

Binarsity: a penalization for one-hot encoded features

Mokhtar Z. Alaya

MOKHTARZAHDI.ALAYA@GMAIL.COM

*Theoretical and applied statistics laboratory
Pierre and Marie Curie University
Paris, France*

Simon Bussy

SIMON.BUSSY@GMAIL.COM

*Theoretical and applied statistics laboratory
Pierre and Marie Curie University
Paris, France*

Stéphane Gaïffas

STEPHANE.GAIFFAS@POLYTECHNIQUE.EDU

*Centre de Mathématiques Appliquées, École Polytechnique
CNRS UMR 7641
91128 Palaiseau, France*

Agathe Guilloux

AGATHE.GUILLOUX@MATH.CNRS.FR

*LaMME, UEVE and UMR 8071
Université Paris Saclay
Evry, France*

Editor:

Abstract

This paper deals with the problem of large-scale linear supervised learning in settings where a large number of continuous features are available. We propose to combine the well-known trick of one-hot encoding of continuous features with a new penalization called *binarsity*. In each group of binary features coming from the one-hot encoding of a single raw continuous feature, this penalization uses total-variation regularization together with an extra linear constraint to avoid collinearity within groups. Non-asymptotic oracle inequalities for generalized linear models are proposed, and numerical experiments illustrate the good performances of our approach on several datasets. It is also noteworthy that our method has a numerical complexity comparable to standard ℓ_1 penalization.

Keywords: Supervised learning, Features binarization, Total-variation, Oracle inequalities, Proximal methods

1. Introduction

In many applications, datasets used for supervised learning contain a large number of continuous features, with a large number of samples. An example is web-marketing, where features are obtained from bag-of-words scaled using tf-idf (Russell, 2013), recorded during the visit of users on websites. A well-known trick (Wu and Coggeshall, 2012; Liu et al., 2002) in this setting is to replace each raw continuous feature by a set of binary features that one-hot encodes the interval containing it, among a list of intervals partitioning the raw feature range. This leads to a non-linear decision function with respect to the raw con-

tinuous features space, and can therefore improve prediction. However, this trick is prone to over-fitting, since it increases significantly the dimension of the problem.

A new penalization. To overcome this problem, we introduce a new penalization called *binarsity*, that penalizes the model weights learned from such grouped one-hot encodings (one group for each raw continuous feature). Since the binary features within these groups are naturally ordered, the binarsity penalization combines a group total-variation penalization, with an extra linear constraint in each group to avoid collinearity between the one-hot encodings. This penalization forces the weights of the model to be as constant (with respect to the order induced by the original feature) as possible within a group, by selecting a minimal number of relevant cut-points. Moreover, if the model weights are all equal within a group, then the full block of weights is zero, because of the extra linear constraint. This allows to perform raw feature selection.

Sparsity. To address the high-dimensionality of features, sparse inference is now an ubiquitous technique for dimension reduction and variable selection, see for instance Bühlmann and van De Geer (2011) and Hastie et al. (2001) among many others. The principle is to induce sparsity (large number of zeros) in the model weights, assuming that only a few features are actually helpful for the label prediction. The most popular way to induce sparsity in model weights is to add a ℓ_1 -penalization (Lasso) term to the goodness-of-fit (Tibshirani, 1996a). This typically leads to sparse parametrization of models, with a level of sparsity that depends on the strength of the penalization. Statistical properties of ℓ_1 -penalization have been extensively investigated, see for instance Knight and Fu (2000); Zhao and Yu (2006); Bunea et al. (2007); Bickel et al. (2009) for linear and generalized linear models and Donoho and Huo (2001); Donoho and Elad (2002); Candès et al. (2008); Candès and Wakin (2008) for compressed sensing, among others.

However, the Lasso ignores ordering of features. In Tibshirani et al. (2005), a structured sparse penalization is proposed, known as fused Lasso, which provides superior performance in recovering the true model in such applications where features are ordered in some meaningful way. It introduces a mixed penalization using a linear combination of the ℓ_1 -norm and the total-variation penalization, thus enforcing sparsity in both the weights and their successive differences. Fused Lasso has achieved great success in some applications such as comparative genomic hybridization (Rapaport et al., 2008), image denoising (Friedman et al., 2007), and prostate cancer analysis (Tibshirani et al., 2005).

Features discretization and cuts. For supervised learning, it is often useful to encode the input features in a new space to let the model focus on the relevant areas (Wu and Coggeshall, 2012). One of the basic encoding technique is *feature discretization* or *feature quantization* (Liu et al., 2002) that partitions the range of a continuous feature into intervals and relates these intervals with meaningful labels. Recent overviews of discretization techniques can be found in Liu et al. (2002) or Garcia et al. (2013).

Obtaining the optimal discretization is a NP-hard problem (Chlebus and Nguyen, 1998), and an approximation can be easily obtained using a greedy approach, as proposed in decision trees: CART (Breiman et al., 1984) and C4.5 (Quinlan, 1993), among others, that sequentially select pairs of features and cuts that minimize some purity measure (intra-variance, Gini index, information gain are the main examples). These approaches build

decision functions that are therefore very simple, by looking only at a single feature at a time, and a single cut at a time. Ensemble methods (boosting (Lugosi and Vayatis, 2004), random forests (Breiman, 2001)) improve this by combining such decisions trees, at the expense of models that are harder to interpret.

Organization of the paper. The main contribution of this paper is the idea to use a total-variation penalization, with an extra linear constraint, on the weights of a model trained on a binarization of the raw continuous features, leading to a procedure that selects multiple cut-points per feature, looking at all features simultaneously. The proposed methodology is described in Section 2. Section 3 establishes an oracle inequality for generalized linear models. Section 4 highlights the results of the method on various datasets and compares its performances to well known classification algorithms. Finally, we discuss the obtained results in Section 5.

Notations. Throughout the paper, for every $q > 0$, we denote by $\|v\|_q$ the usual ℓ_q -quasi norm of a vector $v \in \mathbb{R}^m$, namely $\|v\|_q = (\sum_{k=1}^m |v_k|^q)^{1/q}$, and $\|v\|_\infty = \max_{k=1, \dots, m} |v_k|$. We also denote $\|v\|_0 = |\{k : v_k \neq 0\}|$, where $|A|$ stands for the cardinality of a finite set A . For $u, v \in \mathbb{R}^m$, we denote by $u \odot v$ the Hadamard product $u \odot v = (u_1 v_1, \dots, u_m v_m)^\top$. For any $u \in \mathbb{R}^m$ and any $L \subset \{1, \dots, m\}$, we denote u_L as the vector in \mathbb{R}^m satisfying $(u_L)_k = u_k$ for $k \in L$ and $(u_L)_k = 0$ for $k \in L^c = \{1, \dots, m\} \setminus L$. We write $\mathbf{1}_m$ (resp. $\mathbf{0}_m$) for the vector of \mathbb{R}^m having all coordinates equal to one (resp. zero). Finally, we denote by $\text{sign}(x)$ the set of sub-differentials of the function $x \mapsto |x|$, namely $\text{sign}(x) = \{1\}$ if $x > 0$, $\text{sign}(x) = \{-1\}$ if $x < 0$ and $\text{sign}(0) = [-1, 1]$.

2. The proposed method

Consider a supervised training dataset $(x_i, y_i)_{i=1, \dots, n}$ containing features $x_i = (x_{i,1}, \dots, x_{i,p})^\top \in \mathbb{R}^p$ and labels $y_i \in \mathcal{Y} \subset \mathbb{R}$, that are independent and identically distributed samples of (X, Y) with unknown distribution \mathbb{P} . Let us denote $\mathbf{X} = [x_{i,j}]_{1 \leq i \leq n; 1 \leq j \leq p}$ the $n \times p$ features matrix vertically stacking the n samples of p raw features. Let $\mathbf{X}_{\bullet, j}$ be the j -th feature column of \mathbf{X} .

Binarization. The binarized matrix \mathbf{X}^B is a matrix with an extended number $d > p$ of columns, where the j -th column $\mathbf{X}_{\bullet, j}$ is replaced by $d_j \geq 2$ columns $\mathbf{X}_{\bullet, j, 1}^B, \dots, \mathbf{X}_{\bullet, j, d_j}^B$ containing only zeros and ones. Its i -th row is written

$$x_i^B = (x_{i,1,1}^B, \dots, x_{i,1,d_1}^B, x_{i,2,1}^B, \dots, x_{i,2,d_2}^B, \dots, x_{i,p,1}^B, \dots, x_{i,p,d_p}^B)^\top \in \mathbb{R}^d.$$

In order to simplify presentation of our results, we assume in the paper that all raw features $\mathbf{X}_{\bullet, j}$ are continuous, so that they are transformed using the following one-hot encoding. We consider a full partitioning without overlap: $\cup_{k=1}^{d_j} I_{j,k} = \text{range}(\mathbf{X}_{\bullet, j})$ and $I_{j,k} \cap I_{j,k'} = \emptyset$ for all $k \neq k'$ with $k, k' \in \{1, \dots, d_j\}$, and define

$$x_{i,j,k}^B = \begin{cases} 1 & \text{if } x_{i,j} \in I_{j,k}, \\ 0 & \text{otherwise} \end{cases}$$

for $i = 1, \dots, n$ and $k = 1, \dots, d_j$. A natural choice of intervals is given by quantiles, namely $I_{j,1} = [q_j(0), q_j(\frac{1}{d_j})]$ and $I_{j,k} = (q_j(\frac{k-1}{d_j}), q_j(\frac{k}{d_j})]$ for $k = 2, \dots, d_j$, where $q_j(\alpha)$ denotes a

quantile of order $\alpha \in [0, 1]$ for $\mathbf{X}_{\bullet, j}$. In practice, if there are ties in the estimated quantiles for a given feature, we simply choose the set of ordered unique values to construct the intervals. This principle of binarization is a well-known trick (Garcia et al., 2013), that allows to construct a non-linear decision with respect to the raw feature space. If training data contains also unordered qualitative features, one-hot encoding with ℓ_1 -penalization can be used for instance. Note that, however, not all forms of non-linear decision functions can be approximated using the binarization trick, hence with the approach described in this paper. In particular, it does not work in ‘‘XOR’’ situations, since binarization allows to replace linearities in the decision function by piecewise constant functions, as illustrated in Figure 2.

Goodness-of-fit. Given a loss function $\ell : \mathcal{Y} \times \mathbb{R} \rightarrow \mathbb{R}$, we consider the goodness-of-fit term

$$R_n(\theta) = \frac{1}{n} \sum_{i=1}^n \ell(y_i, m_\theta(x_i)), \tag{1}$$

where $m_\theta(x_i) = \theta^\top x_i^B$ and $\theta \in \mathbb{R}^d$ with $d = \sum_{j=1}^p d_j$. We then have $\theta = (\theta_{1,\bullet}^\top, \dots, \theta_{p,\bullet}^\top)^\top$, with $\theta_{j,\bullet}$ corresponding to the group of coefficients weighting the binarized raw j -th feature. We focus on generalized linear models (Green and Silverman, 1994), where the conditional distribution $Y|X = x$ is assumed to be from a one-parameter exponential family distribution with a density of the form

$$y|x \mapsto f^0(y|x) = \exp\left(\frac{ym^0(x) - b(m^0(x))}{\phi} + c(y, \phi)\right), \tag{2}$$

with respect to a reference measure which is either the Lebesgue measure (e.g. in the Gaussian case) or the counting measure (e.g. in the logistic or Poisson cases), leading to a loss function of the form

$$\ell(y_1, y_2) = -y_1 y_2 + b(y_2).$$

The density described in (2) encompasses several distributions, see Table 1. The functions $b(\cdot)$ and $c(\cdot)$ are known, while the natural parameter function $m^0(\cdot)$ is unknown. The dispersion parameter ϕ is assumed to be known in what follows. It is also assumed that $b(\cdot)$ is three times continuously differentiable. It is standard to notice that

$$\mathbb{E}[Y|X = x] = \int y f^0(y|x) dy = b'(m^0(x)),$$

where b' stands for the derivative of b . This formula explains how b' links the conditional expectation to the unknown m^0 . The results given in Section 3 rely on the following Assumption.

Assumption 1 *Assume that b is three times continuously differentiable, and that there exist constants $C_n > 0$, and $0 < L_n \leq U_n$ such that $C_n = \max_{i=1, \dots, n} |m^0(x_i)| < \infty$ and $L_n \leq \max_{i=1, \dots, n} b''(m^0(x_i)) \leq U_n$.*

This assumption is satisfied for most standard generalized linear models. In Table 1, we list some standard examples that fit in this framework, see also van de Geer (2008); Rigollet (2012).

	ϕ	$b(z)$	$b'(z)$	$b''(z)$	L_n	U_n
Normal	σ^2	$\frac{z^2}{2}$	z	1	1	1
Logistic	1	$\log(1 + e^z)$	$\frac{e^z}{1+e^z}$	$\frac{e^z}{(1+e^z)^2}$	$\frac{e^{C_n}}{(1+e^{C_n})^2}$	$\frac{1}{4}$
Poisson	1	e^z	e^z	e^z	e^{-C_n}	e^{C_n}

Tab. 1: Examples of standard distributions that fit in the considered setting of generalized linear models, with the corresponding constants in Assumption 1.

Binarsity. Several problems occur when using the binarization trick described above:

- (P1) The one-hot-encodings satisfy $\sum_{k=1}^{d_j} \mathbf{X}_{i,j,k}^B = 1$ for $j = 1, \dots, p$, meaning that the columns of each block sum to $\mathbf{1}_n$, making \mathbf{X}^B not of full rank by construction.
- (P2) Choosing the number of intervals d_j for binarization of each raw feature j is not an easy task, as too many might lead to overfitting: the number of model-weights increases with each d_j , leading to a over-parametrized model.
- (P3) Some of the raw features $\mathbf{X}_{\bullet,j}$ might not be relevant for the prediction task, so we want to select raw features from their one-hot encodings, namely induce block-sparsity in θ .

A usual way to deal with (P1) is to impose a linear constraint (Agresti, 2015) in each block. In our penalization term, we impose

$$\sum_{k=1}^{d_j} \theta_{j,k} = 0 \quad (3)$$

for all $j = 1, \dots, p$. Now, the trick to tackle (P2) is to remark that within each block, binary features are ordered. We use a within block total-variation penalization

$$\sum_{j=1}^p \|\theta_{j,\bullet}\|_{\text{TV}, \hat{w}_{j,\bullet}}$$

where

$$\|\theta_{j,\bullet}\|_{\text{TV}, \hat{w}_{j,\bullet}} = \sum_{k=2}^{d_j} \hat{w}_{j,k} |\theta_{j,k} - \theta_{j,k-1}|, \quad (4)$$

with weights $\hat{w}_{j,k} > 0$ to be defined later, to keep the number of different values taken by $\theta_{j,\bullet}$ to a minimal level. Finally, dealing with (P3) is actually a by-product of dealing with (P1) and (P2). Indeed, if the raw feature j is not-relevant, then $\theta_{j,\bullet}$ should have all entries constant because of the penalization (4), and in this case all entries are zero, because of (3). We therefore introduce the following penalization, called *binarsity*

$$\text{bina}(\theta) = \sum_{j=1}^p \left(\sum_{k=2}^{d_j} \hat{w}_{j,k} |\theta_{j,k} - \theta_{j,k-1}| + \delta_1(\theta_{j,\bullet}) \right) \quad (5)$$

where the weights $\hat{w}_{j,k} > 0$ are defined in Section 3 below, and where

$$\delta_{\mathbf{1}}(u) = \begin{cases} 0 & \text{if } \mathbf{1}^\top u = 0, \\ \infty & \text{otherwise.} \end{cases}$$

We consider the goodness-of-fit (1) penalized by (5), namely

$$\hat{\theta} \in \operatorname{argmin}_{\theta \in \mathbb{R}^d} \{R_n(\theta) + \operatorname{bina}(\theta)\}. \quad (6)$$

An important fact is that this optimization problem is numerically cheap, as explained in the next paragraph. Figure 1 illustrates the effect of the binarsity penalization with a varying strength on an example.

In Figure 2, we illustrate on a toy example, when $p = 2$, the decision boundaries obtained for logistic regression (LR) on raw features, LR on binarized features and LR on binarized features with the binarsity penalization.

Proximal operator of binarsity. The proximal operator and proximal algorithms are important tools for non-smooth convex optimization, with important applications in the field of supervised learning with structured sparsity (Bach et al., 2012). The proximal operator of a proper lower semi-continuous (Bauschke and Combettes, 2011) convex function $g : \mathbb{R}^d \rightarrow \mathbb{R}$ is defined by

$$\operatorname{prox}_g(v) \in \operatorname{argmin}_{u \in \mathbb{R}^d} \left\{ \frac{1}{2} \|v - u\|_2^2 + g(u) \right\}.$$

Proximal operators can be interpreted as generalized projections. Namely, if g is the indicator of a convex set $C \subset \mathbb{R}^d$ given by

$$g(u) = \delta_C(u) = \begin{cases} 0 & \text{if } u \in C, \\ \infty & \text{otherwise,} \end{cases}$$

then prox_g is the projection operator onto C . It turns out that the proximal operator of binarsity can be computed very efficiently, using an algorithm (Condat, 2013) that we modify in order to include weights $\hat{w}_{j,k}$. It applies in each group the proximal operator of the total-variation since binarsity penalization is block separable, followed by a centering within each block to satisfy the sum-to-zero constraint, see Algorithm 1 below. We refer to Algorithm 2 in Appendix B for the weighted total-variation proximal operator.

Proposition 1 *Algorithm 1 computes the proximal operator of $\operatorname{bina}(\theta)$ given by (5).*

A proof of Proposition 1 is given in Appendix A. Algorithm 1 leads to a very fast numerical routine, see Section 4. The next section provides a theoretical analysis of our algorithm with an oracle inequality for the prediction error.

3. Theoretical guarantees

We now investigate the statistical properties of (6) where the weights in the binarsity penalization have the form

$$\hat{w}_{j,k} = \mathcal{O}\left(\sqrt{\frac{\log d}{n} \hat{\pi}_{j,k}}\right),$$




Fig. 1: Illustration of the binarsity penalization on the “Churn” dataset (see Section 4 for details) using logistic regression. Figure (a) shows the model weights learned by the Lasso method on the continuous raw features. Figure (b) shows the unpenalized weights on the binarized features, where the dotted green lines mark the limits between blocks corresponding to each raw features. Figures (c) and (d) show the weights with medium and strong binarsity penalization respectively. We observe in (c) that some significant cut-points start to be detected, while in (d) some raw features are completely removed from the model, the same features as those removed in (a).

with

$$\hat{\pi}_{j,k} = \frac{|\{i = 1, \dots, n : x_{i,j} \in (q_j(\frac{k}{d_j}), q_j(1)]\}|}{n}$$

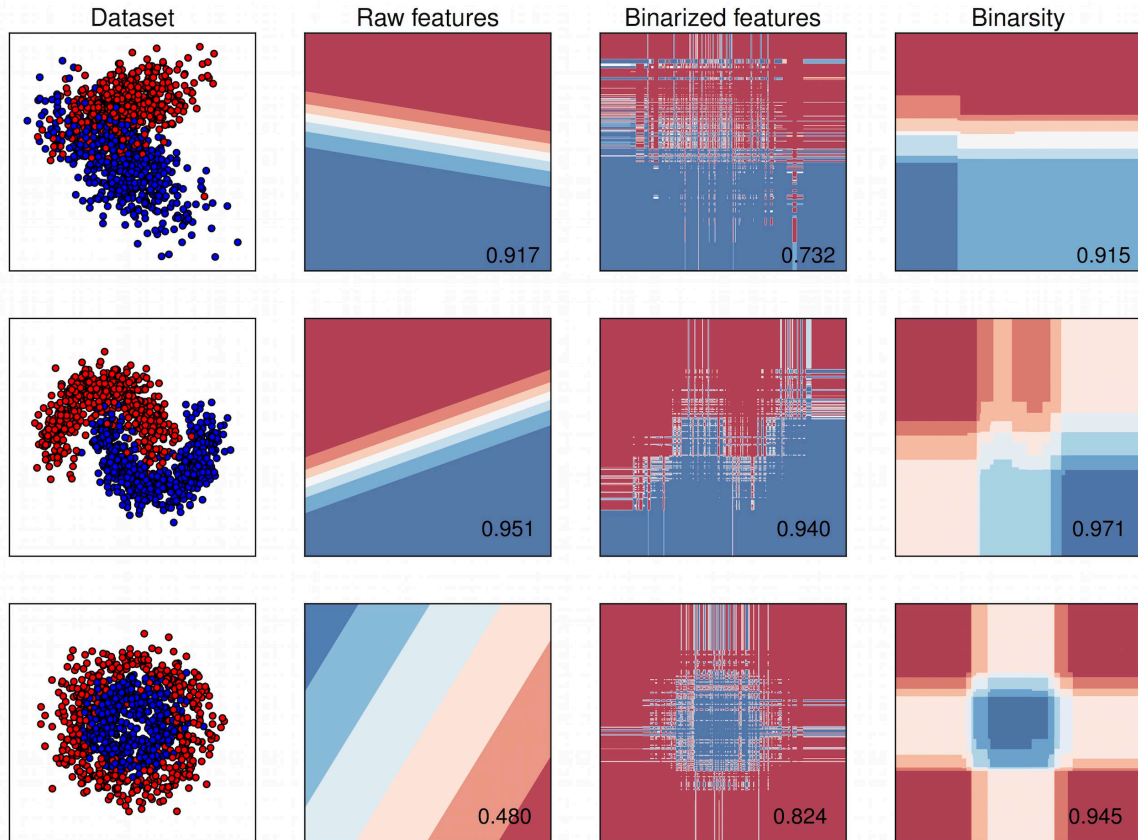


Fig. 2: Illustration of binarsity on 3 simulated toy datasets for binary classification with two classes (blue and red points). We set $n = 1000$, $p = 2$ and $d_1 = d_2 = 100$. In each row, we display the simulated dataset, followed by the decision boundaries for a logistic regression classifier trained on initial raw features, then on binarized features without regularization, and finally on binarized features with binarsity. The corresponding testing AUC score is given on the lower right corner of each figure. Our approach allows to keep an almost linear decision boundary in the first row, while non-linear decision boundaries are learned on the two other examples, without apparent overfitting.

Algorithm 1: Proximal operator of $\text{bina}(\theta)$, see (5)

Input: vector $\theta \in \mathbb{R}^d$ and weights $\hat{w}_{j,k}$ for $j = 1, \dots, p$ and $k = 1, \dots, d_j$

Output: vector $\eta = \text{prox}_{\text{bina}}(\theta)$

for $j = 1$ **to** p **do**

$\beta_{j,\bullet} \leftarrow \text{prox}_{\|\theta_{j,\bullet}\|_{\text{TV}, \hat{w}_{j,\bullet}}}(\theta_{j,\bullet})$ (TV-weighted prox in block j , see (4))
 $\eta_{j,\bullet} \leftarrow \beta_{j,\bullet} - \frac{1}{d_j} \sum_{k=1}^{d_j} \beta_{j,k}$ (within-block centering)

Return: η

for all $k \in \{2, \dots, d_j\}$, see Theorem 2 for a precise definition of $\hat{w}_{j,k}$. Note that $\hat{\pi}_{j,k}$ corresponds to the proportion of ones in the sub-matrix obtained by deleting the first k

columns in the j -th binarized block matrix $\mathbf{X}_{\bullet,j}^B$. In particular, we have $\hat{\pi}_{j,k} > 0$ for all j, k . We consider the risk measure defined by

$$R(m_\theta) = \frac{1}{n} \sum_{i=1}^n \left\{ -b'(m^0(x_i))m_\theta(x_i) + b(m_\theta(x_i)) \right\},$$

which is standard with generalized linear models (van de Geer, 2008). We aim at evaluating how “close” to the minimal possible expected risk our estimated function $m_{\hat{\theta}}$ with $\hat{\theta}$ given by (6) is. To measure this closeness, we establish a non-asymptotic oracle inequality with a fast rate of convergence considering the excess risk of $m_{\hat{\theta}}$, namely $R(m_{\hat{\theta}}) - R(m^0)$. To derive this inequality, we need to impose a restricted eigenvalue assumption on \mathbf{X}^B .

For all $\theta \in \mathbb{R}^d$, let $J(\theta) = [J_1(\theta), \dots, J_p(\theta)]$ be the concatenation of the support sets relative to the total-variation penalization, that is

$$J_j(\theta) = \{k : \theta_{j,k} \neq \theta_{j,k-1}, \text{ for } k = 2, \dots, d_j\}.$$

Similarly, we denote $J^c(\theta) = [J_1^c(\theta), \dots, J_p^c(\theta)]$ the complementary of $J(\theta)$. The restricted eigenvalue condition is defined as follow.

Assumption 2 *Let $K = [K_1, \dots, K_p]$ be a concatenation of index sets. We consider*

$$\kappa(K) \in \inf_{u \in \mathcal{C}_{\text{TV}, \hat{w}}(K) \setminus \{\mathbf{0}_d\}} \left\{ \frac{\|\mathbf{X}^B u\|_2}{\sqrt{n} \|u_K\|_2} \right\}$$

with

$$\mathcal{C}_{\text{TV}, \hat{w}}(K) = \left\{ u \in \mathbb{R}^d : \sum_{j=1}^p \|(u_{j,\bullet})_{K_j^c}\|_{\text{TV}, \hat{w}_{j,\bullet}} \leq 2 \sum_{j=1}^p \|(u_{j,\bullet})_{K_j}\|_{\text{TV}, \hat{w}_{j,\bullet}} \right\}. \quad (7)$$

We suppose that the following condition holds

$$\kappa(K) > 0. \quad (8)$$

The set $\mathcal{C}_{\text{TV}, \hat{w}}(K)$ is a cone composed by all vectors with similar support K . Let us now work locally on

$$B_d(\rho) = \{\theta \in \mathbb{R}^d : \|\theta\|_2 \leq \rho\},$$

the ℓ_2 -ball of radius $\rho > 0$ in \mathbb{R}^d . This restriction has already been considered in the case of high-dimensional generalized linear models (van de Geer, 2008). It allows us to establish a connection, via the notion of self-concordance (Bach, 2010), between the empirical squared ℓ_2 -norm and the empirical Kullback divergence (see Lemma 9 in Appendix C). Theorem 2 gives a risk bound for the estimator $m_{\hat{\theta}}$.

Theorem 2 *Let Assumptions 1 and 2 be satisfied. Fix $A > 0$ and choose*

$$\hat{w}_{j,k} = \sqrt{\frac{2U_n \phi(A + \log d)}{n}} \hat{\pi}_{j,k}. \quad (9)$$

Let $C_n(\rho, p) = 2(C_n + \rho\sqrt{p})$, $\psi(u) = e^u - u - 1$, and consider the following constants

$$C_n(\rho, p, L_n) = \frac{L_n\psi(-C_n(\rho, p))}{C_n^2(\rho, p)}, \quad \epsilon > \frac{2}{C_n(\rho, p, L_n)} \quad \text{and} \quad \zeta = \frac{4}{\epsilon C_n(\rho, p, L_n) - 2}.$$

Then, with probability at least $1 - 2e^{-A}$, any solution $\hat{\theta}$ of problem (6) restricted on $B_d(\rho)$ fulfills the following risk bound

$$R(m_{\hat{\theta}}) - R(m^0) \leq (1 + \zeta) \inf_{\theta \in B_d(\rho)} \left\{ R(m_\theta) - R(m^0) + \frac{\xi |J(\theta)|}{\kappa^2(J(\theta))} \max_{j=1, \dots, p} \|(\hat{w}_{j, \bullet})_{J_j(\theta)}\|_\infty^2 \right\}, \quad (10)$$

where

$$\xi = \frac{512 \epsilon^2 C_n(\rho, p, L_n)}{\epsilon C_n(\rho, p, L_n) - 2}.$$

A proof of Theorem 2 is given in Appendix C. Note that $\hat{w}_{j,k} > 0$, since by construction $\hat{\pi}_{j,k} > 0$ for all j, k . The second term in the right-hand side of (10) can be viewed as a variance term, and its dominant term satisfies

$$\frac{|J(\theta)|}{\kappa^2(J(\theta))} \max_{j=1, \dots, p} \|(\hat{w}_{j, \bullet})_{J_j(\theta)}\|_\infty^2 \leq \frac{\tilde{A} U_n \phi}{\kappa^2(J(\theta))} \frac{|J(\theta)| \log d}{n}, \quad (11)$$

for some positive constant \tilde{A} . The complexity term in (11) depends on both the sparsity and the restricted eigenvalues of the binarized matrix. The value $|J(\theta)|$ characterizes the sparsity of the vector θ , that is the smaller $|J(\theta)|$, the sparser θ . The rate of convergence of the estimator $m_{\hat{\theta}}$ has the expected shape $\log d/n$. Moreover, for the case of least squares regression, the oracle inequality in Theorem 2 is sharp, in the sense that $\zeta = 0$ (see Remark 10 in Appendix C).

4. Numerical experiments

In this section, we first illustrate the fact that the binarsity penalization is roughly only two times slower than basic ℓ_1 -penalization, see the timings in Figure 3. We then compare binarsity to a large number of baselines, see Table 2, using 9 classical binary classification datasets obtained from the UCI Machine Learning Repository (Lichman, 2013), see Table 3.

For each method, we randomly split all datasets into a training and a test set (30% for testing), and all hyper-parameters are tuned on the training set using V -fold cross-validation with $V = 10$. For support vector machine with radial basis kernel (SVM), random forests (RF) and gradient boosting (GB), we use the reference implementations from the `scikit-learn` library (Pedregosa et al., 2011), and we use the `LogisticGAM` procedure from the `pygam` library¹ for the GAM baseline. The binarsity penalization is proposed in the `tick` library (Bacry et al., 2017), we provide sample code for its use in Figure 4. Logistic regression with no penalization or ridge penalization gave similar or lower scores for all considered datasets, and are therefore not reported in our experiments.

1. <https://github.com/dswah/pyGAM>




Fig. 3: Average computing time in second (with the black lines representing \pm the standard deviation) obtained on 100 simulated datasets for training a logistic model with binarsity VS Lasso penalization, both trained on \mathbf{X}^B with $d_j = 10$ for all $j \in 1, \dots, p$. Features are Gaussian with a Toeplitz covariance matrix with correlation 0.5 and $n = 10000$. Note that the computing time ratio between the two methods stays roughly constant and equal to 2.

Name	Description	Reference
Lasso	Logistic regression (LR) with ℓ_1 penalization	Tibshirani (1996b)
Group L1	LR with group ℓ_1 penalization	Meier et al. (2008)
Group TV	LR with group total-variation penalization	
SVM	Support vector machine with radial basis kernel	Schölkopf and Smola (2002)
GAM	Generalized additive model	Hastie and Tibshirani (1990)
RF	Random forest classifier	Breiman (2001)
GB	Gradient boosting	Friedman (2002)

Tab. 2: Baselines considered in our experiments. Note that Group L1 and Group TV are considered on binarized features.

```

# input: features X, labels y
from tick.inference import LogisticRegression
from tick.preprocessing import FeaturesBinarizer
from sklearn.model_selection import train_test_split

# binarize data
binarizer = FeaturesBinarizer(n_cuts=50)
X = binarizer.fit_transform(X)

# shuffle and split training and test sets
X, X_test, y, y_test = train_test_split(X, y, stratify=y)

# fit the model
learner = LogisticRegression(penalty='binarsity', C=1,
                             blocks_start=binarizer.blocks_start,
                             blocks_length=binarizer.blocks_length)

learner.fit(X, y)

# predict on test set
y_pred = learner.predict_proba(X_test)[:, 1]
    
```

Fig. 4: Sample python code for the use of binarsity with logistic regression in the tick library, with the use of the FeaturesBinarizer transformer for features binarization.

Dataset	#Samples	#Features	Reference
Ionosphere	351	34	Sigillito et al. (1989)
Churn	3333	21	Lichman (2013)
Default of credit card	30000	24	Yeh and Lien (2009)
Adult	32561	14	Kohavi (1996)
Bank marketing	45211	17	Moro et al. (2014)
Covertypes	550088	10	Blackard and Dean (1999)
SUSY	5000000	18	Baldi et al. (2014)
HEPMASS	10500000	28	Baldi et al. (2016)
HIGGS	11000000	24	Baldi et al. (2014)

Tab. 3: Basic informations about the 9 considered datasets.

The binarsity penalization does not require a careful tuning of d_j (number of bins for the one-hot encoding of raw feature j). Indeed, past a large enough value, increasing d_j even further barely changes the results since the cut-points selected by the penalization do not change anymore. This is illustrated in Figure 5, where we observe that past 50 bins, increasing d_j even further does not affect the performance, and only leads to an increase of the training time. In all our experiments, we therefore fix $d_j = 50$ for $j = 1, \dots, p$.




Fig. 5: Impact of the number of bins used in each block (d_j) on the classification performance (measured by AUC) and on the training time using the “Adult” and “Default of credit card” datasets. All d_j are equal for $j = 1, \dots, p$, and we consider in all cases the best hyper-parameters selected after cross validation. We observe that past $d_j = 50$ bins, performance is roughly constant, while training time strongly increases.

The results of all our experiments are reported in Figures 6 and 7. In Figure 6 we compare the performance of binarsity with the baselines on all 9 datasets, using ROC curves and the Area Under the Curve (AUC), while we report computing (training) timings in Figure 7. We observe that binarsity consistently outperforms Lasso, as well as Group L1: this highlights the importance of the TV norm within each group. The AUC of Group TV is always slightly below the one of binarsity, and more importantly it involves a much larger training time: convergence is slower for Group TV, since it does not use the linear constraint of binarsity, leading to a ill-conditioned problem (sum of binary features equals 1 in each block). Finally, binarsity outperforms also GAM and its performance is comparable in all considered examples to RF and GB, with computational timings that are orders of magnitude faster, see Figure 7. All these experiments illustrate that binarsity achieves an

extremely competitive compromise between computational time and performance, compared to all considered baselines.

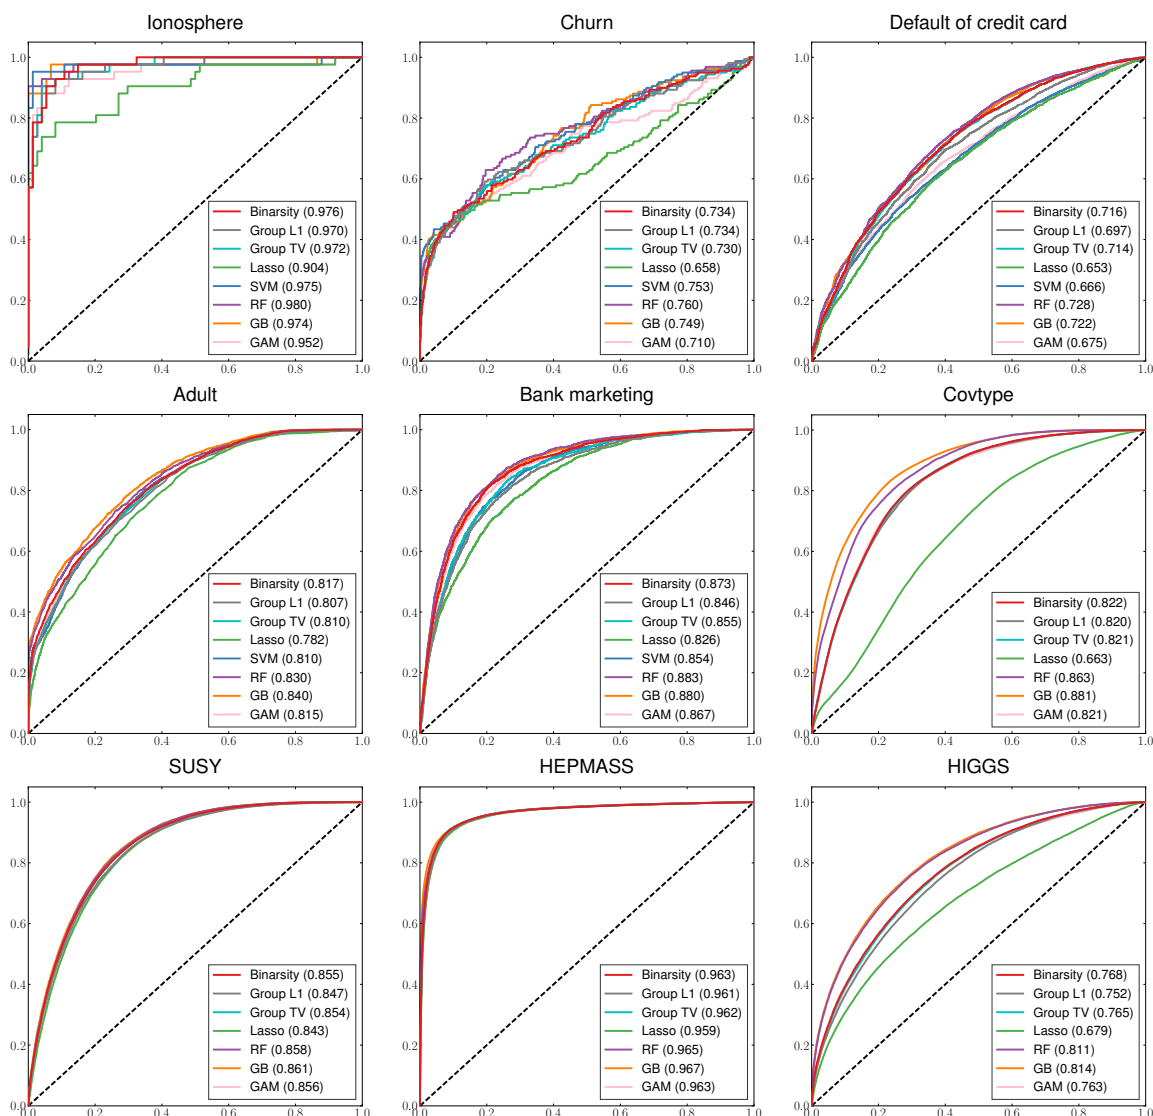


Fig. 6: Performance comparison using ROC curves and AUC scores (given between parenthesis) computed on test sets. The 4 last datasets contain too many examples for SVM (RBF kernel). Binarsity consistently does a better job than Lasso, Group L1, Group TV and GAM. Its performance is comparable to SVM, RF and GB but with computational timings that are orders of magnitude faster, see Figure 7.




Fig. 7: Computing time comparisons (in seconds) between the methods on the considered datasets. Note that the time values are log-scaled. These timings concern the learning task for each model with the best hyper parameters selected, after the cross validation procedure. The 4 last datasets contain too many examples for the SVM with RBF kernel to be trained in a reasonable time. Roughly, binarsity is between 2 and 5 times slower than ℓ_1 penalization on the considered datasets, but is more than 100 times faster than random forests or gradient boosting algorithms on large datasets, such as HIGGS.

5. Conclusion

In this paper, we introduced the binarsity penalization for one-hot encodings of continuous features. We illustrated the good statistical properties of binarsity for generalized linear models by proving non-asymptotic oracle inequalities. We conducted extensive comparisons of binarsity with state-of-the-art algorithms for binary classification on several standard datasets. Experimental results illustrate that binarsity significantly outperforms Lasso, Group L1 and Group TV penalizations and also generalized additive models, while being competitive with random forests and boosting. Moreover, it can be trained orders of magnitude faster than boosting and other ensemble methods. Even more importantly, it provides interpretability. Indeed, in addition to the raw feature selection ability of binarsity, the method pinpoints significant cut-points for all continuous feature. This leads to a much more precise and deeper understanding of the model than the one provided by Lasso on raw features. These results illustrate the fact that binarsity achieves an extremely competitive compromise between computational time and performance, compared to all considered baselines.

Appendix A. Proof of Proposition 1: proximal operator of binarsity

For any fixed $j = 1, \dots, p$, we aim to prove that $\text{prox}_{\|\cdot\|_{\text{TV}, \hat{w}_{j,\bullet}} + \delta_1}$ is the composite proximal operators of $\text{prox}_{\|\cdot\|_{\text{TV}, \hat{w}_{j,\bullet}}}$ and prox_{δ_1} , namely

$$\text{prox}_{\|\cdot\|_{\text{TV}, \hat{w}_{j,\bullet}} + \delta_1}(\theta_{j,\bullet}) = \text{prox}_{\delta_1}(\text{prox}_{\|\cdot\|_{\text{TV}, \hat{w}_{j,\bullet}}}(\theta_{j,\bullet}))$$

for all $\theta_{j,\bullet} \in \mathbb{R}^{d_j}$. Using Theorem 1 in Yu (2013), it is sufficient to show that for all $\theta_{j,\bullet} \in \mathbb{R}^{d_j}$, we have

$$\partial(\|\theta_{j,\bullet}\|_{\text{TV}, \hat{w}_{j,\bullet}}) \subseteq \partial(\|\text{prox}_{\delta_1}(\theta_{j,\bullet})\|_{\text{TV}, \hat{w}_{j,\bullet}}). \quad (12)$$

Clearly, by the definition of the proximal operator, we have $\text{prox}_{\delta_1}(\theta_{j,\bullet}) = \Pi_{\text{span}\{\mathbf{1}_{d_j}\}^\perp}(\theta_{j,\bullet})$, where $\Pi_{\text{span}\{\mathbf{1}_{d_j}\}^\perp}(\cdot)$ stands for the projection onto the hyperplane $\text{span}\{\mathbf{1}_{d_j}\}^\perp$. Besides, we know that

$$\begin{aligned} \Pi_{\text{span}\{\mathbf{1}_{d_j}\}^\perp}(\theta_{j,\bullet}) &= \theta_{j,\bullet} - \Pi_{\{\mathbf{1}_{d_j}\}}(\theta_{j,\bullet}) \\ &= \theta_{j,\bullet} - \frac{\langle \theta_{j,\bullet}, \mathbf{1}_{d_j} \rangle}{\|\mathbf{1}_{d_j}\|_2^2} \mathbf{1}_{d_j} \\ &= \theta_{j,\bullet} - \bar{\theta}_{j,\bullet} \mathbf{1}_{d_j}, \end{aligned}$$

where $\bar{\theta}_{j,\bullet} = \frac{1}{d_j} \sum_{k=1}^{d_j} \theta_{j,k}$. Now, let us define the $d_j \times d_j$ matrix D_j by

$$D_j = \begin{bmatrix} 1 & 0 & & 0 \\ -1 & 1 & & \\ & \ddots & \ddots & \\ 0 & & -1 & 1 \end{bmatrix} \in \mathbb{R}^{d_j} \times \mathbb{R}^{d_j}. \quad (13)$$

We then remark that for all $\theta_{j,\bullet} \in \mathbb{R}^{d_j}$,

$$\|\theta_{j,\bullet}\|_{\text{TV}, \hat{w}_{j,\bullet}} = \sum_{k=2}^{d_j} \hat{w}_{j,k} |\theta_{j,k} - \theta_{j,k-1}| = \|\hat{w}_{j,\bullet} \odot D_j \theta_{j,\bullet}\|_1. \quad (14)$$

Using subdifferential calculus (see details in the proof of Proposition 4 below), one has

$$\partial(\|\theta_{j,\bullet}\|_{\text{TV}, \hat{w}_{j,\bullet}}) = \partial(\|\hat{w}_{j,\bullet} \odot D_j \theta_{j,\bullet}\|_1) = D_j^\top \hat{w}_{j,\bullet} \odot \text{sign}(D_j \theta_{j,\bullet}).$$

Then, the linear constraint $\sum_{k=1}^{d_j} \theta_{j,k} = 0$ entails that

$$D_j^\top \hat{w}_{j,\bullet} \odot \text{sign}(D_j \theta_{j,\bullet}) = D_j^\top \hat{w}_{j,\bullet} \odot \text{sign}(D_j(\theta_{j,\bullet} - \bar{\theta}_{j,\bullet} \mathbf{1}_{d_j})),$$

which leads to (12). Hence, setting $\beta_{j,\bullet} = \text{prox}_{\|\cdot\|_{\text{TV}, \hat{w}_{j,\bullet}}}(\theta_{j,\bullet})$ and $\bar{\beta}_{j,\bullet} = \frac{1}{d_j} \sum_{k=1}^{d_j} \beta_{j,k}$ we get

$$\text{prox}_{\|\cdot\|_{\text{TV}, \hat{w}_{j,\bullet}} + \delta_1}(\theta_{j,\bullet}) = \beta_{j,\bullet} - \bar{\beta}_{j,\bullet} \mathbf{1}_{d_j}$$

which gives Algorithm 1. ■

Appendix B. Algorithm of computing proximal operator of weighted TV penalization

We recall here the algorithm given in Alaya et al. (2015) for computing the proximal operator of weighted total-variation penalization. The latter is defined as follows

$$\beta = \text{prox}_{\|\cdot\|_{\text{TV}, \hat{w}}}(\theta) \in \operatorname{argmin}_{\theta \in \mathbb{R}^m} \left\{ \frac{1}{2} \|\beta - \theta\|_2^2 + \|\theta\|_{\text{TV}, \hat{w}} \right\}. \quad (15)$$

The proposed algorithm consists in running forwardly through the samples $(\theta_1, \dots, \theta_m)$. Using the Karush-Kuhn-Tucker (KKT) optimality conditions for a convex optimization (Boyd and Vandenberghe, 2004), at location k , β_k stays constant where $|u_k| < \hat{w}_{k+1}$. Here u_k is a solution to a dual problem associated to the primal problem (15). If this is not possible, it goes back to the last location where a jump can be introduced in β , validates the current segment until this location, starts a new segment, and continues. This algorithm is described precisely in Algorithm 2.

Appendix C. Proof of Theorem 2: fast oracle inequality under binarsity

The proof relies on some technical properties given below.

Additional notation. Hereafter, we use the following vector notations: $\mathbf{y} = (y_1, \dots, y_n)^\top$, $m^0(\mathbf{X}) = (m^0(x_1), \dots, m^0(x_n))^\top$, $m_\theta(\mathbf{X}) = (m_\theta(x_1), \dots, m_\theta(x_n))^\top$ (recall that $m_\theta(x_i) = \theta^\top x_i^B$), and $b'(m_\theta(\mathbf{X})) = (b'(m_\theta(x_1)), \dots, b'(m_\theta(x_n)))^\top$.

C.1 Empirical Kullback-Leibler divergence.

Let us now define the Kullback-Leibler divergence between the true probability density function f^0 defined in (2) and a candidate f_θ within the generalized linear model ($f_\theta(y|x) = \exp(y m_\theta(x) - b(m_\theta(x)))$) as follows

$$\begin{aligned} KL_n(f^0, f_\theta) &= \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{\mathbb{P}_{\mathbf{y}|X}} \left[\log \frac{f^0(y_i|x_i)}{f_\theta(y_i|x_i)} \right] \\ &:= KL_n(m^0(\mathbf{X}), m_\theta(\mathbf{X})), \end{aligned}$$

where $\mathbb{P}_{\mathbf{y}|X}$ is the joint distribution of $\mathbf{y} = (y_1, \dots, y_n)^\top$ given $\mathbf{X} = (x_1, \dots, x_n)^\top$. We then have the following property.

Lemma 3 *The excess risk verifies $R(m_\theta) - R(m^0) = \phi KL_n(m^0(\mathbf{X}), m_\theta(\mathbf{X}))$.*

Proof. Straightforwardly, one has

$$\begin{aligned} &KL_n(m^0(\mathbf{X}), m_\theta(\mathbf{X})) \\ &= \phi^{-1} \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{\mathbb{P}_{\mathbf{y}|X}} \left[\left(-y_i m_\theta(x_i) + b(m_\theta(x_i)) \right) - \left(-y_i m^0(x_i) + b(m^0(x_i)) \right) \right] \\ &= \phi^{-1} (R(m_\theta) - R(m^0)). \end{aligned}$$

■

Algorithm 2: Proximal operator of weighted TV penalization

Input: vector $\theta = (\theta_1, \dots, \theta_m)^\top \in \mathbb{R}^m$ and weights $\hat{w} = (\hat{w}_1, \dots, \hat{w}_m) \in \mathbb{R}_+^m$.

Output: vector $\beta = \text{prox}_{\|\cdot\|_{\text{TV}, \hat{w}}}(\theta)$

1. **Set** $k = k_0 = k_- = k_+ \leftarrow 1$
 $\beta_{\min} \leftarrow \theta_1 - \hat{w}_2$; $\beta_{\max} \leftarrow \theta_1 + \hat{w}_2$
 $u_{\min} \leftarrow \hat{w}_2$; $u_{\max} \leftarrow -\hat{w}_2$
 2. **if** $k = m$ **then**
 $\beta_m \leftarrow \beta_{\min} + u_{\min}$
 3. **if** $\theta_{k+1} + u_{\min} < \beta_{\min} - \hat{w}_{k+2}$ **then** /* negative jump */
 $\beta_{k_0} = \dots = \beta_{k_-} \leftarrow \beta_{\min}$
 $k = k_0 = k_- = k_+ \leftarrow k_- + 1$
 $\beta_{\min} \leftarrow \theta_k - \hat{w}_{k+1} + \hat{w}_k$; $\beta_{\max} \leftarrow \theta_k + \hat{w}_{k+1} + \hat{w}_k$
 $u_{\min} \leftarrow \hat{w}_{k+1}$; $u_{\max} \leftarrow -\hat{w}_{k+1}$
 4. **else if** $\theta_{k+1} + u_{\max} > \beta_{\max} + \hat{w}_{k+2}$ **then** /* positive jump */
 $\beta_{k_0} = \dots = \beta_{k_+} \leftarrow \beta_{\max}$
 $k = k_0 = k_- = k_+ \leftarrow k_+ + 1$
 $\beta_{\min} \leftarrow \theta_k - \hat{w}_{k+1} - \hat{w}_k$; $\beta_{\max} \leftarrow \theta_k + \hat{w}_{k+1} - \hat{w}_k$
 $u_{\min} \leftarrow \hat{w}_{k+1}$; $u_{\max} \leftarrow -\hat{w}_{k+1}$
 5. **else** /* no jump */
 $\text{set } k \leftarrow k + 1$
 $u_{\min} \leftarrow \theta_k + \hat{w}_{k+1} - \beta_{\min}$
 $u_{\max} \leftarrow \theta_k - \hat{w}_{k+1} - \beta_{\max}$ **if** $u_{\min} \geq \hat{w}_{k+1}$ **then**
 $\beta_{\min} \leftarrow \beta_{\min} + \frac{u_{\min} - \hat{w}_{k+1}}{k - k_0 + 1}$
 $u_{\min} \leftarrow \hat{w}_{k+1}$
 $k_- \leftarrow k$
if $u_{\max} \leq -\hat{w}_{k+1}$ **then**
 $\beta_{\max} \leftarrow \beta_{\max} + \frac{u_{\max} + \hat{w}_{k+1}}{k - k_0 + 1}$
 $u_{\max} \leftarrow -\hat{w}_{k+1}$
 $k_+ \leftarrow k$
 6. **if** $k < m$ **then**
 go to 3.
 7. **if** $u_{\min} < 0$ **then**
 $\beta_{k_0} = \dots = \beta_{k_-} \leftarrow \beta_{\min}$
 $k = k_0 = k_- \leftarrow k_- + 1$
 $\beta_{\min} \leftarrow \theta_k - \hat{w}_{k+1} + \hat{w}_k$
 $u_{\min} \leftarrow \hat{w}_{k+1}$; $u_{\max} \leftarrow \theta_k + \hat{w}_k - u_{\max}$
 go to 2.
 8. **else if** $u_{\max} > 0$ **then**
 $\beta_{k_0} = \dots = \beta_{k_+} \leftarrow \beta_{\max}$
 $k = k_0 = k_+ \leftarrow k_+ + 1$
 $\beta_{\max} \leftarrow \theta_k + \hat{w}_{k+1} - \hat{w}_k$
 $u_{\max} \leftarrow -\hat{w}_{k+1}$; $u_{\min} \leftarrow \theta_k - \hat{w}_k - u_{\min}$
 go to 2.
 9. **else**
 $\beta_{k_0} = \dots = \beta_m \leftarrow \beta_{\min} + \frac{u_{\min}}{k - k_0 + 1}$
-

C.2 Optimality conditions.

To characterize the solution of the problem (6), the following result can be straightforwardly obtained using the Karush-Kuhn-Tucker (KKT) optimality conditions for a convex optimization (Boyd and Vandenberghe, 2004).

Proposition 4 *A vector $\hat{\theta} = (\hat{\theta}_{1,\bullet}^\top, \dots, \hat{\theta}_{p,\bullet}^\top)^\top \in \mathbb{R}^d$ is an optimum of the objective function in (6) if and only if there exists a sequence of subgradients $\hat{h} = (\hat{h}_{j,\bullet})_{j=1,\dots,p} \in \partial(\|\hat{\theta}\|_{\text{TV},\hat{w}})$ and $\hat{g} = (\hat{g}_{j,\bullet})_{j=1,\dots,p} \in \partial(\delta_1(\hat{\theta}_{j,\bullet}))_{j=1,\dots,p}$ such that*

$$\nabla R_n(\hat{\theta}_{j,\bullet}) + \hat{h}_{j,\bullet} + \hat{g}_{j,\bullet} = \mathbf{0}_{d_j},$$

where

$$\begin{cases} \hat{h}_{j,\bullet} = D_j^\top(\hat{w}_{j,\bullet} \odot \text{sign}(D_j \hat{\theta}_{j,\bullet})) & \text{if } j \in J(\hat{\theta}), \\ \hat{h}_{j,\bullet} \in D_j^\top(\hat{w}_{j,\bullet} \odot [-1, +1]^{d_j}) & \text{if } j \in J^c(\hat{\theta}), \end{cases} \quad (16)$$

and where $J(\hat{\theta})$ is the active set of $\hat{\theta}$. The subgradient $\hat{g}_{j,\bullet}$ belongs to

$$\partial(\delta_1(\hat{\theta}_{j,\bullet})) = \{\mu_{j,\bullet} \in \mathbb{R}^{d_j} : \langle \mu_{j,\bullet}, \theta_{j,\bullet} \rangle \leq \langle \mu_{j,\bullet}, \hat{\theta}_{j,\bullet} \rangle \text{ for all } \theta_{j,\bullet} \text{ such that } \mathbf{1}_{d_j}^\top \theta_{j,\bullet} = 0\}.$$

For the generalized linear model, we have

$$\frac{1}{n}(\mathbf{X}_{\bullet,j}^B)^\top (b'(m_{\hat{\theta}}(\mathbf{X})) - \mathbf{y}) + \hat{h}_{j,\bullet} + \hat{g}_{j,\bullet} + \hat{f}_{j,\bullet} = \mathbf{0}_{d_j}, \quad (17)$$

where $\hat{f} = (\hat{f}_{j,\bullet})_{j=1,\dots,p}$ belongs to the normal cone of the ball $B_d(\rho)$.

Proof. We denote by $\partial(\phi)$ the subdifferential mapping of a convex functional ϕ . The function $\theta \mapsto R_n(\theta)$ is differentiable, so the subdifferential of $R_n(\cdot) + \text{bina}(\cdot)$ at a point $\theta = (\theta_{j,\bullet})_{j=1,\dots,p} \in \mathbb{R}^d$ is given by

$$\partial(R_n(\theta) + \text{bina}(\theta)) = \nabla R_n(\theta) + \partial(\text{bina}(\theta)),$$

where $\nabla R_n(\theta) = \left(\frac{\partial(R_n(\theta))}{\partial(\theta_{1,\bullet})}, \dots, \frac{\partial(R_n(\theta))}{\partial(\theta_{p,\bullet})} \right)^\top$ and

$$\partial(\text{bina}(\theta)) = \left(\partial(\|\theta_{1,\bullet}\|_{\text{TV},\hat{w}_{1,\bullet}}) + \partial(\delta_1(\theta_{1,\bullet})), \dots, \partial(\|\theta_{p,\bullet}\|_{\text{TV},\hat{w}_{p,\bullet}}) + \partial(\delta_1(\theta_{p,\bullet})) \right)^\top.$$

We have $\|\theta_{j,\bullet}\|_{\text{TV},\hat{w}_{j,\bullet}} = \|\hat{w}_{j,\bullet} \odot D_j \theta_{j,\bullet}\|_1$ for all $j = 1, \dots, p$. Then, by applying some properties of the subdifferential calculus, we get

$$\partial(\|\theta_{j,\bullet}\|_{\text{TV},\hat{w}_{j,\bullet}}) = \begin{cases} D_j^\top \text{sign}(\hat{w}_{j,\bullet} \odot D_j \theta_{j,\bullet}) & \text{if } D_j \theta_{j,\bullet} \neq \mathbf{0}_{d_j}, \\ D_j^\top(\hat{w}_{j,\bullet} \odot v_j) & \text{otherwise,} \end{cases} \quad (18)$$

where $v_j \in [-1, +1]^{d_j}$, for all $j = 1, \dots, p$. For generalized linear models, we rewrite

$$\hat{\theta} \in \text{argmin}_{\theta \in \mathbb{R}^d} \{R_n(\theta) + \text{bina}(\theta) + \delta_{B_d(\rho)}(\theta)\}, \quad (19)$$

where $\delta_{B_d(\rho)}$ is the indicator function for $B_d(\rho)$. Now, $\hat{\theta} = (\hat{\theta}_{1,\bullet}^\top, \dots, \hat{\theta}_{p,\bullet}^\top)^\top$ is an optimum of Problem (19) if and only if $\mathbf{0}_d \in \nabla R_n(m_{\hat{\theta}}) + \partial(\|\hat{\theta}\|_{\text{TV}, \hat{w}}) + \partial(\delta_{B_d(\rho)}(\hat{\theta}))$. Recall that the subdifferential of $\delta_{B_d(\rho)}(\cdot)$ is the normal cone of $B_d(\rho)$, that is

$$\partial(\delta_{B_d(\rho)}(\hat{\theta})) = \{\eta \in \mathbb{R}^d : \langle \eta, \theta \rangle \leq \langle \eta, \hat{\theta} \rangle \text{ for all } \theta \in B_d(\rho)\}. \quad (20)$$

Straightforwardly, one obtains

$$\frac{\partial(R_n(\theta))}{\partial(\theta_{j,\bullet})} = \frac{1}{n} (\mathbf{X}_{\bullet,j}^B)^\top (b'(m_{\hat{\theta}}(\mathbf{X})) - \mathbf{y}), \quad (21)$$

and equalities (21) and (20) give equation (17), which ends the proof of Proposition 4. \blacksquare

C.3 Compatibility conditions.

Let us define the block diagonal matrix $\mathbf{D} = \text{diag}(D_1, \dots, D_p)$, with D_j , defined in (13), being invertible. We denote its inverse T_j which is defined by the $d_j \times d_j$ lower triangular matrix with entries $(T_j)_{r,s} = 0$ if $r < s$ and $(T_j)_{r,s} = 1$ otherwise. We set $\mathbf{T} = \text{diag}(T_1, \dots, T_p)$. It is clear that $\mathbf{D}^{-1} = \mathbf{T}$. In order to prove Theorem 2, we need, in addition to Assumption 2, the following results which give a compatibility condition (van de Geer, 2008; van de Geer and Lederer, 2013; Dalalyan et al., 2017) satisfied by the matrix \mathbf{T} in Lemma 5 and $\mathbf{X}^B \mathbf{T}$ in Lemma 6. To this end, for any concatenation of subsets $K = [K_1, \dots, K_p]$, we set

$$K_j = \{\tau_j^1, \dots, \tau_j^{b_j}\} \subset \{1, \dots, d_j\} \quad (22)$$

for all $j = 1, \dots, p$ and with the convention that $\tau_j^0 = 0$ and $\tau_j^{b_j+1} = d_j + 1$.

Lemma 5 *Let $\gamma \in \mathbb{R}_+^d$ be a given vector of weights and $K = [K_1, \dots, K_p]$ with K_j given by (22) for all $j = 1, \dots, p$. Then for every $u \in \mathbb{R}^d \setminus \{\mathbf{0}_d\}$, we have*

$$\frac{\|\mathbf{T}u\|_2}{\left| \|u_K \odot \gamma_K\|_1 - \|u_{K^c} \odot \gamma_{K^c}\|_1 \right|} \geq \kappa_{\mathbf{T}, \gamma}(K),$$

where

$$\kappa_{\mathbf{T}, \gamma}(K) = \left\{ 32 \sum_{j=1}^p \sum_{k=1}^{d_j} |\gamma_{j,k+1} - \gamma_{j,k}|^2 + 2|K_j| \|\gamma_{j,\bullet}\|_\infty^2 \Delta_{\min, K_j}^{-1} \right\}^{-1/2},$$

and $\Delta_{\min, K_j} = \min_{r=1, \dots, b_j} |\tau_j^{r_j} - \tau_j^{r_j-1}|$.

Proof. Using Proposition 3 in Dalalyan et al. (2017), we have

$$\begin{aligned} & \|u_K \odot \gamma_K\|_1 - \|u_{K^c} \odot \gamma_{K^c}\|_1 \\ &= \sum_{j=1}^p \|u_{K_j} \odot \gamma_{K_j}\|_1 - \|u_{K_j^c} \odot \gamma_{K_j^c}\|_1 \\ &\leq \sum_{j=1}^p 4 \|T_j u_{j,\bullet}\|_2 \left\{ 2 \sum_{k=1}^{d_j} |\gamma_{j,k+1} - \gamma_{j,k}|^2 + 2(b_j + 1) \|\gamma_{j,\bullet}\|_\infty^2 \Delta_{\min, K_j}^{-1} \right\}^{1/2}. \end{aligned}$$

Applying Hölder's inequality for the right hand side of the last inequality gives

$$\begin{aligned} & \|u_K \odot \gamma_K\|_1 - \|u_{K^c} \odot \gamma_{K^c}\|_1 \\ & \leq \|\mathbf{T}u\|_2 \left\{ 32 \sum_{j=1}^p \sum_{k=1}^{d_j} |\gamma_{j,k+1} - \gamma_{j,k}|^2 + 2|K_j| \|\gamma_{j,\bullet}\|_\infty^2 \Delta_{\min, K_j}^{-1} \right\}^{1/2}. \end{aligned}$$

This completes the proof of Lemma 5. \blacksquare

Now, using Assumption 2 and Lemma 5, we establish a compatibility condition satisfied by the product of matrices $\mathbf{X}^B \mathbf{T}$.

Lemma 6 *Let Assumption 2 holds. Let $\gamma \in \mathbb{R}_+^d$ be a given vector of weights, and $K = [K_1, \dots, K_p]$ such that K_j is given by (22) for all $j = 1, \dots, p$. Then, one has*

$$\inf_{u \in \mathcal{C}_{1, \hat{w}}(K) \setminus \{0_d\}} \left\{ \frac{\|\mathbf{X}^B \mathbf{T}u\|_2}{\sqrt{n} \|u_K \odot \gamma_K\|_1 - \|u_{K^c} \odot \gamma_{K^c}\|_1} \right\} \geq \kappa_{\mathbf{T}, \gamma}(K) \kappa(K), \quad (23)$$

where

$$\mathcal{C}_{1, \hat{w}}(K) = \left\{ u \in \mathbb{R}^d : \sum_{j=1}^p \|(u_{j,\bullet})_{K_j^c}\|_{1, \hat{w}_{j,\bullet}} \leq 2 \sum_{j=1}^p \|(u_{j,\bullet})_{K_j}\|_{1, \hat{w}_{j,\bullet}} \right\}, \quad (24)$$

with $\|\cdot\|_{1,a}$ denoting the weighted ℓ_1 -norm.

Proof. By Lemma 5, we have that

$$\frac{\|\mathbf{X}^B \mathbf{T}u\|_2}{\sqrt{n} \|u_K \odot \gamma_K\|_1 - \|u_{K^c} \odot \gamma_{K^c}\|_1} \geq \kappa_{\mathbf{T}, \gamma}(K) \frac{\|\mathbf{X}^B \mathbf{T}u\|_2}{\sqrt{n} \|\mathbf{T}u\|_2}.$$

Now, we note that if $u \in \mathcal{C}_{1, \hat{w}}(K)$, then $\mathbf{T}u \in \mathcal{C}_{\text{TV}, \hat{w}}(K)$. Hence, by Assumption 2, we get

$$\frac{\|\mathbf{X}^B \mathbf{T}u\|_2}{\sqrt{n} \|u_K \odot \gamma_K\|_1 - \|u_{K^c} \odot \gamma_{K^c}\|_1} \geq \kappa_{\mathbf{T}, \gamma}(K) \kappa(K). \quad \blacksquare$$

C.4 Connection between empirical Kullback-Leibler divergence and the empirical squared norm.

We remark that the binarized matrix \mathbf{X}^B satisfies $\max_{i=1, \dots, n} \|x_i^B\|_2 = \sqrt{p}$. A direct consequence of this remark is given in the next lemma.

Lemma 7 *One has*

$$\max_{i=1, \dots, n} \sup_{\theta \in B_d(\rho)} |\langle x_i^B, \theta \rangle| \leq \rho \sqrt{p}. \quad (25)$$

To compare the empirical Kullback-Leibler divergence and the empirical squared norm, we use Lemma 1 in Bach (2010), that we recall here.

Lemma 8 Let $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ be a convex three times differentiable function such that for all $t \in \mathbb{R}$, $|\varphi'''(t)| \leq M|\varphi''(t)|$ for some $M \geq 0$. Then, for all $t \geq 0$, one has

$$\frac{\varphi''(0)}{M^2} \psi(-Mt) \leq \varphi(t) - \varphi(0) - \varphi'(0)t \leq \frac{\varphi''(0)}{M^2} \psi(Mt),$$

with $\psi(u) = e^u - u - 1$.

Now, we give a version of the previous Lemma in our setting.

Lemma 9 Under Assumption 1 and with $C_n(\rho, p)$ as defined in Theorem 2, one has

$$\begin{aligned} \frac{L_n \psi(-C_n(\rho, p))}{\phi C_n^2(\rho, p)} \frac{1}{n} \|m^0(\mathbf{X}) - m_\theta(\mathbf{X})\|_2^2 &\leq K L_n(m^0(\mathbf{X}), m_\theta(\mathbf{X})), \\ \frac{U_n \psi(C_n(\rho, p))}{\phi C_n^2(\rho, p)} \frac{1}{n} \|m^0(\mathbf{X}) - m_\theta(\mathbf{X})\|_2^2 &\geq K L_n(m^0(\mathbf{X}), m_\theta(\mathbf{X})), \end{aligned}$$

for all $\theta \in B_d(\rho)$.

Proof. Let us consider the function $G_n : \mathbb{R} \rightarrow \mathbb{R}$ defined by $G_n(t) = R_n(m^0 + tm_\eta)$, then

$$G_n(t) = \frac{1}{n} \sum_{i=1}^n b(m^0(x_i) + tm_\eta(x_i)) - \frac{1}{n} \sum_{i=1}^n y_i (m^0(x_i) + tm_\eta(x_i)).$$

By differentiating G_n three times with respect to t , we obtain

$$\begin{aligned} G_n'(t) &= \frac{1}{n} \sum_{i=1}^n m_\eta(x_i) b'(m^0(x_i) + tm_\eta(x_i)) - \frac{1}{n} \sum_{i=1}^n y_i m_\eta(x_i), \\ G_n''(t) &= \frac{1}{n} \sum_{i=1}^n m_\eta^2(x_i) b''(m^0(x_i) + tm_\eta(x_i)), \\ \text{and } G_n'''(t) &= \frac{1}{n} \sum_{i=1}^n m_\eta^3(x_i) b'''(m^0(x_i) + tm_\eta(x_i)). \end{aligned}$$

In all the considered models, we have $|b'''(z)| \leq 2|b''(z)|$, see the following table

Model	ϕ	$b(z)$	$b'(z)$	$b''(z)$	$b'''(z)$	L_n	U_n
Normal	σ^2	$\frac{z^2}{2}$	z	1	0	1	1
Logistic	1	$\log(1 + e^z)$	$\frac{e^z}{1+e^z}$	$\frac{e^z}{(1+e^z)^2}$	$\frac{1-e^z}{1+e^z} b''(z)$	$\frac{e^{C_n}}{(1+e^{C_n})^2}$	$\frac{1}{4}$
Poisson	1	e^z	e^z	e^z	$b''(z)$	e^{-C_n}	e^{C_n}

Then, we get $|G_n'''(t)| \leq 2\|m_\eta\|_\infty |G_n''(t)|$ where $\|m_\eta\|_\infty := \max_{i=1, \dots, n} |m_\eta(x_i)|$. Applying Lemma 8 with $M = 2\|m_\eta\|_\infty$, we obtain

$$G_n''(0) \frac{\psi(-2\|m_\eta\|_\infty t)}{4\|m_\eta\|_\infty^2} \leq G_n(t) - G_n(0) - tG_n'(0) \leq G_n''(0) \frac{\psi(2\|m_\eta\|_\infty t)}{4\|m_\eta\|_\infty^2}.$$

for all $t \geq 0$. Taking $t = 1$ leads to

$$\begin{aligned} G_n''(0) \frac{\psi(-2\|m_\eta\|_\infty)}{4\|m_\eta\|_\infty^2} &\leq R_n(m^0 + m_\eta) - R_n(m^0) - G_n'(0), \\ G_n''(0) \frac{\psi(2\|m_\eta\|_\infty)}{4\|m_\eta\|_\infty^2} &\geq R_n(m^0 + m_\eta) - R_n(m^0) - G_n'(0). \end{aligned}$$

A short calculation gives that

$$-G_n'(0) = \frac{1}{n} \sum_{i=1}^n m_\eta(x_i) (y_i - b'(m^0(x_i))), \text{ and } G_n''(0) = \frac{1}{n} \sum_{i=1}^n m_\eta^2(x_i) b''(m_\theta(x_i)).$$

It is clear that $\mathbb{E}_{\mathbb{P}_{\mathbf{y}|\mathbf{X}}}[-G_n'(0)] = 0$. Then

$$G_n''(0) \frac{\psi(-2\|m_\eta\|_\infty)}{4\|m_\eta\|_\infty^2} \leq R(m^0 + m_\eta) - R(m^0) \leq G_n''(0) \frac{\psi(2\|m_\eta\|_\infty)}{4\|m_\eta\|_\infty^2}.$$

Now choose $m_\eta = m_\theta - m^0$, and using Assumption 1 and (25) in Lemma 7, we have

$$2\|m_\eta\|_\infty \leq 2 \max_{i=1, \dots, n} (|\langle x_i^B, \theta \rangle| + |m^0(x_i)|) \leq 2(\rho\sqrt{p} + C_n) = C_n(\rho, p).$$

Hence, we obtain

$$\begin{aligned} G_n''(0) \frac{\psi(-C_n(\rho, p))}{C_n^2(\rho, p)} &\leq R(m_\theta) - R(m_0) = \phi K L_n(m^0(\mathbf{X}), m_\theta(\mathbf{X})), \\ G_n''(0) \frac{\psi(C_n(\rho, p))}{C_n^2(\rho, p)} &\geq R(m_\theta) - R(m_0) = \phi K L_n(m^0(\mathbf{X}), m_\theta(\mathbf{X})), \end{aligned}$$

with $G_n''(0) = n^{-1} \sum_{i=1}^n (m_\theta(x_i) - m^0(x_i))^2 b''(m^0(x_i))$. It entails that

$$\begin{aligned} \frac{L_n \psi(-C_n(\rho, p))}{\phi C_n^2(\rho, p)} \frac{1}{n} \|m^0(\mathbf{X}) - m_\theta(\mathbf{X})\|_2^2 &\leq K L_n(m^0(\mathbf{X}), m_\theta(\mathbf{X})) \\ &\leq \frac{U_n \psi(C_n(\rho, p))}{\phi C_n^2(\rho, p)} \frac{1}{n} \|m^0(\mathbf{X}) - m_\theta(\mathbf{X})\|_2^2. \end{aligned}$$

■

C.5 Proof of Theorem 2.

Recall that for all $\theta \in \mathbb{R}^d$,

$$R_n(m_\theta) = \frac{1}{n} \sum_{i=1}^n b(m_\theta(x_i)) - \frac{1}{n} \sum_{i=1}^n y_i m_\theta(x_i)$$

and

$$\hat{\theta} \in \operatorname{argmin}_{\theta \in B_d(\rho)} \{R_n(\theta) + \operatorname{bina}(\theta)\}. \quad (26)$$

According to Proposition 4, Equation (26) involves that there is $\hat{h} = (\hat{h}_{j,\bullet})_{j=1,\dots,p} \in \partial(\|\hat{\theta}\|_{\text{TV},\hat{w}})$, $\hat{g} = (\hat{g}_{j,\bullet})_{j=1,\dots,p} \in \partial(\delta_1(\hat{\theta}_{j,\bullet}))_{j=1,\dots,p}$ and $\hat{f} = (\hat{f}_{j,\bullet})_{j=1,\dots,p} \in \partial(\delta_{B_d(\rho)}(\hat{\theta}))$ such that

$$\left\langle \frac{1}{n}(\mathbf{X}^B)^\top (b'(m_{\hat{\theta}}(\mathbf{X})) - \mathbf{y}) + \hat{h} + \hat{g} + \hat{f}, \hat{\theta} - \theta \right\rangle = 0$$

for all $\theta \in \mathbb{R}^d$, which can be written

$$\begin{aligned} & \frac{1}{n} \langle b'(m_{\hat{\theta}}(\mathbf{X})) - b'(m^0(\mathbf{X})), m_{\hat{\theta}}(\mathbf{X}) - m_\theta(\mathbf{X}) \rangle \\ & - \frac{1}{n} \langle \mathbf{y} - b'(m^0(\mathbf{X})), m_{\hat{\theta}}(\mathbf{X}) - m_\theta(\mathbf{X}) \rangle + \langle \hat{h} + \hat{g} + \hat{f}, \hat{\theta} - \theta \rangle = 0. \end{aligned}$$

For any $\theta \in B_d(\rho)$ such that $\mathbf{1}^\top \theta = 0$, and $h \in \partial(\|\theta\|_{\text{TV},\hat{w}})$, the monotony of the subdifferential mapping implies $\langle \hat{h}, \theta - \hat{\theta} \rangle \leq \langle h, \theta - \hat{\theta} \rangle$, $\langle \hat{g}, \theta - \hat{\theta} \rangle \leq 0$, and $\langle \hat{f}, \theta - \hat{\theta} \rangle \leq 0$. Therefore

$$\frac{1}{n} \langle b'(m_{\hat{\theta}}(\mathbf{X})) - b'(m^0(\mathbf{X})), m_{\hat{\theta}}(\mathbf{X}) - m_\theta(\mathbf{X}) \rangle \leq \frac{1}{n} \langle \mathbf{y} - b'(m^0(\mathbf{X})), m_{\hat{\theta}}(\mathbf{X}) - m_\theta(\mathbf{X}) \rangle - \langle h, \hat{\theta} - \theta \rangle. \quad (27)$$

We consider now the function $H_n : \mathbb{R} \rightarrow \mathbb{R}$, defined by

$$H_n(t) = \frac{1}{n} \sum_{i=1}^n b(m_{\hat{\theta}+t\eta}(x_i)) - \frac{1}{n} \sum_{i=1}^n b'(m^0(x_i)) m_{\hat{\theta}+t\eta}(x_i)$$

By differentiating H_n three times with respect t , we obtain

$$\begin{aligned} H'_n(t) &= \frac{1}{n} \sum_{i=1}^n m_\eta(x_i) b'(m_{\hat{\theta}+t\eta}(x_i)) - \frac{1}{n} \sum_{i=1}^n b'(m^0(x_i)) m_\eta(x_i), \\ H''_n(t) &= \frac{1}{n} \sum_{i=1}^n m_\eta^2(x_i) b''(m_{\hat{\theta}+t\eta}(x_i)), \\ \text{and } H'''_n(t) &= \frac{1}{n} \sum_{i=1}^n m_\eta^3(x_i) b'''(m_{\hat{\theta}+t\eta}(x_i)). \end{aligned}$$

Using Lemma 7, we have $|H'''_n(t)| \leq 2\rho\sqrt{p}|H''_n(t)|$. Applying now Lemma 8 with $M(\rho, p) = 2\rho\sqrt{p}$, we obtain

$$H''_n(0) \frac{\psi(-tM(\rho, p))}{M^2(\rho, p)} \leq H_n(t) - H_n(0) - tH'_n(0) \leq H''_n(0) \frac{\psi(tM(\rho, p))}{M^2(\rho, p)},$$

for all $t \geq 0$. Taking $t = 1$ and $\eta = \theta - \hat{\theta}$ implies

$$\begin{aligned} H_n(1) &= \frac{1}{n} \sum_{i=1}^n b(m_\theta(x_i)) - \frac{1}{n} \sum_{i=1}^n b'(m^0(x_i)) m_\theta(x_i) = R(m_\theta), \\ \text{and } H_n(0) &= \frac{1}{n} \sum_{i=1}^n b(m_{\hat{\theta}}(x_i)) - \frac{1}{n} \sum_{i=1}^n b'(m^0(x_i)) m_{\hat{\theta}}(x_i) = R(m_{\hat{\theta}}). \end{aligned}$$

Moreover, we have

$$\begin{aligned} H'_n(0) &= \frac{1}{n} \sum_{i=1}^n \langle x_i^B, \theta - \hat{\theta} \rangle b'(m_{\hat{\theta}}(x_i)) - \frac{1}{n} \sum_{i=1}^n b'(m^0(x_i)) \langle x_i^B, \hat{\theta} - \theta \rangle \\ &= \frac{1}{n} \langle b'(m_{\hat{\theta}}(\mathbf{X})) - b'(m^0(\mathbf{X})), \mathbf{X}^B(\theta - \hat{\theta}) \rangle, \\ \text{and } H''_n(0) &= \frac{1}{n} \sum_{i=1}^n \langle x_i^B, \hat{\theta} - \theta \rangle^2 b''(m_{\hat{\theta}}(x_i)). \end{aligned}$$

Then, we deduce that

$$\begin{aligned} H''_n(0) \frac{\psi(-M(\rho, p))}{M(\rho, p)^2} &\leq R(m_\theta) - R(m_{\hat{\theta}}) - \frac{1}{n} \langle b'(m_{\hat{\theta}}(\mathbf{X})) - b'(m^0(\mathbf{X})), \mathbf{X}^B(\theta - \hat{\theta}) \rangle \\ &= \phi KL_n(m^0(\mathbf{X}), m_\theta(\mathbf{X})) - \phi KL_n(m^0(\mathbf{X}), m_{\hat{\theta}}(\mathbf{X})) \\ &\quad + \frac{1}{n} \langle b'(m_{\hat{\theta}}(\mathbf{X})) - b'(m^0(\mathbf{X})), m_{\hat{\theta}}(\mathbf{X}) - m_\theta(\mathbf{X}) \rangle. \end{aligned}$$

Then, with Equation (27), one has

$$\begin{aligned} \phi KL_n(m^0(\mathbf{X}), m_{\hat{\theta}}(\mathbf{X})) + H''_n(0) \frac{\psi(-M(\rho, p))}{M^2(\rho, p)} \\ \leq \phi KL_n(m^0(\mathbf{X}), m_\theta(\mathbf{X})) + \frac{1}{n} \langle \mathbf{y} - b'(m^0(\mathbf{X})), m_{\hat{\theta}}(\mathbf{X}) - m_\theta(\mathbf{X}) \rangle - \langle h, \hat{\theta} - \theta \rangle. \end{aligned} \quad (28)$$

As $H''_n(0) \geq 0$, it implies that

$$\begin{aligned} \phi KL_n(m^0(\mathbf{X}), m_{\hat{\theta}}(\mathbf{X})) &\leq \phi KL_n(m^0(\mathbf{X}), m_\theta(\mathbf{X})) \\ &\quad + \frac{1}{n} \langle \mathbf{y} - b'(m^0(\mathbf{X})), m_{\hat{\theta}}(\mathbf{X}) - m_\theta(\mathbf{X}) \rangle - \langle h, \hat{\theta} - \theta \rangle. \end{aligned} \quad (29)$$

If $\frac{1}{n} \langle \mathbf{y} - b'(m^0(\mathbf{X})), \mathbf{X}^B(\hat{\theta} - \theta) \rangle - \langle h, \hat{\theta} - \theta \rangle < 0$, it follows that

$$KL_n(m^0(\mathbf{X}), m_{\hat{\theta}}(\mathbf{X})) \leq KL_n(m^0(\mathbf{X}), m_\theta(\mathbf{X})),$$

then Theorem 2 holds. From now on, let us assume that

$$\frac{1}{n} \langle \mathbf{y} - b'(m^0(\mathbf{X})), m_{\hat{\theta}}(\mathbf{X}) - m_\theta(\mathbf{X}) \rangle - \langle h, \hat{\theta} - \theta \rangle \geq 0. \quad (30)$$

We first derive a bound on $\frac{1}{n} \langle \mathbf{y} - b'(m^0(\mathbf{X})), m_{\hat{\theta}}(\mathbf{X}) - m_\theta(\mathbf{X}) \rangle$. Using $\mathbf{D}^{-1} = \mathbf{T}$, we focus on finding out a bound of $\frac{1}{n} \langle (\mathbf{X}^B \mathbf{T})^\top (\mathbf{y} - b'(m^0(\mathbf{X}))), \mathbf{D}(\hat{\theta} - \theta) \rangle$. In one hand, one has

$$\begin{aligned} \frac{1}{n} \langle (\mathbf{X}^B)^\top (\mathbf{y} - b'(m^0(\mathbf{X}))), \hat{\theta} - \theta \rangle &= \frac{1}{n} \langle (\mathbf{X}^B \mathbf{T})^\top (\mathbf{y} - b'(m^0(\mathbf{X}))), \mathbf{D}(\hat{\theta} - \theta) \rangle \\ &\leq \frac{1}{n} \sum_{j=1}^p \sum_{k=1}^{d_j} \left| \langle (\mathbf{X}_{\bullet, j}^B T_j)_{\bullet, k}, \mathbf{y} - b'(m^0(\mathbf{X})) \rangle \right| \left| (D_j(\hat{\theta}_{j, \bullet} - \theta_{j, \bullet}))_k \right|, \end{aligned}$$

where $(\mathbf{X}_{\bullet,j}^B T_j)_{\bullet,k} = ((\mathbf{X}_{\bullet,j}^B T_j)_{1,k}, \dots, (\mathbf{X}_{\bullet,j}^B T_j)_{n,k})^\top \in \mathbb{R}^n$ is the k -th column of the matrix $(\mathbf{X}_{\bullet,j}^B T_j)$. Let us consider the event

$$\mathcal{E}_n = \bigcap_{j=1}^p \bigcap_{k=2}^{d_j} \mathcal{E}_{n,j,k}, \text{ where } \mathcal{E}_{n,j,k} = \left\{ \frac{1}{n} \left| \langle (\mathbf{X}_{\bullet,j}^B T_j)_{\bullet,k}, \mathbf{y} - b'(m^0(\mathbf{X})) \rangle \right| \leq \hat{w}_{j,k} \right\}.$$

Then, on \mathcal{E}_n , we have

$$\begin{aligned} \frac{1}{n} \langle (\mathbf{X}^B)^\top (\mathbf{y} - b'(m^0(\mathbf{X}))), \hat{\theta} - \theta \rangle &\leq \sum_{j=1}^p \sum_{k=1}^{d_j} \hat{w}_{j,k} | (D_j(\hat{\theta}_{j,\bullet} - \theta_{j,\bullet}))_k | \\ &\leq \sum_{j=1}^p \|\hat{w}_{j,\bullet} \odot D_j(\hat{\theta}_{j,\bullet} - \theta_{j,\bullet})\|_1. \end{aligned} \quad (31)$$

In another hand, from the definition of the subgradient $(h_{j,\bullet})_{j=1,\dots,p} \in \partial(\|\theta\|_{\text{TV},\hat{w}})$ (see Equation (16)), one can choose h such that

$$h_{j,k} = \left(D_j^\top (\hat{w}_{j,\bullet} \odot \text{sign}(D_j \theta_{j,\bullet})) \right)_k$$

for all $k = 1, \dots, J_j(\theta)$ and

$$h_{j,k} = \left(D_j^\top (\hat{w}_{j,\bullet} \odot \text{sign}(D_j \hat{\theta}_{j,\bullet})) \right)_k = \left(D_j^\top (\hat{w}_{j,\bullet} \odot \text{sign}(D_j(\hat{\theta}_{j,\bullet} - \theta_{j,\bullet}))) \right)_k$$

for all $k = 1, \dots, J_j^c(\theta)$. Using a triangle inequality and the fact that $\langle \text{sign}(x), x \rangle = \|x\|_1$, we obtain

$$\begin{aligned} -\langle h, \hat{\theta} - \theta \rangle &\leq \sum_{j=1}^p \left\| (\hat{w}_{j,\bullet})_{J_j(\theta)} \odot D_j(\hat{\theta}_{j,\bullet} - \theta_{j,\bullet})_{J_j(\theta)} \right\|_1 \\ &\quad - \sum_{j=1}^p \left\| (\hat{w}_{j,\bullet})_{J_j^c(\theta)} \odot D_j(\hat{\theta}_{j,\bullet} - \theta_{j,\bullet})_{J_j^c(\theta)} \right\|_1 \\ &\leq \sum_{j=1}^p \left\| (\hat{\theta}_{j,\bullet} - \theta_{j,\bullet})_{J_j(\theta)} \right\|_{\text{TV},\hat{w}_{j,\bullet}} - \sum_{j=1}^p \left\| (\hat{\theta}_{j,\bullet} - \theta_{j,\bullet})_{J_j^c(\theta)} \right\|_{\text{TV},\hat{w}_{j,\bullet}}. \end{aligned} \quad (32)$$

Combining inequalities (31) and (32), we get

$$\sum_{j=1}^p \left\| (\hat{\theta}_{j,\bullet} - \theta_{j,\bullet})_{J_j^c(\theta)} \right\|_{\text{TV},\hat{w}_{j,\bullet}} \leq 2 \sum_{j=1}^p \left\| (\hat{\theta}_{j,\bullet} - \theta_{j,\bullet})_{J_j(\theta)} \right\|_{\text{TV},\hat{w}_{j,\bullet}}$$

on \mathcal{E}_n . Hence

$$\sum_{j=1}^p \left\| (\hat{w}_{j,\bullet})_{J_j^c(\theta)} \odot D_j(\hat{\theta}_{j,\bullet} - \theta_{j,\bullet})_{J_j^c(\theta)} \right\|_1 \leq 2 \sum_{j=1}^p \left\| (\hat{w}_{j,\bullet})_{J_j(\theta)} \odot D_j(\hat{\theta}_{j,\bullet} - \theta_{j,\bullet})_{J_j(\theta)} \right\|_1.$$

This means that

$$\hat{\theta} - \theta \in \mathcal{C}_{\text{TV}, \hat{w}}(J(\theta)) \text{ and } \mathbf{D}(\hat{\theta} - \theta) \in \mathcal{C}_{1, \hat{w}}(J(\theta)), \quad (33)$$

see (7) and (24). Now, going back to (29) and taking into account (33), the compatibility of $\mathbf{X}^B \mathbf{T}$ (see (23)), on \mathcal{E}_n the following holds

$$\phi KL_n(m^0(\mathbf{X}), m_{\hat{\theta}}(\mathbf{X})) \leq \phi KL_n(m^0(\mathbf{X}), m_{\theta}(\mathbf{X})) + 2 \sum_{j=1}^p \|(\hat{w}_{j, \bullet})_{J_j(\theta)} \odot D_j(\hat{\theta}_{j, \bullet} - \theta_{j, \bullet})_{J_j(\theta)}\|_1.$$

Then

$$KL_n(m^0(\mathbf{X}), m_{\hat{\theta}}(\mathbf{X})) \leq KL_n(m^0(\mathbf{X}), m_{\theta}(\mathbf{X})) + \frac{\|m_{\hat{\theta}}(\mathbf{X}) - m_{\theta}(\mathbf{X})\|_2}{\sqrt{n} \phi \kappa_{\mathbf{T}, \hat{\gamma}}(J(\theta)) \kappa(J(\theta))}, \quad (34)$$

where $\hat{\gamma} = (\hat{\gamma}_{1, \bullet}^\top, \dots, \hat{\gamma}_{p, \bullet}^\top)^\top$ such that

$$\hat{\gamma}_{j, k} = \begin{cases} 2\hat{w}_{j, k} & \text{if } k \in J_j(\theta), \\ 0 & \text{if } k \in J_j^c(\theta), \end{cases}$$

for all $j = 1, \dots, p$ and

$$\kappa_{\mathbf{T}, \hat{\gamma}}(J(\theta)) = \left\{ 32 \sum_{j=1}^p \sum_{k=1}^{d_j} |\hat{\gamma}_{j, k+1} - \hat{\gamma}_{j, k}|^2 + 2|J_j(\theta)| \|\hat{\gamma}_{j, \bullet}\|_\infty^2 \Delta_{\min, J_j(\theta)}^{-1} \right\}^{-1/2}.$$

Next, we find an upper bound for $1/\kappa_{\mathbf{T}, \hat{\gamma}}^2(J(\theta))$. We have

$$\frac{1}{\kappa_{\mathbf{T}, \hat{\gamma}}^2(J(\theta))} = 32 \sum_{j=1}^p \sum_{k=1}^{d_j} |\hat{\gamma}_{j, k+1} - \hat{\gamma}_{j, k}|^2 + 2|J_j(\theta)| \|\hat{\gamma}_{j, \bullet}\|_\infty^2 \Delta_{\min, J_j(\theta)}^{-1}.$$

Note that $\|\hat{\gamma}_{j, \bullet}\|_\infty \leq 2\|\hat{w}_{j, \bullet}\|_\infty$. We write the set $J_j(\theta) = \{k_j^1, \dots, k_j^{|J_j(\theta)|}\}$ and we set $B_r = \llbracket k_j^{r-1}, k_j^r \rrbracket = \{k_j^{r-1}, k_j^{r-1} + 1, \dots, k_j^r - 1\}$ for $r = 1, \dots, |J_j(\theta)| + 1$ with the convention that $k_j^0 = 0$ and $k_j^{|J_j(\theta)|+1} = d_j + 1$. Then

$$\begin{aligned} \sum_{k=1}^{d_j} |\hat{\gamma}_{j, k+1} - \hat{\gamma}_{j, k}|^2 &= \sum_{r=1}^{|J_j(\theta)|+1} \sum_{k \in B_r} |\hat{\gamma}_{j, k+1} - \hat{\gamma}_{j, k}|^2 \\ &= \sum_{r=1}^{|J_j(\theta)|+1} |\hat{\gamma}_{j, k_j^{r-1}+1} - \hat{\gamma}_{j, k_j^{r-1}}|^2 + |\hat{\gamma}_{j, k_j^r} - \hat{\gamma}_{j, k_j^r-1}|^2 \\ &= \sum_{r=1}^{|J_j(\theta)|+1} \hat{\gamma}_{j, k_j^{r-1}}^2 + \hat{\gamma}_{j, k_j^r}^2 \\ &= \sum_{r=1}^{|J_j(\theta)|} 2 \hat{\gamma}_{j, k_j^r}^2 \\ &\leq 8 |J_j(\theta)| \|(\hat{w}_{j, \bullet})_{J_j(\theta)}\|_\infty^2. \end{aligned}$$

Therefore

$$\begin{aligned}
 \frac{1}{\kappa_{\mathbf{T},\hat{\gamma}}^2(J(\theta))} &\leq 32 \sum_{j=1}^p \left\{ 8 |J_j(\theta)| \|\hat{w}_{j,\bullet}\|_{J_j(\theta)}^2 \right\} + 8 |J_j(\theta)| \|\hat{w}_{j,\bullet}\|_{J_j(\theta)}^2 \Delta_{\min, J_j(\theta)}^{-1} \\
 &\leq (32 \times 8) \sum_{j=1}^p \left\{ 1 + \frac{1}{\Delta_{\min, J_j(\theta)}} \right\} |J_j(\theta)| \|\hat{w}_{j,\bullet}\|_{J_j(\theta)}^2 \\
 &\leq 512 |J(\theta)| \max_{j=1, \dots, p} \|\hat{w}_{j,\bullet}\|_{J_j(\theta)}^2.
 \end{aligned}$$

Remark 10 For the case of least squares regression where $y_i|x_i$ has Gaussian distribution with mean $m^0(x_i)$ and variance $\phi = \sigma^2$. Using inequalities (28) and (34), we get

$$\begin{aligned}
 \phi K L_n(m^0(\mathbf{X}), m_{\hat{\theta}}(\mathbf{X})) &+ \frac{\psi(-M(\rho, p))}{M^2(\rho, p)} \frac{1}{n} \|m_{\hat{\theta}}(\mathbf{X}) - m_{\theta}(\mathbf{X})\|_2^2 \\
 &\leq \phi K L_n(m^0(\mathbf{X}), m_{\theta}(\mathbf{X})) + \frac{\|m_{\hat{\theta}}(\mathbf{X}) - m_{\theta}(\mathbf{X})\|_2}{\sqrt{n} \kappa_{\mathbf{T},\hat{\gamma}}(J(\theta)) \kappa(J(\theta))} \\
 &\leq \phi K L_n(m^0(\mathbf{X}), m_{\theta}(\mathbf{X})) \\
 &\quad + 2 \frac{\sqrt{\psi(-M(\rho, p))}}{M(\rho, p)} \frac{1}{\sqrt{n}} \|m_{\hat{\theta}}(\mathbf{X}) - m_{\theta}(\mathbf{X})\|_2 \frac{M(\rho, p)}{\sqrt{\psi(-M(\rho, p)) \kappa_{\mathbf{T},\hat{\gamma}}(J(\theta)) \kappa(J(\theta))}}
 \end{aligned}$$

Using the fact that $2uv \leq u^2 + v^2$ it yields

$$\phi K L_n(m^0(\mathbf{X}), m_{\hat{\theta}}(\mathbf{X})) \leq \phi K L_n(m^0(\mathbf{X}), m_{\theta}(\mathbf{X})) + \frac{M^2(\rho, p)}{\psi(-M(\rho, p)) \kappa_{\mathbf{T},\hat{\gamma}}^2(J(\theta)) \kappa^2(J(\theta))}$$

Hence, we derive the following sharp oracle inequality

$$R(m_{\hat{\theta}}) - R(m^0) \leq \inf_{\theta \in B_d(\rho)} \left\{ R(m_{\theta}) - R(m^0) + \frac{\xi |J(\theta)|}{\kappa^2(J(\theta))} \max_{j=1, \dots, p} \|\hat{w}_{j,\bullet}\|_{J_j(\theta)}^2 \right\},$$

where

$$\xi = \frac{512 M^2(\rho, p)}{\psi(-M(\rho, p))}.$$

Now for generalized linear models, we use the connection between the empirical norm and the Kullback-Leibler divergence. First, We have

$$\begin{aligned}
 &\frac{\|m_{\hat{\theta}}(\mathbf{X}) - m_{\theta}(\mathbf{X})\|_2}{\sqrt{n} \phi \kappa_{\mathbf{T},\hat{\gamma}}(J(\theta)) \kappa(J(\theta))} \\
 &\leq \frac{1}{\phi \kappa_{\mathbf{T},\hat{\gamma}}(J(\theta)) \kappa(J(\theta))} \left(\frac{1}{\sqrt{n}} \|m_{\hat{\theta}}(\mathbf{X}) - m^0(\mathbf{X})\|_2 + \frac{1}{\sqrt{n}} \|m^0(\mathbf{X}) - m_{\theta}(\mathbf{X})\|_2 \right).
 \end{aligned}$$

Therefore, by Lemma 9, we get

$$\begin{aligned}
 &\frac{\|m_{\hat{\theta}}(\mathbf{X}) - m_{\theta}(\mathbf{X})\|_2}{\sqrt{n} \phi \kappa_{\mathbf{T},\hat{\gamma}}(J(\theta)) \kappa(J(\theta))} \\
 &\leq \frac{2}{\sqrt{\phi} \kappa_{\mathbf{T},\hat{\gamma}}(J(\theta)) \kappa(J(\theta))} \left(\sqrt{C_n(\rho, p, L_n)^{-1} K L_n(m^0(\mathbf{X}), m_{\hat{\theta}}(\mathbf{X}))} \right. \\
 &\quad \left. + \sqrt{C_n(\rho, p, L_n)^{-1} K L_n(m^0(\mathbf{X}), m_{\theta}(\mathbf{X}))} \right).
 \end{aligned}$$

We now use the elementary inequality $2uv \leq \epsilon u^2 + v^2/\epsilon$ with $\epsilon > 0$. Therefore (34) becomes

$$\begin{aligned} KL_n(m^0(\mathbf{X}), m_{\hat{\theta}}(\mathbf{X})) &\leq KL_n(m^0(\mathbf{X}), m_{\theta}(\mathbf{X})) + \frac{\epsilon}{\phi \kappa_{\mathbf{T}, \hat{\gamma}}^2(J(\theta)) \kappa^2(J(\theta))} \\ &\quad + 2(\epsilon C_n(\rho, p, L_n))^{-1} KL_n(m^0(\mathbf{X}), m_{\hat{\theta}}(\mathbf{X})) \\ &\quad + 2(\epsilon C_n(\rho, p, L_n))^{-1} KL_n(m^0(\mathbf{X}), m_{\theta}(\mathbf{X})). \end{aligned}$$

By choosing $2(\epsilon C_n(\rho, p, L_n))^{-1} < 1$, we get

$$\begin{aligned} KL_n(m^0(\mathbf{X}), m_{\hat{\theta}}(\mathbf{X})) &\leq \frac{1 + 2(\epsilon C_n(\rho, p, L_n))^{-1}}{1 - 2(\epsilon C_n(\rho, p, L_n))^{-1}} KL_n(m^0(\mathbf{X}), m_{\theta}(\mathbf{X})) \\ &\quad + \frac{\epsilon^2}{(1 - 2(\epsilon C_n(\rho, p, L_n))^{-1}) \phi \kappa_{\mathbf{T}, \hat{\gamma}}^2(J(\theta)) \kappa^2(J(\theta))} \\ &\leq \frac{\epsilon C_n(\rho, p, L_n) + 2}{\epsilon C_n(\rho, p, L_n) - 2} KL_n(m^0(\mathbf{X}), m_{\theta}(\mathbf{X})) \\ &\quad + \frac{\epsilon^2 C_n(\rho, p, L_n)}{(\epsilon C_n(\rho, p, L_n) - 2) \phi \kappa_{\mathbf{T}, \hat{\gamma}}^2(J(\theta)) \kappa^2(J(\theta))}. \end{aligned}$$

Setting

$$\frac{\epsilon C_n(\rho, p, L_n) + 2}{\epsilon C_n(\rho, p, L_n) - 2} = 1 + \frac{4}{\epsilon C_n(\rho, p, L_n) - 2} = 1 + \zeta,$$

we get the desired result in (10).

Finally, we have to compute the probability of the complementary of the event \mathcal{E}_n . This is given by the following:

$$\begin{aligned} \mathbb{P}[\mathcal{E}_n^c] &\leq \sum_{j=1}^p \sum_{k=2}^{d_j} \mathbb{P}\left[\frac{1}{n} \left| \langle (\mathbf{X}_{\bullet, j}^B T_j)_{\bullet, k}, \mathbf{y} - b'(m^0(\mathbf{X})) \rangle \right| \geq \hat{w}_{j, k}\right] \\ &\leq \sum_{j=1}^p \sum_{k=2}^{d_j} \mathbb{P}\left[\sum_{i=1}^n |(\mathbf{X}_{\bullet, j}^B T_j)_{i, k} (y_i - b'(m^0(x_i)))| \geq n \hat{w}_{j, k}\right]. \end{aligned}$$

Let $\xi_{i, j, k} = (\mathbf{X}_{\bullet, j}^B T_j)_{i, k}$, and $Z_i = y_i - b'(m^0(x_i))$. Note that conditionally on x_i , the random variables (Z_i) are independent. It can be easily shown (see Theorem 5.10 in Lehmann and Casella (1998)) that the moment generating function of Z (copy of Z_i) is given by

$$\mathbb{E}[\exp(tZ)] = \exp\left(\phi^{-1}\{b(m^0(x) + t) - tb'(m^0(x) - b(m^0(x)))\}\right). \quad (35)$$

Applying Lemma 6.1 in Rigollet (2012), using (35) and Assumption 1, we can derive the following Chernoff-type bounds

$$\mathbb{P}\left[\sum_{i=1}^n |\xi_{i, j, k} Z_i| \geq n \hat{w}_{j, k}\right] \leq 2 \exp\left(-\frac{n^2 \hat{w}_{j, k}^2}{2U_n \phi \|\xi_{\bullet, j, k}\|_2^2}\right), \quad (36)$$

where $\xi_{\bullet,j,k} = (\xi_{1,j,k}, \dots, \xi_{n,j,k})^\top \in \mathbb{R}^n$. We have

$$\mathbf{X}_{\bullet,j}^B T_j = \begin{bmatrix} 1 & \sum_{k=2}^{d_j} x_{1,j,k}^B & \sum_{k=3}^{d_j} x_{1,j,k}^B & \cdots & \sum_{k=d_{j-1}}^{d_j} x_{1,j,k}^B & x_{1,j,d_j}^B \\ \vdots & \vdots & \vdots & & \vdots & \vdots \\ 1 & \sum_{k=2}^{d_j} x_{n,j,k}^B & \sum_{k=3}^{d_j} x_{n,j,k}^B & \cdots & \sum_{k=d_{j-1}}^{d_j} x_{n,j,k}^B & x_{n,j,d_j}^B \end{bmatrix}.$$

Therefore,

$$\|\xi_{\bullet,j,k}\|_2^2 = \sum_{i=1}^n (\mathbf{X}_{\bullet,j}^B T_j)_{\bullet,k}^2 = \#\left(\left\{i \in [n] : x_{i,j} \in \bigcup_{r=k}^{d_j} I_{j,r}\right\}\right) = n\hat{\pi}_{j,k}. \quad (37)$$

Using weights $\hat{w}_{j,k}$ (see (9) in Theorem 2), and (36) together with (37), we find that the probability of the complementary event \mathcal{E}_n^c is smaller than $2e^{-A}$. This concludes the proof of Theorem 2. \blacksquare

References

- A. Agresti. *Foundations of Linear and Generalized Linear Models*. John Wiley & Sons, 2015.
- M. Z. Alaya, S. Gaïffas, and A. Guillaou. Learning the intensity of time events with change-points. *Information Theory, IEEE Transactions on*, 61(9):5148–5171, 2015.
- F. Bach. Self-concordant analysis for logistic regression. *Electron. J. Statist.*, 4:384–414, 2010.
- F. Bach, R. Jenatton, J. Mairal, and G. Obozinski. Optimization with sparsity-inducing penalties. *Foundations and Trends® in Machine Learning*, 4(1):1–106, 2012.
- E. Bacry, M. Bompain, S. Gaïffas, and S. Poulsen. tick: a Python library for statistical learning, with a particular emphasis on time-dependent modeling. *ArXiv e-prints*, July 2017.
- P. Baldi, P. Sadowski, and D. Whiteson. Searching for exotic particles in high-energy physics with deep learning. *Nature communications*, 5, 2014.
- P. Baldi, K. Cranmer, T. Faucett, P. Sadowski, and D. Whiteson. Parameterized neural networks for high-energy physics. *The European Physical Journal C*, 76(5):1–7, Apr 2016.
- H. H. Bauschke and P. L. Combettes. *Convex analysis and monotone operator theory in Hilbert spaces*. CMS Books in Mathematics/Ouvrages de Mathématiques de la SMC. Springer, New York, 2011.
- P. J. Bickel, Y. Ritov, and A. B. Tsybakov. Simultaneous analysis of Lasso and Dantzig selector. *Ann. Statist.*, 37(4):1705–1732, 2009. ISSN 0090-5364.
- J. A. Blackard and D. J. Dean. Comparative accuracies of artificial neural networks and discriminant analysis in predicting forest cover types from cartographic variables. *Computers and electronics in agriculture*, 24(3):131–151, 1999.

- S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge University Press, Cambridge, 2004. ISBN 0-521-83378-7.
- L. Breiman. Random forests. *Mach. Learn.*, 45(1):5–32, 2001.
- L. Breiman, J. Friedman, R. Olshen, and C. Stone. *Classification and Regression Trees*. Wadsworth and Brooks, Monterey, CA, 1984.
- P. Bühlmann and S. van De Geer. *Statistics for high-dimensional data*. Springer Series in Statistics. Springer, Heidelberg, 2011.
- F. Bunea, A. Tsybakov, and M. Wegkamp. Sparsity oracle inequalities for the Lasso. *Electron. J. Statist.*, 1:169–194, 2007.
- E. J. Candès and M. B. Wakin. An Introduction To Compressive Sampling. *Signal Processing Magazine, IEEE*, 25(2):21–30, 2008.
- E. J. Candès, M. B. Wakin, and S. P. Boyd. Enhancing sparsity by reweighted ℓ_1 minimization. *Journal of Fourier Analysis and Applications*, 14(5):877–905, 2008.
- B. Chlebus and S. H. Nguyen. On finding optimal discretizations for two attributes. In Lech Polkowski and Andrzej Skowron, editors, *Rough Sets and Current Trends in Computing*, volume 1424 of *Lecture Notes in Computer Science*, pages 537–544. Springer Berlin Heidelberg, 1998.
- L. Condat. A Direct Algorithm for 1D Total Variation Denoising. *IEEE Signal Processing Letters*, 20(11):1054–1057, 2013.
- A. S. Dalalyan, M. Hebiri, and J. Lederer. On the prediction performance of the Lasso. *Bernoulli*, 23(1):552–581, 2017.
- D. L. Donoho and M. Elad. Optimally sparse representation in general (non-orthogonal) dictionaries via ℓ_1 minimization. In *PROC. NATL ACAD. SCI. USA 100 2197202*, 2002.
- D. L. Donoho and X. Huo. Uncertainty principles and ideal atomic decomposition. *Information Theory, IEEE Transactions on*, 47(7):2845–2862, 2001.
- J. Friedman, T. Hastie, H. Höfling, and R. Tibshirani. Pathwise coordinate optimization. *Ann. Appl. Stat.*, 1(2):302–332, 2007.
- J. H. Friedman. Stochastic gradient boosting. *Computational Statistics & Data Analysis*, 38(4):367–378, 2002.
- S. Garcia, J. Luengo, J. A. Saez, V. Lopez, and F. Herrera. A survey of discretization techniques: Taxonomy and empirical analysis in supervised learning. *IEEE Transactions on Knowledge and Data Engineering*, 25(4):734–750, 2013.
- P. J. Green and B. W. Silverman. *Nonparametric regression and generalized linear models: a roughness penalty approach*. Chapman and Hall, London, 1994.
- T. Hastie and R. Tibshirani. *Generalized additive models*. Wiley Online Library, 1990.

- T. Hastie, R. Tibshirani, and J. Friedman. *The elements of statistical learning*. Springer Series in Statistics. Springer-Verlag, New York, 2001.
- K. Knight and W. Fu. Asymptotics for Lasso-type estimators. *Ann. Statist.*, 28(5):1356–1378, 2000.
- R. Kohavi. Scaling up the accuracy of naive-Bayes classifiers: A decision-tree hybrid. In *KDD*, volume 96, pages 202–207, 1996.
- E. L. Lehmann and G. Casella. *Theory of point estimation*. Springer texts in statistics. Springer, New York, 1998.
- M. Lichman. UCI Machine Learning Repository, 2013.
- H. Liu, F. Hussain, C. L. Tan, and M. Dash. Discretization: an enabling technique. *Data Min. Knowl. Discov.*, 6(4):393–423, 2002.
- G. Lugosi and N. Vayatis. On the Bayes-risk consistency of regularized boosting methods. *Annals of Statistics*, pages 30–55, 2004.
- L. Meier, S. van De Geer, and P. Bühlmann. The group lasso for logistic regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70(1):53–71, 2008.
- S. Moro, P. Cortez, and P. Rita. A data-driven approach to predict the success of bank telemarketing. *Decision Support Systems*, 62:22–31, 2014.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- J. R. Quinlan. *C4.5: Programs for Machine Learning (Morgan Kaufmann Series in Machine Learning)*. Morgan Kaufmann, 1 edition, 1993.
- F. Rapaport, E. Barillot, and J. P. Vert. Classification of arraycgh data using fused SVM. *Bioinformatics*, 24(13):i375–i382, 2008.
- P. Rigollet. Kullback Leibler aggregation and misspecified generalized linear models. *Ann. Statist.*, 40(2):639–665, 2012.
- M. A. Russell. *Mining the Social Web: Data Mining Facebook, Twitter, LinkedIn, Google+, GitHub, and More*. O’Reilly Media, 2013.
- B. Schölkopf and A. J. Smola. *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press, 2002.
- V. G. Sigillito, S. P. Wing, L. V. Hutton, and K. B. Baker. Classification of radar returns from the ionosphere using neural networks. *Johns Hopkins APL Technical Digest*, 10(3):262–266, 1989.

- R. Tibshirani. Regression shrinkage and selection via the Lasso. *J. Roy. Statist. Soc. Ser. B*, 58(1):267–288, 1996a.
- R. Tibshirani, M. Saunders, S. Rosset, J. Zhu, and K. Knight. Sparsity and smoothness via the fused Lasso. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 67(1):91–108, 2005.
- Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996b.
- S. van de Geer. High-dimensional generalized linear models and the Lasso. *Ann. Statist.*, 36(2):614–645, 2008.
- S. van de Geer and J. Lederer. *The Lasso, correlated design, and improved oracle inequalities*, volume Volume 9 of *Collections*, pages 303–316. Institute of Mathematical Statistics, Beachwood, Ohio, USA, 2013.
- J. Wu and S. Coggeshall. *Foundations of Predictive Analytics (Chapman & Hall/CRC Data Mining and Knowledge Discovery Series)*. Chapman & Hall/CRC, 1st edition, 2012.
- I. C. Yeh and C. H. Lien. The comparisons of data mining techniques for the predictive accuracy of probability of default of credit card clients. *Expert Systems with Applications*, 36(2):2473–2480, 2009.
- Y. L. Yu. On decomposing the proximal map. In C.J.C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 91–99. 2013.
- P. Zhao and B. Yu. On model selection consistency of Lasso. *J. Mach. Learn. Res.*, 7: 2541–2563, 2006.