conversely, Proposition 4 below shows that discrete strong stability is also a necessary condition for continuous strong stability, in the case where $p = \infty$ and F is equal to its LCE.

Proposition 4. If $F = F_-$ and J is strongly stable w.r.t Ω_∞ , then J is strongly stable w.r.t F . Similarly, for any

monotone F , if J is strongly stable w.r.t Θ_∞ , then J is strongly stable w.r.t F .

Finally, in the special case of monotone submodular function, the following Corollary 3, along with Proposition 3 demonstrates that all definitions of stability become equivalent. We thus recover the result in [1] showing that discrete weakly stable supports correspond to the set of allowed sparsity patterns for monotone submodular functions.

Corollary 3. If F is monotone submodular and J is weakly stable w.r.t $\Omega_\infty = \Theta_\infty$ then J is weakly stable w.r.t F .

4.3 Examples

We highlight in this section what are the supports recovered by the adaptive estimator (AE) (6) with the homogeneous convex relaxation Ω_p and non-homogeneous convex relaxation Θ_p of some examples of structure priors. For simplicity, we will focus on the case $p = \infty$. Also in all the examples we consider below, weak and strong discrete stability are equivalent, so we omit the weak/strong specification. Note that it is desirable that the regularizer used enforces the recovery of only the non-zero patterns satisfying the desired structure.

Monotone submodular functions: As discussed above, for this class of functions, all stability definitions are equivalent and $\Omega_\infty = \Theta_\infty = f_L$. As a result, AE recovers any discrete stable non-zero pattern. This includes the following examples (c.f., [31] for further examples).

- **Cardinality:** As a strictly monotone function, all supports are stable w.r.t to it. Thus AE recovers *all* non-zero patterns with Ω_∞ and Θ_∞ , given by the ℓ_1 -norm.
- **Overlap count function:** $F_\cap(A) = \sum_{G \in \mathcal{G}, G \cap A \neq \emptyset} d_G$ where \mathcal{G} is a collection of predefined groups G and d_G their associated weights. Ω_∞ and Θ_∞ are given by the ℓ_1/ℓ_∞ -group Lasso norm, and stable patterns are complements of union of groups. For example, for hierarchical groups (i.e., groups consisting of each node and its descendants on a tree), AE recovers rooted connected tree supports.
- **Modified range function:** The range function can be transformed into a submodular function, if scaled by a constant as suggested in [1], yielding the monotone submodular function $F^{\text{mr}}(A) = d - 1 + \text{range}(A)$, $\forall A \neq \emptyset$ and $F^{\text{mr}}(\emptyset) = 0$. This can actually be written as an instance of F_\cap with groups defined as $\mathcal{G} = \{[1, k] : 1 \leq k \leq d\} \cup \{[k, d] : 1 \leq k \leq d\}$. This norm was proposed to induce interval patterns by [23], and indeed its stable patterns are interval supports. We will compare this function in the experiments with the direct convex relaxations of the range function.

Range function: The range function is $\frac{1}{d-1}$ -submodular,




Figure 3: (Left column) Best Hamming distance and (Right column) best least square error to the true vector w^* , along the regularization path, averaged over 5 runs.

thus its discrete strongly and weakly stable supports are identical and they correspond to interval supports. As a result, AE recovers interval supports with Θ_∞ . On the other hand, since the homogeneous LCE of the range function is the cardinality, AE recovers all supports with Ω_∞ .

Down monotone structures: Functions of the form $F(A) = |A| + \iota_{A \in \mathcal{M}}(A)$, where \mathcal{M} is down-monotone, also have their discrete strongly and weakly stable supports identical and given by the feasible set \mathcal{M} . These structures include the dispersive and graph models discussed in examples 4 and 5. Since their homogeneous LCE is also the cardinality, then AE recovers all supports with Ω_∞ , and only feasible supports with Θ_∞ .

5 Numerical Illustration

To illustrate the results presented in this paper, we consider the problem of estimating the support of a parameter vector $w^* \in \mathbb{R}^d$ whose support is an interval. It is natural then to choose as combinatorial penalty the range function whose stable supports are intervals. We aim to study the effect of adaptive weights, as well as the effect of the choice of homogeneous vs. non-homogeneous convex relaxation for regularization, on the quality of support recovery.

As discussed in Section 4.3, the ℓ_∞ -homogeneous convex envelope of the range is simply the ℓ_1 -norm. Its ℓ_∞ -non-homogeneous convex envelope Θ_∞^r can be computed using the formulation (3), where only interval sets need to be considered in the constraints, leading to a quadratic number of constraints. We also consider the ℓ_1/ℓ_∞ -norm that corresponds to the convex relaxation of the modified range

function F^{mr} .

We consider a simple regression setting in which $w^* \in \mathbb{R}^d$ is a constant signal whose support is an interval. The choice of $p = \infty$ is well suited for constant valued signals. The design matrix $X \in \mathbb{R}^{d \times n}$ is either drawn as (1) an i.i.d Gaussian matrix with normalized columns, or (2) a correlated Gaussian matrix with normalized columns, with the off-diagonal values of the covariance matrix set to a value $\rho = 0.5$. We observe noisy linear measurements $y = Xw^* + \epsilon$, where the noise vector is i.i.d. with variance σ^2 , where σ is varied between 10^{-5} and 1. We solve problem (6) with and without adaptive weights $|w^0|^{\alpha-1}$, where w^0 is taken to be the least squares solution and $\alpha = 0.3$.

We assess the estimators obtained through the different regularizers both in terms of support recovery and in terms of estimation error. Figure 5 plots (in logscale) these two criteria against the noise level σ . We plot the best achieved error on the regularization path, where the regularization parameter λ was varied between 10^{-6} and 10^3 . We set the parameters to $d = 250, k = 100, n = 500$.

We observe that the adaptive weight scheme helps in support recovery, especially in the correlated design setting. Indeed, Lasso is only guaranteed to recover the support under an ‘‘irrepresentability condition’’ [38]. This is satisfied with high probability only in the non-correlated design. On the other hand, adaptive weights allow us to recover any strongly stable support, without any additional condition, as shown in Theorem 2. The ℓ_1/ℓ_∞ -norm performs poorly in this setup. In fact, the modified range function F^{mr} , introduced a gap of d between non-empty sets and the empty set. This leads to the undesirable behavior, already documented in [1, 23] of adding all the variables in one step, as opposed to gradually. Adaptive weights seem to correct for this effect, as seen by the significant improvement in performance. Finally, note that choosing the ‘‘tighter’’ non-homogeneous convex relaxation leads to better support recovery. Indeed, Θ_∞^r performs better than ℓ_1 -norm in all setups.

6 Conclusion

We presented an analysis of homogeneous and non-homogeneous convex relaxations of ℓ_p -regularized combinatorial penalties. Our results show that structure encoded by submodular priors can be equally well expressed by both relaxations, while the non-homogeneous relaxation is able to express the structure of more general monotone set functions. We also identified necessary and sufficient stability conditions on the supports to be correctly recovered. We proposed an adaptive weight scheme that is guaranteed to recover supports that satisfy the sufficient stability conditions, in the asymptotic setting, even under correlated design matrix.

Acknowledgements

We thank Ya-Ping Hsieh for helpful discussions. This work was supported in part by the European Commission under ERC Future Proof, SNF 200021-146750, SNF CRSII2-147633, NCCR Marvel. Francis Bach acknowledges support from the chaire Economie des nouvelles données with the data science joint research initiative with the fonds AXA pour la recherche, and the Initiative de Recherche “Machine Learning for Large-Scale Insurance” from the Institut Louis Bachelier.

References

- [1] F. Bach. Structured sparsity-inducing norms through submodular functions. In *NIPS*, pages 118–126, 2010.
- [2] F. Bach. Learning with submodular functions: A convex optimization perspective. *arXiv preprint arXiv:1111.6453*, 2011.
- [3] Francis Bach, Rodolphe Jenatton, Julien Mairal, Guillaume Obozinski, et al. Optimization with sparsity-inducing penalties. *Foundations and Trends® in Machine Learning*, 4(1):1–106, 2012.
- [4] R.G. Baraniuk, V. Cevher, M.F. Duarte, and C. Hegde. Model-based compressive sensing. *Information Theory, IEEE Transactions on*, 56(4):1982–2001, 2010.
- [5] Emmanuel J Candes, Michael B Wakin, and Stephen P Boyd. Enhancing sparsity by reweighted minimization. *Journal of Fourier analysis and applications*, 14(5):877–905, 2008.
- [6] V. Chandrasekaran, B. Recht, P.A. Parrilo, and A.S. Willsky. The convex geometry of linear inverse problems. *Foundations of Computational Mathematics*, 12:805–849, 2012.
- [7] Abhimanyu Das and David Kempe. Submodular meets spectral: Greedy algorithms for subset selection, sparse approximation and dictionary selection. *arXiv preprint arXiv:1102.3975*, 2011.
- [8] Ingrid Daubechies, Ronald DeVore, Massimo Fornasier, and C Sinan Güntürk. Iteratively reweighted least squares minimization for sparse recovery. *Communications on Pure and Applied Mathematics*, 63(1):1–38, 2010.
- [9] Shaddin Dughmi. Submodular functions: Extensions, distributions, and algorithms. a survey. *arXiv preprint arXiv:0912.0322*, 2009.
- [10] M. El Halabi and V. Cevher. A totally unimodular view of structured sparsity. *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics*, pp. 223–231, 2015.
- [11] Ethan R Elenberg, Rajiv Khanna, Alexandros G Dimakis, and Sahand Negahban. Restricted strong convexity implies weak submodularity. *arXiv preprint arXiv:1612.00804*, 2016.
- [12] Jianqing Fan and Runze Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American statistical Association*, 96(456):1348–1360, 2001.
- [13] Mário AT Figueiredo, José M Bioucas-Dias, and Robert D Nowak. Majorization–minimization algorithms for wavelet-based image restoration. *IEEE Transactions on Image processing*, 16(12):2980–2991, 2007.
- [14] LLdiko E Frank and Jerome H Friedman. A statistical view of some chemometrics regression tools. *Technometrics*, 35(2):109–135, 1993.
- [15] S. Fujishige. *Submodular Functions and Optimization*. Elsevier, 2005.
- [16] Gilles Gasso, Alain Rakotomamonjy, and Stéphane Canu. Recovering sparse signals with a certain family of nonconvex penalties and dc programming. *IEEE Transactions on Signal Processing*, 57(12):4686–4698, 2009.
- [17] Hamed Hassani, Mahdi Soltanolkotabi, and Amin Karbasi. Gradient methods for submodular maximization. In *Advances in Neural Information Processing Systems*, pages 5843–5853, 2017.
- [18] Chinmay Hegde, Marco F Duarte, and Volkan Cevher. Compressive sensing recovery of spike trains using a structured sparsity model. In *SPARS’09-Signal Processing with Adaptive Sparse Structured Representations*, 2009.
- [19] Chinmay Hegde, Piotr Indyk, and Ludwig Schmidt. A nearly-linear time framework for graph-structured sparsity. In *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, pages 928–937, 2015.
- [20] J. Huang, T. Zhang, and D. Metaxas. Learning with structured sparsity. *The Journal of Machine Learning Research*, 12:3371–3412, 2011.
- [21] Laurent Jacob, Guillaume Obozinski, and Jean-Philippe Vert. Group lasso with overlap and graph lasso. In *Proceedings of the 26th annual international conference on machine learning*, pages 433–440. ACM, 2009.
- [22] R. Jenatton, J. Mairal, G. Obozinski, and F. Bach. Proximal methods for hierarchical sparse coding. *Journal of Machine Learning Research*, 12:2297–2334, 2011.

- [23] Rodolphe Jenatton, Jean-Yves Audibert, and Francis Bach. Structured variable selection with sparsity-inducing norms. *Journal of Machine Learning Research*, 12(Oct):2777–2824, 2011.
- [24] Rodolphe Jenatton, Guillaume Obozinski, and Francis Bach. Structured sparse principal component analysis. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 366–373, 2010.
- [25] Vladimir Jojic, Suchi Saria, and Daphne Koller. Convex envelopes of complexity controlling penalties: the case against premature envelopment. In *International Conference on Artificial Intelligence and Statistics*, pages 399–406, 2011.
- [26] Keith Knight and Wenjiang Fu. Asymptotics for lasso-type estimators. *Annals of statistics*, pages 1356–1378, 2000.
- [27] Andreas Krause and Carlos E Guestrin. Near-optimal nonmyopic value of information in graphical models. *arXiv preprint arXiv:1207.1394*, 2012.
- [28] L. Lovász. Submodular functions and convexity. In *Mathematical Programming The State of the Art*, pages 235–257. Springer, 1983.
- [29] Charles A Micchelli, Jean M Morales, and Massimiliano Pontil. Regularizers for structured sparsity. *Advances in Computational Mathematics*, pages 1–35, 2013.
- [30] Sahand Negahban, Pradeep Ravikumar, Martin J Wainwright, and Bin Yu. A unified framework for high-dimensional analysis of m-estimators with decomposable regularizers. In *Adv. Neural Inf. Proc. Sys.(NIPS)*, 2011.
- [31] G. Obozinski and F. Bach. Convex relaxation for combinatorial penalties. *arXiv preprint arXiv:1205.1240*, 2012.
- [32] Guillaume Obozinski, Laurent Jacob, and Jean-Philippe Vert. Group lasso with overlaps: the latent group lasso approach. *arXiv preprint arXiv:1110.0413*, 2011.
- [33] Art B Owen. A robust hybrid of lasso and ridge regression. *Contemporary Mathematics*, 443:59–72, 2007.
- [34] Maurice Sion et al. On general minimax theorems. *Pacific J. Math*, 8(1):171–176, 1958.
- [35] Jan Vondrák. Continuous extensions of submodular functions. CS 369P: Polyhedral techniques in combinatorial optimization, <http://theory.stanford.edu/~jvondrak/CS369P-files/lec17.pdf>, November 2010.
- [36] Ming Yuan and Yi 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.
- [37] Peng Zhao, Guilherme Rocha, and Bin Yu. Grouped and hierarchical model selection through composite absolute penalties. *Department of Statistics, UC Berkeley, Tech. Rep.*, 703, 2006.
- [38] Hui Zou. The adaptive lasso and its oracle properties. *Journal of the American statistical association*, 101(476):1418–1429, 2006.
- [39] Hui Zou and Runze Li. One-step sparse estimates in nonconcave penalized likelihood models. *Annals of statistics*, 36(4):1509, 2008.