



Partitioned least squares

Roberto Esposito¹ · Mattia Cerrato² · Marco Locatelli³

Received: 21 December 2023 / Revised: 28 March 2024 / Accepted: 7 June 2024
© The Author(s) 2024

Abstract

Linear least squares is one of the most widely used regression methods in many fields. The simplicity of the model allows this method to be used when data is scarce and allows practitioners to gather some insight into the problem by inspecting the values of the learnt parameters. In this paper we propose a variant of the linear least squares model allowing practitioners to partition the input features into groups of variables that they require to contribute similarly to the final result. We show that the new formulation is not convex and provide two alternative methods to deal with the problem: one non-exact method based on an alternating least squares approach; and one exact method based on a reformulation of the problem. We show the correctness of the exact method and compare the two solutions showing that the exact solution provides better results in a fraction of the time required by the alternating least squares solution (when the number of partitions is small). We also provide a branch and bound algorithm that can be used in place of the exact method when the number of partitions is too large as well as a proof of NP-completeness of the optimization problem.

Keywords Least squares · Interpretability · Convex optimization · NP completeness

Editor: Rita P. Ribeiro.

Mattia Cerrato, Marco Locatelli have contributed equally to this work.

✉ Roberto Esposito
roberto.esposito@unito.it

Mattia Cerrato
mcerrato@uni-mainz.de

Marco Locatelli
marco.locatelli@unipr.it

¹ Computer Science Department, University of Turin, Corso Svizzera 185, 10125 Turin, Italy

² Computer Science Department, Johannes Gutenberg Universität Mainz, Staudingerweg 9, 55112 Mainz, Germany

³ Engineering and Architecture Department, University of Parma, Parco Area delle Scienze 181/A, 43124 Parma, Italy

1 Introduction

Linear regression models are among the most extensively employed statistical methods in science and industry alike (Bro et al., 2002; Intriligator et al., 1978; Isobe et al., 1990; Nievergelt, 2000; Reeder et al., 2004). Their simplicity, ease of use and performance in low-data regimes enables their usage in various prediction tasks. As the number of observations usually exceeds the number of variables, a practitioner has to resort to approximating the solution of an overdetermined system. Least squares approximation benefits from a closed-form solution and is perhaps the most well known approach in linear regression analysis. Among the benefits of linear regression models there is the possibility of easily interpreting how much each variate is contributing to the approximation of the dependent variable by means of observing the magnitudes and signs of the associated parameters.

In some application domains, partitioning the variables in non-overlapping subsets is beneficial either as a way to insert human knowledge into the regression analysis task or to further improve model interpretability. When considering high-dimensionality data, grouping variables together is also a natural way to make it easier to reason about the data and the regression result. As an example, consider a regression task where the dependent variable is the score achieved by students in an University or College exam. A natural way to group the dependent variables is to divide them into two groups where one contains the variables which represent a student's effort in the specific exam (hours spent studying, number of lectures attended...), while another contains the variables related to previous effort and background (number of previous exams passed, number of years spent at University or College, grade average...). Assuming all these variables could be measured accurately, it might be interesting to know how much each group of variables contributes to the student's score. As a further example, when analyzing complex chemical compounds, it is possible to group together fine-grained features to obtain a partition which refers to high-level properties of the compound (such as structural, interactive and bond-forming among others), and knowing how much each high-level property contributes to the result of the analysis is often of great practical value (Caron et al., 2013). The LIMPET dataset that we introduce in Sect. 5 is a clear-cut example of problems with such structure. In the LIMPET dataset, we have a large number of features that can be grouped in well understood high-level structures and where variables in each group necessarily have to contribute in the same direction (i.e., positively or negatively) to the prediction of lipophilicity of the compound under study.

In this paper, we present a novel variation of linear regression that incorporates feature partitioning into discernible groups. This adapted formulation empowers the analyst to exclude unwanted, unrealistic solutions wherein features within a group are assigned parameters of contrary signs. Thus, the analyst is able to inject domain-specific knowledge into the model. Furthermore, the parameters obtained by solving the problem allow one to easily assess the contribution of each group to the dependent variable as well as the importance of each element of the group.

The newly introduced problem is not easy to solve and indeed we will prove the non-convexity of the objective, and the NP-completeness of the problem itself. In Sect. 3 we introduce two possible algorithms to solve the problem. One is based on an Alternate Convex Search method (Wendell & Hurter, 1976), where the optimization of the parameters is iterative and can get trapped into local minima; the other is based on a reformulation of the original problem into an exponential number of sub-problems, where the exponent is the cardinality K of the partition. We prove convergence of the alternating least square

algorithm and the global optimality of the result returned by the second approach. We also provide guidance for building a branch and bound (Lawler & Wood, 1966) solution that might be useful when the cardinality of the partition is too large to use the exact algorithm.

We test the two algorithms on several datasets. Our experiments include data extracted from the analysis of chemical compounds (Caron et al., 2013) in a particular setting where this kind of analysis already proved to be of value to practitioners, and a number of datasets having a large amount of features which we selected from the UCI repository (Dua & Graff, 2017): in this latter case the number, size, and composition of the partition has been decided arbitrarily just to experiment with the provided algorithms. Our experimental results show that the exact algorithm is usually a good choice, the non-exact algorithm being preferable when high accuracy is not required and/or the cardinality of the partition is too large. Finally, we present and discuss the application of our algorithms to the problem of predicting house prices, showing that the solution provided by our approach leads to more interpretable and actionable results with respect to a least squares model.

While to the best of our knowledge the regression problem and the algorithms we present are novel, there has been previous work dealing with alternative formulations to the linear regression problem. Some of them have shown to be of great practical use and have received attention from both researchers and practitioners.

Partial Least Squares (PLS) Regression (Wold et al., 2001) is a very popular method in hard sciences such as chemistry and chemometrics. PLS has been designed to address the undesirable behavior of ordinary least squares when the dataset is small, especially when the number of features is large in comparison. In such cases, one can try to select a smaller set of features allowing a better behavior. A very popular way to select important features is to use Principal Component Analysis (PCA) to select the features that contributes most to the variation in the dataset. However, since PCA is based on the data matrix alone, one risks to filter out features that are highly correlated with the target variables in \mathbf{Y} . PLS has been explicitly designed to solve this problem by decomposing \mathbf{X} and \mathbf{Y} simultaneously and in such a way to explain as much as possible of the covariance between \mathbf{X} and \mathbf{Y} (Abdi, 2010). Our work is substantially different from these approaches since we are not concerned at all with the goal of removing variables. On the contrary, we group them so to inject domain knowledge in the model, make the result more interpretable, and to provide valuable information about the importance of each group.

Yet another set of techniques that resembles our work are those where a partition of the variables is used to *select* groups of features. Very well known members of this family of algorithms are group lasso methods (Bakin, 1999; Yuan & Lin, 2006; Huang et al., 2012) provide a review of such methodologies). In these works, the authors tackle the problem of selecting grouped variables for accurate prediction. In this case, as in ours, the groups for the variables are defined by the user, but in their case the algorithm needs to predict which subset of the groups will lead to better performances (i.e., either all variables in a group will be used as part of the solution or none of them will be). This is a rather different problem with respect to the one that we introduce here. In our case, we shall assume that all groups are relevant to the analysis. However, in our case we seek a solution where all variables in the same group contributes in the same direction (i.e., with the same sign) to the solution. We argue that this formulation allows for an easier interpretation of the contribution of the whole group as well as of the variables included in each group.

Other techniques that bear some resemblance to our proposal are latent class models McCutcheon (1987). Latent class models are a categorical extension to factor analysis, trying to relate a set of observed variables to a set of latent variables. The value taken by these latter (usually discrete McCutcheon 1987) variables should explain much of the variance in

the former ones. Our problem formulation, on the other hand, constrains the prediction of a continuous target variable by grouping together sets of observed variables. While the solutions found by the algorithms we propose in this paper may reveal interesting patterns (see Sect. 5), our method is not unsupervised and cannot be straightforwardly used to describe the variation of the data via discrete, unobservable factors. Our proposal assumes the availability of a dependent, continuous variable which an analyst is interested in predicting.

In this paper we introduce a new least squares problem and provide algorithms to solve it. We note that we presented the original problem formulation for PartitionedLS in a 2019 paper Esposito et al. (2019). In this follow-up paper, we provide the following new results:

- A revised definition for the PartitionedLS-b problem (see Sect. 3), which allows for an improved optimality proof;
- A complete proof of optimality for the optimal algorithm PartLS-opt, only sketched in previous work Esposito et al. (2019);
- A proof of NP-completeness for the PartitionedLS problem (not present in previous work);
- A new branch-and-bound algorithm that may be used in conjunction with PartLS-opt when the number of partitions is high;
- Information about how to update the algorithms to regularize the solutions;
- Information about how to leverage the non-negative least squares algorithm (Lawson & Hanson, 1995) to improve numerical stability;
- An experimentation of the optimal and the approximate algorithms over three new datasets;
- An experiment showing how the branch-and-bound algorithm compares with the enumerative one;
- A new experiment and a discussion of the interpretability of the results obtained by our approach when applied to the problem of predicting house prices;
- A comparison of the generalization performances of our method with Least Squares, Partial Least Squares, and Principal Component Regression.

2 Model description

In this work we denote matrices with capital bold letters such as \mathbf{X} and vectors with lower-case bold letters as \mathbf{v} . In the text we use a regular (non-bold) font weight when we refer to the *name* of the vector or when we refer to scalar values contained in the vector. In other words, we use the bold face only when we refer to the vector itself. For instance, we might say that the values in the α vector are those contained in the vector $\boldsymbol{\alpha}$, which contains in position i the scalar α_i . We consistently define each piece of notation as soon as we use it, but we also report it in Table 1, where the reader can more easily access the whole notation employed throughout the paper.

Let us consider the problem of inferring a linear least squares model to predict a real variable y given a vector $\mathbf{x} \in \mathbb{R}^M$. We will assume that the examples are available at learning time as an $N \times M$ matrix \mathbf{X} and $N \times 1$ column vector \mathbf{y} . We will also assume that the problem is expressed in homogeneous coordinates, i.e., that \mathbf{X} has an additional column containing values equal to 1, and that the intercept term of the affine function is included into the weight vector \mathbf{w} to be computed.

Table 1 Notation

Symbol(s)	Definition
a_i	i -th component of vector \mathbf{a}
$(\cdot)_n$	Shorthand to specify vectors (or matrices) in terms of their components. For instance $(i)_i$ shall denote a vector \mathbf{v} such that $v_i = i$
k, K	k is the index for iterating over the K subsets belonging to the partition
m, M	m is the index for iterating over the M variables
\mathbf{X}	An $N \times M$ matrix containing the descriptions of the training instances
$\mathbf{A} \times \mathbf{B}$	Matrix multiplication operation (we also simply write it \mathbf{AB} when the notation appears clearer)
\mathbf{y}	A vector of length N containing the labels assigned to the examples in \mathbf{X}
\bullet	Wildcard used in subscriptions to denote whole columns or whole rows: e.g., $\mathbf{X}_{:,k}$ denotes the k -th column of matrix \mathbf{X} and $\mathbf{X}_{m,\bullet}$ denotes its m -th row
\star	Denotes an optimal solution, e.g., p^\star denotes the optimal solution of the PartitionedLS problem, while p_b^\star denotes the optimal solution of the PartitionedLS-b problem
\mathbf{P}	A $M \times K$ partition matrix, $P_{m,k} \in \{0, 1\}$, with $P_{m,k} = 1$ iff variable α_m belongs to the k -th element of the partition
P_k	The set of all indices in the k -th element of the partition: $\{m P_{k,m} = 1\}$
$k[m]$	Index of the partition element to which α_m belongs, i.e.: $k[m]$ is such that $m \in P_{k[m]}$
\circ	Hadamard (i.e., element-wise) product. When used to multiply a matrix by a column vector, it is intended that the columns of the matrix are each one multiplied (element-wise) by the column vector
\oslash	Hadamard (i.e., element-wise) division
\geq	Element-wise larger-than operator: $\alpha \geq 0$ is equivalent to $\alpha_m \geq 0$ for $m \in 1..M$

The standard least squares formulation for the problem at hand is to minimize the quadratic loss over the residuals, i.e.:

$$\text{minimize}_w \|\mathbf{X}\mathbf{w} - \mathbf{y}\|_2^2.$$

This is a problem that has the closed form solution $\mathbf{w} = (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{y}$. As mentioned in Sect. 1, in many application contexts where M is large, the resulting model is hard to interpret. However, it is often the case that domain experts can partition the elements in the weights vector into a small number of groups and that a model built on this partition would provide more accurate results (by incorporating domain knowledge) or/and be much easier to interpret. Then, let \mathbf{P} be a “partition” matrix for the problem at hand (this is not a partition matrix in the linear algebra sense, it is simply a matrix containing the information needed to partition the features of the problem). More formally, let \mathbf{P} be a $M \times K$ matrix where $P_{m,k} \in \{0, 1\}$ is equal to 1 iff feature number m belongs to the k -th partition element. We will also write P_k to denote the set $\{m | P_{m,k} = 1\}$ of all the features belonging to the k -th partition element.

Here we introduce the Partitioned Least Squares (PartitionedLS) problem, a model where we introduce K additional variables and express the whole regression problem in terms of these new variables (and in terms of how the original variables contribute to the predictions made using them). The simplest way to describe the new model is to consider its regression function (to make the discussion easier, we start with the data matrix \mathbf{X} expressed in non-homogenous coordinates and switch to homogenous coordinates afterwards):

$$f(\mathbf{X}) = \left(\sum_{k=1}^K \beta_k \sum_{m \in P_k} \alpha_m x_{n,m} + t \right)_n, \quad (1)$$

i.e., $f(\mathbf{X})$ computes a vector whose n -th component is the one reported within parenthesis (see Table 1 for details on the notation). The first summation is over the K sets in the partition that domain experts have identified as relevant, while the second one iterates over all variables in that set. We note that the m -th α weight contributes to the k -th element of the partition only if feature number m belongs to it. As we shall see, we require that all α values are nonnegative, and that $\forall k : \sum_{m \in P_k} \alpha_m = 1$. Consequently, the expression returns a vector of predictions calculated in terms of two sets of weights: the β weights, which are meant to capture the magnitude and the sign of the contribution of the k -th element of the partition, and the α weights, which are meant to capture how each feature in the k -th set contributes to it. We note that the α weight vector is of the same length as the vector \mathbf{w} in the least squares formulation. Despite this similarity, we prefer to use a different symbol because the interpretation of (and the constraints on) the α weights are different with respect to the w weights.

It is easy to verify that the definition of f in (1) can be rewritten in matrix notation as:

$$\begin{aligned} f(\mathbf{X}) &= \left(\sum_{k=1}^K \beta_k \sum_m P_{m,k} \alpha_m x_{n,m} + t \right)_n \\ &= \mathbf{X} \times (\mathbf{P} \circ \boldsymbol{\alpha}) \times \boldsymbol{\beta} + t, \end{aligned} \quad (2)$$

where \circ is the Hadamard product extended to handle column-wise products. More formally, if \mathbf{Z} is a $A \times B$ matrix, $\mathbf{1}$ is a B dimensional vector with all entries equal to 1, and \mathbf{a} is a column vector of length A , then $\mathbf{Z} \circ \mathbf{a} \triangleq \mathbf{Z} \circ (\mathbf{a} \times \mathbf{1}^T)$; where the \circ symbol on the right hand side of the definition is the standard Hadamard product. Equation (2) can be rewritten in homogeneous coordinates as:

$$f(\mathbf{X}) = \mathbf{X} \times (\mathbf{P} \circ \boldsymbol{\alpha}) \times \boldsymbol{\beta}, \quad (3)$$

where \mathbf{X} incorporates a column with all entries equal to 1, and we consider an additional group (with index $K+1$) having a single α_{M+1} variable in it. Given the constraints on α variables, α_{M+1} is forced to assume a value equal to 1 and the value of t is then totally incorporated into β_{K+1} . In the following we will assume for ease of notation that the problem is given in homogeneous coordinates and that the constants M and K already account for the additional single-variable group.

Definition 1 The partitioned least square (PartitionedLS) problem is formulated as:

$$\begin{aligned} &\text{minimize}_{\boldsymbol{\alpha}, \boldsymbol{\beta}} \|\mathbf{X} \times (\mathbf{P} \circ \boldsymbol{\alpha}) \times \boldsymbol{\beta} - \mathbf{y}\|_2^2 \\ &\text{s.t. } \boldsymbol{\alpha} \geq 0 \\ &\quad \mathbf{P}^T \times \boldsymbol{\alpha} = \mathbf{1}. \end{aligned}$$

In summary, we want to minimize the squared residuals of $f(\mathbf{X})$, as defined in (3), under the constraint that for each subset k in the partition, the set of weights form a distribution: they need to be all nonnegative as imposed by $\boldsymbol{\alpha} \geq 0$ constraint and they need to sum to 1 as imposed by constraint $\mathbf{P}^T \times \boldsymbol{\alpha} = \mathbf{1}$.

Unfortunately we do not know a closed form solution for this problem. Furthermore, the problem is not convex and hence hard to solve to global optimality using standard out-of-the-box solvers. Even worse, later on we shall prove that the problem is actually NP-complete. The following theorem states the non-convexity of the objective function formally.

Theorem 1 *The PartitionedLS problem is not convex.*

Proof It suffices to show that the Hessian of the objective function is not positive semidefinite. The details of the proof can be found in Esposito et al. (2019). □

In the following we will provide two algorithms that solve the above problem. One is an alternating least squares approach which scales well with K , but it is not guaranteed to provide the globally optimal solution. The other one is a reformulation of the problem through a (possibly) large number of convex problems whose minimum is guaranteed to be the globally optimal solution of the original problem. Even though the second algorithm does not scale well with K , we believe that this *should not be a problem* since the PartitionedLS is by design well suited for a small group of interpretable groups. However, we do sketch a possible branch and bound strategy to mitigate this problem in Sect. 3.4.

Remark 1 The PartitionedLS model presented so far has no regularization mechanism in place and, as such, it risks overfitting the training set. Since the α values are normalized by definition, the only parameters that need regularization are those collected in the β vector. Then, the regularized version of the objective function simply adds a penalty on the size of the β vector:

$$\|\mathbf{X} \times (\mathbf{P} \circ \boldsymbol{\alpha}) \times \boldsymbol{\beta}\|_2^2 + \eta \|\boldsymbol{\beta}\|_2^2, \quad (4)$$

where the squared euclidean norm could be substituted with the L1 norm in case a LASSO-like regularization is preferred.

3 Algorithms

3.1 Alternating least squares approach

In the PartitionedLS problem we aim at minimizing a non-convex objective, where the non-convexity depends on the multiplicative interaction between α and β variables in the expression $\|\mathbf{X} \times (\mathbf{P} \circ \boldsymbol{\alpha}) \times \boldsymbol{\beta} - \mathbf{y}\|_2^2$. Interestingly, if one fixes $\boldsymbol{\alpha}$, the expression $\mathbf{X} \times (\mathbf{P} \circ \boldsymbol{\alpha})$ results in a matrix \mathbf{X}' that does not depend on any variable. Then, the whole expression can be rewritten as a problem p_α whose objective function $\|\mathbf{X}'\boldsymbol{\beta} - \mathbf{y}\|_2^2$ depends on the parameter vector $\boldsymbol{\alpha}$ and is the convex objective function of a standard least squares problem in the β variables. In a similar way, it can be shown that by fixing $\boldsymbol{\beta}$ one also ends up with a convex optimization problem p_β . Indeed, after fixing $\boldsymbol{\beta}$, the objective function is the squared norm of a vector whose components are affine functions of vector $\boldsymbol{\alpha}$ (see Sect. 3.3 for more details). These observations naturally lead to the formulation of an alternating least squares solution where one alternates between solving p_α and p_β . In Algorithm 1 we formalize this intuition into the PartLS-alt function where, after initializing $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ randomly, we iterate

until some stopping criterion is satisfied (in our experiments we fixed a number T of iterations, but one may want to stop the algorithm as soon as \mathbf{a} and \mathbf{c} do not change between two iterations). At each iteration we take the latest estimate for the α variables and solve the p_α problem based on that estimate, we then keep the newly found β variables and solve the p_β problem based on them. At each iteration the overall objective is guaranteed not to increase in value and, indeed, we prove that, if the algorithm is never stopped, the sequence of α and β vectors found by PartLS-alt has at least one accumulation point and that all accumulation points are partial optima¹ with the same function value.

Theorem 2 *Let $\zeta_i = (\alpha_i, \beta_i)$ be the sequence of α and β vectors found by PartLS-alt to the PartitionedLS problem and assume that the objective function is regularized as described in (4), then:*

- (1) *The sequence of ζ_i has at least one accumulation point, and*
- (2) *All accumulation points are partial optima attaining the same value of the objective function.*

Proof The PartitionedLS problem is actually a biconvex optimization problem and Algorithm 1 is actually a specific instantiation of the *Alternating Convex Search* strategy (Gorski et al., 2007) to solve biconvex problems. Theorem 4.9 in (Gorski et al., 2007) implies that:

- If the sequence ζ_i is contained in a compact set then it has at least one accumulation point, and
- If for each accumulation point ζ^* of the sequence ζ_i , either the optimal solution of the problem with fixed α is unique, or the optimal solution of the problem with fixed β is unique; then all accumulation points are partial optima and have the same function value.

The first requirement is fulfilled in our case since α is constrained by definition into $[0, 1]^M$, while the regularization term prevents β from growing indefinitely. The second requirement is fulfilled since for fixed α the optimization function is quadratic and strictly convex in β . Hence, the solution is unique. \square

Algorithm 1 Alternating least squares solution to the PartitionedLS problem. The notation $\text{const}(\alpha)$ (respectively $\text{const}(\beta)$) is just to emphasize that the current value of α (respec-

¹ A partial optima of a function $f(\alpha, \beta)$ is a point (α^*, β^*) such that $\forall \alpha : f(\alpha^*, \beta^*) \leq f(\alpha, \beta^*)$ and $\forall \beta : f(\alpha^*, \beta^*) \leq f(\alpha^*, \beta)$.

tively β) will be used as a constant in the following step

```

1   function PartLS-alt (X, y, P)
2     alpha = random(M)
3     beta = random(K)
4
5     for t in 1...T
6       a = const(alpha)
7       p* = minimize_beta (||X x (P o a) x beta - y||_2^2)
8
9
10      c = const(beta)
11      p* = minimize_alpha (||X x (P o alpha) x c - y||_2^2,
12                          alpha >= 0,
13                          P^T x alpha = 1)
14    end
15
16    return (p*, alpha, beta)
17  end

```

3.2 Reformulation as a set of convex subproblems

Here we show how the PartitionedLS problem can be reformulated as a new problem with binary variables which, in turn, can be split into a set of convex problems such that the smallest objective function value among all local (and global) minimizers of these convex problems is also the global optimum value of the PartitionedLS problem.

Definition 2 The *PartitionedLS-b problem* is a PartitionedLS problem in which the β variables are substituted by a binary variable vector $\mathbf{b} \in \{-1, 1\}^K$, and the normalization constraints over the α variables are dropped:

$$\begin{aligned}
 & \text{minimize}_{\alpha, \mathbf{b}} \|\mathbf{X} \times (\mathbf{P}^\circ \alpha) \times \mathbf{b} - \mathbf{y}\|_2^2 \\
 & \text{s.t. } \alpha \geq 0 \\
 & \quad \mathbf{b} \in \{-1, 1\}^K.
 \end{aligned}$$

The PartitionedLS-b problem turns out to be a Mixed Integer Nonlinear Programming (MINLP) problem with a peculiar structure. More specifically, we note that the above definition actually defines 2^K minimization problems, one for each of the possible instances of vector \mathbf{b} . Interestingly, each one of the minimization problems can be shown to be convex by the same argument used in Sect. 3.1 (for fixed β variables) and we will prove that the minimum attained by minimizing all sub-problems corresponds to the global minimum of the original problem. We also show that by simple algebraic manipulation of the result found by a PartitionedLS-b solution, it is possible to write a corresponding PartitionedLS solution attaining the same objective.

The main breakthrough here derives from noticing that in the original formulation the β variables are used to keep track of two facets of the solution: (i) The magnitude and (ii) The sign of the contribution of each subset in the partition of the variables. With the \mathbf{b} vector

keeping track of the signs, one only needs to reconstruct the magnitude of the β contributions to recover the solution of the original problem.

The following theorem states the equivalence between the PartitionedLS and the PartitionedLS-b problem. More precisely, we will prove that for any feasible solution of one of the two problems, one can build a feasible solution of the other problem with the same objective function value, from which equality between the optimal values of the two problems immediately follows.

Theorem 3 *Let (α, \mathbf{b}) be a feasible solution of the PartitionedLS-b problem. Then, there exists a feasible solution $(\hat{\alpha}, \hat{\beta})$ of the PartitionedLS problem such that:*

$$\|\mathbf{X} \times (\mathbf{P} \circ \alpha) \times \mathbf{b} - \mathbf{y}\|_2^2 = \|\mathbf{X} \times (\mathbf{P} \circ \hat{\alpha}) \times \hat{\beta} - \mathbf{y}\|_2^2. \quad (5)$$

Analogously, for each feasible solution $(\hat{\alpha}, \hat{\beta})$ of the PartitionedLS problem, there exists a feasible solution (α, \mathbf{b}) of the PartitionedLS-b problem such that (5) holds. Finally, $p^* = p_b^*$, where p^* and p_b^* denote, respectively, the optimal value of the PartitionedLS problem and of the PartitionedLS-b problem.

Proof Let (α, \mathbf{b}) be a feasible solution of the PartitionedLS-b problem and let $\bar{\beta}$ be a normalization vector containing in $\bar{\beta}_k$ the normalization factor for variables in partition subset k :

$$\bar{\beta} = \left(\sum_{m \in P_k} \alpha_m \right)_k = \mathbf{P}^\top \times \alpha.$$

Then, for each m such that $\bar{\beta}_{k[m]} \neq 0$, we define $\hat{\alpha}_m$ as follows:

$$\hat{\alpha}_m = \frac{\alpha_m}{\bar{\beta}_{k[m]}},$$

while for any m such that $\bar{\beta}_{k[m]} = 0$ we can define $\hat{\alpha}_m$, e.g., as follows:

$$\hat{\alpha}_m = \frac{1}{|P_{k[m]}|}.$$

In fact, for any k such that $\bar{\beta}_k = 0$, any definition of $\hat{\alpha}_m$ for $m \in P_k$ such that $\sum_{m \in P_k} \hat{\alpha}_m = 1$ would be acceptable. The $\hat{\beta}$ vector can be reconstructed simply by taking the Hadamard product of \mathbf{b} and $\bar{\beta}$:

$$\hat{\beta} = \mathbf{b} \circ \bar{\beta}.$$

In order to prove (5), we only need to prove that

$$\mathbf{X} \times (\mathbf{P} \circ \alpha) \times \mathbf{b} = \mathbf{X} \times (\mathbf{P} \circ \hat{\alpha}) \times \hat{\beta}.$$

The equality is proved as follows:

$$\begin{aligned}
 \mathbf{X} \times (\mathbf{P} \circ \hat{\boldsymbol{\alpha}}) \times \hat{\boldsymbol{\beta}} &= \mathbf{X} \times \left(\mathbf{P} \circ \left(\frac{\alpha_m}{\bar{\beta}_{k[m]}} \right)_m \right) \times (b_k \bar{\beta}_k)_k \\
 &= \left(\sum_k b_k \bar{\beta}_k \sum_{m \in P_k} \frac{\alpha_m}{\bar{\beta}_{k[m]}} x_{n,m} \right)_n \\
 &= \left(\sum_k b_k \bar{\beta}_k \sum_{m \in P_k} \frac{\alpha_m}{\bar{\beta}_k} x_{n,m} \right)_n \\
 &= \left(\sum_k b_k \sum_{m \in P_k} \alpha_m x_{n,m} \right)_n \\
 &= \mathbf{X} \times (\mathbf{P} \circ \boldsymbol{\alpha}) \times \mathbf{b},
 \end{aligned}$$

where in between row 2 and row 3 we used the fact that $\bar{\beta}_k$ and $\bar{\beta}_{k[m]}$ are two ways to write the same thing (the former using directly the partition number k , and the latter using the notation $k[m]$ to get the partition number from the feature number m). To be more precise, we only considered the case when $\bar{\beta}_{k[m]} \neq 0$ for all m . But the result can be easily extended to the case when $\bar{\beta}_{k[m]} = 0$ for some m , by observing that in this case the corresponding terms give a null contribution to both sides of the equality.

Now, let $(\hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\beta}})$ be a feasible solution of the PartitionedLS problem. Then, we can build a feasible solution $(\boldsymbol{\alpha}, \mathbf{b})$ for the PartitionedLS-b problem as follows. For any $k \in \{1, \dots, K\}$ let:

$$b_k = \begin{cases} -1 & \text{if } \hat{\beta}_k < 0 \\ +1 & \text{otherwise,} \end{cases}$$

while for each m , let:

$$\alpha_m = b_{k[m]} \hat{\beta}_{k[m]} \hat{\alpha}_m.$$

Equivalence between the objective function values at $(\hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\beta}})$ and $(\boldsymbol{\alpha}, \mathbf{b})$ is proved in a way completely analogous to what we have seen before.

Finally, the equivalence between the optimal values of the two problems is an immediate corollary of the previous parts of the proof. In particular, it is enough to observe that for any optimal solution of one of the two problems, there exists a feasible solution of the other problem with the same objective function value, so that both $p^* \geq p_b^*$ and $p^* \leq p_b^*$ holds, and, thus, $p^* = p_b^*$. □

The complete algorithm, which detects and returns the best solution of the PartitionedLS-b problems by iterating over all possible vectors \mathbf{b} , is implemented by the function PartLS-opt reported in Algorithm 2.

Algorithm 2 PartitionedLS-b solution to the PartitionedLS problem. The function extract_min retrieves the $(\hat{p}, \hat{\alpha}, \hat{\beta})$ tuple in the results array attaining the lowest \hat{p} value

```

1 function PartLS-opt (X, y, P)
2   results = []
3
4   for  $\hat{\mathbf{b}}$  in  $\{1, -1\}^K$ 
5      $\hat{p} = \text{minimize}_{\hat{\alpha}} (\|(\mathbf{X} \times (\mathbf{P} \circ \hat{\alpha}) \times \hat{\mathbf{b}} - \mathbf{y})\|_2^2), \hat{\alpha} \succeq 0)$ 
6
7     results += ( $\hat{p}, \hat{\alpha}, \hat{\mathbf{b}}$ )
8   end
9
10   $p^*, \alpha, \mathbf{b} = \text{extract\_best}(\text{results})$ 
11
12
13   $\bar{\beta} = \mathbf{P}^\top \times \alpha$ 
14   $\hat{\alpha} = (\mathbf{P} \circ \alpha \circ \bar{\beta}^\top) \times \mathbf{1}$ 
15   $\hat{\beta} = \mathbf{b} \circ \bar{\beta}$ 
16
17  return ( $p^*, \hat{\alpha}, \hat{\beta}$ )
18 end

```

Remark 2 When dealing with the PartitionedLS-b problem, the regularization term introduced for the objective function of the PartitionedLS problem, reported in (4), needs to be slightly updated so to accommodate the differences in the objective function when used in Algorithm 2. In this second case, since the β variables do not appear in the optimization problems obtained after fixing the different binary vectors \mathbf{b} , the regularization term $\|\beta\|_2^2$ is replaced by $\|\mathbf{P}^\top \times \alpha\|_2^2$. We notice that since the new regularization term is still convex, it does not hinder the convexity of the optimization problems.

3.3 Numerical stability

The optimization problems solved within Algorithms 1 and 2, despite being convex, are sometimes hard to solve due to numerical problems. General-purpose solvers often find the data matrix to be ill-conditioned and return sub-optimal results Björck (1996); Cucker et al. (2007). In this section we show how to rewrite the problems so to mitigate these difficulties. The main idea is to recast the minimization problems as standard least squares and non-negative least squares problems, and to employ efficient solvers for these specific problems rather than the general-purpose ones.

We start by noticing that the minimization problem at line 7 of Algorithm 1 can be easily solved by a standard least square algorithm since the expression $\mathbf{X} \times (\mathbf{P} \circ \alpha)$ computes to a constant matrix \mathbf{X}' and the original problem simplifies to the ordinary least squares problem: $\text{minimize}_{\beta} (\|\mathbf{X}'\beta - \mathbf{y}\|_2^2)$.

For what concerns the minimization problem at line 13 of the same algorithm, we notice that we can initially ignore the constraint $\mathbf{P}^\top \times \alpha = \mathbf{1}$. Without such constraint, the problem turns out to be a non-negative least squares problem. Indeed, we note that expression $\mathbf{X} \times (\mathbf{P} \circ \alpha) \times \mathbf{c}$ can be rewritten as the constant matrix $\mathbf{X}\mathbf{O}(\mathbf{P}\mathbf{O}\mathbf{c}^\top \times \mathbf{1})^\top$

multiplied by the vector α , so that the whole minimization problem could be rewritten as:

$$\begin{aligned} &\text{minimize}_{\alpha} \|\mathbf{X} \circ (\mathbf{P} \circ \mathbf{c}^T \times \mathbf{1})^T \times \alpha - \mathbf{y}\|_2^2 \\ &s.t. \quad \alpha \geq 0. \end{aligned}$$

After such problem has been solved, the solution of the problem including the constraint $\mathbf{P}^T \times \alpha = 1$ can be easily obtained by dividing each α subset by a normalizing factor and multiplying the corresponding β variable by the same normalizing factor (it is the same kind of operations we exploited in Sect. 3.2; in that context the normalizing factors were denoted with $\bar{\beta}$).

In a completely analogous way we can rewrite the minimization problem at line 5 of Algorithm 2 as:

$$\begin{aligned} &\text{minimize}_{\hat{\alpha}} \|\mathbf{X} \circ (\mathbf{P} \circ \hat{\mathbf{b}}^T \times \mathbf{1})^T \times \hat{\alpha} - \mathbf{y}\|_2^2 \\ &s.t. \quad \hat{\alpha} \geq 0, \end{aligned} \tag{6}$$

which, again, is a non-negative least squares problem.

As previously mentioned, by rewriting the optimization problems as described above and by employing special-purpose solvers for the least squares and the non-negative least squares problems, solutions appear to be more stable and accurate.

Remark 3 Many non-negative least squares solvers do not admit an explicit regularization term. An l_2 -regularization term equivalent to $\rho \|\beta\|_2^2 = \rho \|\mathbf{P}^T \times \alpha\|_2^2 = \rho \sum_k (\sum_{m \in P_k} \alpha_m)^2$ can be implicitly added by augmenting the data matrix \mathbf{X} with K additional rows. The trick is done by setting all the additional y to 0 and the k -th additional row as follows:

$$x_{N+k,m} = \begin{cases} \sqrt{\rho} & \text{if } m \in P_k \\ 0 & \text{otherwise.} \end{cases}$$

When the additional k -th row and the additional y are plugged into the expression inside the norm in (6), the expression evaluates to:

$$\sum_{m \in P_k} \sqrt{\rho} b_k \alpha_m - 0 = b_k \sqrt{\rho} \sum_{m \in P_k} \alpha_m,$$

which reduces to $\rho \sum_k (\sum_{m \in P_k} \alpha_m)^2$ when squared and summed over all the k as a result of the evaluation of the norm.

3.4 An alternative branch-and-bound approach

Algorithm 2 is based on a *complete enumeration* of all possible 2^K vectors \mathbf{b} . Of course, such an approach becomes too expensive as soon as K gets large. As already previously commented, PartLS-opt is by design well suited for small K values, so that complete enumeration should be a valid option most of the times. However, for the sake of completeness, in this section we discuss a branch-and-bound approach, based on *implicit enumeration*, which could be employed as K gets large. Pseudo-code detailing the approach is reported in Algorithm 3.

Algorithm 3 Pseudo code for a depth-first implementation of the branch and bound optimization of PartLS-opt. μ is the current upper bound of the optimal value, Σ is the set of constraints associated to the current node. For the sake of simplicity, the algorithm returns only the optimal values. It is easy to modify it to keep track of the best solution as well. `lower_bound` computes the relaxation of either (8) or (9) subject to the constraints in Σ and returns the lower bound lb (the lower bound itself) and α (the values of the variables attaining it)

```

1  function PartLS-bnb( $\mathbf{X}, \mathbf{y}, \mathbf{P}, \mu, \Sigma$ )
2       $\mathbf{Q} = \mathbf{X}^\top \mathbf{X}$ 
3       $\mathbf{q} = -2\mathbf{X}^\top \mathbf{y}$ 
4       $q_0 = \mathbf{y}^\top \mathbf{y}$ 
5
6       $lb, \alpha = \text{lower\_bound}(\mathbf{Q}, \mathbf{q}, q_0, \Sigma)$ 
7      if  $lb \geq \mu$ 
8          # No solution better than the upper
9          # bound can be found in the current
10         # branch
11         return  $\mu$ 
12     end
13
14      $\nu = \left( \sum_{i,j \in P_k} \max\{0, -\alpha_i \alpha_j\} \right)_k$ 
15     if  $\nu = 0$ 
16         # optimal solution found for this branch
17         # we can avoid further splitting
18         # constraints
19         return  $lb$ 
20     end
21
22      $k = \arg \max_k \nu_k$ 
23
24      $\Sigma_+ = \Sigma \cup \{\forall i \in P_k : \alpha_i \geq 0\}$ 
25      $\Sigma_- = \Sigma \cup \{\forall i \in P_k : \alpha_i \leq 0\}$ 
26
27      $\mu_+ = \text{PartLS-bnb}(\mathbf{X}, \mathbf{y}, \mathbf{P}, \mu, \Sigma_+)$ 
28      $\mu_- = \text{PartLS-bnb}(\mathbf{X}, \mathbf{y}, \mathbf{P}, \min(\mu, \mu_+), \Sigma_-)$ 
29
30     return  $\min(\mu, \mu_+, \mu_-)$ 
31 end

```

First, we remark that the PartitionedLS-b problem can be reformulated as follows

$$\begin{aligned}
 & \underset{\alpha}{\text{minimize}} \sum_n \left(\sum_k \sum_{m \in P_k} \alpha_m x_{n,m} - y_n \right)^2 \\
 & \text{s.t.} \quad \alpha_i \alpha_j \geq 0 \quad \forall i, j \in P_k, \quad \forall k \in \{1, \dots, K\},
 \end{aligned} \tag{7}$$

where we notice that vector \mathbf{b} and the nonnegativity constraints $\alpha \geq 0$ have been eliminated, and replaced by the new constraints, which impose that for any k , all variables α_m such that $m \in P_k$ must have the same sign. The new problem is a quadratic one with a convex quadratic objective function and simple (but non-convex) bilinear constraints. We note that, having removed the \mathbf{b} variables, the scalar objective do not need the distinction

between groups anymore and it can be rewritten as $\sum_n (\sum_m \alpha_m x_{n,m} - y_n)^2$ or, in matrix form, as $\|\mathbf{X}\boldsymbol{\alpha} - \mathbf{y}\|^2 = (\mathbf{X}\boldsymbol{\alpha} - \mathbf{y})^\top (\mathbf{X}\boldsymbol{\alpha} - \mathbf{y})$. Hence, we can reformulate the problem as follows

$$\begin{aligned} &\text{minimize}_{\boldsymbol{\alpha}} \boldsymbol{\alpha}^\top \mathbf{Q}\boldsymbol{\alpha} + \mathbf{q}^\top \boldsymbol{\alpha} + q_0 \\ &\text{s.t.} \quad \alpha_i \alpha_j \geq 0 \quad \forall i, j \in P_k, \quad \forall k \in \{1, \dots, K\}, \end{aligned} \tag{8}$$

where $\mathbf{Q} = \mathbf{X}^\top \mathbf{X}$, $\mathbf{q} = -2\mathbf{X}^\top \mathbf{y}$, and $q_0 = \mathbf{y}^\top \mathbf{y}$. Different lower bounds for this problem can be computed. The simplest one is obtained by simply removing all the constraints, which results in an unconstrained convex quadratic problem. A stronger, but more costly, lower bound can be obtained by solving the classical semidefinite relaxation of quadratic programming problems. First, we observe that problem (8) can be rewritten as follows (see Shor, 1987)

$$\begin{aligned} &\text{minimize}_{\boldsymbol{\alpha}, \mathbf{A}} \mathbf{Q} \bullet \mathbf{A} + \mathbf{q}^\top \boldsymbol{\alpha} + q_0 \\ &\text{s.t.} \quad \mathbf{A}_k = \boldsymbol{\alpha}_{P_k} \boldsymbol{\alpha}_{P_k}^\top \quad \forall k \in \{1, \dots, K\} \\ &\quad \mathbf{A}_k \geq \mathbf{O} \quad \forall k \in \{1, \dots, K\}, \end{aligned} \tag{9}$$

where $\mathbf{Q} \bullet \mathbf{A} = \sum_{i,j} Q_{ij} A_{ij}$, and $\boldsymbol{\alpha}_{P_k}$ is the restriction of $\boldsymbol{\alpha}$ to the entries in P_k , $k \in \{1, \dots, K\}$. Next, we observe that the equality constraint $\mathbf{A} = \boldsymbol{\alpha}\boldsymbol{\alpha}^\top$ is equivalent to requiring that \mathbf{A} is a psd (positive semidefinite) matrix and is of rank one. If we remove the (non-convex) rank one requirement, we end up with the following convex relaxation of (8) requiring the solution of a semidefinite programming problem:

$$\begin{aligned} &\text{minimize}_{\boldsymbol{\alpha}, \mathbf{A}} \mathbf{Q} \bullet \mathbf{A} + \mathbf{q}^\top \boldsymbol{\alpha} + q_0 \\ &\text{s.t.} \quad \mathbf{A}_k - \boldsymbol{\alpha}_{P_k} \boldsymbol{\alpha}_{P_k}^\top \text{ is psd} \quad \forall k \in \{1, \dots, K\} \\ &\quad \mathbf{A}_k \geq \mathbf{O} \quad \forall k \in \{1, \dots, K\}. \end{aligned}$$

Note that by Schur complement, constraint “ $\mathbf{A}_k - \boldsymbol{\alpha}_{P_k} \boldsymbol{\alpha}_{P_k}^\top$ is psd” is equivalent to the following semidefinite constraint:

$$\begin{pmatrix} 1 & \boldsymbol{\alpha}_{P_k}^\top \\ \boldsymbol{\alpha}_{P_k} & \mathbf{A}_k \end{pmatrix} \text{ is psd.}$$

No matter which problem we solve to get a lower bound, after having solved it we can consider the vector $\boldsymbol{\alpha}^*$ of the optimal values of the α variables at its optimal solution and we can compute the following quantity for each $k \in \{1, \dots, K\}$

$$v_k = \sum_{i,j \in P_k} \max\{0, -\alpha_i^* \alpha_j^*\}.$$

If $v_k = 0$ for all k , then the optimal solution of the relaxed problem is feasible and also optimal for the original problem (8) and we are done. Otherwise, we can select an index k such that $v_k > 0$ (e.g., the largest one, corresponding to the largest violation of the constraints), and split the original problem into two subproblems, one where we impose that all variables α_m , $m \in P_k$, are nonnegative, and the other where we impose that all variables α_m , $m \in P_k$, are nonpositive. Lower bounds for the new subproblems can be easily computed by the same convex relaxations employed for the original problem (8), but with the additional constraints. The violations v_k are computed also for the subproblems and, in case one of them is strictly positive, the corresponding subproblem may be further split into two further subproblems, unless its lower bound becomes at some point larger than or equal to

the current global upper bound of the problem, which is possibly updated each time a new feasible solution of (8) is detected. As previously commented, Algorithm 3 provides a possible implementation of the branch-and-bound approach. More precisely, Algorithm 3 is an implementation where nodes of the branch-and-bound tree are visited in a depth-first manner. An alternative implementation is, e.g., the one where nodes are visited in a lowest-first manner, i.e., the first node to be visited is the one with the lowest lower bound.

4 Complexity

In this section we establish the theoretical complexity of the PartitionedLS-b problem. In view of reformulation (7), it is immediately seen that the cases where $|P_k| = 1$ for all $k = 1, \dots, K$, are polynomially solvable. Indeed, in this situation problem (7) becomes unconstrained and has a convex quadratic objective function. Here we prove that as soon as we move from $|P_k| = 1$ to $|P_k| = 2$, the problem becomes NP-complete. We prove this by showing that each instance of the NP-complete problem SUBSET SUM (see, e.g., Garey and Johnson, 1979) can be transformed in polynomial time into an instance of problem (7). We recall that problem SUBSET SUM is defined as follows. Let s_1, \dots, s_k be a collection of K positive integers. We want to establish whether there exists a partition of this set of integers into two subsets such that the sums of the integers belonging to the two subsets is equal, i.e., whether there exist $I_1, I_2 \subseteq \{1, \dots, K\}$ such that:

$$I_1 \cup I_2 = \{1, \dots, K\}, I_1 \cap I_2 = \emptyset, \sum_{k \in I_1} s_k = \sum_{k \in I_2} s_k. \quad (10)$$

Now, let us consider an instance of problem (7) with K partitions and two variables $\alpha_{m_1,k}$ and $\alpha_{m_2,k}$ for each partition k (implying $M = 2K$). The data matrix \mathbf{X} and vector \mathbf{y} have $N = 3K + 1$ rows defined as follows (when k and m are not restricted, they are assumed to vary on $\{1 \dots K\}$ and $\{1 \dots M\}$ respectively):

$$\begin{aligned} x_{k,m_1,k} &= 1, \\ x_{k,m_2,k} &= -1, \\ x_{k,m} &= 0, y_k = -s_k && m \notin \{m_{1,k}, m_{2,k}\} \\ x_{K+k,m_1,k} &= \sqrt{\rho}, x_{K+k,m_2,k} = 0, \\ x_{K+k,m} &= 0, \\ y_{K+k} &= 0 && m \notin \{m_{1,k}, m_{2,k}\} \\ x_{2K+k,m_1,k} &= 0, \\ x_{2K+k,m_2,k} &= \sqrt{\rho}, \\ x_{2K+k,m} &= 0, y_{2K+k} = 0 && m \notin \{m_{1,k}, m_{2,k}\} \\ x_{3K+1,m} &= 1, y_{3K+1} = 0. \end{aligned}$$

When the values so defined are plugged into problem (7) we obtain:

$$\begin{aligned}
 &\text{minimize}_{\alpha} \sum_{k=1}^K (\alpha_{m_{1,k}} - \alpha_{m_{2,k}} - s_k)^2 + \rho \sum_{k=1}^K (\alpha_{m_{1,k}})^2 + \\
 &\quad \rho \sum_{k=1}^K (\alpha_{m_{2,k}})^2 + \left[\sum_{k=1}^K (\alpha_{m_{1,k}} + \alpha_{m_{2,k}}) \right]^2 \\
 &\text{s.t. } \alpha_{m_{1,k}} \alpha_{m_{2,k}} \geq 0 \quad \forall k \in \{1, \dots, K\},
 \end{aligned} \tag{11}$$

with $\rho > 0$.

We prove the following theorem, which states that an instance of the SUBSET SUM problem (10) can be solved by solving the corresponding instance (11) of problem (7), and, thus, establishes NP-completeness of the PartitionedLS-b problem.

Theorem 4 *The optimal value of (11) is equal to*

$$\frac{\rho \sum_{k=1}^K s_k^2}{1 + \rho},$$

if and only if there exist I_1, I_2 such that (10) holds, i.e., if and only if the SUBSET SUM problem admits a solution.

Proof As a first step we derive the optimal solutions of the following restricted two-dimensional problems for $k \in \{1, \dots, K\}$:

$$\begin{aligned}
 &\text{minimize}_{\alpha_{m_{1,k}}, \alpha_{m_{2,k}}} (\alpha_{m_{1,k}} - \alpha_{m_{2,k}} - s_k)^2 + \\
 &\quad \rho (\alpha_{m_{1,k}})^2 + \rho (\alpha_{m_{2,k}})^2 \\
 &\text{s.t. } \alpha_{m_{1,k}} \alpha_{m_{2,k}} \geq 0.
 \end{aligned} \tag{12}$$

This problem admits at least a global minimizer since its objective function is strictly convex quadratic. Global minimizers should be searched for among *regular* KKT points and *irregular* points. Regular points are those who fulfill a *constraint qualification*. In particular, in this problem all feasible points, except the origin, fulfill the constraint qualification based on the linear independence of the gradients of the active constraints. This is trivially true since there is a single constraint and the gradient of such constraint is null only at the origin. Thus, the only irregular point is the origin. In order to detect the KKT points, we first write down the KKT conditions:

$$\begin{aligned}
 &2(\alpha_{m_{1,k}} - \alpha_{m_{2,k}} - s_k) + 2\rho\alpha_{m_{1,k}} - \mu\alpha_{m_{2,k}} = 0 \\
 &-2(\alpha_{m_{1,k}} - \alpha_{m_{2,k}} - s_k) + 2\rho\alpha_{m_{2,k}} - \mu\alpha_{m_{1,k}} = 0 \\
 &\alpha_{m_{1,k}} \alpha_{m_{2,k}} \geq 0 \\
 &\mu\alpha_{m_{1,k}} \alpha_{m_{2,k}} = 0 \\
 &\mu \geq 0,
 \end{aligned}$$

where μ is the Lagrange multiplier of the constraint. We can enumerate all KKT points of problem (12). By summing up the first two equations, we notice that

$$(\mu - 2\rho)(\alpha_{m_{1,k}} + \alpha_{m_{2,k}}) = 0,$$

must hold. This equation is satisfied if:

- Either $\alpha_{m_{1,k}} + \alpha_{m_{2,k}} = 0$, which implies $\alpha_{m_{1,k}} = \alpha_{m_{2,k}} = 0$, in view of $\alpha_{m_{1,k}} \alpha_{m_{2,k}} \geq 0$. As previously mentioned, the origin is the unique irregular point. So, it is not a KKT point but when searching for the global minimizer, we need to compute the objective function value also at such point and this is equal to s_k^2 ;
- Or $\mu = 2\rho > 0$, which implies, in view of the complementarity condition, that $\alpha_{m_{1,k}} \alpha_{m_{2,k}} = 0$, and, after substitution in the first two equations, we have the two KKT points

$$\left(\frac{s_k}{1+\rho}, 0 \right), \quad \left(0, -\frac{s_k}{1+\rho} \right).$$

The objective function value at both these KKT points is equal to $\frac{\rho}{1+\rho} s_k^2$, lower than the objective function value at the origin, and, thus, these KKT points are the two global minima of the restricted problem (12).

Based on the above result, we have that problem

$$\begin{aligned} & \text{minimize}_{\alpha} \quad \sum_{k=1}^K (\alpha_{m_{1,k}} - \alpha_{m_{2,k}} - s_k)^2 + \\ & \quad \quad \quad \rho \sum_{k=1}^K (\alpha_{m_{1,k}})^2 + \rho \sum_{k=1}^K (\alpha_{m_{2,k}})^2 \\ & \text{s.t.} \quad \quad \alpha_{m_{1,k}} \alpha_{m_{2,k}} \geq 0 \quad \forall k \in \{1, \dots, K\}, \end{aligned}$$

which is the original one (11) without the last term $\left[\sum_{k=1}^K (\alpha_{m_{1,k}} + \alpha_{m_{2,k}}) \right]^2$, and which can be split into the K subproblems (12), has global minimum value equal to $\frac{\rho \sum_{k=1}^K s_k^2}{1+\rho}$ and $2K$ global minima defined as follows: for each $I_1, I_2 \subseteq \{1, \dots, K\}$ such that $I_1 \cap I_2 = \emptyset$ and $I_1 \cup I_2 = \{1, \dots, K\}$,

$$\alpha_{m_{1,k}}^* = \begin{cases} \frac{s_k}{1+\rho} & k \in I_1 \\ 0 & k \notin I_1, \end{cases} \quad \alpha_{m_{2,k}}^* = \begin{cases} -\frac{s_k}{1+\rho} & k \in I_2 \\ 0 & k \notin I_2. \end{cases}$$

Now, if we replace these coordinates in the omitted term $\left[\sum_{k=1}^K (\alpha_{m_{1,k}} + \alpha_{m_{2,k}}) \right]^2$, we have the following

$$\left[\sum_{k=1}^K (\alpha_{m_{1,k}}^* + \alpha_{m_{2,k}}^*) \right]^2 = \frac{1}{(1+\rho)^2} \left[\sum_{k \in I_1} s_k - \sum_{k \in I_2} s_k \right]^2,$$

which is equal to 0 for some I_1, I_2 if and only if the SUBSET SUM problem admits a solution. As a consequence the optimal value of problem (11) is equal to $\frac{\rho \sum_{k=1}^K s_k^2}{1+\rho}$ if and only if the SUBSET SUM problem admits a solution, as we wanted to prove. \square

5 Experiments

In this section, we present the experimental findings obtained through the application of the algorithms proposed in this paper over several commonly used datasets (see Table 2).

Table 2 Summary of dataset statistics

Dataset	Rows	Columns	# Feature groups
Limpet	82	44	6
Facebook comment volume	40000	53	5
Superconductivity	10000	81	7
YearPredictionMSD	10000	90	9
Ames house regression	2931	79	10

We stress that the number of feature groups may be adapted depending on the analyst's needs and the constraints they wish to impose on the regressor

In Sect. 5.1, we investigate the properties in terms of regression performance and runtime of PartLS-opt, PartLS-alt, and PartLS-bnb, providing insights about when one should be preferred over the other.

In Sect. 5.2, we provide an example of interpreting the solution provided by our approach. Unfortunately, interpretability is not easily measurable and is, in general, highly task-dependent Doshi-Velez and Kim (2017). Nonetheless, previous research has discussed interpretability of models across multiple dimensions, such as simulatability, decomposability and algorithmic transparency (Lipton, 2016). To show the benefits of framing a regression task as a partitioned least squares problem, we report an experiment analyzing the solution found by the PartLS-opt algorithm on an additional dataset (the Ames House Prices dataset). In particular, we will show that the grouped solution found via the Partitioned Least Squares formulation is arguably more simulatable and decomposable compared to the more commonly employed “feature-by-feature” linear regression solutions. Finally, in Sect. 5.3, we compare the generalization performances of our approach with those of least squares and two established variants: Partial Least Squares (PLS) and Principal Component Regression (PCR).

5.1 Runtime vs. solution quality

We start by experimenting with PartLS-opt and PartLS-alt and on four regression problems on the following datasets: *Limpet*, *Facebook Comment Volume*, *Superconductivity*, and *YearPredictionMSD*. Details about these datasets may be found in the Appendix. We choose these datasets because of their relatively high number of features. In particular, the *Limpet* dataset had already been the subject of a block-relevance analysis in previous literature (Ermondi & Caron, 2012; Caron et al., 2013). We ran PartLS-alt (Algorithm 1) in a multi-start fashion with 100 randomly generated starting points. The four panels in Fig. 1 report the best objective value obtained during these random restarts along with the cumulative time needed to obtain that value (so the rightmost point will plot the cumulative time of the 100 restarts versus the best objective obtained in the whole experiment). We repeated the experiment using two different values of parameter T (number of iterations), setting it to 20 and 100, respectively. So for a single random restart with $T = 20$ (or $T = 100$), Algorithm 1 will alternate 20 (100) times before returning. As one would expect, we see that increasing the value of parameter T slows down the algorithm, but allows it to converge to better solutions.

The experiments confirm that PartLS-opt retrieves more accurate solutions, as expected due to its global optimality property established in Sect. 3.2. Depending on

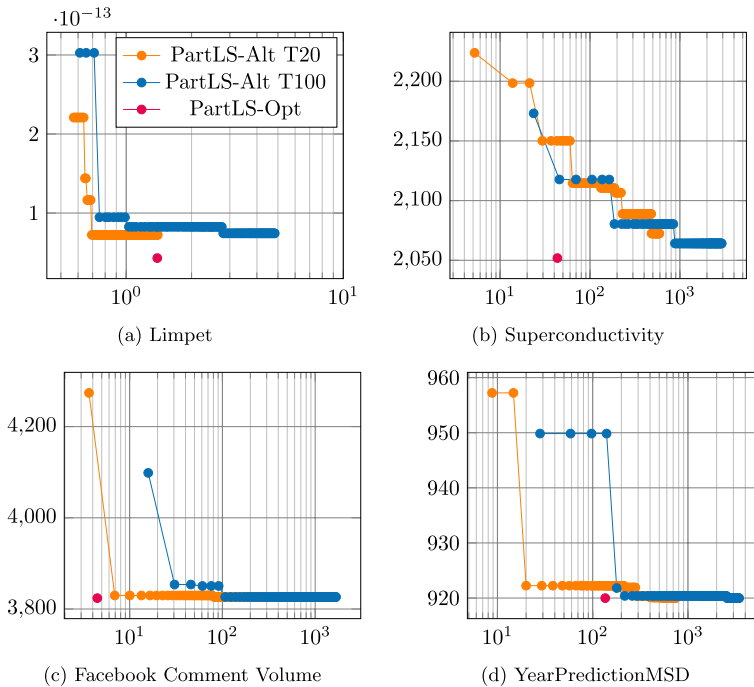


Fig. 1 Plot of the behavior of the two proposed algorithms on four datasets. PartLS-alt has been repeated 100 times following a multi-start strategy and in two settings ($T = 20$ and $T = 100$). Each point on the orange and blue lines reports the cumulative time and best objective found during these 100 restarts. PartLS-opt outputs a single solution, drawn in green (Color figure online)

the dataset, this solution may be either cheaper or more costly to compute compared to the approximate solution obtained by PartLS-alt. Notably, in typical scenarios, the alternating least squares approach, PartLS-alt, outperforms PartLS-opt in terms of running time only when the total number of iterations (and thus the total number of convex sub-problems to be solved) is smaller than 2^K . However, in our experimentation, this often results in solutions that grossly approximate the optimal one. Consequently, we find that PartLS-opt is likely preferable in most cases, providing an optimal solution within a reasonable timeframe, often even quicker than PartLS-alt. Furthermore, although the alternating algorithm can occasionally deliver a solution faster than PartLS-opt, which might be deemed “good enough”, its iterative nature introduces a degree of uncertainty.

Clearly there are cases where the number of groups or where the time required to solve a single convex problem is very large. In these cases, when approximate solutions are acceptable for the application at hand, PartLS-alt could be a very compelling solution. We conclude by noting that a use case with a large number of groups appears to us not very plausible. In fact, it could be argued that the reduced interpretability of the results defies one of the main motivations behind employing the Partitioned Least Squares model in the first place.

It is worth mentioning that, in case a problem with a large K were to arise, the PartLS-bnb algorithm (see Algorithm 3) is likely to allow users to retrieve the optimal solution more efficiently than PartLS-opt. We propose here a further experiment with synthetic data, through which we show when it is convenient to switch from PartLS-opt to the Branch-and-Bound approach implemented in PartLS-bnb and discussed in

Table 3 Computing times (in seconds) for PartLS-opt and Algorithm 3, for different values of K, N, M

K	N	M	Time PartLS-opt	Time BB
10	100	400	3.62	0.81
10	100	600	3.74	0.58
10	150	400	9.50	1.15
10	150	600	9.63	2.52
10	200	400	14.90	2.09
10	200	600	15.19	7.05
15	100	400	141.71	1.70
15	100	600	151.29	1.65
15	150	400	346.65	5.95
15	150	600	403.24	4.31
15	200	400	525.27	8.24
15	200	600	595.25	18.26
20	100	400	–	1.71
20	100	600	–	1.42
20	150	400	–	13.26
20	150	600	–	3.67
20	200	400	–	22.57
20	200	600	–	6.60

Sect. 3.4. In all the previously discussed experiments the cardinality K of the partition is relatively small. Thus, PartLS-opt is able to solve the related problems efficiently. However, the computing times of PartLS-opt quickly increases exponentially as K increases. In these cases, a Branch-and-Bound approach is a much better choice. To better clarify this point, we report in Table 3 the results on synthetic data obtained by randomly generating in the interval $[-10, 10]$ the entries of \mathbf{X} , by generating \mathbf{y} by adding some random noise generated in the interval $[-50, 50]$ to each entry of a target solution $\mathbf{y}_{\text{ref}} = \mathbf{X}\mathbf{w}_{\text{ref}}$, and by randomly generating the K sets in the partition. In the table we compare the computing times (in seconds), for different values K, N, M , of PartLS-opt and of Algorithm 3 (with lower bounds at branch-and-bound nodes computed through the solution of least squares problems with additional non-negativity constraints). A – denotes a computing time exceeding 1,000 s. The results clearly show that, as K increases, a branch-and-bound approach is much more efficient than PartLS-opt.

5.2 Interpretability on ames house prices

We present here an analysis of a solution found by PartLS-opt on the Ames House Prices dataset, which is publicly available via Kaggle Anna Montoya (2016). This dataset has a relatively high number of columns—79 in total—each detailing one particular characteristic of housing properties in Ames, Iowa. The task is to predict the selling price of each house.

We propose a grouping of the features into 10 groups, each one representing a high-level characteristic of the property (see Table 5). As an example, we collect 6 columns




Fig. 2 Feature groups and associated β values as found by PartLS-opt on the Ames House Prices dataset Anna Montoya (2016)




Fig. 3 α values for group “Outside Facilities” found by PartLS-opt on the Ames House Prices dataset Anna Montoya (2016)




Fig. 4 Feature weights for group “Outside Facilities” of a regularized linear regression on the Ames House Prices dataset Anna Montoya (2016)

referring to the availability and quality of air conditioning systems, electrical system, heating and fireplaces in a “Power and Temperature” group. Other feature groups refer to overall quality of the construction work and materials employed (“Building Quality”), external facilities such as garages or swimming pools (“Outside Facilities”). We show the feature groups we designed and the β values found by PartLS-opt² in Fig. 2. We note that the grouped solution enabled by the partitioned least squares formulation is able to give a high-level summary of the regression result. An analyst is therefore able to communicate easily to, e.g. an individual selling their house, that the price is mostly determined by the building quality and the attractiveness of the lot. A deeper analysis is of course possible by investigating the α values found by the algorithm. For instance, we report the α s contributions for the “Outside Facilities” group in Fig. 3. Here, one is able to notice that garage quality has the biggest impact on the property’s price, which is potentially actionable knowledge.

In Fig. 4, we report the weights of the features in the “Outside Facilities” group as learnt by the least squares algorithm.

We argue that the group- and feature-level analysis made possible by our contributions improves on the interpretability of ungrouped linear regression. While linear regression is a relatively simple model and therefore intuitively satisfies some notion of transparency, previous work has established that this is not necessarily the case. Lipton (2016) discusses interpretability of models around three separate dimensions: simulatability, decomposability and algorithmic transparency. Simulatability may be achieved when a person is able to

² In this the regularization parameter has been set to $\rho = 10$

contemplate all the model at once in reasonable amount of time. While the amount of time is of course subjective, Lipton stresses the fact that linear models may not be simulatable if a high number of features are involved. On the other hand, the partitioned least squares formulation we propose finds a higher-level, grouped solution via the β values. Thus, a practitioner would be able to build a simpler mental model of the solution by focusing on the groups rather than the individual features.

5.3 Quality of the inferred model

While one of the major benefits of the Partitioned Least Squares problem is in simplifying the interpretation of the results, it should be self-evident that this would be a pointless exercise if the returned hypothesis were not at least comparable to other widely used techniques in terms of generalization capabilities. In this section, we investigate generalization performances of regressors learnt by PartLS-opt and compare them with Least Squares (LS), Principal Component Regression (PCR) and Partial Least Squares (PLS). All experiments are repeated 100 times on different train/test splits.

We experiment on the four datasets earlier in this section and on an additional dataset *Artificial* which we created for this specific test. The main goal of this dataset is to showcase a situation where we have complete and accurate domain knowledge about the partition. The artificial dataset contains 70 training samples and 930 test samples. Samples contains feature values randomly sampled from a normal distribution. This dataset's target variable may be computed without cross-partition feature interactions: Specifically, the target is computed as $\mathbf{y} = \mathbf{X} \times (\mathbf{P} \circ \boldsymbol{\alpha}) \times \boldsymbol{\beta} + t$, where \mathbf{P} is a partition in 5 sets having cardinalities 5, 10, 4, 12, 6. The \mathbf{X} matrix has been perturbed with gaussian noise with mean 0 and standard deviation 0.05 after generating the target column. t and $\boldsymbol{\alpha}$ have been generated using a uniform distribution in $[0, 1]$. $\boldsymbol{\alpha}$ has then been normalized so that $P^T \times \boldsymbol{\alpha} = \mathbf{1}$. The $\boldsymbol{\beta}$ are the normalization factors used to ensure $P^T \times \boldsymbol{\alpha} = \mathbf{1}$ multiplied by 11, 4, 2, 1, and 3. Signs of the groups have been set to $-1, 1, 1, -1, 1$. These latter parameters have been set arbitrarily and without tuning.

Results are reported in Table 4. When the test error of a method is significantly better³ than the competitors, it is shown in bold. If more than one result is in bold, then the bold-faced results are not significantly better with respect to each other, but are significantly better than all the remaining results.

The PCR algorithm has been run setting the maximal number of principal components equal to the number of groups in the P matrix. We conducted experiments with PartLS-opt three times, utilizing three distinct partitioning methods: one based on the partitions devised by ourselves ("P arbitrary" or "P from DK"), one where features were grouped based on their signs in the solution found by LS in the experiment with the same train/test split ("P from LS"), and one ("P opt") using this same methodology but on the results of LS run on the full dataset (i.e., before the train/test split). This latter experiment aims at simulating a situation where the domain knowledge closely matches the "natural" partitioning of the columns of the dataset. The settings "P from LS" and "P opt" aim to demonstrate that, with the correct partitioning, the algorithm converges to an optimal solution.

³ According to a paired t-test at the 99% confidence level.

Table 4 Train/Test mean squared errors of Least Squares (LS), Principal Component Regression (PCR), Partial Least Squares (PLS) and Partitioned Least Squares (PartLS). Partitioned Least Squares has been run with several configurations of the partition matrix: in “P from DK” experiments the partition matrix is built from domain knowledge; in “P arbitrary” the partition matrix is arbitrary (used instead of DK when domain knowledge was not available); in “P from LS” we use a binary partition matrix based on the signs of the attributes in the solution of the LS problem computed on the training set; in “P opt” the partition matrix is built as in “P from LS”, but the LS is run over the complete dataset. We report the average and the standard deviation of results over 100 repetitions based on different train/test splits. Bold faced averages are statistically better at the 99% significance level (using a dpaired t-test)

Dataset	Method	Train	Test
Artificial	LS	0.18 ±0.04	0.81 ±0.16
	PCR	23.56 ±4.48	30.33 ±2.26
	PLS	0.51 ±0.10	2.06 ±0.64
	PartLS: P from LS	0.18 ±0.04	0.81 ±0.16
	PartLS: P from DK	0.19 ±0.05	0.68 ±0.11
	PartLS: P opt	0.19 ±0.05	0.68 ±0.11
Limpet	LS	0.00 ±0.00	4.14 ±3.95
	PCR	0.71 ±0.13	1.05 ±0.39
	PLS	0.64 ±0.76	1.27 ±1.65
	PartLS: P from LS	0.00 ±0.00	6.91 ±12.56
	PartLS: P from DK	0.00 ±0.00	46.34 ±47.30
	PartLS: P opt	0.00 ±0.00	4.44 ±4.20
Facebook	LS	2785.66 ±450.79	3012.92 ±725.57
	PCR	3672.98 ±503.71	3687.18 ±789.64
	PLS	3765.28 ±509.81	3784.41 ±798.20
	PartLS: P from LS	2785.66 ±450.79	3012.92 ±725.57
	PartLS: P arbitrary	2850.49 ±457.82	3023.13 ±733.19
	PartLS: P opt	2791.91 ±450.93	2988.19 ±722.48
Year prediction	LS	89.96 ±1.97	91.87 ±2.20
	PCR	115.24 ±2.39	115.52 ±2.71
	PLS	3.42E5 ±6.8E3	3.5E5 ±1.1E4
	PartLS: P from LS	89.96 ±1.97	91.87 ±2.20
	PartLS: P arbitrary	100.30 ±2.02	101.83 ±2.28
	PartLS: P opt	90.18 ±1.97	91.70 ±2.18
Superconductivity	LS	307.64 ±3.68	310.87 ±5.76
	PCR	762.74 ±6.77	764.47 ±11.35
	PLS	566.86 ±5.25	567.64 ±8.32
	PartLS: P from LS	307.64 ±3.68	310.87 ±5.76
	PartLS: P arbitrary	376.82 ±3.91	379.36 ±5.87
	PartLS: P opt	307.73 ±3.68	310.77 ±5.76

We start discussing these results by noting that, for all datasets except Limpet where the problem is heavily under-determined, LS and PartLS-opt on “P from LS” yield identical results. This is expected as the problems are equivalent from the optimizer perspective. Furthermore, results of PartLS-opt on “P opt” show that, when the provided partition is accurate, the inductive bias allows for better generalization in most situations.

On the Artificial dataset, PartLS-opt on “P from DK” significantly outperforms the competitors. This shows once more that, when the correct partitioning is provided, PartLS-opt exhibits an inductive bias that enhances generalization. On this dataset PartLS-opt attains the same results when using “P from DK” and “P opt”. This is expected since “P from DK” is built using perfect knowledge of the P matrix, and we verified that the signs of the features found by LS on the complete dataset induce a partition that can be formally shown to be equivalent to the one that have been used to generate the data.

The PCR and the PLS algorithms are the clear winners on the Limpet dataset. The dataset matrix is under-determined and collinear, which is the ideal case for these techniques. In all the other cases, their inductive biases significantly hinder the algorithms performances, most likely because the number of principal components guessed on the basis of the P matrix was not sufficient to explain enough variance in the data. Setting the number of maximal number of principal components to be equal to the number of features does not seem to change much the results: either they converge to the LS solution, or they obtain result not too distant from the ones presented in Table 4.

For Facebook, Year Prediction, and Superconductivity datasets, PartLS-opt yields the best performances when equipped with the “P opt” partition. It lags a little behind LS when equipped with the partitions we used in our previous experiments (“P arbitrary”), which is totally reasonable since those partitions were chosen to showcase the difference in the time performances between the approaches, rather than the quality of the generalization results.

The experiments overall demonstrate that the constraints imposed by the Partitioned Least Squares approach can serve as a strong inductive bias when the partition knowledge is accurate. However, the technique encounters difficulties when analyzing datasets with many collinear features. Indeed, the current formulation of Partitioned Least Squares does not address this specific issue, suggesting that further research is needed to tackle this challenge.

6 Conclusions

In this paper we presented an alternative least squares linear regression formulation. Our model enables scientists and practitioners to group features together into partitions, hence allowing the modeling of higher level abstractions which are easier to reason about. We provided rigorous proofs of the non-convexity of the problem and presented PartLS-alt and PartLS-opt, two algorithms to cope with the problem.

PartLS-alt is an iterative algorithm based on the alternating least squares method. The algorithm is proved to converge, but there is no guarantee that the accumulation point results in a globally optimal solution. On the contrary, as experiments have shown, the algorithm can be trapped in a local minimizer and return an approximate solution. Experiments suggest that it could be faster and preferable to the exact algorithm PartLS-opt in some circumstances (e.g., when the time needed to solve a single sub-problem is large and the application allows for sub-optimal answers).

PartLS-opt is an enumerative, exact, algorithm and our contribution includes a formal optimality proof. In our experimentation, we confirmed that it behaves very well under several different settings, although its time complexity grows exponentially with the number of groups. We argue that this exponential growth in time complexity should not impede its adoption: a large number of groups seems implausible in practical scenarios since it would

undermine interpretability of the results and hence the attractiveness of the problem formulation. However, for the sake of completeness and to provide guidance to the interested reader, we provided a branch-and-bound solution that shares the same optimality guarantees of PartLS-opt. This latter formulation, depending on the actual structure of the problem as implied by the data, might save computation by pruning the search space, possibly avoiding to solve a large number of sub-problems. In Sect. 5.1 we have shown the benefits of this strategy when the number of partition sets increases, but we intend to further investigate this issue in future work.

In Sect. 5.3, we explore how the constraints introduced by the Partitioned Least Squares formulation impact the generalization properties of the inferred model. Our findings indicate that when the partition knowledge aligns with the underlying data distribution, the Partitioned Least Squares algorithms are very effective in leveraging this information. However, the results obtained from the Limpet dataset clearly demonstrate that collinearity can pose a challenge for the proposed technique and, indeed, neither the problem formulation, nor the proposed algorithms try to address this issue. We believe that addressing collinearity problems represents an interesting avenue for future research.

One topic for further research is about how to evaluate the partitions created by a domain expert. In this work, we have taken feature partitions “at face value” or otherwise assumed that an agreed-upon partitioning was developed by an expert. Investigating the challenges of the (human) partitioning process, possibly by performing an interactive user study as suggested by Doshi-Velez and Kim (2017), is a possible avenue for future developments.

Table 5 Summary of the groups of features used in the Ames House Prices experiment. See the Kaggle Anna Montoya (2016) repository for detailed information about the meaning of each feature label

Group	Features
LotDescription	MSSubClass, MSZoning, LotFrontage, LotArea, Street, Alley, LotShape, LandContour, LotConfig, LandSlope
BuildingPlacement	Utilities, Neighborhood, Condition1, Condition2
BuildingAge	YearBuilt, YearRemodAdd
BuildingQuality	BldgType, HouseStyle, OverallQual, OverallCond, RoofStyle, RoofMatl, Exterior1st, Exterior2nd, MasVnrType, MasVnrArea, ExterQual, ExterCond, Foundation, Functional
Basement	BsmtQual, BsmtCond, BsmtExposure, BsmtFinType1, BsmtFinSF1, BsmtFinType2, BsmtFinSF2, BsmtUnfSF, TotalBsmtSF
PowerAndTemperature	Heating, HeatingQC, CentralAir, Electrical, Fireplaces, FireplaceQu
Sizes	1stFlrSF, 2ndFlrSF, LowQualFinSF, GrLivArea
Rooms	BsmtFullBath, BsmtHalfBath, FullBath, HalfBath, BedroomAbvGr, KitchenAbvGr, KitchenQual, TotRmsAbvGrd
OutsideFacilities	GarageType, GarageYrBlt, GarageFinish, GarageCars, GarageArea, GarageQual, GarageCond, PavedDrive, WoodDeckSF, OpenPorchSF, EnclosedPorch, 3SsnPorch, ScreenPorch, PoolArea, PoolQC, Fence
Various	MiscFeature, MiscVal, MoSold, YrSold, SaleType, SaleCondition

Appendix

Dataset descriptions

Limpet dataset

This dataset (Caron et al., 2016) contains 82 features describing measurements over simulated (VolSurf+ (Goodford, 1985)) models of 44 drugs. The regression task is the prediction of the lipophilicity of the 44 compounds. The 82 features are partitioned into 6 groups according to the kind of property they describe. The six groups have been identified by domain experts and are characterized in (Ermondi & Caron, 2012) as follows:

- *Size/Shape*: 7 features describing the size and shape of the solute;
- *OH2*: 19 features expressing the solute's interaction with water molecules;
- *NI*: 5 features describing the solute's ability to form hydrogen bond interactions with the donor group of the probe;
- *O*: 5 features expressing the solute's ability to form hydrogen bond interactions with the acceptor group of the probe;
- *DRY*: 28 features describing the solute's propensity to participate in hydrophobic interactions;
- *Others*: 18 descriptors describing mainly the imbalance between hydrophilic and hydrophobic regions.

This dataset, while not high-dimensional in the broadest sense of the term, can be partitioned into well-defined, interpretable groups of variables. Moreover and perhaps more importantly, this is a clear case where the Partitioned Least Squares formulation is important to correctly handle the structure of the problem: each group contains variables describing physical properties of the compound that are theoretically bound to act in the same "direction" on the target variable (its lipophilicity). Previous literature which employed this dataset has indeed focused on leveraging the data's structure to obtain explainable results (Caron et al., 2013). We used as training/test split the same one proposed in (Caron et al., 2016).

For this particular problem, the number of groups is 6 and PartLS-opt needs to solve just $2^6 = 64$ convex problems. It terminates in ~ 1.4 seconds reaching a value of the objective function of about $4.3 \cdot 10^{-14}$ (note that the annotation " $1e - 13$ " at the top of the plot denotes that all values on the y axis are to be multiplied by 10^{-13}). PartLS-alt (Algorithm 1) in this particular case is doing very well. Even though the plot shows that PartLS-opt reaches a better loss value, PartLS-alt starts already at a very low value of about $3 \cdot 10^{-13}$ requiring a fraction of the time needed by its optimal counterpart. It is also worth noting that, despite the small changes in the objective value reached by the two algorithms, the configuration of the α and β variables are substantially different.

Facebook comment volume dataset

The Facebook Comment Volume dataset (Singh, 2016) contains more than 40 thousand training vectors along with 53 features. Each sample represents a post published on the social media service by a "Facebook Page", an entity which other users can follow and

“like” so to receive updates on their Facebook activity. Features range from the number of users which “like” and follow the page to the number of comments the post received during different time frames. We removed the column which indicated whether a post was a paid advertisement, as this feature only contained 0 values, i.e., no advertisements were collected. Then, we divided the features into 5 blocks, each containing 10 features save for the last one which contained 11 features. The task here is to predict how many comments the same post will receive in the next few hours. The dataset is hosted at the UCI repository (Dua & Graff, 2017). To keep training time and memory usage low, we limited the training samples to the first 15000 examples of the training set. On this dataset, PartLS-opt is able to find the highest quality solution in less than 5 s. PartLS-alt with $T = 20$ finds a similar quality solution after about 7 s. PartLS-alt with $T = 100$ takes more than 3 min to converge to a comparable objective value.

Superconductivity dataset

The Superconductivity dataset contains 81 features representing characteristics of superconductors. The dataset contains 21264 examples. In our experiment we trained the model over the first 10000 examples. The task is to predict a material’s critical temperature. The features are derived from a superconductor’s atomic mass, density and fusion heat among others. We refer the reader to the original paper (Hamidieh, 2018) for the specific details about the process. In our experiment, we created 7 feature blocks with 10 features each and an additional one which contained 11 features. PartLS-opt takes ~ 47 seconds reaching an objective value of ~ 2051 . At about the same computational cost, PartLS-alt with $T = 20$ reaches an objective of ~ 2150 . It will take the algorithm about ~ 440 seconds to lower that figure to a loss objective value (~ 2072) comparable to the one obtained by PartLS-opt. Setting $T = 100$ slightly improves the situation: after about 40 seconds the loss objective is ~ 2117 , which lowers to ~ 2080 after ~ 186 seconds and to ~ 2064 after ~ 881 seconds.

YearPredictionMSD dataset

We also propose an experimentation on the YearPredictionMSD dataset. It is a subset of the Million Songs dataset (Bertin-Mahieux et al., 2011). When compared with the original dataset, it has about half the examples (around 500 thousands) and instead of the raw audio and metadata 90 timbre-related features are included. As for the Superconductivity dataset, we limited our experimentation to the first 10000 examples. The target variable represents the year a song has been released in. In this dataset we experimented with 9 blocks of 10 to 12 features. PartLS-opt takes ~ 130 seconds to reach the optimal loss at ~ 920 . PartLS-alt with $T = 20$ is instead able to find a solution which is reasonably close (~ 922) to the optimal one in a much shorter time (around 20 s). When $T = 100$ is used instead, PartLS-alt reaches a reasonable approximation only after ~ 178 seconds.

Author contributions Roberto Esposito, Marco Locatelli, and Mattia Cerrato all contributed to the conception and design of the work. Roberto Esposito and Marco Locatelli designed the original version of the algorithms. Roberto Esposito, Marco Locatelli and Mattia Cerrato designed the experiments and contributed to the writing of the manuscript. All authors read and approved the final manuscript.

Funding Open access funding provided by Università degli Studi di Torino within the CRUI-CARE Agreement.

Data availability All datasets, with the exception of the LIMPET dataset, are publicly available. The LIMPET dataset can be obtained by contacting the authors of the original paper (Caron et al., 2013). The repository for the experiments contains code to download and to pre-process the datasets (or the datasets themselves when not available for downloading) as well as the scripts to actually launch the experiments. Pre-processing consists in packing the data in a format suitable for the algorithms and to partition the data as mentioned in Sect. 5.

Code availability A Julia (Bezanson et al., 2012) implementation of algorithms PartLS-opt, PartLS-alt, and PartLS-bnb is available at <https://github.com/ml-unito/PartitionedLS>; the code for the experiments is available at: <https://github.com/ml-unito/PartitionedLS-experiments-2>.

Declarations

Conflict of interest The authors declared that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Abdi, H. (2010). Partial least squares regression and projection on latent structure regression (PLS regression). *WIREs Computational Statistics*, 2(1), 97–106.
- Anna Montoya, D. (2016). House Prices - Advanced Regression Techniques. Kaggle (2016). <https://kaggle.com/competitions/house-prices-advanced-regression-techniques>
- Bakin, S. (1999). Adaptive regression and model selection in data mining problems. PhD thesis, School of Mathematical Sciences, Australian National University.
- Bertin-Mahieux, T., Ellis, D.P.W., Whitman, B., & Lamere, P. (2011). The million song dataset. In: Proceedings of the 12th International Conference on Music Information Retrieval (ISMIR 2011)
- Bezanson, J., Karpinski, S., Shah, V.B., & Edelman, A. (2012). Julia: A fast dynamic language for technical computing. CoRR [arXiv:1209.5145](https://arxiv.org/abs/1209.5145)
- Björck, Å. (1996). Numerical methods for least squares problems.
- Bro, R., Sidiropoulos, N. D., & Smilde, A. K. (2002). Maximum likelihood fitting using ordinary least squares algorithms. *Journal of Chemometrics: A Journal of the Chemometrics Society*, 16(8–10), 387–400.
- Caron, G., Vallaro, M., & Ermondi, G. (2013). The block relevance (BR) analysis to aid medicinal chemists to determine and interpret lipophilicity. *MedChemCommun*, 10, 1376–1381.
- Caron, G., Vallaro, M., Ermondi, G., Goetz, G. H., Abramov, Y. A., Philippe, L., & Shalaeva, M. (2016). A fast chromatographic method for estimating lipophilicity and ionization in nonpolar membrane-like environment. *Molecular Pharmaceutics*, 13(3), 1100–1110.
- Cucker, F., Diao, H., & Wei, Y. (2007). On mixed and componentwise condition numbers for moore-penrose inverse and linear least squares problems. *Mathematics of Computation*, 76(258), 947–963.
- Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. arXiv preprint [arXiv:1702.08608](https://arxiv.org/abs/1702.08608).
- Dua, D., & Graff, C. (2017). UCI Machine Learning Repository (2017). <http://archive.ics.uci.edu/ml>
- Ermondi, G., & Caron, G. (2012). Molecular interaction fields based descriptors to interpret and compare chromatographic indexes. *Journal of Chromatography A*, 1252, 84–89.
- Esposito, R., Cerrato, M., & Locatelli, M. (2019). Partitioned least squares. In: AI*IA 2019 – Advances in Artificial Intelligence.

- Garey, M.R., & Johnson, D.S. (1979). *Computers and intractability: A guide to the theory of NP-completeness* (Series of Books in the Mathematical Sciences).
- Goodford, P. J. (1985). A computational procedure for determining energetically favorable binding sites on biologically important macromolecules. *Journal of Medicinal Chemistry*, 28(7), 849–857.
- Gorski, J., Pfeuffer, F., & Klamroth, K. (2007). Biconvex sets and optimization with biconvex functions: A survey and extensions. *Mathematical Methods of Operations Research*, 66(3), 373–407.
- Hamidieh, K. (2018). A data-driven statistical model for predicting the critical temperature of a superconductor. *Computational Materials Science*, 154, 346–354.
- Huang, J., Breheny, P., & Ma, S. (2012). A selective review of group selection in high-dimensional models. *Statistical Science: A Review Journal of the Institute of Mathematical Statistics*, 27(4), 392. <https://doi.org/10.1214/12-STS392>
- Intriligator, M.D., Bodkin, R.G., & Hsiao, C. (1978). *Econometric Models, Techniques, and Applications*.
- Isobe, T., Feigelson, E. D., Akritas, M. G., & Babu, G. J. (1990). Linear regression in astronomy. *The Astrophysical Journal*, 364, 104–113.
- Lawler, E. L., & Wood, D. E. (1966). Branch-and-bound methods: A survey. *Operations Research*, 14(4), 699–719.
- Lawson, C.L., & Hanson, R.J. (1995). *Solving Least Squares Problems* vol.15.
- Lipton, Z. (2016). The mythos of model interpretability. *Communications of the ACM*, 61, 31–57.
- McCutcheon, A.L. (1987). Latent class analysis.
- Nievergelt, Y. (2000). A tutorial history of least squares with applications to astronomy and geodesy. *Journal of Computational and Applied Mathematics*, 121(1–2), 37–72.
- Reeder, S. B., Wen, Z., Yu, H., Pineda, A. R., Gold, G. E., Markl, M., & Pelc, N. J. (2004). Multicoil dixon chemical species separation with an iterative least-squares estimation method. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 51(1), 35–45.
- Shor, N. Z. (1987). Quadratic optimization problems. *Soviet Journal of Computer and Systems Sciences*, 25, 1–11.
- Singh, K. (2016). Facebook comment volume prediction. *International Journal of Simulation- Systems, Science and Technology- IJSSST*, 16(5), 16.
- Wendell, R. E., & Hurter, A. P., Jr. (1976). Minimization of a non-separable objective function subject to disjoint constraints. *Operations Research*, 24(4), 643–657.
- Wold, S., Sjöström, M., & Eriksson, L. (2001). Pls-regression: a basic tool of chemometrics. *Chemometrics and Intelligent Laboratory Systems*, 58(2), 109–130.
- Yuan, M., & Lin, Y. (2006). Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society, Series B*, 68, 49–67.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.