

LEARNING WITH CLUSTERING STRUCTURE

VINCENT ROULET, FAJWEL FOGEL, ALEXANDRE D’ASPREMONT, AND FRANCIS BACH

ABSTRACT. We study supervised learning problems using clustering constraints to impose structure on either features or samples, seeking to help both prediction and interpretation. The problem of clustering features arises naturally in text classification for instance, to reduce dimensionality by grouping words together and identify synonyms. The sample clustering problem on the other hand, applies to multiclass problems where we are allowed to make multiple predictions and the performance of the best answer is recorded. We derive a unified optimization formulation highlighting the common structure of these problems and produce algorithms whose core iteration complexity amounts to a k-means clustering step, which can be approximated efficiently. We extend these results to combine sparsity and clustering constraints, and develop a new projection algorithm on the set of clustered sparse vectors. We prove convergence of our algorithms on random instances, based on a union of subspaces interpretation of the clustering structure. Finally, we test the robustness of our methods on artificial data sets as well as real data extracted from movie reviews.

1. INTRODUCTION

Adding structural information to supervised learning problems can significantly improve prediction performance. Sparsity for example has been proven to improve statistical and practical performance [Bach et al., 2012]. Here, we study clustering constraints that seek to group either features or samples, to both improve prediction and provide additional structural insights on the data.

When there exists some groups of highly correlated features for instance, reducing dimensionality by assigning uniform weights inside each distinct group of features can be beneficial both in terms of prediction and interpretation [Bondell and Reich, 2008] by significantly reducing dimension. This often occurs in text classification for example, where it is natural to group together words having the same meaning for a given task [Dhillon et al., 2003; Jiang et al., 2011].

On the other hand, learning a unique predictor for all samples can be too restrictive. For recommendation systems for example, users can be partitioned in groups, each having different tastes. Here, we study how to learn a partition of the samples that achieves the best within-group prediction [Guzman-Rivera et al., 2014; Zhang, 2003]

These problems can of course be tackled by grouping synonyms or clustering samples in an unsupervised preconditioning step. However such partitions might not be optimized or relevant for the prediction task. Prior hypotheses on the partition can also be added as in Latent Dirichlet Allocation [Blei et al., 2003] or Mixture of Experts [Jordan, 1994]. We present here a unified framework that highlights the clustered structure of these problems without adding prior information on these clusters. While constraining the predictors, our framework allows the use of any loss function for the prediction task. We propose several optimization schemes to solve these problems efficiently.

First, we formulate an explicit convex relaxation which can be solved efficiently using the conditional gradient algorithm [Frank and Wolfe, 1956; Jaggi, 2013], where the core inner step amounts to solving a clustering problem. We then study an approximate projected gradient scheme similar to the Iterative Hard Thresholding (IHT) algorithm [Blumensath and Davies, 2009] used in compressed sensing. While

Date: September 19, 2016.

Key words and phrases. Clustering, Multitask, Dimensionality Reduction, Supervised Learning, Supervised Clustering.

A shorter, preliminary version of this paper appeared at the NIPS 2015 workshop “Transfer and Multi-Task Learning: Trends and New Perspectives”.

constraints are non-convex, projection on the feasible set reduces to a clustering subproblem akin to k-means. In the particular case of feature clustering for regression, the k-means steps are performed in dimension one, and can therefore be solved exactly by dynamic programming [Bellman, 1973; Wang and Song, 2011]. When a sparsity constraint is added to the feature clustering problem for regression, we develop a new dynamic program that gives the exact projection on the set of sparse and clustered vectors.

We provide a theoretical convergence analysis of our projected gradient scheme generalizing the proof made for IHT. Although our structure is similar to sparsity, we show that imposing a clustered structure, while helping interpretability, does not allow us to significantly reduce the number of samples, as in the sparse case for example.

Finally, we describe experiments on both synthetic and real datasets involving large corpora of text from movie reviews. The use of k-means steps makes our approach fast and scalable while comparing very favorably with standard benchmarks and providing meaningful insights on the data structure.

2. LEARNING & CLUSTERING FEATURES OR SAMPLES

Given n sample points represented by the matrix $X = (x_1, \dots, x_n)^T \in \mathbb{R}^{n \times d}$ and corresponding labels $y = (y_1, \dots, y_n)$, real or nominal depending on the task (classification or regression), we seek to compute linear predictors represented by W . Clustering features or samples is done by constraining W and our problems take the generic form

$$\begin{aligned} & \text{minimize} && \mathbf{Loss}(y, X, W) + R(W) \\ & \text{subject to} && W \in \mathcal{W}, \end{aligned}$$

in the prediction variable W , where $\mathbf{Loss}(y, X, W)$ is a learning loss (for simplicity, we consider only squared or logistic losses in what follows), $R(W)$ is a classical regularizer and \mathcal{W} encodes the clustering structure.

The clustering constraint partitions features or samples into Q groups $\mathcal{G}_1, \dots, \mathcal{G}_Q$ of size s_1, \dots, s_Q by imposing that all features or samples within a cluster \mathcal{G}_q share a common predictor vector or coefficient v_q , solving the supervised learning problem. To define it algebraically we use a matrix Z that assigns the features or the samples to the Q groups, *i.e.* $Z_{iq} = 1$ if feature or sample i is in group \mathcal{G}_q and 0 otherwise. Denoting $V = (v_1, \dots, v_Q)$, the prediction variable is decomposed as $W = ZV$ leading to the supervised learning problem with clustering constraint

$$\begin{aligned} & \text{minimize} && \mathbf{Loss}(y, X, W) + R(W) \\ & \text{subject to} && W = ZV, Z \in \{0, 1\}^{m \times Q}, Z\mathbf{1} = \mathbf{1}, \end{aligned} \tag{1}$$

in variables W , V and Z whose dimensions depend on whether features ($m = d$) or samples ($m = n$) are clustered.

Although this formulation is non-convex, we observe that the core non-convexity emerges from a clustering problem on the predictors W , which we can deal with using k-means approximations, as detailed in Section 3. We now present in more details two key applications of our formulation: dimensionality reduction by clustering features and learning experts by grouping samples. We only detail regression formulations, extensions for classification are given in the Appendix 7.1. Our framework also applies to clustered multitask as a regularization hypothesis, and we refer the reader to the Appendix 7.2 for more details on this formulation.

2.1. Dimensionality reduction: clustering features. Given a prediction task, we want to reduce dimensionality by grouping together features which have a similar influence on the output [Bondell and Reich, 2008], *e.g.* synonyms in a text classification problem. The predictor variable W is here reduced to a single vector, whose coefficients take only a limited number of values. In practice, this amounts to a quantization of the classifier vector, supervised by a learning loss.

Our objective is to form Q groups of features $\mathcal{G}_1, \dots, \mathcal{G}_Q$, assigning a unique weight v_q to all features in group \mathcal{G}_q . In other words, we search a predictor $w \in \mathbb{R}^d$ such that $w_j = v_q$ for all $j \in \mathcal{G}_q$. This problem can be written

$$\begin{aligned} & \text{minimize} && \frac{1}{n} \sum_{i=1}^n \text{loss}(y_i, w^T x_i) + \frac{\lambda}{2} \|w\|_2^2 \\ & \text{subject to} && w = Zv, Z \in \{0, 1\}^{d \times Q}, Z\mathbf{1} = \mathbf{1}, \end{aligned} \quad (2)$$

in the variables $w \in \mathbb{R}^d$, $v \in \mathbb{R}^Q$ and Z . In what follows, $\text{loss}(y_i, w^T x_i)$ will be a squared or logistic loss that measures the quality of prediction for each sample. Regularization can either be seen as a standard l_2 regularization on w with $R(w) = \frac{\lambda}{2} \|w\|_2^2$, or a weighted regularization on v , $R(v) = \frac{\lambda}{2} \sum_{q=1}^Q s_q \|v_q\|_2^2$.

Note that fused lasso [Tibshirani et al., 2005] in dimension one solves a similar problem that also quantizes the regression vector using an ℓ_1 penalty on coefficient differences. The crucial difference with our setting is that fused lasso assumes that the variables are ordered and minimizes the total variation of the coefficient vector. Here we do not make any ordering assumption on the regression vector.

2.2. Learning experts: clustering samples. Mixture of experts [Jordan, 1994] is a standard model for prediction that seeks to learn Q predictors called “experts”, each predicting labels for a different group of samples. For a new sample x the prediction is then given by a weighted sum of the predictions of all experts $\hat{y} = \sum_{q=1}^Q p_q v_q^T x$. The weights p_q are given by a prior probability depending on x . Here, we study a slightly different setting where we also learn Q experts, but assignments to groups are only extracted from the labels y and not based on the feature variables x as illustrated by the graphical model in Figure 1.

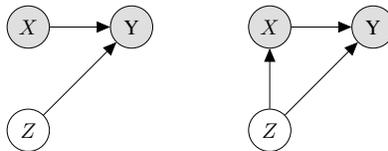


FIGURE 1. Learning multiple diverse experts (*left*), mixture of experts model (*right*). The assignment matrix Z gives the assignment to groups, grey variables are observed, arrows represent dependance of variables.

This means that while we learn several experts (classifier vectors), the information contained in the features x is not sufficient to select the best experts. Given a new point x we can only give Q diverse answers or an approximate weighted prediction $\hat{y} = \sum_{q=1}^Q \frac{s_q}{n} v_q^T x$. Our algorithm will thus return several answers and minimizes the loss of the best of these answers. This setting was already studied by Zhang [2003] for general predictors, it is also related to subspace clustering [Elhamifar and Vidal, 2009], however here we already know in which dimension the data points lie.

Given a prediction task, our objective is to find Q groups $\mathcal{G}_1, \dots, \mathcal{G}_Q$ of sample points to maximize within-group prediction performance. Within each group \mathcal{G}_q , samples are predicted using a common linear predictor v_q . Our problem can be written

$$\text{minimize} \frac{1}{n} \sum_{q=1}^Q \sum_{i \in \mathcal{G}_q} \text{loss}(y_i, v_q^T x_i) + \frac{\lambda}{2} \sum_{q=1}^Q s_q \|v_q\|_2^2 \quad (3)$$

in the variables $V = (v_1, \dots, v_Q) \in \mathbb{R}^{d \times Q}$ and $\mathcal{G} = (\mathcal{G}_1, \dots, \mathcal{G}_Q)$ such that \mathcal{G} is a partition of the n samples. As in the problem of clustering features above, $\text{loss}(y_i, v_q^T x_i)$ measures the quality of prediction for each sample and $R(V) = \frac{\lambda}{2} \sum_{q=1}^Q s_q \|v_q\|_2^2$ is a weighted regularization. Using an assignment matrix $Z \in \{0, 1\}^{n \times Q}$ and an auxiliary variable $W = (w_1, \dots, w_n) \in \mathbb{R}^{d \times n}$ such that $W = VZ^T$, which means

$w_i = v_q$ if $i \in \mathcal{G}_q$, problem (3) can be rewritten

$$\begin{aligned} & \text{minimize} && \frac{1}{n} \sum_{i=1}^n \mathbf{loss}(y_i, w_i^T x_i) + \frac{\lambda}{2} \sum_{i=1}^n \|w_i\|_2^2 \\ & \text{subject to} && W^T = ZV^T, Z \in \{0, 1\}^{n \times Q}, Z\mathbf{1} = \mathbf{1}, \end{aligned}$$

in the variables $W \in \mathbb{R}^{d \times n}$, $V \in \mathbb{R}^{d \times Q}$ and Z . Once again, our problem fits in the general formulation given in (1) and in the sections that follows, we describe several algorithms to solve this problem efficiently.

3. APPROXIMATION ALGORITHMS

We now present optimization strategies to solve learning problems with clustering constraints. We begin by simple greedy procedures and a more refined convex relaxation solved using approximate conditional gradient. We will show that this latter relaxation is exact in the case of feature clustering because the inner one dimensional clustering problem can be solved exactly by dynamic programming.

3.1. Greedy algorithms. For both clustering problems discussed above, greedy algorithms can be derived to handle the clustering objective. A straightforward strategy to group features is to first train predictors as in a classical supervised learning problem, and then cluster weights together using k-means. In the same spirit, when clustering sample points, one can alternate minimization on the predictors of each group and assignment of each point to the group where its loss is smallest. These methods are fast but unstable and highly dependent on initialization. However, alternating minimization can be used to refine the solution of the more robust algorithms proposed below.

3.2. Convex relaxation using conditional gradient algorithm. Another approach is to relax the problem by considering the convex hull of the feasible set and use the conditional gradient method (a.k.a. Frank-Wolfe, [Frank and Wolfe, 1956; Jaggi, 2013]) on the relaxed convex problem. Provided that an affine minimization oracle can be computed efficiently, the key benefit of using this method when minimizing a convex objective over a non-convex set is that it automatically solves a convex relaxation, i.e. minimizes the convex objective over the convex hull of the feasible set, without ever requiring this convex hull to be formed explicitly.

In our case, the convex hull of the set $\{W : W = ZV, Z \in \{0, 1\}^{m \times Q}, Z\mathbf{1} = \mathbf{1}\}$ is the entire space so the relaxed problem loses the initial clustering structure. However in the special case of a squared loss, i.e. $\mathbf{loss}(y, \hat{y}) = \frac{1}{2}(y - \hat{y})^2$, minimization in V can be performed analytically and our problem reduces to a clustering problem for which this strategy is relevant. We illustrate this simplification in the case of clustering features for a regression task, detailed computations and explicit procedures for other settings are given in Appendix 7.3.

Replacing $w = Zv$ in (2), the objective function in problem (2) becomes

$$\begin{aligned} \phi(v, Z) &= \frac{1}{2n} \sum_{i=1}^n (y_i - (Zv)^T x_i)^2 + \frac{\lambda}{2} \|Zv\|_2^2 \\ &= \frac{1}{2n} v^T Z^T X^T X Z v + \frac{\lambda}{2} v^T Z^T Z v - \frac{1}{n} y^T X Z v + \frac{1}{2n} y^T y. \end{aligned}$$

Minimizing in v and using the Sherman-Woodbury-Morrison formula we then get

$$\begin{aligned} \min_v \phi(v, Z) &= \frac{1}{2n} y^T (\mathbf{I} - XZ(Z^T X^T X Z + \lambda n Z^T Z)^{-1} Z^T X^T) y \\ &= \frac{1}{2n} y^T \left(\mathbf{I} + \frac{1}{n\lambda} XZ(Z^T Z)^{-1} Z^T X^T \right)^{-1} y, \end{aligned}$$

and the resulting clustering problem is then formulated in terms of the normalized equivalence matrix

$$M = Z(Z^T Z)^{-1} Z^T$$

such that $M_{ij} = 1/s_q$ if item i and j are in the same group \mathcal{G}_q and 0 otherwise.

Writing $\mathcal{M} = \{M : M = Z(Z^T Z)^{-1} Z^T, Z \in \{0, 1\}^{d \times Q}, Z\mathbf{1} = \mathbf{1}\}$ the set of equivalence matrices for partitions into at most Q groups, our partitioning problem can be written

$$\begin{aligned} & \text{minimize} && \psi(M) \triangleq y^T \left(\mathbf{I} + \frac{1}{n\lambda} X M X^T \right)^{-1} y \\ & \text{subject to} && M \in \mathcal{M}. \end{aligned}$$

in the matrix variable $M \in \mathbf{S}_n$. We now relax this last problem by solving it (implicitly) over the convex hull of the set of equivalence matrices using the conditional gradient method. Its generic form is described in Algorithm (1), where the scalar product is the canonical one on matrices, *i.e.* $\langle A, B \rangle = \text{Tr}(A^T B)$. At each iteration, the algorithm requires solving an linear minimization oracle over the feasible set. This gives the direction for the next step and an estimated gap to the optimum which is used as stopping criterion.

Algorithm 1 Conditional gradient algorithm

Initialize $M_0 \in \mathcal{M}$

for $t = 0, \dots, T$ **do**

 Solve linear minimization oracle

$$\Delta_t = \underset{N \in \text{hull}(\mathcal{M})}{\text{argmin}} \langle N, \nabla \psi(M_t) \rangle \quad (4)$$

if $\text{gap}(M_t, M_*) \leq \epsilon$ **then**

return M_t

else

 Set $M_{t+1} = M_t + \alpha_t(\Delta_t - M_t)$

end if

end for

The estimated gap is given by the linear oracle as

$$\text{gap}(M_t, M_*) \triangleq -\langle \Delta_t - M_t, \nabla \psi(M_t) \rangle.$$

By definition of the oracle and convexity of the objective function, we have

$$-\langle \Delta_t - M_t, \nabla \psi(M_t) \rangle \geq -\langle M_* - M_t, \nabla \psi(M_t) \rangle \geq \psi(M_t) - \psi(M_*).$$

Crucially here, the linear minimization oracle in (4) is equivalent to a projection step. This projection step is itself equivalent to a k-means clustering problem which can be solved exactly in the feature clustering case and well approximated in the other scenarios detailed in the appendix. For a fixed matrix $M \in \text{hull}(\mathcal{M})$, we have that

$$P \triangleq -\nabla \psi(M) = \frac{1}{2n^2\lambda} X^T \left(\mathbf{I} + \frac{1}{n\lambda} X M X^T \right)^{-1} y y^T \left(\mathbf{I} + \frac{1}{n\lambda} X M X^T \right)^{-1} X$$

is positive semidefinite (this is the case for all the settings considered in this paper). Writing $P^{\frac{1}{2}}$ its matrix square root we get

$$\begin{aligned} \underset{N \in \text{hull}(\mathcal{M})}{\text{argmin}} \langle N, \nabla \psi(M) \rangle &= \underset{N \in \mathcal{M}}{\text{argmin}} \text{Tr}(N^T \nabla \psi(M)) \\ &= \underset{N \in \mathcal{M}}{\text{argmin}} -\text{Tr}(N P^{\frac{1}{2}} P^{\frac{1}{2}T}) \\ &= \underset{N \in \mathcal{M}}{\text{argmin}} \text{Tr}((\mathbf{I} - N) P^{\frac{1}{2}} P^{\frac{1}{2}T}) \\ &= \underset{N \in \mathcal{M}}{\text{argmin}} \|P^{\frac{1}{2}} - N P^{\frac{1}{2}}\|_F^2 \\ &= \underset{Z}{\text{argmin}} \min_V \|P^{\frac{1}{2}} - ZV\|_F^2, \end{aligned}$$

because N is an orthonormal projection ($N^2 = N$, $N^T = N$) and so is $(I - N)$. Given a matrix W , we also have

$$\operatorname{argmin}_{Z, V} \|W - ZV\|_F^2 = \operatorname{argmin} \sum_{q=1}^Q \sum_{i \in \mathcal{G}_q} \|w_i - v_q\|_2^2, \quad (5)$$

where the minimum is taken over centroids v_q and partition $(\mathcal{G}_1, \dots, \mathcal{G}_Q)$. This means that computing the linear minimization oracle on $\nabla\psi(M)$ is equivalent to solving a k-means clustering problem on $P^{1/2}$. This k-means problem can itself be solved approximately using the k-means++ algorithm which performs alternate minimization on the assignments and the centroids after an appropriate random initialization. Although this is a non-convex subproblem, k-means++ guarantees a constant approximation ratio on its solution [Arthur and Vassilvitskii, 2007]. We write $\text{k-means}(V, Q)$ the approximate solution of the projection. Overall, this means that the linear minimization oracle (4) can therefore be computed approximately. Moreover, in the particular case of grouping features for regression, the k-means subproblem is one-dimensional and can be solved exactly using dynamic programming [Bellman, 1973; Wang and Song, 2011] so that convergence of the algorithm is ensured.

The complete method is described as Algorithm 2 where we use the classical stepsize for conditional gradient $\alpha_t = \frac{2}{t+2}$. A feasible solution for the original non-convex problem is computed from the solution of the relaxed problem using Frank-Wolfe rounding, *i.e.* output the last linear oracle.

Algorithm 2 Conditional gradient on the equivalence matrix

Input: X, y, Q, ϵ

Initialize $M_0 \in \mathcal{M}$

for $t = 0, \dots, T$ **do**

 Compute the matrix square root $P^{\frac{1}{2}}$ of $-\nabla\psi(M_0)$

 Get oracle $\Delta_t = \text{k-means}(P^{\frac{1}{2}}, Q)$

if $-\text{Tr}(\Delta_t - M_t)^T \nabla\psi(M_t) \leq \epsilon$ **then**

return M_t

else

 Set $M_{t+1} = M_t + \alpha_t(\Delta_t - M_t)$

end if

end for

Z^* is given by the last k-means

V^* is given by the analytic solution of the minimization for Z^* fixed

Output: V^*, Z^*

3.3. Complexity. The core complexity of Algorithm 2 is concentrated in the inner k-means subproblem, which standard alternating minimization approximates at cost $O(t_K Q p)$, where t_K is the number of alternating steps, Q is the number of clusters, and p is the product of the dimensions of V . However, computation of the gradient requires to invert matrices and to compute a matrix square root of the gradient at each iteration, which can slow down computations for large datasets. The choice of the number of clusters can be done given an a priori on the problem (*e.g.* knowing the number of hidden groups in the sample points), or cross-validation, idem for the other regularization parameters.

4. PROJECTED GRADIENT ALGORITHM

In practice, convergence of the conditional gradient method detailed above can be quite slow and we also study a projected gradient algorithm to tackle the generic problem in (1). Although simple and non-convex in general, this strategy used in the context of sparsity can produce scalable and convergent algorithms in certain scenarios, as we will see below.

4.1. **Projected gradient.** We can exploit the fact that projecting a matrix W on the feasible set

$$\{\tilde{W} : \tilde{W} = ZV, Z \in \{0, 1\}^{m \times Q}, Z\mathbf{1} = \mathbf{1}\}$$

is equivalent to a clustering problem, with

$$\operatorname{argmin}_{Z, V} \|W - ZV\|_F^2 = \operatorname{argmin} \sum_{q=1}^Q \sum_{i \in \mathcal{G}_q} \|w_i - v_q\|_2^2,$$

where the minimum is taken over centroids v_q and partition $(\mathcal{G}_1, \dots, \mathcal{G}_Q)$. The k-means problem can be solved approximately with the k-means++ algorithm as mentioned in Section 3.2. We will analyze this algorithm for clustering features for regression in which the projection can be found exactly. Writing $\text{k-means}(V, Q)$ the approximate solution of the projection, ϕ the objective function and α_t the stepsize, the full method is summarized as Algorithm 3 and its implementation is detailed in Section 4.3.

Algorithm 3 Proj. Gradient Descent

Input: X, y, Q, ϵ

Initialize $W_0 = 0$

while $|\phi(W_t) - \phi(W_{t-1})| \geq \epsilon$ **do**

$$W_{t+\frac{1}{2}} = W_t - \alpha_t(\nabla \text{Loss}(y, X, W_t) + \nabla R(W_t))$$

$$[Z_{t+1}, V_{t+1}] = \text{k-means}(W_{t+\frac{1}{2}}, Q)$$

$$W_{t+1} = Z_{t+1}V_{t+1}$$

end while

Z^* and V^* are given through k-means

Output: W^*, Z^*, V^*

4.2. **Convergence.** We now analyze the convergence of the projected gradient algorithm with a constant stepsize $\alpha_t = 1$, applied to the feature clustering problem for regression. We focus on a problem with squared loss without regularization term, which reads

$$\begin{aligned} & \text{minimize} && \frac{1}{2n} \|Xw - y\|_2^2 \\ & \text{subject to} && w = Zv, Z \in \{0, 1\}^{d \times Q}, Z\mathbf{1} = \mathbf{1} \end{aligned}$$

in the variables $w \in \mathbb{R}^d$, $v \in \mathbb{R}^Q$ and Z . We assume that the regression values y are generated by a linear model whose coefficients w^* satisfy the constraints above, up to additive noise, with

$$y = Xw^* + \eta$$

where $\eta \sim \mathcal{N}(0, \sigma^2)$. Hence we study convergence of our algorithm to w^* , *i.e.* to the partition \mathcal{G}^* of its coefficients and its Q values.

We will exploit the fact that each partition \mathcal{G} defines a subspace of vectors w , so the feasible set can be written as a union of subspaces. Let \mathcal{G} be a partition and define

$$\mathcal{U}_{\mathcal{G}} = \{w : w = Zv, Z \in \mathcal{Z}(\mathcal{G})\},$$

where $\mathcal{Z}(\mathcal{G})$ is the set of assignment matrices corresponding to \mathcal{G} . Since permuting the columns of Z together with the coefficients of v has no impact on w , the matrices in $\mathcal{Z}(\mathcal{G})$ are identical up to a permutation of their columns. So, for $Z \in \mathcal{Z}(\mathcal{G})$, $\mathcal{Z}(\mathcal{G}) = \{Z\Pi, \Pi \text{ permutation matrix}\}$, therefore $\mathcal{U}_{\mathcal{G}}$ is a subspace and the corresponding assignment matrices are its different basis.

To a feasible vector w , we associate the partition \mathcal{G} of its values that has the least number of groups. This partition and its corresponding subspace are uniquely defined and, denoting \mathcal{P} the set of partitions in at most Q clusters, our problem (6) can thus be written

$$\begin{aligned} & \text{minimize} && \frac{1}{2n} \|Xw - y\|_2^2 \\ & \text{subject to} && w \in \bigcup_{\mathcal{G} \in \mathcal{P}} \mathcal{U}_{\mathcal{G}}. \end{aligned}$$

where the variable $w \in \mathbb{R}^d$ belongs to a union of subspaces $\mathcal{U}_{\mathcal{G}}$.

We will write the projected gradient algorithm for (6) as a fixed point algorithm whose contraction factor depends on the singular values of the design matrix X on collections of subspaces generated by the partitions \mathcal{G} . We only need to consider largest subspaces in terms of inclusion order, which are the ones generated by the partitions into exactly Q groups. Denoting \mathcal{P}_Q this set of partitions, the collections of subspaces are defined as

$$\begin{aligned}\mathcal{E}_1 &= \{\mathcal{U}_{\mathcal{G}}, \mathcal{G} \in \mathcal{P}_Q\}, \\ \mathcal{E}_2 &= \{\mathcal{U}_{\mathcal{G}_1} + \mathcal{U}_{\mathcal{G}_2}, (\mathcal{G}_1, \mathcal{G}_2) \in \mathcal{P}_Q\}, \\ \mathcal{E}_3 &= \{\mathcal{U}_{\mathcal{G}_1} + \mathcal{U}_{\mathcal{G}_2} + \mathcal{U}_{\mathcal{G}_3}, (\mathcal{G}_1, \mathcal{G}_2, \mathcal{G}_3) \in \mathcal{P}_Q\}.\end{aligned}$$

Our main convergence result follows. Provided that the contraction factor is sufficient, it states the convergence of the projected gradient scheme to the original vector up to a constant error of the order of the noise.

Proposition 4.1. *Given that projection on $\bigcup_{\mathcal{G} \in \mathcal{P}} \mathcal{U}_{\mathcal{G}}$ is well defined, the projected gradient algorithm applied to (3) converges to the original w^* as*

$$\|w^* - w_t\|_2 \leq \rho^t \|w^*\|_2 + \frac{1 - \rho^t}{1 - \rho} \nu \|\eta\|_2,$$

where

$$\begin{aligned}\rho &\triangleq 2 \max_{\mathcal{U} \in \mathcal{E}_3} \|I - \frac{1}{n} \Pi_{\mathcal{U}}^T X^T X \Pi_{\mathcal{U}}\|_2 \\ \nu &\triangleq \frac{2}{n} \max_{\mathcal{U} \in \mathcal{E}_2} \|X \Pi_{\mathcal{U}}\|_2\end{aligned}$$

and $\Pi_{\mathcal{U}}$ is any orthonormal basis of the subspace \mathcal{U} .

Proof. To describe the algorithm we define \mathcal{G}_t and \mathcal{G}_* as the partitions associated respectively with w_t and w^* containing the least number of groups and

$$\begin{cases} w_{t+1/2} &= w_t - \nabla \mathbf{Loss}(X, y, w_t) = w_t - \frac{1}{n} X^T X (w_t - w^*) + \frac{1}{n} X^T \eta \\ w_{t+1} &= \operatorname{argmin}_{w \in \bigcup_{\mathcal{G} \in \mathcal{P}} \mathcal{U}_{\mathcal{G}}} \|w - w_{t+1/2}\|_2^2 \\ \mathcal{U}_t &= \mathcal{U}_{\mathcal{G}_t} \\ \mathcal{U}_{t,*} &= \mathcal{U}_{\mathcal{G}_t} + \mathcal{U}_{\mathcal{G}_*} \\ \mathcal{U}_{t,t+1,*} &= \mathcal{U}_{\mathcal{G}_t} + \mathcal{U}_{\mathcal{G}_{t+1}} + \mathcal{U}_{\mathcal{G}_*}.\end{cases}$$

Orthonormal projections on \mathcal{U}_t , $\mathcal{U}_{t,*}$ and $\mathcal{U}_{t,t+1,*}$ are given respectively by $P_t, P_{t,*}, P_{t,t+1,*}$. Therefore by definition $w_t \in \mathcal{U}_t$, $(w_t, w^*) \in \mathcal{U}_{t,*}$ and $(w_t, w_{t+1}, w^*) \in \mathcal{U}_{t,t+1,*}$.

We can now control convergence, with

$$\begin{aligned}\|w^* - w_{t+1}\|_2 &= \|P_{t+1,*}(w^* - w_{t+1})\|_2 \\ &\leq \|P_{t+1,*}(w^* - w_{t+1/2})\|_2 + \|P_{t+1,*}(w_{t+1/2} - w_{t+1})\|_2.\end{aligned}\tag{6}$$

In the second term, as $w^* \in \bigcup_{\mathcal{G} \in \mathcal{P}} \mathcal{U}_{\mathcal{G}}$ and $w_{t+1} = \operatorname{argmin}_{w \in \bigcup_{\mathcal{G} \in \mathcal{P}} \mathcal{U}_{\mathcal{G}}} \|w - w_{t+1/2}\|_2$, we have

$$\|w_{t+1} - w_{t+1/2}\|_2^2 \leq \|w^* - w_{t+1/2}\|_2^2$$

which is equivalent to

$$\|P_{t+1,*}(w_{t+1} - w_{t+1/2})\|_2^2 + \|(I - P_{t+1,*})w_{t+1/2}\|_2^2 \leq \|P_{t+1,*}(w^* - w_{t+1/2})\|_2^2 + \|(I - P_{t+1,*})w_{t+1/2}\|_2^2$$

and this last statement implies

$$\|P_{t+1,*}(w_{t+1} - w_{t+1/2})\|_2 \leq \|P_{t+1,*}(w^* - w_{t+1/2})\|_2.$$

This means that we get from (6)

$$\begin{aligned}
\|w^* - w_{t+1}\|_2 &\leq 2\|P_{t+1,*}(w^* - w_{t+1/2})\|_2 \\
&= 2\|P_{t+1,*}(w^* - w_t - \frac{1}{n}X^T X(w^* - w_t) - \frac{1}{n}X^T \eta)\|_2 \\
&\leq 2\|P_{t+1,*}(I - \frac{1}{n}X^T X)(w^* - w_t)\|_2 + \frac{2}{n}\|P_{t+1,*}(X^T \eta)\|_2 \\
&= 2\|P_{t+1,*}(I - \frac{1}{n}X^T X)P_{t,*}(w^* - w_t)\|_2 + \frac{2}{n}\|P_{t+1,*}(X^T \eta)\|_2 \\
&\leq 2\|P_{t+1,*}(I - \frac{1}{n}X^T X)P_{t,*}\|_2\|w^* - w_t\|_2 + \frac{2}{n}\|P_{t+1,*}X^T\|_2\|\eta\|_2.
\end{aligned}$$

Now, assuming

$$2\|P_{t+1,*}(I - \frac{1}{n}X^T X)P_{t,*}\|_2 \leq \rho \quad (7)$$

$$\frac{2}{n}\|P_{t+1,*}X^T\|_2 \leq \nu \quad (8)$$

and summing the latter inequality over t , using that $w_0 = 0$, we get

$$\|w^* - w_t\|_2 \leq \rho^t\|w^*\|_2 + \frac{1 - \rho^t}{1 - \rho}\nu\|\eta\|_2.$$

We bound ρ and ν using the information of X on all possible subspaces of \mathcal{E}_2 or \mathcal{E}_3 . For a subspace $\mathcal{U} \in \mathcal{E}_2$ or \mathcal{E}_3 , we define $P_{\mathcal{U}}$ the orthonormal projection on it and $\Pi_{\mathcal{U}}$ any orthonormal basis of it. For (8) we get

$$\|P_{t+1,*}X^T\|_2 = \|XP_{t+1,*}\|_2 \leq \max_{\mathcal{U} \in \mathcal{E}_2} \|XP_{\mathcal{U}}\|_2 = \max_{\mathcal{U} \in \mathcal{E}_2} \|X\Pi_{\mathcal{U}}\|_2,$$

which is independent of the choice of $\Pi_{\mathcal{U}}$.

For (7), using that $\mathcal{U}_{t,*} \subset \mathcal{U}_{t,t+1,*}$ and $\mathcal{U}_{t+1,*} \subset \mathcal{U}_{t,t+1,*}$, we have

$$\begin{aligned}
\|P_{t+1,*}(I - X^T X)P_{t,*}\|_2 &\leq \|P_{t,t+1,*}(I - \frac{1}{n}X^T X)P_{t,t+1,*}\|_2 \\
&\leq \max_{\mathcal{U} \in \mathcal{E}_3} \|P_{\mathcal{U}}(I - \frac{1}{n}X^T X)P_{\mathcal{U}}\|_2 \\
&= \max_{\mathcal{U} \in \mathcal{E}_3} \|\Pi_{\mathcal{U}}(I - \frac{1}{n}\Pi_{\mathcal{U}}^T X^T X \Pi_{\mathcal{U}})\Pi_{\mathcal{U}}^T\|_2 \\
&= \max_{\mathcal{U} \in \mathcal{E}_3} \|I - \frac{1}{n}\Pi_{\mathcal{U}}^T X^T X \Pi_{\mathcal{U}}\|_2,
\end{aligned}$$

which is independent of the choice of $\Pi_{\mathcal{U}}$ and yields the desired result. ■ We now show that ρ and ν derive

from bounds on the singular values of X on the collections \mathcal{E}_2 and \mathcal{E}_3 . Denoting $s_{\min}(A)$ and $s_{\max}(A)$ respectively the smallest and largest singular values of a matrix A , we have

$$\max_{\mathcal{U} \in \mathcal{E}_2} \|X\Pi_{\mathcal{U}}\|_2 = \max_{\mathcal{U} \in \mathcal{E}_2} s_{\max}(X\Pi_{\mathcal{U}}),$$

and assuming $\mathcal{U} \in \mathcal{E}_3$ and that

$$1 - \delta \leq s_{\min}\left(\frac{X\Pi_{\mathcal{U}}}{\sqrt{n}}\right) \leq s_{\max}\left(\frac{X\Pi_{\mathcal{U}}}{\sqrt{n}}\right) \leq 1 + \delta,$$

for some $\delta > 0$, then [Vershynin, 2010, Lemma 5.38] shows

$$\|I - \frac{1}{n}\Pi_{\mathcal{U}}^T X^T X \Pi_{\mathcal{U}}\|_2 \leq 3 \max\{\delta, \delta^2\}.$$

We now show that for isotropic independent sub-Gaussian data x_i , these singular values depend on the number of subspaces of \mathcal{E}_1 , N , their dimension D and the number of samples n . This proposition reformulates results of [Vershynin \[2010\]](#) to exploit the union of subspace structure.

Proposition 4.2. *Let $\mathcal{E}_1, \mathcal{E}_2, \mathcal{E}_3$ be the finite collections of subspaces defined above, let $D = \max_{\mathcal{U} \in \mathcal{E}_1} \dim(\mathcal{U})$ and $N = \text{Card}(\mathcal{E}_1)$. Assuming that the rows x_i of the design matrix are n isotropic independent sub-gaussian, we have*

$$\frac{1}{\sqrt{n}} \max_{\mathcal{U} \in \mathcal{E}_2} \|X\Pi_{\mathcal{U}}\|_2 \leq 1 + \delta_2 + \epsilon \quad \text{and} \quad \max_{\mathcal{U} \in \mathcal{E}_3} \left\| I - \frac{1}{n} \Pi_{\mathcal{U}}^T X^T X \Pi_{\mathcal{U}} \right\|_2 \leq 3 \max\{\delta_3 + \epsilon, (\delta_3 + \epsilon)^2\},$$

with probability larger than $1 - \exp(-c\epsilon^2 n)$, where $\delta_p = C_0 \sqrt{\frac{pD}{n}} + \sqrt{\frac{1+p \log(N)}{cn}}$, $\Pi_{\mathcal{U}}$ is any orthonormal basis of \mathcal{U} and C_0, c depend only on the sub-gaussian norm of the x_i .

Proof. Let us fix $\mathcal{U} \in \mathcal{E}_p$, with $p = 2$ or 3 and $\Pi_{\mathcal{U}}$ one of its orthonormal basis. By definition of \mathcal{E}_p , $\dim(\mathcal{U}) \leq pD$. The rows of $X\Pi_{\mathcal{U}}$ are orthogonal projections of the rows of X onto \mathcal{U} , so they are still independent sub-gaussian isotropic random vectors. We can therefore apply [[Vershynin, 2010](#), Theorem 5.39] on $X\Pi_{\mathcal{U}} \in \mathbb{R}^{n \times \dim(\mathcal{U})}$. Hence for any $s \geq 0$, with probability at least $1 - 2 \exp(-cs^2)$, the smallest and largest singular values of the rescaled matrix $\frac{X\Pi_{\mathcal{U}}}{\sqrt{n}}$ written respectively $s_{\min}(\frac{X\Pi_{\mathcal{U}}}{\sqrt{n}})$ and $s_{\max}(\frac{X\Pi_{\mathcal{U}}}{\sqrt{n}})$ are bounded by

$$1 - C_0 \sqrt{\frac{pD}{n}} - \frac{s}{\sqrt{n}} \leq s_{\min} \left(\frac{X\Pi_{\mathcal{U}}}{\sqrt{n}} \right) \leq s_{\max} \left(\frac{X\Pi_{\mathcal{U}}}{\sqrt{n}} \right) \leq 1 + C_0 \sqrt{\frac{pD}{n}} + \frac{s}{\sqrt{n}}, \quad (9)$$

where c and C_0 depend only on the sub-gaussian norm of the x_i . Now taking the union bound on all subsets of \mathcal{E}_p , (9) holds for any $\mathcal{U} \in \mathcal{E}_p$ with probability

$$\begin{aligned} 1 - 2 \binom{N}{p} \exp(-cs^2) &\geq 1 - 2 \left(\frac{eN}{p} \right)^p \exp(-cs^2) \\ &\geq 1 - 2 \exp(1 + p \log(N) - cs^2). \end{aligned}$$

Taking $s = \sqrt{\frac{1+p \log(N)}{c}} + \epsilon \sqrt{n}$, we get for all $\mathcal{U} \in \mathcal{E}_p$,

$$1 - \delta_p - \epsilon \leq s_{\min} \left(\frac{X\Pi_{\mathcal{U}}}{\sqrt{n}} \right) \leq s_{\max} \left(\frac{X\Pi_{\mathcal{U}}}{\sqrt{n}} \right) \leq 1 + \delta_p + \epsilon,$$

with probability at least $1 - \exp(-c\epsilon^2 n)$, where $\delta_p = C_0 \sqrt{\frac{pD}{n}} + \sqrt{\frac{1+p \log(N)}{cn}}$. Therefore

$$\frac{1}{\sqrt{n}} \max_{\mathcal{U} \in \mathcal{E}_2} \|X\Pi_{\mathcal{U}}\|_2 \leq 1 + \delta_2 + \epsilon.$$

Then [[Vershynin, 2010](#), Theorem 5.39] yields

$$\max_{\mathcal{U} \in \mathcal{E}_3} \left\| I - \frac{1}{n} \Pi_{\mathcal{U}}^T X^T X \Pi_{\mathcal{U}} \right\|_2 \leq 3 \max\{\delta_3 + \epsilon, (\delta_3 + \epsilon)^2\},$$

hence the desired result. ■

Overall here, Proposition 4.1 shows that the projected gradient method converges when the contraction factor ρ is strictly less than one. When observations x_i are isotropic independent sub-gaussian, this means

$$C_0 \sqrt{\frac{3D}{n}} < \frac{1}{3} \quad \text{and} \quad \sqrt{\frac{1+3 \log(N)}{cn}} < \frac{1}{3}$$

which is also

$$n = \Omega(D) \quad \text{and} \quad n = \Omega(\log(N)) \quad (10)$$

The first condition in (10) means that subspaces must be low-dimensional, in our case $D = 3Q$ and we naturally want the number of groups Q to be small. The second condition in (10) means that the structure (clustering here) is restrictive enough, *i.e.* that the number of possible configurations, N , is small enough.

As we show below, in the simple clustering case however, this number of subspaces is quite large, growing essentially as Q^d .

Proposition 4.3. *The number of subspaces N in \mathcal{E}_1 is lower bounded by*

$$N \geq Q^{d-Q}$$

Proof. \mathcal{E}_1 is indexed by the number of partitions in exactly Q clusters, *i.e.* the Stirling number of second kind $\left\{ \begin{smallmatrix} d \\ Q \end{smallmatrix} \right\}$. Standard bounds on the Stirling number of the second kind give

$$\frac{1}{2}(Q^2 + Q + 2)Q^{d-Q-1} - 1 \leq \left\{ \begin{smallmatrix} d \\ Q \end{smallmatrix} \right\} \leq \frac{1}{2}(ed/Q)^Q Q^{d-Q}. \quad (11)$$

hence $N \geq Q^{d-Q}$. ■

This last proposition means that although the intrinsic dimension of our variables is of order $D = 3Q$, the number of subspaces N is such that we need roughly $n \geq 3d \log(Q)$, *i.e.* approximately as many samples as features, so the clustering structure is not specific enough to reduce the number of samples required by our algorithm to converge. On the other hand, given this many samples, the algorithm provably converges to a clustered output, which helps interpretation.

As a comparison, classical sparse recovery problems have the same structure [Rao et al., 2012], as k -sparse vectors for instance can be described as $\{w : w = Zv, Z^T \mathbf{1} = 1\}$ and so are part of a “union of subspaces”. However in the case of sparse vectors the number of subspaces grows as d^k which means recovery requires much less samples than features.

4.3. Implementation and complexity. In our implementation we use a backtracking line search on the stepsize α_t that guarantees decreasing of the objective. At each iteration if

$$\hat{W}_{t+1} = \text{k-means}(W_t - \alpha_t(\nabla \text{Loss}(y, X, W_t) + \nabla R(W_t)), Q)$$

decreases the objective value we take $W_{t+1} = \hat{W}_{t+1}$ and we increase the stepsize by a constant factor $\alpha_{t+1} = a\alpha_t$ with $a > 1$. If \hat{W}_{t+1} increases the objective value we decrease the stepsize by a constant factor $\alpha_t \leftarrow b\alpha_t$, with $b < 1$, output a new \hat{W}_{t+1} and iterate this scheme until \hat{W}_{t+1} decreases the objective value or the stepsize reaches a stopping value ϵ . We observed better results with this line search than with constant stepsize, in particular when the number of samples for clustering features is small.

Complexity of Algorithm 3 is measured by the cost of the projection and the number of iterations until convergence. If approximated by k-means++ the projection step costs $O(t_K Q p)$, where t_K is the number of alternating steps, Q is the number of clusters, and p is the product of the dimensions of V . When clustering features for regression, the dynamic program of Zhang [2003] solving exactly the projection step is in $O(d^2 Q)$ and ours for k -sparse vectors, detailed in Section 5.1, is in $O(k^2 Q)$. We observed convergence of the projected gradient algorithm in less than 100 iterations which makes it highly scalable. As for the convex relaxation the choice of the number of clusters is done given a priori on the problem.

5. SPARSE AND CLUSTERED LINEAR MODELS

Algorithm 3 can also be applied when a sparsity constraint is added to the linear model, provided that the projection is still defined. This scenario arises for instance in text prediction when we want both to select a few relevant words and to group them to reduce dimensionality. Formally the problem of clustering features (2) becomes then

$$\begin{aligned} & \text{minimize} && \frac{1}{n} \sum_{i=1}^n \text{loss}(y_i, w^T x_i) + \frac{\lambda}{2} \|w\|_2^2 \\ & \text{subject to} && w = SZv, S \in \{0, 1\}^{d \times k}, Z \in \{0, 1\}^{k \times Q}, S^T \mathbf{1} = \mathbf{1}, Z \mathbf{1} = \mathbf{1}, \end{aligned}$$

in the variables $w \in \mathbb{R}^d$, $v \in \mathbb{R}^Q$, S and Z , where S is a matrix of k canonical vectors which assigns nonzero coefficients and Z an assignment matrix of k variables in Q clusters.

We develop a new dynamic program to get the projection on k -sparse vectors whose non-zero coefficients are clustered in Q groups and apply our previous theoretical analysis to prove convergence of the projected gradient scheme on random instances.

5.1. Projection on k -sparse Q -clustered vectors. Let \mathcal{W} be the set of k -sparse vectors whose non-zero values can be partitioned in at most Q groups. Given $x \in \mathbb{R}^d$, we are interested in its projection on \mathcal{W} , which we formulate as a partitioning problem. For a feasible $w \in \mathcal{W}$, with $\mathcal{G}_0 = \{i : w_i = 0\}$ and $\mathcal{G}_1, \dots, \mathcal{G}_{Q'}$, with $Q' \leq Q$, the partition of its non-zero values such that $w_i = v_q$ if and only if $i \in \mathcal{G}_q$, the distance between x and w is given by

$$\|x - w\|_2^2 = \sum_{i \in \mathcal{G}_0} x_i^2 + \sum_{q=1}^{Q'} \sum_{i \in \mathcal{G}_q} (x_i - v_q)^2.$$

The projection is solution of

$$\begin{aligned} & \text{minimize} && \sum_{i \in \mathcal{G}_0} x_i^2 + \sum_{q=1}^{Q'} \sum_{i \in \mathcal{G}_q} (x_i - v_q)^2 \\ & \text{subject to} && \text{Card} \left(\bigcup_{q=1}^{Q'} \mathcal{G}_q \right) \leq k, \quad 0 \leq Q' \leq Q, \end{aligned} \quad (12)$$

in the number of groups Q' , the partition $\mathcal{G} = (\mathcal{G}_0, \dots, \mathcal{G}_{Q'})$ of $\{1, \dots, d\}$ and $v \in \mathbb{R}^{Q'}$. For a fixed number of non-zero values k' , the objective is clearly decreasing in the number of groups Q' , which measures the degrees of freedom of the projection, however it cannot exceed k' . We will use this argument below to get the best parameter Q' . For a fixed partition \mathcal{G} , minimization in v gives the barycenters of the Q' groups, $\mu_q = \frac{1}{s_q} \sum_{i \in \mathcal{G}_q} x_i$. Inserting them in (12), the objective can be developed as

$$\sum_{i \in \mathcal{G}_0} x_i^2 + \sum_{q=1}^{Q'} \sum_{i \in \mathcal{G}_q} x_i^2 + \mu_q^2 - 2\mu_q x_i = \sum_{i=1}^d x_i^2 - \sum_{q=1}^{Q'} s_q \mu_q^2.$$

Splitting this objective between positive and negative barycenters, we get that the minimizer of (12) solves

$$\begin{aligned} & \text{maximize} && \sum_{q: \mu_q < 0} s_q \mu_q^2 + \sum_{q: \mu_q > 0} s_q \mu_q^2 \\ & \text{subject to} && \text{Card} \left(\bigcup_{q=1}^{Q'} \mathcal{G}_q \right) \leq k, \quad 0 \leq Q' \leq Q, \end{aligned} \quad (13)$$

in the number of groups Q' and the partition $\mathcal{G} = (\mathcal{G}_0, \dots, \mathcal{G}_{Q'})$ of $\{1, \dots, d\}$, where $\mu_q = \frac{1}{s_q} \sum_{i \in \mathcal{G}_q} x_i$.

We tackle this problem by finding the best balance between the two terms of the objective. We define $f_-(j, q)$ the optimal value of $\sum s_p \mu_p^2$ when picking j points clustered in q groups forming only negative barycenters, *i.e.* the solution of the problem

$$\begin{aligned} & \text{maximize} && \sum_{p=1}^q s_p \mu_p^2 \\ & \text{subject to} && \mu_p = \frac{1}{s_p} \sum_{i \in \mathcal{G}_p} x_i < 0 \\ & && \text{Card} \left(\bigcup_{p=1}^q \mathcal{G}_p \right) = j, \end{aligned} \quad (P_-(j, q))$$

in the partition $\mathcal{G} = \{\mathcal{G}_0, \dots, \mathcal{G}_q\}$ of $\{1, \dots, d\}$. We define $f_+(j, q)$ similarly except that it constraints barycenters to be positive. Using remark above on the parameter Q' , problem (13) is then equivalent to

$$\begin{aligned} & \text{maximize} && f_-(j, q) + f_+(k' - j, Q' - q) \\ & \text{subject to} && 0 \leq k' \leq k, \quad 0 \leq j \leq k', \quad Q' = \min(k', Q), \quad 0 \leq q \leq Q', \end{aligned} \quad (14)$$

in variables j , k' and q .

Now we show that f_- and f_+ can be computed by dynamic programming, we begin with f_- . Remark that $(P_-(j, q))$ is a partitioning problem on the j smallest values of x . To see this, let $S_- \subset \{1, \dots, d\}$ be the optimal subset of indexes taken for $(P_-(j, q))$ and $i \in S_-$. If there exists $j \notin S_-$ such that $x_j \leq x_i$, then

swapping j and i would increase the magnitude of the barycenter of the group that i belongs to and so the objective. Now for $(P_-(j, q))$ a feasible problem, let $\mathcal{G}_1, \dots, \mathcal{G}_q$ be its optimal partition whose corresponding barycenters are in ascending order and x_i be the smallest value of x in \mathcal{G}_q , then necessarily $\mathcal{G}_1, \dots, \mathcal{G}_{q-1}$ is optimal to solve $P_-(i-1, q-1)$. We order therefore the values of x in ascending order and use the following dynamic program to compute f_- ,

$$f_-(j, q) = \max_{\substack{q \leq i \leq j \\ \mu(x_i, \dots, x_j) < 0}} f_-(i-1, q-1) + (j-i+1)\mu(x_i, \dots, x_j)^2, \quad (15)$$

where $\mu(x_i, \dots, x_j) = \frac{1}{j-i+1} \sum_{l=i}^j x_l$ can be computed in constant time using that

$$\mu(x_i, \dots, x_j) = \frac{x_i + (j-i)\mu(x_{i+1}, \dots, x_j)}{j-i+1}.$$

By convention $f_-(j, q) = -\infty$ if (15) and so $(P_-(j, q))$ are not feasible. f_- is initialized as a grid of $k+1$ and $Q+1$ columns such that $f_-(0, q) = 0$ for any q , $f_-(j, 0) = 0$ and $f_-(j, 1) = j\mu(x_1, \dots, x_j)^2$ for any $j \geq 1$. Values of f_- are stored to compute (14). Two auxiliary variables I_- and μ_- store respectively the indexes of the smallest value of x in group \mathcal{G}_q and the barycenter of the group \mathcal{G}_q , defined by

$$\begin{aligned} I_-(j, q) &= \operatorname{argmax}_{\substack{q \leq i \leq j \\ \mu(x_i, \dots, x_j) < 0}} f_-(i-1, q-1) + (j-i+1)\mu(x_i, \dots, x_j)^2, \\ \mu_-(j, q) &= \mu(x_{I_-(j, q)}, \dots, x_j), \quad i = I_-(j, q). \end{aligned}$$

I_- and μ_- are initialized by $I_-(j, 1) = 1$ and $\mu_-(j, 1) = \mu(x_1, \dots, x_j)$. The same dynamic program can be used to compute f_+ , I_+ and μ_+ , defined similarly as I_- and μ_- , by reversing the order of the values of x . A grid search on $f(j, q, k') = f_-(j, q) + f_+(k' - j, Q' - q)$, with $Q' = \min(k', Q)$, gives the optimal balance between positive and negative barycenters. A backtrack on I_- and I_+ finally gives the best partition and the projection with the associated barycenters given in μ_- and μ_+ .

Each dynamic program needs only to build the best partitions for the k smallest or largest partitions so their complexity is in $O(k^2Q)$. The complexity of the grid search is $O(k^2Q)$ and the complexity of the backtrack is $O(Q)$. The overall complexity of the projection is therefore $O(k^2Q)$.

5.2. Convergence. Our theoretical convergence analysis can directly be applied to this setting for a problem with squared loss without regularization. The feasible set is again a union of subspaces

$$w \in \bigcup_{\substack{S \in \{0,1\}^{d \times k}, S^T \mathbf{1} = \mathbf{1} \\ Z \in \{0,1\}^{k \times Q}, Z \mathbf{1} = \mathbf{1}}} \{w : w = SZv\}.$$

However the number of largest subspaces in terms of inclusion order is smaller. They are defined by selecting k features among d and partitioning these k features into Q groups so that their number is $N = \binom{d}{k} \left\{ \begin{matrix} k \\ Q \end{matrix} \right\}$. Using classical bounds on the binomial coefficient and (11), we have for $k \geq 3$, $Q \geq 3$,

$$N \leq d^k k^Q Q^{k-Q}.$$

Our analysis thus predicts that only

$$n \geq 36 \max \left\{ QC_0^2, \frac{1}{c} (k \log d + Q \log(k) + (k-Q) \log(Q)) \right\}$$

isotropic independent sub-Gaussian samples are sufficient for the projected gradient algorithm to converge. It produces $Q+1$ cluster of features, one being a cluster of zero features, reducing dimensionality, while needing roughly as many samples as non-zero features.

6. NUMERICAL EXPERIMENTS

We now test our methods, first on artificial datasets to check their robustness to noisy data. We then test our algorithms for feature clustering on real data extracted from movie reviews. While our approach is general and applies to both features and samples, we observe that our algorithms compare favorably with specialized algorithms for these tasks.

6.1. Synthetic dataset.

6.1.1. *Clustering constraint on sample points.* We test the robustness of our method when the information of the regression problem leading to the partition of the samples lies in a few features. We generate n data points (x_i, y_i) for $i = 1, \dots, n$, with $x_i \in \mathbb{R}^d$, $d = 8$, and $y_i \in \mathbb{R}$, divided in $Q = 3$ clusters corresponding to regression tasks with weight vectors v_q . Regression labels for points x_i in cluster \mathcal{G}_q are given by $y_i = v_q^T x_i + \eta_y$, where $\eta_y \sim \mathcal{N}(0, \sigma_y^2)$. We test the robustness of the algorithms to the addition of noisy dimensions by completing x_i with d_n dimensions of noise $\eta_d \sim \mathcal{N}(0, \sigma_d)$. For testing the models we take the difference between the true label and the best prediction such that the Mean Square Error (MSE) is given by

$$\text{Loss}(y, X, W) = \frac{1}{2n} \sum_{i=1}^n \min_{q=1, \dots, Q} (y_i - v_q^T x_i)^2. \quad (16)$$

The results are reported in Table 1 where the intrinsic dimension is 10 and the proportion of dimensions of noise $d_n/(d + d_n)$ increases. On the algorithmic side, ‘‘Oracle’’ refers to the least-squares fit given the true assignments, which can be seen as the best achievable error rate, AM refers to alternate minimization, PG refers to projected gradient with squared loss, CG refers to conditional gradient and RC to regression clustering as proposed by Zhang [2003], implemented using the Harmonic K-means formulation. PG, CG and RC were followed by AM refinement. 1000 points were used for training, 100 for testing. The regularization parameters were 5-fold cross-validated using a logarithmic grid. Noise on labels is $\sigma_y = 10^{-1}$ and noise on added dimensions is $\sigma_d = 1$. Results were averaged over 50 experiments with figures after the \pm sign corresponding to one standard deviation.

	p = 0	p = 0.25	p = 0.5	p = 0.75	p = 0.9	p = 0.95
Oracle	0.52±0.08	0.55±0.07	0.55±0.10	0.58±0.09	0.71±0.11	1.17±0.18
AM	0.52 ±0.08	0.55 ±0.07	5.57±4.11	6.93±14.39	101.08±55.49	133.48±52.20
PG	1.53±7.13	3.98±17.65	3.20 ±13.23	5.64±20.50	91.33±39.32	131.48 ±50.90
CG	0.87±2.45	1.16±4.29	3.64±11.02	5.43 ±14.33	91.19±53.00	136.57±58.60
RC	0.52 ±0.08	0.55 ±0.07	5.59±20.27	13.45±28.76	59.19 ±37.97	135.77±66.96

TABLE 1. Test MSE given by (16) along proportion of added dimensions of noise $p = d_n/(d + d_n)$.

All algorithms perform similarly, RC and AM get better results without added noise. None of the present algorithms get a significantly better behavior with a majority of noisy dimensions.

6.1.2. *Clustering constraint on features.* We test the robustness of our method when with the number of training samples or the level of noise in the labels. We generate n data points (x_i, y_i) for $i = 1, \dots, n$ with $x_i \in \mathbb{R}^d$, $d = 100$, and $y_i \in \mathbb{R}$. Regression weights w have only 5 different values v_q for $q = 1, \dots, 5$, uniformly distributed around 0. Regression labels are given by $y_i = w^T x_i + \eta$, where $\eta \sim \mathcal{N}(0, \sigma^2)$. We vary the number of samples n or the level of noise σ and measure $\|w_* - \hat{w}\|_2$, the l_2 norm of the difference between the true vector of weights w_* and the estimated ones \hat{w} .

In Table 2 and 3, we compare the proposed algorithms to Least Squares (LS), Least Squares followed by K-means on the weights (using associated centroids as predictors) (LSK) and OSCAR [Bondell and Reich, 2008]. For OSCAR we used a submodular approach [Bach et al., 2012] to compute the corresponding

proximal algorithm, which makes it scalable. “Oracle” refers to the Least Square solution given the true assignments of features and can be seen as the best achievable error rate. Here too, PG refers to projected gradient with squared loss (initialized with the solution of Least Square followed by k-means), CG refers to conditional gradient, CGPG refers to conditional gradient followed by PG. When varying the number of samples, noise on labels is set to $\sigma = 0.5$ and when varying level of noise σ number of samples is set to $n = 150$. Parameters of the algorithms were all cross-validated using a logarithmic grid. Results were averaged over 50 experiments and figures after the \pm sign correspond to one standard deviation.

	$n = 50$	$n = 75$	$n = 100$	$n = 125$	$n = 150$
Oracle	0.16±0.06	0.14±0.04	0.10±0.04	0.10±0.04	0.09±0.03
LS	61.94±17.63	51.94±16.01	21.41±9.40	1.02±0.18	0.70±0.09
LSK	62.93±18.05	57.78±17.03	10.18±14.96	0.31±0.19	0.19±0.12
PG	63.31±18.24	52.72±16.51	5.52±14.33	0.14±0.09	0.09±0.04
CG	61.81±17.78	52.59±16.58	17.24±13.87	1.20±1.38	1.05±1.37
CGPG	62.29±18.15	50.15±17.43	0.64±2.03	0.15±0.19	0.17±0.53
OS	61.54±17.59	52.87±15.90	11.32±7.03	1.25±0.28	0.71±0.10

TABLE 2. Measure of $\|w_* - \hat{w}\|_2$, the l_2 norm of the difference between the true vector of weights w_* and the estimated ones \hat{w} along number of samples n .

	$\sigma = 0.05$	$\sigma = 0.1$	$\sigma = 0.5$	$\sigma = 1$
Oracle	0.86±0.27	1.72±0.54	8.62±2.70	17.19±5.43
LS	7.04±0.92	14.05±1.82	70.39±9.20	140.41±18.20
LSK	1.44±0.46	2.88±0.91	19.10±12.13	48.09±27.46
PG	0.87±0.27	1.74±0.52	9.11±4.00	26.23±18.00
CG	23.91±36.51	122.31±145.77	105.45±136.79	155.98±177.69
CGPG	1.52±3.13	140.83±710.32	17.34±53.31	24.80±16.32
OS	14.43±2.45	18.89±3.46	71.00±10.12	140.33±18.83

TABLE 3. Measure of $\|w_* - \hat{w}\|_2$, the l_2 norm of the difference between the true vector of weights w_* and the estimated ones \hat{w} along level of noise σ .

We observe that both PG and CGPG give significantly better results than other methods and even reach the performance of the Oracle for $n > d$ and for small σ , while for $n \leq d$ results are in the same range.

6.2. Real data.

6.2.1. *Predicting ratings from reviews using groups of words.* We perform “sentiment” analysis of newspaper movie reviews. We use the publicly available dataset introduced by Pang and Lee [2005] which contains movie reviews paired with star ratings. We treat it as a regression problem, taking responses for y in $(0, 1)$ and word frequencies as covariates. The corpus contains $n = 5006$ documents and we reduced the initial vocabulary to $d = 5623$ words by eliminating stop words, rare words and words with small TF-IDF mean on whole corpus. We evaluate our algorithms for regression with clustered features against standard regression approaches: Least-Squares (LS), and Least-Squares followed by k-means on predictors (LSK), Lasso and Iterative Hard Thresholding (IHT). We also tested our projected gradient with sparsity constraint, initialized by the solution of LSK (PGS) or by the solution of CG (CGPGS). Number of clusters, sparsity constraints and regularization parameters were 5-fold cross-validated using respectively grids going from 5 to 15, $d/2$ to $d/5$ and logarithmic grids. Cross validation and training were made on 80% on the dataset and tested on the remaining 20% it gave $Q = 15$ number of clusters and $d/2$ sparsity constraint for our algorithms.

Results are reported in Table 4, figures after the \pm sign correspond to one standard deviation when varying the training and test sets on 20 experiments.

All methods perform similarly except IHT and Lasso whose hypotheses does not seem appropriate for the problem. Our approaches have the benefit to reduce dimensionality from 5623 to 15 and provide meaningful cluster of words. The clusters with highest absolute weights are also the ones with smallest number of words, which confirms the intuition that only a few words are very discriminative. We illustrate this in Table 5, picking randomly words of the four clusters within which associated predictor weights v_q have largest magnitude.

LS	LSK	PG	CG	CGPG	OS
1.51 \pm 0.06	1.53 \pm 0.06	1.52 \pm 0.06	1.58 \pm 0.07	1.49 \pm 0.08	1.47 \pm 0.07
	PGS	CGPGS	IHT	Lasso	
	1.53 \pm 0.06	1.49 \pm 0.07	2.19 \pm 0.12	3.77 \pm 0.17	

TABLE 4. $100 \times$ mean square errors for predicting movie ratings associated with reviews.

First and Second Cluster (negative) sizes 1 and 7	bad, awful, worst, boring, ridiculous, watchable, suppose, disgusting,
Last and Before Last Cluster (positive) sizes 4 and 40	perfect, hilarious, fascinating, great wonderfully, perfectly, goodspirited, world, intelligent, wonderfully, unexpected, gem, recommendation, excellent, rare, unique, marvelous, good-spirited, mature, send, delightful, funniest

TABLE 5. Clustering of words on movie reviews. We show clusters of words within which associated predictor weights v_q have largest magnitude. First and second one are associated to a negative coefficient and therefore bad feelings about movies, last and before last ones to a positive coefficient and good feelings about movies.

Acknowledgements. AA is at CNRS, at the Département d’Informatique at École Normale Supérieure, 2 rue Simone Iff, 75012 Paris, France. FB is at the Département d’Informatique at École Normale Supérieure and INRIA, Sierra project-team, PSL Research University. The authors would like to acknowledge support from a starting grant from the European Research Council (ERC project SIPA), an AMX fellowship, the MSR-Inria Joint Centre, as well as support from the chaire *Économie des nouvelles données*, the *data science* joint research initiative with the *fonds AXA pour la recherche* and a gift from Société Générale Cross Asset Quantitative Research.

REFERENCES

- Andreas Argyriou, Theodoros Evgeniou, and Massimiliano Pontil. Convex multi-task feature learning. *Machine Learning*, 73(3):243–272, 2008.
- David Arthur and Sergei Vassilvitskii. k-means++: The advantages of careful seeding. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 1027–1035. Society for Industrial and Applied Mathematics, 2007.
- Francis Bach, Rodolphe Jenatton, Julien Mairal, and Guillaume Obozinski. Optimization with sparsity-inducing penalties. *Found. Trends Mach. Learn.*, 4(1):1–106, January 2012. ISSN 1935-8237. doi: 10.1561/22000000015.
- Richard Bellman. A note on cluster analysis and dynamic programming. *Mathematical Biosciences*, 18(3): 311–312, 1973.

- David M. Blei, Andrew Y. Ng, and Michael I. Jordan. Latent dirichlet allocation. *J. Mach. Learn. Res.*, 3: 993–1022, March 2003. ISSN 1532-4435.
- Thomas Blumensath and Mike E. Davies. Iterative hard thresholding for compressed sensing. *Applied and Computational Harmonic Analysis*, 27(3):265–274, 2009.
- Howard D. Bondell and Brian J. Reich. Simultaneous regression shrinkage, variable selection, and supervised clustering of predictors with oscar. *Biometrics*, 64(1):115–123, 2008.
- Carlo Ciliberto, Youssef Mroueh, Tomaso Poggio, and Lorenzo Rosasco. Convex learning of multiple tasks and their structure. In *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, pages 1548–1557, 2015.
- Inderjit S. Dhillon, Subramanyam Mallela, and Rahul Kumar. A divisive information theoretic feature clustering algorithm for text classification. *The Journal of Machine Learning Research*, 3:1265–1287, 2003.
- Ehsan Elhamifar and Rene Vidal. Sparse subspace clustering. In *In CVPR*, 2009.
- Marguerite Frank and Philip Wolfe. An algorithm for quadratic programming. *Naval research logistics quarterly*, 3(1-2):95–110, 1956.
- Abner Guzman-Rivera, Pushmeet Kohli, Dhruv Batra, and Rob Rutenbar. Efficiently enforcing diversity in multi-output structured prediction. In *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics*, pages 284–292, 2014.
- Laurent Jacob, Jean-Philippe Vert, and Francis Bach. Clustered multi-task learning: A convex formulation. In *Advances in Neural Information Processing Systems 21*, pages 745–752. 2009.
- Martin Jaggi. Revisiting frank-wolfe: Projection-free sparse convex optimization. In *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, pages 427–435, 2013.
- Jung-Yi Jiang, Ren-Jia Liou, and Shie-Jue Lee. A fuzzy self-constructing feature clustering algorithm for text classification. *Knowledge and Data Engineering, IEEE Transactions on*, 23(3):335–349, 2011.
- Michael I. Jordan. Hierarchical mixtures of experts and the em algorithm. *Neural Computation*, 6:181–214, 1994.
- Bo Pang and Lillian Lee. Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales. In *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics*, pages 115–124. Association for Computational Linguistics, 2005.
- N. Rao, B. Recht, and R. Nowak. Signal Recovery in Unions of Subspaces with Applications to Compressive Imaging. *ArXiv e-prints*, September 2012.
- Robert Tibshirani, Michael Saunders, Saharon Rosset, Ji Zhu, and Keith Knight. Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(1): 91–108, 2005.
- Roman Vershynin. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*, 2010.
- Haizhou Wang and Mingzhou Song. Ckmeans. 1d. dp: optimal k-means clustering in one dimension by dynamic programming. *The R Journal*, 3(2):29–33, 2011.
- Bin Zhang. Regression clustering. In *ICDM*, pages 451–. IEEE Computer Society, 2003. ISBN 0-7695-1978-4.

7. APPENDIX

7.1. Formulations for classification . We present here formulations of clustering either features or samples when our task is to classify samples into K classes. For both settings we assume that n sample points are given, represented by the matrix $X = (x_1, \dots, x_n)^T \in \mathbb{R}^{n \times d}$ and corresponding labels $Y = (y_1, \dots, y_K) \in \{0, 1\}^{n \times K}$.

7.1.1. Clustering features for classification. Here we search K predictors $W = (w_1, \dots, w_K)$, each of them having features clustered in Q groups $\{\mathcal{G}_1, \dots, \mathcal{G}_Q\}$ such that for any k , $w_k^j = v_k^q$ if feature j is in group q . Partition of the features is shared by all predictors but each has different centroids represented in the vector v_k . Using an assignment matrix Z and the matrix of centroids $V = (v_1, \dots, v_K)$, our problem can therefore be written

$$\begin{aligned} & \text{minimize} && \frac{1}{n} \sum_{i=1}^n \mathbf{loss}(y_i, x_i^T W) + \frac{\lambda}{2} \|W\|_F^2 \\ & \text{subject to} && W = ZV, Z \in \{0, 1\}^{d \times Q}, Z\mathbf{1} = \mathbf{1} \end{aligned}$$

in variables $W \in \mathbb{R}^{d \times K}$, $V \in \mathbb{R}^{Q \times K}$ and Z . $\mathbf{loss}(y_i, x_i^T W)$ is a squared or logistic multiclass loss and regularization can either be seen as a standard ℓ_2 regularization on the w_k or a weighted regularization on the centroids v_k .

7.1.2. Clustering samples for classification. Here our objective is to form Q groups $\{\mathcal{G}_1, \dots, \mathcal{G}_Q\}$ of sample points to maximize the within-group prediction performance. For classification, within each group \mathcal{G}_q , samples are predicted using a common matrix of predictors $V^q = (v_1^q, \dots, v_K^q)$. Our problem can be written

$$\text{minimize} \frac{1}{n} \sum_{i \in \mathcal{G}_q} \mathbf{loss}(y_i, x_i^T V^q) + \frac{\lambda}{2} \sum_{q=1}^Q s_q \|V^q\|_F^2 \tag{17}$$

in the variables $V = (V^1, \dots, V^Q) \in \mathbb{R}^{d \times K \times Q}$ and $\mathcal{G} = (\mathcal{G}_1, \dots, \mathcal{G}_Q)$ such that \mathcal{G} is a partition of the n samples. $\mathbf{loss}(y_i, x_i^T W)$ is a squared or logistic multiclass and $\frac{\lambda}{2} \sum_{q=1}^Q s_q \|V^q\|_F^2$ is a weighted regularization. Using an assignment matrix $Z \in \{0, 1\}^{n \times Q}$ and auxiliary variables $(W^1, \dots, W^n) \in \mathbb{R}^{d \times K \times n}$ such that $W^i = V^q$ if $i \in \mathcal{G}_q$, problem (17) can be rewritten

$$\begin{aligned} & \text{minimize} && \frac{1}{n} \sum_{i=1}^n \mathbf{loss}(y_i, W^i x_i) + \frac{\lambda}{2} \sum_{i=1}^n \|W^i\|_F^2 \\ & \text{subject to} && \tilde{W}^T = Z\tilde{V}^T, Z \in \{0, 1\}^{n \times Q}, Z\mathbf{1} = \mathbf{1}, \end{aligned}$$

in the variables $W \in \mathbb{R}^{d \times K \times n}$, $V \in \mathbb{R}^{d \times K \times Q}$ and Z , where $\tilde{W} = (\text{Vec}(W^1), \dots, \text{Vec}(W^n))$, $\tilde{V} = (\text{Vec}(V^1), \dots, \text{Vec}(V^Q))$ and for a matrix A , $\text{Vec}(A)$ concatenates its columns into one vector.

Remark that in that case we must have $K > Q$ otherwise we output more possible answers than classes (in that case the problem is ill-posed).

7.2. Clustered multitask . Our framework applies also to transfer learning by clustering similar tasks. Given a set of K supervised tasks like regression or binary classification, transfer learning aims at jointly solving these tasks, hoping that each task can benefit from the information given by other tasks. For simplicity, we illustrate the case of multi-category classification, which can be extended to the general multitask setting. When performing classification with one-versus-all majority vote, we train one binary classifier for each class vs. all others. Using a regularizing penalty such as the squared ℓ_2 norm, the problem of multitask learning can be cast as

$$\text{minimize} \frac{1}{n} \sum_{k=1}^K \sum_{i=1}^n \mathbf{loss}(y_i^k, w_k^T x_i) + \lambda \sum_{i=1}^K \|w_k\|_2^2. \tag{18}$$

in the matrix variable $W = (w_1, \dots, w_K) \in \mathbb{R}^{d \times K}$ of classifier vectors (one column per task). We write $\mathbf{Loss}(y, X, W)$ and $R(W)$ the first and second term of this problem. Various strategies are used to leverage the information coming from related tasks, such as low rank [Argyriou et al. \[2008\]](#) or structured norm

penalties Ciliberto et al. [2015] on the matrix of classifiers W . Here we follow the clustered multitask setting introduced in Jacob et al. [2009]. Namely we add a penalty Ω on the classifiers (w_1, \dots, w_K) which enforce them to be clustered in Q groups $\mathcal{G}_1, \dots, \mathcal{G}_Q$ around centroids $V = (v_1, \dots, v_Q) \in \mathbb{R}^{d \times Q}$. This penalty can be decomposed in

- A measure of the **norm of the barycenter of centers** $\bar{v} = \frac{1}{K} \sum_{q=1}^Q s_q v_q$

$$\Omega_{mean}(V) = \frac{\lambda_m}{2} K \|\bar{v}\|_2^2$$

- A measure of the **variance between clusters**

$$\Omega_{between}(V) = \frac{\lambda_b}{2} \sum_{q=1}^Q s_q \|v_q - \bar{v}\|_2^2$$

- A measure of the **variance within clusters**

$$\Omega_{within}(W, V) = \frac{\lambda_w}{2} \sum_{q=1}^Q \sum_{i \in \mathcal{G}_q} \|w_i - v_q\|_2^2$$

The total penalty $\Omega(W, V) = \Omega_{mean}(V) + \Omega_{between}(V) + \Omega_{within}(W, V)$ is illustrated in Figure 2.

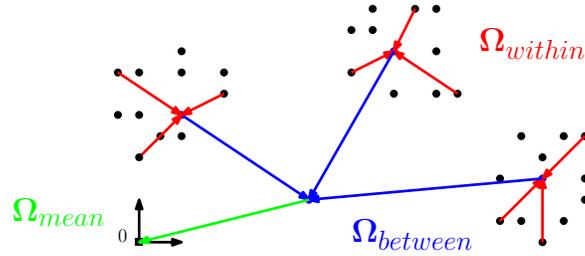


FIGURE 2. Decomposed clustering penalty on K classes in the space of classifier vectors.

The clustered multitask learning problem can then be written using an assignment matrix Z and an auxiliary variable W . Denoting $\Pi = \mathbf{I} - \frac{\mathbf{1}\mathbf{1}^T}{K}$ the centering matrix of the K classes, we develop each term of the penalty,

$$\begin{aligned} \Omega_{mean}(V, Z) &= \frac{\lambda_M}{2} \mathbf{Tr}(VZ^T(\mathbf{I} - \Pi)ZV^T), \\ \Omega_{between}(V, Z) &= \frac{\lambda_B}{2} \mathbf{Tr}(VZ^T\Pi ZV^T), \\ \Omega_{within}(W, V, Z) &= \frac{\lambda_W}{2} \|W - VZ^T\|_F^2. \end{aligned}$$

Using $\tilde{W} = VZ^T$ the total penalty can then be written

$$\Omega(W, \tilde{W}) = \frac{\lambda_M}{2} \mathbf{Tr}(\tilde{W}(\mathbf{I} - \Pi)\tilde{W}^T) + \frac{\lambda_B}{2} \mathbf{Tr}(\tilde{W}\Pi\tilde{W}^T) + \frac{\lambda_W}{2} \|W - \tilde{W}\|_F^2,$$

and the problem is

$$\begin{aligned} &\text{minimize } \mathbf{Loss}(y, X, W) + R(W) + \Omega(W, \tilde{W}) \\ &\text{s.t. } \tilde{W}^T = ZV^T, \quad Z \in \{0, 1\}^{K \times Q}, \quad Z\mathbf{1} = \mathbf{1}, \end{aligned}$$

in variables $W \in \mathbb{R}^{d \times K}$, $\tilde{W} \in \mathbb{R}^{d \times K}$, $V \in \mathbb{R}^{d \times Q}$ and Z .

7.3. Convex relaxations formulations.

7.3.1. *Clustering samples for regression task.* We use a squared loss $l(\hat{y}, y) = \frac{1}{2}(y - \hat{y})^2$ in (3) and minimize in V to get a clustering problem that we can tackle using Frank-Wolfe method. We fix a partition \mathcal{G} and define for each group $\mathcal{G}_q = \{k_1, \dots, k_{s_q}\} \subset \{1, \dots, d\}$, the matrix $E \in \{0, 1\}^{s_q \times n}$ that picks the s_q points of \mathcal{G}_q , i.e. $(E_q)_{ij} = 1$ if $j = k_i$ and 0 otherwise. Therefore $y_q = E_q y \in \mathbb{R}^{s_q}$ and $X_q = E_q X \in \mathbb{R}^{s_q \times d}$ are respectively the vector of labels and the matrix of sample vectors of the group \mathcal{G}_q . We naturally have $E_q E_q^T = \mathbf{I}$ as rows of E_q are orthonormal and $E_q^T E_q$ is a diagonal matrix where $Z_q = \mathbf{diag}(E_q^T E_q) \in \{0, 1\}^n$ is the assignment vector in group \mathcal{G}_q , i.e. $(Z_q)_j = 1$ if $j \in \mathcal{G}_q$ and 0 otherwise. $Z = (Z_1, \dots, Z_Q)$ is therefore an assignment matrix for the partition \mathcal{G} .

Minimizing in v and using the Sherman-Woodbury-Morrison formula, we obtain a function of the partition

$$\begin{aligned} \tilde{\psi}(\mathcal{G}) &= \min_{v_1, \dots, v_Q} \frac{1}{2n} \sum_{q=1}^Q \|y_q - X_q v_q\|_2^2 + \frac{\lambda}{2} \sum_{q=1}^Q s_q \|v_q\|_2^2 \\ &= \frac{1}{2n} \sum_{q=1}^Q \|y_q\|_2^2 - y_q^T X_q (s_q \lambda n \mathbf{I} + X_q^T X_q)^{-1} X_q^T y_q \\ &= \frac{1}{2n} \sum_{q=1}^Q y_q^T \left(\mathbf{I} + \frac{1}{s_q \lambda n} X_q X_q^T \right)^{-1} y_q. \end{aligned}$$

Formulating terms of the sum as solutions of an optimization problem, we get

$$\begin{aligned} \tilde{\psi}(\mathcal{G}) &= \frac{1}{2n} \sum_{q=1}^Q \max_{\alpha_q \in \mathbb{R}^{s_q}} -\alpha_q^T \left(\mathbf{I} + \frac{1}{s_q \lambda n} X_q X_q^T \right) \alpha_q + 2y_q^T \alpha_q \\ &= \frac{1}{2n} \max_{\substack{\alpha = (\alpha_1; \dots; \alpha_Q) \\ \alpha_q \in \mathbb{R}^{s_q}}} \sum_{q=1}^Q -\alpha_q^T \left(\mathbf{I} + \frac{1}{s_q \lambda n} X_q X_q^T \right) \alpha_q + 2y_q^T \alpha_q, \end{aligned}$$

where $(\alpha_1; \dots; \alpha_Q) = (\alpha_1^T, \dots, \alpha_Q^T)^T$ stacks vectors α_q in one vector of size $\sum_{q=1}^Q s_q = n$. Using that $E = (E_1; \dots; E_Q) = (E_1^T, \dots, E_Q^T)^T \in \{0, 1\}^{n \times n}$ is an orthonormal matrix, we make the change of variable $\beta = E^T \alpha$ (and so $\alpha = E\beta$) such that for $\alpha = (\alpha_1; \dots; \alpha_Q)$, $\alpha_q \in \mathbb{R}^{s_q}$, $\alpha_q = E_q \beta$. Decomposing X_q and y_q and using $E_q^T E_q = \mathbf{diag}(Z_q)$, we get

$$\begin{aligned} \tilde{\psi}(\mathcal{G}) &= \frac{1}{2n} \max_{\beta \in \mathbb{R}^n} \sum_{q=1}^Q -\beta^T E_q^T \left(\mathbf{I} + \frac{1}{s_q \lambda n} X_q X_q^T \right) E_q \beta + 2y_q^T E_q \beta \\ &= \frac{1}{2n} \max_{\beta \in \mathbb{R}^n} \sum_{q=1}^Q -\beta^T E_q^T \left(\mathbf{I} + \frac{1}{s_q \lambda n} E_q X X^T E_q^T \right) E_q \beta + 2y^T E_q^T E_q \beta \\ &= \frac{1}{2n} \max_{\beta \in \mathbb{R}^n} \sum_{q=1}^Q -\beta^T \mathbf{diag}(Z_q) \beta - \frac{1}{s_q \lambda n} \beta^T \mathbf{diag}(Z_q) X X^T \mathbf{diag}(Z_q) \beta + 2y^T \mathbf{diag}(Z_q) \beta. \end{aligned}$$

For q fixed, $\left(\frac{1}{s_q} \mathbf{diag}(Z_q) X X^T \mathbf{diag}(Z_q) \right)_{ij} = \frac{1}{s_q} x_i^T x_j$ if $(i, j) \in \mathcal{G}_q$ and 0 otherwise. So

$$\sum_{q=1}^Q \frac{1}{s_q} \mathbf{diag}(Z_q) X X^T \mathbf{diag}(Z_q) = X X^T \circ M,$$

where $M = Z(Z^T Z)^{-1} Z^T$ is the normalized equivalence matrix of the partition \mathcal{G} and \circ denotes the Hadamard product. Using $\sum_{q=1}^Q \mathbf{diag}(Z_q) = \mathbf{I}$, we finally get a function of the equivalence matrix

$$\begin{aligned}\psi(M) &= \frac{1}{2n} \max_{\beta \in \mathbb{R}^n} -\beta^T \left(\mathbf{I} + \frac{1}{\lambda n} X X^T \circ M \right) \beta + 2y^T \beta \\ &= \frac{1}{2n} y^T \left(\mathbf{I} + \frac{1}{\lambda n} X X^T \circ M \right)^{-1} y.\end{aligned}$$

Its gradient is given by

$$\nabla \psi(M) = -\frac{1}{2\lambda n^2} X X^T \circ \left(\left(\mathbf{I} + \frac{1}{\lambda n} X X^T \circ M \right)^{-1} y y^T \left(\mathbf{I} + \frac{1}{\lambda n} X X^T \circ M \right)^{-1} \right).$$

Algorithm 2 can be applied to minimize ψ . The linear oracle can indeed be computed with k-means using that the gradient is negative semi-definite. For a fixed Z , the linear predictors v_q for each cluster of points are given by

$$\begin{aligned}v_q &= (n\lambda s_q \mathbf{I} + X_q^T X_q)^{-1} X_q^T y_q \\ &= (n\lambda s_q \mathbf{I} + X^T E_q^T E_q X)^{-1} X^T E_q^T E_q y \\ &= (n\lambda s_q \mathbf{I} + X^T \mathbf{diag}(Z_q) X)^{-1} X^T \mathbf{diag}(Z_q) y.\end{aligned}$$

7.3.2. Convex relaxations for classification. We observe that convex relaxations for classification derive from computations of the convex relaxations for regression by replacing vector of labels y by the corresponding matrix of labels Y .

INRIA - SIERRA PROJECT TEAM & D.I., UMR 8548,
ÉCOLE NORMALE SUPÉRIEURE, PARIS, FRANCE.
E-mail address: vincent.roulet@inria.fr

C.M.A.P., ÉCOLE POLYTECHNIQUE, UMR CNRS 7641
E-mail address: fajwel.fogel@cmap.polytechnique.fr

CNRS & D.I., UMR 8548,
ÉCOLE NORMALE SUPÉRIEURE, PARIS, FRANCE.
E-mail address: aspremon@ens.fr

INRIA - SIERRA PROJECT TEAM & D.I., UMR 8548,
ÉCOLE NORMALE SUPÉRIEURE, PARIS, FRANCE.
E-mail address: francis.bach@inria.fr