

Fast Screening Rules for Optimal Design via Quadratic Lasso Reformulation

Guillaume Sagnol

*Institut für Mathematik, Sekr. MA 5-2
Technische Universität Berlin, Straße des 17. Juni 136
10623 Berlin, Germany*

SAGNOL@MATH.TU-BERLIN.DE

Luc Pronzato

*Université Côte d’Azur, CNRS, Laboratoire I3S
2000 route des lucioles
06900 Sophia Antipolis, France*

LUC.PRONZATO@CNRS.FR

Editor: Christophe Giraud

Abstract

The problems of Lasso regression and optimal design of experiments share a critical property: their optimal solutions are typically *sparse*, i.e., only a small fraction of the optimal variables are non-zero. Therefore, the identification of the support of an optimal solution reduces the dimensionality of the problem and can yield a substantial simplification of the calculations. It has recently been shown that linear regression with a *squared* ℓ_1 -norm sparsity-inducing penalty is equivalent to an optimal experimental design problem. In this work, we use this equivalence to derive safe screening rules that can be used to discard inessential samples. Compared to previously existing rules, the new tests are much faster to compute, especially for problems involving a parameter space of high dimension, and can be used dynamically within any iterative solver, with negligible computational overhead. Moreover, we show how an existing homotopy algorithm to compute the regularization path of the lasso method can be reparametrized with respect to the squared ℓ_1 -penalty. This allows the computation of a Bayes c -optimal design in a finite number of steps and can be several orders of magnitude faster than standard first-order algorithms. The efficiency of the new screening rules and of the homotopy algorithm are demonstrated on different examples based on real data.

Keywords: Design of experiments, Screening rules, Sparsity-inducing penalty, Lasso, L -optimality

1. Introduction

Estimation or prediction of unknown quantities from experimental data are among the most classical problems in statistics and machine learning. *Optimal design of experiments* plays a central role in this process, as it questions *which* data should be collected in order to make estimation/prediction as accurate as possible. In a regression problem, the quality of an experimental design is usually measured through the covariance matrix of the estimator it produces, and a scalar criterion, function of this matrix, is then minimized to provide an A -, c -, D -, E - or L -optimal design. One can refer, e.g., to Fedorov (1972); Silvey (1980); Pukelsheim (1993) for a thorough exposition of the theory of optimal experiments and a

discussion of optimality criteria. Similarly, supervised learning in a classification problem requires labelling samples that are originally unlabelled. This operation may involve taking physical measurements, conducting a poll, running computer simulations or consulting an expert. Here, the aim of optimal design theory is to select *the best possible* subset of the data to be labelled, subject to a budget constraint. In machine learning, optimal design is also known as *active learning* (Cohn et al., 1996).

In this paper, we focus on optimal design for Bayesian estimation. The design is computed off-line, before the collection of any data. However, we assume that prior information on the quantity to be estimated is available, which means that the approach could also be used in a sequential setting: in that case, the selection of *batches of additional samples* relies on the information gathered through data collection at previous iterations. A central objective of this work is to derive screening rules that can be used to safely discard *useless samples*, and hence speed-up the computation of Bayesian optimal designs, by exploiting the recently discovered connection between c -optimal design and a problem similar to Lasso-regression (Sagnol and Pauwels, 2019).

For most optimality criteria, the computation of an optimal *exact design*, i.e., a multiset of samples of given cardinality minimizing the optimality criterion, is NP-hard (Welch, 1982; Civril and Magdon-Ismail, 2009; Černý and Hladík, 2012). To circumvent this issue, standard approaches rely on *approximate design theory*; see, e.g., Silvey (1980); Pukelsheim (1993). When the set of admissible experimental conditions (the full data set available) is finite, it consists in solving a continuous relaxation of the problem in which the *design* is represented by a vector \mathbf{w} lying in the probability simplex, such that the weight w_i corresponds to the fraction of experimental resources allocated to the i -th sample. Then, various rounding techniques can be used to turn \mathbf{w} into an exact design satisfying performance guarantees; see, e.g., Pukelsheim and Rieder (1992); Singh and Xie (2020).

In general, only a few samples contribute to an optimal design \mathbf{w}^* ; that is, most of the optimal weights w_i^* equal zero. We say that a sample with $w_i^* = 0$ is *inessential*. Identifying inessential samples is crucial to algorithms that compute optimal designs, since their removal from the set of candidates speeds up subsequent iterations: indeed, the complexity per iteration of design algorithms that operate on a finite set of samples S is typically proportional to the cardinality of S . The idea originated from Pronzato (2003) and Harman and Pronzato (2007) for D -optimal design (and for the construction of the minimum-volume ellipsoid containing a set of points). It was further extended to Kiefer’s (1974) φ_p -criteria (Pronzato, 2013), to E -optimal design (Harman and Rosa, 2019), and to the elimination of inessential points in the smallest enclosing ball problem (Pronzato, 2019a). A related idea was recently used for the computation of *exact designs*, for criteria used in the field of active learning (Anstreicher, 2020; Li et al., 2022). There, convex relaxations of the problem with some variables fixed at 0 or 1 are solved at internal nodes of a branch-and-bound tree to obtain certificates that some design points are inessential, or must be part of the exact optimal design. This permits to accelerate the pruning of the branch-and-bound tree.

The recent work of Sagnol and Pauwels (2019) sheds more light on the sparsity of optimal designs, as it shows the equivalence between c -optimal design and a variant of the Lasso regression problem in which the sparsity-inducing ℓ_1 -norm penalty is squared. Analogously, L -optimal design is equivalent to group-Lasso regression with a squared $\ell_{1,2}$ -norm penalty; see Section 3.1 for more details on the precise meaning of “equivalent”. In the field of Lasso-

regression, the idea of deriving screening criteria, i.e., inequalities satisfied by points that do not support any optimal solution and can thus be safely eliminated, dates back to El Ghaoui et al. (2012). These screening rules have been refined in subsequent work (Xiang and Ramadge, 2012; Fercoq et al., 2015; Xiang et al., 2016; Ndiaye et al., 2017), with important developments concerning their usage to speed up the solution of a Lasso problem: Bonnefoy et al. (2015) developed dynamical tests that can be run during the progress of any Lasso algorithm (rather than only at the beginning); Wang et al. (2015) proposed sequential tests in which a sequence of Lasso problems are solved for decreasing values of the regularization parameter α , the optimal solution $\mathbf{x}^*(\alpha_k)$ of the problem for the regularization parameter $\alpha = \alpha_k$ being used to screen-out features in the Lasso problem with regularization parameter $\alpha_{k+1} < \alpha_k$.

Our contribution in this work is the development of efficient screening rules for c - and L -optimal designs in the context of Bayesian estimation. In contrast to the elimination rules presented in our previous work (Pronzato and Sagnol, 2021), which also work in the absence of prior but require heavy algebraic computations such as taking matrix square-roots, the present paper leverages the equivalence with a Lasso problem and the geometry of its dual. This yields lightweight screening rules that only require the computation of matrix-vector products and can be used periodically with any algorithm computing an optimal design. While the techniques we use are similar to those used for example in (Fercoq et al., 2015) for the Lasso problem, the adaptation is not straightforward since the sparsity-inducing penalty is squared, which changes the geometric nature of the dual problem. We also show that a homotopy algorithm used to compute the regularization path of a lasso problem (Osborne et al., 2000; Efron et al., 2004) can be adapted to construct a (Bayesian) c -optimal design, sometimes much faster than with state-of-the-art algorithms. Last but not least, the code used in our experiments has been published in the form of a python package (`qlasso`) and is available at <https://gitlab.com/gsnagnol/qlasso>.

The rest of the paper is organized as follows: Section 2 introduces the c - and L -optimal design problems considered in this paper. Section 3 presents the quadratic lasso problem and gives precise statements about its equivalence with an optimal design problem. The dual quadratic lasso problem is derived and we show that it can be interpreted as a projection over a polyhedral cone. Our main result is stated in Theorem 3.6 in the form of a general screening rule, which is declined in two corollaries giving inequalities adapted for algorithms iterating either on lasso variables (Corollary 3.7) or on design weights (Corollary 3.8). These results are adapted to the case of L -optimality in Appendix A. The adaptation of the homotopy algorithm for the lasso to the case of Bayesian c -optimality is derived in Section 4. Finally, numerical examples demonstrating the performance of the new screening rules and the homotopy algorithm are presented in Section 5.

2. Background on optimal designs and notation

Let $\mathbb{H} = \{\mathbf{H}_i, i = 1, \dots, p\}$ denote a set of $m \times m$ symmetric positive definite matrices, which we shall call elementary information matrices. For any vector of weights \mathbf{w} in the

probability simplex

$$\mathcal{P}_p = \{\mathbf{w} \in \mathbb{R}^p, \mathbf{w} \geq \mathbf{0}, \sum_{i=1}^p w_i = 1\},$$

we denote by $\mathbf{M}(\mathbf{w})$ the information matrix

$$\mathbf{M}(\mathbf{w}) = \sum_{i=1}^p w_i \mathbf{H}_i. \quad (2.1)$$

For \mathbf{c} a given vector in \mathbb{R}^m , a c -optimal design is a solution of the optimization problem

$$\min_{\mathbf{w} \in \mathcal{P}_p} \phi_c(\mathbf{w}) = \Phi_c[\mathbf{M}(\mathbf{w})] = \mathbf{c}^\top \mathbf{M}^{-1}(\mathbf{w}) \mathbf{c}, \quad (2.2)$$

where the decision variable \mathbf{w} defines a probability measure over the finite space $[p] = \{1, \dots, p\}$ and is called *design*. We shall denote \mathbf{w}^* a c -optimal design. Note that nor \mathbf{w}^* neither $\mathbf{M}(\mathbf{w}^*)$ are necessarily unique, but Theorem 3.2 shows that $\mathbf{M}^{-1}(\mathbf{w}^*)\mathbf{c}$ is unique. More generally, we also consider L -optimal designs that minimize a linear optimality criterion: given an $m \times m$ positive semidefinite matrix $\mathbf{C} = \mathbf{K}\mathbf{K}^\top \succeq 0$, an L -optimal design solves

$$\min_{\mathbf{w} \in \mathcal{P}_p} \phi_L(\mathbf{w}) = \Phi_L[\mathbf{M}(\mathbf{w})] = \text{trace}[\mathbf{C}\mathbf{M}^{-1}(\mathbf{w})] = \text{trace}[\mathbf{K}^\top \mathbf{M}^{-1}(\mathbf{w}) \mathbf{K}]. \quad (2.3)$$

Note that L -optimality contains c -optimality (for $\mathbf{K} = \mathbf{c}$, see (2.2)), and A -optimality (for $\mathbf{K} = \mathbf{C} = \mathbf{I}_m$) as special cases (hence L -optimality is also called $A_{\mathbf{K}}$ -optimality by some authors).

These problems arise in the context of Bayesian estimation for the linear model

$$Y_i = \mathbf{a}_i^\top \boldsymbol{\theta} + \varepsilon_i \quad \forall i \in [p], \quad (2.4)$$

where $\mathbf{a}_i \in \mathbb{R}^m$ is a sample of m features, Y_i is the (noisy) observation for the i th sample, and the errors ε_i are mutually independent and normal $\mathcal{N}(0, \sigma^2)$. We further assume that a prior $\mathcal{N}(\boldsymbol{\theta}^0, \boldsymbol{\Sigma})$ on the vector of unknown parameters $\boldsymbol{\theta}$ is available, with $\boldsymbol{\Sigma}$ an $m \times m$ -positive definite matrix (which we write $\boldsymbol{\Sigma} \succ 0$).

Suppose that n observations are collected according to the design $\mathbf{w} \in \mathcal{P}_p$, with \mathbf{w} such such that $n w_i \in \mathbb{Z}_{\geq 0}$ for all i ; that is, for each i we collect $n_i = n w_i$ independent observations $Y_{i,j}$, $j \in [n_i]$. Then, the posterior variance of $\mathbf{c}^\top \boldsymbol{\theta}$ is $\frac{\sigma^2}{n} \cdot \mathbf{c}^\top \left(\sum_i w_i \mathbf{a}_i \mathbf{a}_i^\top + \frac{\sigma^2}{n} \boldsymbol{\Sigma}^{-1} \right)^{-1} \mathbf{c}$. Minimizing this posterior variance is equivalent to minimizing $\mathbf{c}'^\top \mathbf{M}^{-1}(\mathbf{w}) \mathbf{c}'$ where $\mathbf{c}' = \boldsymbol{\Sigma}^{1/2} \mathbf{c}$, $\mathbf{M}(\mathbf{w}) = \sum_{i=1}^p w_i \mathbf{H}_i$ and $\mathbf{H}_i = \mathbf{a}'_i \mathbf{a}'_i{}^\top + \lambda \mathbf{I}_m$ for all i , with $\mathbf{a}'_i = \boldsymbol{\Sigma}^{1/2} \mathbf{a}_i$ and $\lambda = \sigma^2/n$. We can thus assume without any loss of generality that $\boldsymbol{\Sigma} = \mathbf{I}_m$ and that \mathbf{H}_i is of the form

$$\mathbf{H}_i = \mathbf{a}_i \mathbf{a}_i^\top + \lambda \mathbf{I}_m, \quad \lambda > 0. \quad (2.5)$$

Following the same lines as above, we see that minimizing $\phi_L(\mathbf{w})$ is equivalent to minimizing the sum of posterior variances of the vector $\mathbf{K}^\top \boldsymbol{\theta}$ in the regression model with Gaussian prior. One may refer to Pilz (1983) for a book-length exposition on optimal design

for Bayesian estimation (Bayesian optimal design). It is noteworthy that L -optimality finds other applications than in design for parameter estimation. For instance, A -optimality can be used to construct space-filling designs, by kernel-based (Gauthier and Pronzato, 2017) or geometrical (Pronzato and Zhigljavsky, 2019) approaches. In (Pronzato, 2019b), L -optimal designs are also used to estimate Sobol’ indices for sensitivity analysis. An example of space-filling based on L -optimality with $p = 2^{13} = 8192$ and $m = 50$ is presented in Section 5.2.

In the rest of the paper we adopt the machinery of *approximate design theory* and ignore the integrality requirements $n w_i \in \mathbb{Z}_{\geq 0}$. We thus obtain the convex optimization problems (2.2) and (2.3) for c - and L -optimality, respectively. Various algorithms have been proposed to solve these problems, starting from traditional methods such as vertex-direction algorithms (Wynn, 1970; Fedorov, 1972), which are adaptations of the celebrated Frank-Wolfe method, and multiplicative weight update algorithms (Fellman, 1974; Yu, 2010). Another approach is to reformulate these problems as second-order cone programs, which can be handled by interior-point solvers (Sagnol, 2011). Recent progress has been obtained through randomization (Harman et al., 2020), or through a reformulation as an unconstrained problem with squared lasso penalty, for which algorithms such as FISTA or block coordinate descent can be used (Sagnol and Pauwels, 2019). The screening rules presented in this paper can be used to speed-up any of the aforementioned algorithms. These rules are in the form of inequalities that are satisfied by any inessential sample \mathbf{a}_i such that $w_i^* = 0$ for all optimal designs \mathbf{w}^* . When this occurs, we also say that the design point \mathbf{a}_i or the matrix \mathbf{H}_i cannot *support* an optimal design.

We denote respectively by $\|\mathbf{x}\| = (\sum_{i=1}^p x_i^2)^{1/2}$, $\|\mathbf{x}\|_1 = \sum_{i=1}^p |x_i|$ and $\|\mathbf{v}\|_\infty = \max_{i \in [p]} |x_i|$ the ℓ_2 -, ℓ_1 - and ℓ_∞ -norm of a vector $\mathbf{x} \in \mathbb{R}^p$; \mathbf{e}_i denotes the i -th canonical basis vector of \mathbb{R}^p . For a $p \times r$ matrix \mathbf{X} , we denote by $\|\mathbf{X}\|_F = [\text{trace}(\mathbf{X}^\top \mathbf{X})]^{1/2}$ the Frobenius norm of \mathbf{X} and by $\|\mathbf{X}\|_{1,2} = \sum_{i=1}^p \|\mathbf{X}_{i,\cdot}\| = \sum_{i=1}^p (\sum_{j=1}^r \mathbf{X}_{i,j}^2)^{1/2}$ its $\ell_{1,2}$ -norm, with $\mathbf{X}_{i,\cdot}$ denoting the i th row of \mathbf{X} ; $\mathcal{B}_m(\mathbf{x}, \rho)$ denotes the closed Euclidean ball with center \mathbf{x} and radius ρ in \mathbb{R}^m .

3. Optimal design and (quadratic) Lasso

3.1 Equivalence of Optimal Design and Quadratic Lasso

Sagnol and Pauwels (2019) have shown the equivalence between c - (respectively, L -) optimal design and a quadratic Lasso (respectively, group-Lasso) problem.

We present below a simplified proof of this result for the case of c -optimality, which also leads to a more precise statement. The extension of this result to the case of L -optimality is carried out in Appendix A.

Theorem 3.1. *Let \mathbf{H}_i be given by (2.5) and denote by \mathbf{A} the $m \times p$ matrix with columns $(\mathbf{a}_i)_{i \in [p]}$. Then, the c -optimal design problem (2.2) is equivalent to the following problem, which we call quadratic Lasso:*

$$\min_{\mathbf{x} \in \mathbb{R}^p} \mathcal{L}_\lambda(\mathbf{x}) = \|\mathbf{A}\mathbf{x} - \mathbf{c}\|^2 + \lambda \|\mathbf{x}\|_1^2, \tag{3.1}$$

in the following sense:

- (i) the optimal value of (3.1) is equal to $\lambda \phi_c(\mathbf{w}^*)$, where \mathbf{w}^* is a c -optimal design;

(ii) If \mathbf{x}^* solves the quadratic Lasso problem and $\mathbf{A}^\top \mathbf{c} \neq \mathbf{0}$, then $\mathbf{x}^* \neq \mathbf{0}$ and $\widehat{\mathbf{w}}^* = \widehat{\mathbf{w}}(\mathbf{x}^*)$ is c -optimal, where

$$\widehat{w}_i(\mathbf{x}) = \frac{|x_i|}{\|\mathbf{x}\|_1}, \quad \forall i \in [p], \quad \mathbf{x} \neq \mathbf{0}; \quad (3.2)$$

(iii) If $\mathbf{w}^* \in \mathcal{P}_p$ is c -optimal, then $\widehat{\mathbf{x}}^* = \widehat{\mathbf{x}}(\mathbf{w}^*)$ is optimal for (3.1), where

$$\widehat{x}_i(\mathbf{w}) = w_i \mathbf{a}_i^\top \mathbf{M}^{-1}(\mathbf{w}) \mathbf{c}, \quad \forall i \in [p]. \quad (3.3)$$

(iv) In the pathological case $\mathbf{A}^\top \mathbf{c} = \mathbf{0}$, the unique optimal solution to the quadratic lasso is $\mathbf{x}^* = \mathbf{0}$, while every design $\mathbf{w} \in \mathcal{P}_p$ is c -optimal.

Proof We introduce the function $v : \mathbb{R}^p \times \mathcal{P}_p \rightarrow \mathbb{R} \cup \{\infty\}$ by

$$v(\mathbf{x}, \mathbf{w}) = \|\mathbf{A}\mathbf{x} - \mathbf{c}\|^2 + \lambda \sum_{i=1}^p \frac{x_i^2}{w_i}, \quad (3.4)$$

where the function is defined by continuity when $w_i = 0$, that is, we assume that $\frac{x_i^2}{0} = 0$ if $x_i = 0$ and $\frac{x_i^2}{0} = \infty$ otherwise¹. We note that v is convex, as $\mathbf{x} \mapsto \|\mathbf{A}\mathbf{x} - \mathbf{c}\|^2$ is convex and $(x_i, w_i) \mapsto \frac{x_i^2}{w_i} = w_i \cdot \left(\frac{x_i}{w_i}\right)^2$ is the *perspective function* of $x_i \mapsto x_i^2$; see, e.g., Boyd and Vandenberghe (2004).

We next use the fact that for all $\mathbf{x} \in \mathbb{R}^p \setminus \{\mathbf{0}\}$, the function $\mathbf{w} \mapsto \sum_i x_i^2/w_i$ is minimized over \mathcal{P}_p for $\mathbf{w} = \widehat{\mathbf{w}}(\mathbf{x})$ given by (3.2). Therefore, it holds

$$\min_{\mathbf{w} \in \mathcal{P}_p} v(\mathbf{x}, \mathbf{w}) = v(\mathbf{x}, \widehat{\mathbf{w}}(\mathbf{x})) = \|\mathbf{A}\mathbf{x} - \mathbf{c}\|^2 + \lambda \sum_{i=1}^p \frac{x_i^2}{|x_i|} \cdot \|\mathbf{x}\|_1 = \mathcal{L}_\lambda(\mathbf{x}), \quad (3.5)$$

and the equation $\min_{\mathbf{w} \in \mathcal{P}_p} v(\mathbf{x}, \mathbf{w}) = \mathcal{L}_\lambda(\mathbf{x})$ remains valid for $\mathbf{x} = \mathbf{0}$ (with $\mathcal{L}_\lambda(\mathbf{0}) = \|\mathbf{c}\|^2$).

On the other hand, we can also minimize $v(\mathbf{x}, \mathbf{w})$ with respect to \mathbf{x} for a fixed $\mathbf{w} \in \mathcal{P}_p$. If $w_i = 0$, then all minimizers $\widehat{\mathbf{x}}$ of $v(\mathbf{x}, \mathbf{w})$ must satisfy $x_i = 0$, as otherwise $v(\mathbf{x}, \mathbf{w}) = \infty$. We thus restrict to optimizing the other coordinates of x_i , which corresponds to minimizing a function of the same form as $\mathbf{x} \mapsto v(\mathbf{x}, \mathbf{w})$, but for some $\mathbf{w} \neq \mathbf{0}$. We thus assume w.l.o.g. that $\mathbf{w} \neq \mathbf{0}$. Denote by $\mathbf{D}(\mathbf{w}) = \text{diag}\{w_1, \dots, w_p\}$. The problem to solve is in fact a least square problem, as

$$v(\mathbf{x}, \mathbf{w}) = \|\mathbf{A}\mathbf{x} - \mathbf{c}\|^2 + \lambda \mathbf{x}^\top \mathbf{D}^{-1}(\mathbf{w}) \mathbf{x}.$$

The unique minimizer is thus

$$\begin{aligned} \widehat{\mathbf{x}} &= \left[\mathbf{A}^\top \mathbf{A} + \lambda \mathbf{D}^{-1}(\mathbf{w}) \right]^{-1} \mathbf{A}^\top \mathbf{c} = \frac{1}{\lambda} \left[\mathbf{I}_m - \mathbf{D}(\mathbf{w}) \mathbf{A}^\top \mathbf{M}^{-1}(\mathbf{w}) \mathbf{A} \right] \mathbf{D}(\mathbf{w}) \mathbf{A}^\top \mathbf{c} \\ &= \mathbf{D}(\mathbf{w}) \mathbf{A}^\top \mathbf{M}^{-1}(\mathbf{w}) \mathbf{c}, \end{aligned}$$

1. As shown in Sagnol and Pauwels (2019), $v(\mathbf{x}, \mathbf{w})$ actually represents the variance of the unbiased linear estimator $\widehat{\boldsymbol{\xi}} = \mathbf{x}^\top Y + (\mathbf{c} - \mathbf{A}^\top \mathbf{x})^\top Z$ for $\boldsymbol{\xi} = \mathbf{c}^\top \boldsymbol{\theta}$, in the model with observations $Y = \mathbf{A}^\top \boldsymbol{\theta} + \boldsymbol{\epsilon}$, $Z = \boldsymbol{\theta} + \boldsymbol{\nu}$ and errors $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \lambda \text{Diag}\{w_1, \dots, w_p\}^{-1})$, $\boldsymbol{\nu} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_m)$, with $\boldsymbol{\epsilon}$ and $\boldsymbol{\nu}$ mutually independent.

where we have used the Sherman-Morrison-Woodbury identity for the second equality and the definition of $\mathbf{M}(\mathbf{w})$ which gives $\mathbf{A}\mathbf{D}(\mathbf{w})\mathbf{A}^\top = \mathbf{M}(\mathbf{w}) - \lambda\mathbf{I}$ for the last one. Therefore, $\hat{x}_i = w_i \mathbf{a}_i^\top \mathbf{M}^{-1}(\mathbf{w})\mathbf{c}$. Note that this formula is also valid for the coordinates $i \in [p]$ with $w_i = 0$, as we set $\hat{x}_i = 0$ in this case. Thus, $\hat{\mathbf{x}} = \hat{\mathbf{x}}(\mathbf{w})$, as given by (3.3).

Substitution in $v(\mathbf{x}, \mathbf{w})$ yields

$$\begin{aligned} \min_{\mathbf{x} \in \mathbb{R}^p} v(\mathbf{x}, \mathbf{w}) &= v(\hat{\mathbf{x}}(\mathbf{w}), \mathbf{w}) = \|\mathbf{A}\mathbf{D}(\mathbf{w})\mathbf{A}^\top \mathbf{M}^{-1}(\mathbf{w})\mathbf{c} - \mathbf{c}\|^2 + \lambda \sum_{i=1}^p w_i (\mathbf{a}_i^\top \mathbf{M}^{-1}(\mathbf{w})\mathbf{c})^2 \\ &= \|\lambda \mathbf{M}^{-1}(\mathbf{w})\mathbf{c}\|^2 + \lambda \sum_{i=1}^p w_i \cdot \mathbf{c}^\top \mathbf{M}^{-1}(\mathbf{w}) \mathbf{a}_i \mathbf{a}_i^\top \mathbf{M}^{-1}(\mathbf{w}) \mathbf{c} \\ &= \lambda \cdot [\mathbf{c}^\top \mathbf{M}^{-1}(\mathbf{w}) (\lambda \mathbf{I}_m + \sum_{i=1}^p w_i \mathbf{a}_i \mathbf{a}_i^\top) \mathbf{M}^{-1}(\mathbf{w}) \mathbf{c}] \\ &= \lambda \phi_c(\mathbf{w}), \end{aligned} \tag{3.6}$$

where we have used $\mathbf{A}\mathbf{D}(\mathbf{w})\mathbf{A}^\top + \lambda \mathbf{I}_m = \sum_{i=1}^p w_i \mathbf{a}_i \mathbf{a}_i^\top + \lambda \mathbf{I}_m = \mathbf{M}(\mathbf{w})$.

Now, the theorem easily follows from the above observations. For (i), we use that

$$\min_{\mathbf{x} \in \mathbb{R}^p} \mathcal{L}_\lambda(\mathbf{x}) = \min_{\mathbf{x} \in \mathbb{R}^p, \mathbf{w} \in \mathcal{P}_p} v(\mathbf{x}, \mathbf{w}) = \min_{\mathbf{w} \in \mathcal{P}_p} \lambda \phi_c(\mathbf{w}), \tag{3.7}$$

where the first equality results from (3.5) and the second one from (3.6). For (ii), we first prove that $\mathbf{x}^* \neq \mathbf{0}$ holds for every optimal solution to the quadratic lasso whenever $\mathbf{A}^\top \mathbf{c} \neq \mathbf{0}$. Let i be any index such that $\mathbf{a}_i^\top \mathbf{c} \neq 0$. Then, we define $\mathbf{x} = \frac{\mathbf{a}_i^\top \mathbf{c}}{\|\mathbf{a}_i\|^2 + \lambda} \mathbf{e}_i$. It holds

$$\mathcal{L}_\lambda(\mathbf{x}) = \left\| \frac{\mathbf{a}_i^\top \mathbf{c}}{\|\mathbf{a}_i\|^2 + \lambda} \mathbf{a}_i - \mathbf{c} \right\|^2 + \lambda \frac{(\mathbf{a}_i^\top \mathbf{c})^2}{(\|\mathbf{a}_i\|^2 + \lambda)^2} = \|\mathbf{c}\|^2 - \frac{(\mathbf{a}_i^\top \mathbf{c})^2}{\|\mathbf{a}_i\|^2 + \lambda} < \|\mathbf{c}\|^2 = \mathcal{L}_\lambda(\mathbf{0}),$$

which shows that $\mathbf{0}$ cannot be optimal for the quadratic lasso. Now, let \mathbf{x}^* be an optimal solution to the quadratic lasso. We have, for any $\mathbf{w} \in \mathcal{P}_p$,

$$\begin{aligned} \lambda \phi_c(\mathbf{w}) &\stackrel{(3.6)}{=} \min_{\mathbf{x} \in \mathbb{R}^p} v(\mathbf{x}, \mathbf{w}) \geq \min_{\mathbf{x} \in \mathbb{R}^p} \min_{\mathbf{w} \in \mathcal{P}_p} v(\mathbf{x}, \mathbf{w}) \stackrel{(3.5)}{=} \min_{\mathbf{x} \in \mathbb{R}^p} \mathcal{L}_\lambda(\mathbf{x}) = \mathcal{L}_\lambda(\mathbf{x}^*) \\ &\stackrel{(3.5)}{=} v(\mathbf{x}^*, \hat{\mathbf{w}}(\mathbf{x}^*)) \geq \min_{\mathbf{x}' \in \mathbb{R}^p} v(\mathbf{x}', \hat{\mathbf{w}}(\mathbf{x}^*)) \stackrel{(3.6)}{=} \lambda \phi_c(\hat{\mathbf{w}}(\mathbf{x}^*)), \end{aligned}$$

which shows that $\hat{\mathbf{w}}^* = \hat{\mathbf{w}}(\mathbf{x}^*)$ is c -optimal. Similarly, for (iii), let \mathbf{w}^* be a c -optimal design. Then, for any $\mathbf{x} \in \mathbb{R}^p$,

$$\begin{aligned} \mathcal{L}_\lambda(\mathbf{x}) &\stackrel{(3.5)}{=} \min_{\mathbf{w} \in \mathcal{P}_p} v(\mathbf{x}, \mathbf{w}) \geq \min_{\mathbf{w} \in \mathcal{P}_p} \min_{\mathbf{x} \in \mathbb{R}^p} v(\mathbf{x}, \mathbf{w}) \stackrel{(3.6)}{=} \min_{\mathbf{w} \in \mathcal{P}_p} \lambda \phi_c(\mathbf{w}) = \lambda \phi_c(\mathbf{w}^*) \\ &\stackrel{(3.6)}{=} v(\hat{\mathbf{x}}(\mathbf{w}^*), \mathbf{w}^*) \geq \min_{\mathbf{w}' \in \mathcal{P}_p} v(\hat{\mathbf{x}}(\mathbf{w}^*), \mathbf{w}') \stackrel{(3.5)}{=} \mathcal{L}_\lambda(\hat{\mathbf{x}}(\mathbf{w}^*)), \end{aligned}$$

and $\hat{\mathbf{x}}^* = \hat{\mathbf{x}}(\mathbf{w}^*)$ is thus an optimal solution of the quadratic lasso problem.

It only remains to handle the pathological case (iv), that is, $\mathbf{A}^\top \mathbf{c} = \mathbf{0}$. For any $\mathbf{x} \neq \mathbf{0}$, we have $\mathcal{L}_\lambda(\mathbf{x}) = \|\mathbf{Ax} - \mathbf{c}\|^2 + \lambda \|\mathbf{x}\|_1^2 = \|\mathbf{c}\|^2 + \|\mathbf{Ax}\|^2 + \lambda \|\mathbf{x}\|_1^2 > \|\mathbf{c}\|^2 = \mathcal{L}_\lambda(\mathbf{0})$, so the unique optimal solution to (3.1) is $\mathbf{x}^* = \mathbf{0}$. On the other hand, by (3.7), c -optimal designs are minimizers of $\mathbf{w} \mapsto v(\mathbf{x}^*, \mathbf{w}) = v(\mathbf{0}, \mathbf{w})$, which is the constant function $\mathbf{w} \mapsto \|\mathbf{c}\|^2$. ■

For all $\mathbf{x} \in \mathbb{R}^p$, observe that

$$\mathcal{L}_\lambda(\mathbf{x}) = v(\mathbf{x}, \widehat{\mathbf{w}}(\mathbf{x})) \geq \min_{\mathbf{x}' \in \mathbb{R}^p} v(\mathbf{x}', \widehat{\mathbf{w}}(\mathbf{x})) = \lambda \phi_c(\widehat{\mathbf{w}}(\mathbf{x})), \quad (3.8)$$

with equality when \mathbf{x} is such that $\widehat{\mathbf{w}}(\mathbf{x})$ corresponds to c -optimal weights \mathbf{w}^* , and therefore when $x_i = \widehat{x}_i(\mathbf{w}^*) = w_i^* \mathbf{a}_i^\top \mathbf{M}^{-1}(\mathbf{w}^*) \mathbf{c}$, $\forall i \in [p]$. In other words, both minimization problems $\min_{\mathbf{x}} \mathcal{L}_\lambda(\mathbf{x})$ and $\min_{\mathbf{x}} \lambda \phi_c(\widehat{\mathbf{w}}(\mathbf{x}))$ share the same optimal value $\lambda \phi_c(\mathbf{w}^*)$, but the former objective function dominates the latter.

Remark 1. By alternating minimization of (3.4) with respect to \mathbf{w} and \mathbf{x} , with $\mathbf{w}^k \rightarrow \mathbf{x}^{k+1} = \widehat{\mathbf{x}}(\mathbf{w}^k) \rightarrow \mathbf{w}^{k+1} = \widehat{\mathbf{w}}(\mathbf{x}^{k+1}) \rightarrow \dots$, where $\widehat{\mathbf{w}}(\mathbf{x})$ and $\widehat{\mathbf{x}}(\mathbf{w})$ are respectively defined by (3.2) and (3.3), we obtain that $\mathbf{w}^{k+1} = \widehat{\mathbf{w}}(\widehat{\mathbf{x}}(\mathbf{w}^k))$ is given by

$$w_i^{k+1} = \frac{w_i^k |\mathbf{a}_i^\top \mathbf{M}^{-1}(\mathbf{w}^k) \mathbf{c}|}{\sum_{j=1}^p w_j^k |\mathbf{a}_j^\top \mathbf{M}^{-1}(\mathbf{w}^k) \mathbf{c}|}.$$

Since we only consider weights that sum to one, we may rewrite $\phi_c(\mathbf{w})$ as $\phi_c(\mathbf{w}) = \phi'_c(\mathbf{w}) = \mathbf{c}^\top [(\sum_{i=1}^p w_i \mathbf{a}_i \mathbf{a}_i^\top) + \lambda \mathbf{I}_m]^{-1} \mathbf{c}$. Denote by $\nabla \phi'_c(\mathbf{w})$ the gradient of $\phi'_c(\cdot)$ at \mathbf{w} . Since its i -th component $\{\nabla \phi'_c(\mathbf{w})\}_i$ is equal to $-\mathbf{c}^\top \mathbf{M}^{-1}(\mathbf{w}) \mathbf{a}_i \mathbf{a}_i^\top \mathbf{M}^{-1}(\mathbf{w}) \mathbf{c} = -(\mathbf{a}_i^\top \mathbf{M}^{-1}(\mathbf{w}) \mathbf{c})^2$, the alternate minimization algorithm above corresponds to

$$w_i^{k+1} = \frac{w_i^k |\{\nabla \phi'_c(\mathbf{w}^k)\}_i|^{1/2}}{\sum_{j=1}^p w_j^k |\{\nabla \phi'_c(\mathbf{w}^k)\}_j|^{1/2}}, \quad (3.9)$$

which coincides with a variant of the multiplicative weight update algorithm of Fellman (1974) for the minimization of $\phi'_c(\mathbf{w})$; see also Yu (2010). ◁

3.2 The dual quadratic Lasso

The primal corresponds to (3.1); that is, $\min_{\mathbf{x} \in \mathbb{R}^p} \mathcal{L}_\lambda(\mathbf{x})$. Introducing an auxiliary variable $\mathbf{z} = \mathbf{Ax} - \mathbf{c}$, we can write this problem as a saddle point problem

$$\min_{\mathbf{x} \in \mathbb{R}^p, \mathbf{z} \in \mathbb{R}^m} \max_{\mathbf{y} \in \mathbb{R}^m} \lambda \|\mathbf{x}\|_1^2 + \|\mathbf{z}\|^2 + 2\mathbf{y}^\top (\mathbf{z} - \mathbf{Ax} + \mathbf{c}), \quad (3.10)$$

and therefore the Lagrangian dual problem reads

$$\max_{\mathbf{y} \in \mathbb{R}^m} \left[2\mathbf{y}^\top \mathbf{c} + \min_{\mathbf{z} \in \mathbb{R}^m} \left(\|\mathbf{z}\|^2 + 2\mathbf{y}^\top \mathbf{z} \right) + \min_{\mathbf{x} \in \mathbb{R}^p} \left(\lambda \|\mathbf{x}\|_1^2 - 2\mathbf{y}^\top \mathbf{Ax} \right) \right]. \quad (3.11)$$

Now, recall that the convex conjugate of a function $f : \mathbb{R}^p \rightarrow \mathbb{R}$ is $f^* : \mathbf{y} \mapsto \sup_{\mathbf{x}} \mathbf{x}^\top \mathbf{y} - f(\mathbf{x})$, so the above problem can be rewritten as

$$\max_{\mathbf{y} \in \mathbb{R}^m} \left[2\mathbf{y}^\top \mathbf{c} - 2g_2^*(-\mathbf{y}) - 2\lambda g_1^*\left(\frac{\mathbf{A}^\top \mathbf{y}}{\lambda}\right) \right],$$

where $g_1(\mathbf{x}) = \|\mathbf{x}\|_1^2/2$ and $g_2(\mathbf{x}) = \|\mathbf{x}\|^2/2$. The dual problem then takes a simple form, noting that for any norm $\|\cdot\|$, the convex conjugate of $g : \mathbf{x} \mapsto \|\mathbf{x}\|^2/2$ is $g^* : \mathbf{y} \mapsto \|\mathbf{y}\|_*^2/2$, where $\|\cdot\|_*$ is the dual norm of $\|\cdot\|$; see Boyd and Vandenberghe (2004, Chapter 3). Hence the expression to be maximized in (3.11) is $2\mathbf{y}^\top \mathbf{c} - \|\mathbf{y}\|^2 - \|\mathbf{A}^\top \mathbf{y}\|_\infty^2/\lambda$. To summarize, since $2\mathbf{y}^\top \mathbf{c} - \|\mathbf{y}\|^2 = \|\mathbf{c}\|^2 - \|\mathbf{y} - \mathbf{c}\|^2$, the dual (quadratic) Lasso is

$$\max_{\mathbf{y} \in \mathbb{R}^m} \mathcal{D}_\lambda(\mathbf{y}) = \|\mathbf{c}\|^2 - \|\mathbf{y} - \mathbf{c}\|^2 - \frac{\|\mathbf{A}^\top \mathbf{y}\|_\infty^2}{\lambda}. \quad (3.12)$$

Moreover, there is no duality gap, since both the primal and dual problems are unconstrained, and Slater's constraint qualification trivially holds (see, e.g., Boyd and Vandenberghe (2004)):

$$\min_{\mathbf{x} \in \mathbb{R}^p} \mathcal{L}_\lambda(\mathbf{x}) = \max_{\mathbf{y} \in \mathbb{R}^m} \mathcal{D}_\lambda(\mathbf{y}). \quad (3.13)$$

3.3 Dual-optimal solution

The following theorem gives a necessary and sufficient condition for a pair $(\mathbf{x}^*, \mathbf{y}^*)$ to be primal-dual optimal for the quadratic Lasso.

Theorem 3.2. *A pair $(\mathbf{x}^*, \mathbf{y}^*) \in \mathbb{R}^p \times \mathbb{R}^m$ is primal-dual optimal, i.e., $\mathcal{D}_\lambda(\mathbf{y}^*) = \mathcal{L}_\lambda(\mathbf{x}^*)$, if and only if*

$$\mathbf{y}^* = \mathbf{c} - \mathbf{A}\mathbf{x}^*, \quad \|\mathbf{A}^\top \mathbf{y}^*\|_\infty = \lambda \|\mathbf{x}^*\|_1, \quad \text{and } \forall i \in [p], \begin{cases} |\mathbf{a}_i^\top \mathbf{y}^*| \leq \lambda \|\mathbf{x}^*\|_1 & \text{if } x_i^* = 0; \\ \mathbf{a}_i^\top \mathbf{y}^* = \lambda \text{sign}(x_i^*) \|\mathbf{x}^*\|_1 & \text{otherwise.} \end{cases} \quad (3.14)$$

In particular,

$$x_i^* (\|\mathbf{A}^\top \mathbf{y}^*\|_\infty - |\mathbf{a}_i^\top \mathbf{y}^*|) = 0 \text{ for all } i \in [p]. \quad (3.15)$$

Moreover, the dual optimal point \mathbf{y}^* is unique and satisfies $\mathbf{y}^* = \lambda \mathbf{M}_*^{-1} \mathbf{c}$, where $\mathbf{M}_* = \mathbf{M}(\widehat{\mathbf{w}}(\mathbf{x}^*))$, with $\widehat{\mathbf{w}}(\mathbf{x})$ given by (3.2).

Proof (i) *Uniqueness of the dual optimal solution.* The dual problem $\max_{\mathbf{y} \in \mathbb{R}^m} \mathcal{D}_\lambda(\mathbf{y})$ is equivalent to

$$\min_{\mathbf{y} \in \mathbb{R}^m, \|\mathbf{A}^\top \mathbf{y}\|_\infty \leq \sqrt{\lambda} u} \|\mathbf{y} - \mathbf{c}\|^2 + u^2 = \min_{\bar{\mathbf{y}} \in \mathcal{P}_\lambda(\mathbf{A})} \|\bar{\mathbf{y}} - \bar{\mathbf{c}}\|^2, \quad (3.16)$$

where $\bar{\mathbf{y}} = (\mathbf{y}^\top, u)^\top$ and $\bar{\mathbf{c}} = (\mathbf{c}^\top, 0)^\top$ belong to \mathbb{R}^{m+1} and $\mathcal{P}_\lambda(\mathbf{A})$ is the polyhedral cone

$$\mathcal{P}_\lambda(\mathbf{A}) = \{\bar{\mathbf{y}} = (\mathbf{y}^\top, u)^\top \in \mathbb{R}^{m+1} : |\mathbf{a}_i^\top \mathbf{y}| \leq \sqrt{\lambda} u \text{ for all } i = 1, \dots, m\}, \quad (3.17)$$

showing that the optimal solution for the dual corresponds to the unique orthogonal projection of $\bar{\mathbf{c}}$ onto $\mathcal{P}_\lambda(\mathbf{A})$; see Figure 1 for an illustration.

(ii) *Optimality conditions.* The optimization problem is convex and Slater's condition holds, so the Karush-Kuhn-Tucker (KKT) conditions are necessary and sufficient to characterize a pair $(\mathbf{x}^*, \mathbf{y}^*)$ of primal-dual optimal solutions. The primal problem consists of minimizing $2(\lambda g_1(\mathbf{x}) + g_2(\mathbf{z}))$ under the constraint $\mathbf{z} = \mathbf{A}\mathbf{x} - \mathbf{c}$, so the KKT system reduces to primal feasibility, i.e., $\mathbf{z}^* = \mathbf{A}\mathbf{x}^* - \mathbf{c}$, and stationarity of the Lagrangian (3.10), i.e.,

$$\mathbf{0} \in \partial \left(\mathbf{x}, \mathbf{z} \mapsto 2(\lambda g_1(\mathbf{x}) + g_2(\mathbf{z}) + \mathbf{y}^\top (\mathbf{z} - \mathbf{A}\mathbf{x} + \mathbf{c})) \right) \Big|_{\mathbf{x}=\mathbf{x}^*, \mathbf{z}=\mathbf{z}^*}.$$

Differentiating with respect to \mathbf{z} gives $\nabla_{\mathbf{z}}(\frac{1}{2}\|\mathbf{z}\|^2 + \mathbf{y}^{*\top}\mathbf{z})|_{\mathbf{z}=\mathbf{z}^*} = \mathbf{z}^* + \mathbf{y}^* = \mathbf{0}$, which already shows that $\mathbf{y}^* = -\mathbf{z}^* = \mathbf{c} - \mathbf{A}\mathbf{x}^*$.

As the Lagrangian (3.10) is not differentiable with respect to \mathbf{x} , the stationarity condition with respect to \mathbf{x} is slightly more complicated. We must solve

$$\mathbf{0} \in \partial(\mathbf{x} \mapsto \lambda g_1(\mathbf{x}) - \mathbf{y}^{*\top} \mathbf{A}\mathbf{x})|_{\mathbf{x}=\mathbf{x}^*},$$

which is equivalent to $\mathbf{A}^\top \mathbf{y}^* / \lambda \in \partial g_1(\mathbf{x}^*)$. Since $\partial g_1(\mathbf{x}) = \{\|\mathbf{x}\|_1 \cdot \mathbf{u} : \mathbf{u} \in \partial\|\mathbf{x}\|_1\}$ (see, e.g., Sagnol and Pauwels (2019)) and since $\mathbf{u} \in \partial\|\mathbf{x}\|_1$ if and only if $u_i \in [-1, 1]$ whenever $x_i = 0$ and $u_i = \text{sign}(x_i)$ otherwise, we get the following equivalent condition:

$$\begin{aligned} & \exists \mathbf{u} \in \partial\|\mathbf{x}^*\|_1 : \quad \mathbf{A}^\top \mathbf{y}^* = \lambda \|\mathbf{x}^*\|_1 \mathbf{u} \\ \iff & \forall i \in [p], \quad (x_i^* = 0 \wedge |\mathbf{a}_i^\top \mathbf{y}^*| \leq \lambda \|\mathbf{x}^*\|_1) \quad \text{or} \quad (x_i^* \neq 0 \wedge \mathbf{a}_i^\top \mathbf{y}^* = \lambda \text{sign}(x_i^*) \|\mathbf{x}^*\|_1). \end{aligned}$$

(iii) *Alternative expression of \mathbf{y}^* .* We prove that $\mathbf{M}_*(\mathbf{c} - \mathbf{A}\mathbf{x}^*) = \lambda \mathbf{c}$, which implies $\mathbf{c} - \mathbf{A}\mathbf{x}^* = \lambda \mathbf{M}_*^{-1} \mathbf{c}$, as desired. Let $\mathbf{D}_* = \mathbf{D}(\widehat{\mathbf{w}}(\mathbf{x}^*)) = (1/\|\mathbf{x}^*\|_1) \text{diag}\{x_1^*, \dots, x_p^*\}$. We have


$$\begin{aligned} \mathbf{M}_*(\mathbf{c} - \mathbf{A}\mathbf{x}^*) - \lambda \mathbf{c} &= (\mathbf{M}_* - \lambda \mathbf{I})(\mathbf{c} - \mathbf{A}\mathbf{x}^*) - \lambda \mathbf{A}\mathbf{x}^* = \mathbf{A} \mathbf{D}_* \mathbf{A}^\top \mathbf{y}^* - \lambda \mathbf{A}\mathbf{x}^* \\ &= \mathbf{A}[\mathbf{D}_* \mathbf{A}^\top \mathbf{y}^* - \lambda \mathbf{x}^*]. \end{aligned}$$

Finally, the optimality conditions imply $|\mathbf{a}_i^*| \mathbf{a}_i^\top \mathbf{y}^* = \lambda x_i \|\mathbf{x}\|_1^*$, $\forall i \in [p]$, hence $\mathbf{D}_* \mathbf{A}^\top \mathbf{y}^* - \lambda \mathbf{x}^* = \mathbf{0}$. \blacksquare

We also recall a well known analogous optimality condition for the problem of c -optimal design. Based on a stationarity condition, the so-called *Equivalence Theorem* from optimal design theory states that $\mathbf{w}^* \in \mathcal{S}_p$ is c -optimal if and only if the directional derivative of $\phi_c(\cdot)$ in the direction of any vertex of \mathcal{S}_p is non-negative; see, e.g., Silvey (1980); Pukelsheim (1993).

Theorem 3.3. *Suppose that the \mathbf{H}_i 's in \mathbb{H} satisfy (2.5). Then the vector of weights \mathbf{w}^* is c -optimal if and only if $\mathbf{c}^\top \mathbf{M}^{-1}(\mathbf{w}^*)(\mathbf{a}_i \mathbf{a}_i^\top + \lambda \mathbf{I}_m) \mathbf{M}^{-1}(\mathbf{w}^*) \mathbf{c} \leq \phi_c(\mathbf{w}^*)$ for all $i \in [p]$, with equality for all i such that $w_i^* > 0$.*

Note that the theorem implies that $w_i^*(|\mathbf{a}_i^\top \mathbf{M}^{-1}(\mathbf{w}^*) \mathbf{c}| - \|\mathbf{A}^\top \mathbf{M}^{-1}(\mathbf{w}^*) \mathbf{c}\|_\infty) = 0$ for all $i \in [p]$ when \mathbf{w}^* is c -optimal. In particular, $|\mathbf{a}_i^\top \mathbf{M}^{-1}(\mathbf{w}^*) \mathbf{c}|$ does not depend on $i \in [p]$ as soon as $w_i^* > 0$, and the elements of $\widehat{\mathbf{x}}(\mathbf{w}^*)$ must be of the form $\widehat{x}_i(\mathbf{w}^*) = \pm w_i^* \|\mathbf{A}^\top \mathbf{M}^{-1}(\mathbf{w}^*) \mathbf{c}\|_\infty$.


 Figure 1: Uniqueness of the dual optimal solution $\bar{\mathbf{y}}^*$.

Hence, $|\hat{x}_i(\mathbf{w}^*)|$ is proportional to w_i^* and we have $\widehat{\mathbf{w}}(\widehat{\mathbf{x}}(\mathbf{w}^*)) = \mathbf{w}^*$, i.e., \mathbf{w}^* is a fixed point of the mapping $\mathbf{w} \mapsto \widehat{\mathbf{w}}(\widehat{\mathbf{x}}(\mathbf{w}))$.

We next prove an important property about $\widehat{\mathbf{x}}(\mathbf{w})$. For $\mathbf{w} \in \mathcal{P}_p$ and $\mathbf{x} \in \mathbb{R}^p$, we define

$$\mathbf{y}_1(\mathbf{x}) = \mathbf{c} - \mathbf{A}\mathbf{x} \quad \text{and} \quad \mathbf{y}_2(\mathbf{w}) = \lambda \mathbf{M}^{-1}(\mathbf{w})\mathbf{c}. \quad (3.18)$$

Note that Theorem 3.2 implies that $\mathbf{y}_1(\mathbf{x}^*) = \mathbf{y}_2(\widehat{\mathbf{w}}(\mathbf{x}^*))$ for any optimal solution \mathbf{x}^* of the quadratic lasso problem. The next lemma states that the equation obtained by inverting the role of \mathbf{w} and \mathbf{x} holds for any design $\mathbf{w} \in \mathcal{P}_p$.

Lemma 3.4. *Let $\mathbf{w} \in \mathcal{P}_p$. Then,*

$$\mathbf{y}_1(\widehat{\mathbf{x}}(\mathbf{w})) = \mathbf{y}_2(\mathbf{w}), \quad (3.19)$$

where $\widehat{\mathbf{x}}(\mathbf{w})$ is defined by (3.3), and $\mathbf{y}_1, \mathbf{y}_2$ are defined by (3.18).

Proof Denote $\mathbf{D} = \mathbf{D}(\mathbf{w}) = \text{diag}\{w_1, \dots, w_p\}$. By definition, $\mathbf{M} = \mathbf{M}(\mathbf{w}) = \mathbf{A}\mathbf{D}\mathbf{A}^\top + \lambda \mathbf{I}_m$ and (3.3) gives $\widehat{\mathbf{x}}(\mathbf{w}) = \mathbf{D}\mathbf{A}^\top \mathbf{M}^{-1}\mathbf{c}$. It follows that

$$\mathbf{y}_1(\widehat{\mathbf{x}}(\mathbf{w})) = \mathbf{c} - \mathbf{A}\widehat{\mathbf{x}}(\mathbf{w}) = \mathbf{c} - \mathbf{A}\mathbf{D}\mathbf{A}^\top \mathbf{M}^{-1}\mathbf{c} = \left[\mathbf{M} - \mathbf{A}\mathbf{D}\mathbf{A}^\top \right] \mathbf{M}^{-1}\mathbf{c} = \lambda \mathbf{M}^{-1}\mathbf{c} = \mathbf{y}_2(\mathbf{w}).$$

■

We conclude this section by making explicit the connection between the quadratic and standard lasso problems. This connection will be exploited in Section 4 to derive a new algorithm for the computation of c -optimal designs. The (standard) lasso problem reads

$$\min_{\mathbf{x} \in \mathbb{R}^p} \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{c}\|^2 + \alpha \|\mathbf{x}\|_1, \quad (3.20)$$

that is, (3.20) is the Lagrange relaxation of the minimization of $\frac{1}{2}\|\mathbf{Ax} - \mathbf{c}\|^2$ under the constraint $\|\mathbf{x}\|_1 < t$, or equivalently $\|\mathbf{x}\|_1^2 < t^2$, for some $t > 0$. The formulation (3.1) with a squared penalty $\|\mathbf{x}\|_1^2$ thus corresponds to the Lagrangian relaxation of the same problem, with the precise relationship between t^2 and λ depending on \mathbf{c} and \mathbf{A} . (Note that there is a factor $\frac{1}{2}$ in front of $\|\mathbf{Ax} - \mathbf{c}\|^2$ in the standard lasso problem (3.20), but not in (3.1): this factor is commonly introduced in the standard lasso to have full symmetry with the dual problem; we have refrained from introducing such a factor for the quadratic lasso problem (3.1), as this would make the formulas of the next section more complex.)

Theorem 3.5. *Let $\alpha < \|\mathbf{A}^\top \mathbf{c}\|_\infty$ and \mathbf{x}^* be an optimal solution to the standard lasso problem (3.20). Then, $\mathbf{x}^* \neq \mathbf{0}^*$ and \mathbf{x}^* is an optimal solution to the quadratic lasso (3.1) with $\lambda = \lambda(\alpha) = \alpha/\|\mathbf{x}^*\|_1$. Conversely, for all $\lambda > 0$, if $\mathbf{A}^\top \mathbf{c} \neq \mathbf{0}$ and \mathbf{x}^* is an optimal solution to the quadratic lasso (3.1), then \mathbf{x}^* is also an optimal solution to the standard lasso (3.20) with $\alpha = \alpha(\lambda) = \lambda\|\mathbf{x}^*\|_1$.*

Proof This property is a simple consequence of the KKT conditions of the standard lasso problem (3.20), which can be found, e.g., in Xiang et al. (2016), and read as follows: \mathbf{x}^* solves (3.20) if and only if there exists a dual vector $\boldsymbol{\theta}^*$ such that

$$\mathbf{c} = \mathbf{Ax}^* + \alpha\boldsymbol{\theta}^* \quad \text{and} \quad \forall i \in [p], \mathbf{a}_i^\top \boldsymbol{\theta}^* \in \begin{cases} \{\text{sign}(x_i^*)\} & \text{if } x_i^* \neq 0; \\ [-1, 1] & \text{if } x_i^* = 0. \end{cases} \quad (3.21)$$

First note that $\mathbf{x}^* = \mathbf{0}$ can only satisfy (3.21) if $\boldsymbol{\theta}^* = \mathbf{c}/\alpha$, which implies $\|\mathbf{A}^\top \mathbf{c}\|_\infty = \alpha\|\mathbf{A}^\top \boldsymbol{\theta}^*\|_\infty \leq \alpha$.

Now, let $\alpha < \|\mathbf{A}^\top \mathbf{c}\|_\infty$ and $\mathbf{x}^* \neq \mathbf{0}$ be an optimal solution to the standard lasso problem (3.20), and denote by $\boldsymbol{\theta}^*$ a corresponding optimal dual vector. Define $\lambda = \frac{\alpha}{\|\mathbf{x}^*\|_1}$, and $\mathbf{y}^* = \alpha\boldsymbol{\theta}^* = \lambda\|\mathbf{x}^*\|_1\boldsymbol{\theta}^*$. Clearly, (3.21) and $\mathbf{x}^* \neq \mathbf{0}$ imply that $(\mathbf{x}^*, \mathbf{y}^*)$ fulfills the KKT system (3.14), hence \mathbf{x}^* is an optimal solution to the quadratic lasso (3.1). Conversely, if $\mathbf{A}^\top \mathbf{c} \neq \mathbf{0}$ and \mathbf{x}^* solves the quadratic lasso, we know from Theorem 3.1 that $\mathbf{x}^* \neq \mathbf{0}$, so $\alpha = \alpha(\lambda) = \lambda\|\mathbf{x}^*\|_1 > 0$. Further, if $(\mathbf{x}^*, \mathbf{y}^*)$ solves the KKT equations (3.14), then $(\mathbf{x}^*, \boldsymbol{\theta}^*)$ solves (3.21) with $\boldsymbol{\theta}^* = \mathbf{y}^*/\alpha$. Thus, \mathbf{x}^* solves the standard lasso with regularizer α . \blacksquare

3.4 Screening rules for inessential \mathbf{a}_i 's

The objective is to derive an inequality that must be satisfied by any \mathbf{a}_i that may support an optimal design. To this end, we first use information about the current iterate \mathbf{x} for Problem (3.1) to construct a dual solution $\mathbf{y}(\mathbf{x})$. Then, we derive a bound on $\|\mathbf{y}(\mathbf{x}) - \mathbf{y}^*\|$ and exploit the optimality conditions of the dual problem (3.16).

Theorem 3.6. *Let $\mathbf{y} \in \mathbb{R}^m$ be an ϵ -suboptimal dual solution, i.e., $\mathcal{D}_\lambda(\mathbf{y}) \geq \mathcal{D}_\lambda(\mathbf{y}^*) - \epsilon$. Then, any \mathbf{a}_i satisfying*

$$D_0(\mathbf{a}_i; \mathbf{y}, \epsilon) = \|\mathbf{A}^\top \mathbf{y}\|_\infty - |\mathbf{a}_i^\top \mathbf{y}| - \sqrt{\epsilon(\|\mathbf{a}_i\|^2 + \lambda)} > 0$$

cannot support an optimal design.

Proof We use the same construction as in part (i) of the proof of Theorem 3.2, and denote $\bar{\mathbf{y}} = (\mathbf{y}^\top, u)^\top$, $\bar{\mathbf{y}}^* = (\mathbf{y}^{*\top}, u^*)^\top$ and $\bar{\mathbf{c}} = (\mathbf{c}^\top, 0)^\top$, with \mathbf{y}^* the optimal solution of the dual problem (3.12), $u = \|\mathbf{A}^\top \mathbf{y}\|_\infty / \sqrt{\lambda}$ and $u^* = \|\mathbf{A}^\top \mathbf{y}^*\|_\infty / \sqrt{\lambda}$. Then, $\mathcal{D}_\lambda(\mathbf{y}) = \|\mathbf{c}\|^2 - \|\bar{\mathbf{y}} - \bar{\mathbf{c}}\|^2$ and $\mathcal{D}_\lambda(\mathbf{y}^*) = \|\mathbf{c}\|^2 - \|\bar{\mathbf{y}}^* - \bar{\mathbf{c}}\|^2$. The quantity $(\bar{\mathbf{y}} - \bar{\mathbf{y}}^*)^\top (\bar{\mathbf{y}}^* - \bar{\mathbf{c}})$ is nonnegative, as $\bar{\mathbf{y}} \in \mathcal{P}_\lambda(A)$ is feasible for the conic formulation of the dual problem (3.16), and by optimality of $\bar{\mathbf{y}}^*$ the gradient $\nabla \|\bar{\mathbf{y}} - \bar{\mathbf{c}}\|_{\bar{\mathbf{y}} = \bar{\mathbf{y}}^*}^2 = 2(\bar{\mathbf{y}}^* - \bar{\mathbf{c}})$ defines a supporting hyperplane to $\mathcal{P}_\lambda(A)$.

If we denote $r = \|\bar{\mathbf{y}} - \bar{\mathbf{c}}\|$ and $r_* = \|\bar{\mathbf{y}}^* - \bar{\mathbf{c}}\|$, then we can write $r^2 = \|\bar{\mathbf{y}} - \bar{\mathbf{y}}^* + \bar{\mathbf{y}}^* - \bar{\mathbf{c}}\|^2 = \|\bar{\mathbf{y}} - \bar{\mathbf{y}}^*\|^2 + r_*^2 + 2(\bar{\mathbf{y}} - \bar{\mathbf{y}}^*)^\top (\bar{\mathbf{y}}^* - \bar{\mathbf{c}})$, hence $r^2 - r_*^2 \geq \|\bar{\mathbf{y}} - \bar{\mathbf{y}}^*\|^2$. In addition,

$$r^2 - r_*^2 = \mathcal{D}_\lambda(\mathbf{y}^*) - \mathcal{D}_\lambda(\mathbf{y}) \leq \epsilon.$$

We thus obtain $\|\bar{\mathbf{y}} - \bar{\mathbf{y}}^*\|^2 \leq r^2 - r_*^2 \leq \epsilon$; i.e. $\bar{\mathbf{y}}^* \in \mathcal{B}_{m+1}(\bar{\mathbf{y}}, \sqrt{\epsilon})$.

Due to the KKT-condition (3.14) for Problem (3.16), it holds $x_i^*(|\mathbf{a}_i^\top \mathbf{y}^*| - \sqrt{\lambda} u^*) = 0$ for all $i \in [p]$. Therefore, any \mathbf{a}_i that supports an optimal design corresponds to an active constraint at the optimum, and thus satisfies $|\mathbf{a}_i^\top \mathbf{y}^*| = \sqrt{\lambda} u^*$. But $\bar{\mathbf{y}}^* \in \mathcal{B}_{m+1}(\bar{\mathbf{y}}, \sqrt{\epsilon})$ implies

$$|\mathbf{a}_i^\top \mathbf{y}^*| - \sqrt{\lambda} u^* \leq |\mathbf{a}_i^\top \mathbf{y}| - \sqrt{\lambda} u + \sqrt{\epsilon(\|\mathbf{a}_i\|^2 + \lambda)},$$

showing that any \mathbf{a}_i such that $|\mathbf{a}_i^\top \mathbf{y}| < \sqrt{\lambda} u - \sqrt{\epsilon(\|\mathbf{a}_i\|^2 + \lambda)}$ cannot support an optimal design. \blacksquare

Specializing the above result for the dual solutions $\mathbf{y}_1(\mathbf{x}) = \mathbf{c} - \mathbf{A}\mathbf{x}$ and $\mathbf{y}_2(\hat{\mathbf{w}}(\mathbf{x})) = \lambda \mathbf{M}^{-1}(\hat{\mathbf{w}}(\mathbf{x}))\mathbf{c}$ (which are reasonable choices by Theorem 3.2), we obtain the screening criteria given in Corollaries 3.7 and 3.8.

Corollary 3.7. *Let $\mathbf{x} \in \mathbb{R}^p$, and define*

$$\begin{aligned} D_1(\mathbf{a}_i; \mathbf{x}) &= D_0(\mathbf{a}_i; \mathbf{y}_1, \mathcal{L}_\lambda(\mathbf{x}) - \mathcal{D}_\lambda(\mathbf{y}_1)) \\ &= \|\mathbf{v}\|_\infty - |v_i| - \left(\frac{\|\mathbf{v}\|_\infty^2}{\lambda} + \lambda \|\mathbf{x}\|_1^2 - 2\mathbf{x}^\top \mathbf{v} \right)^{1/2} \sqrt{\|\mathbf{a}_i\|^2 + \lambda}, \end{aligned} \quad (3.22)$$

where $\mathbf{y}_1 = \mathbf{y}_1(\mathbf{x}) = \mathbf{c} - \mathbf{A}\mathbf{x}$ and $\mathbf{v} = \mathbf{v}(\mathbf{x}) = \mathbf{A}^\top \mathbf{y}_1(\mathbf{x})$. Then, any \mathbf{a}_i satisfying $D_1(\mathbf{a}_i; \mathbf{x}) > 0$ cannot support an optimal design.

Proof The result directly follows from Theorem 3.6. The expression $\mathcal{L}_\lambda(\mathbf{x}) - \mathcal{D}_\lambda[\mathbf{y}_1(\mathbf{x})]$ is an upper bound on the duality gap at \mathbf{x} , as $\mathcal{D}_\lambda(\mathbf{y}^*) - \mathcal{D}_\lambda[\mathbf{y}_1(\mathbf{x})] = \mathcal{L}_\lambda(\mathbf{x}^*) - \mathcal{D}_\lambda[\mathbf{y}_1(\mathbf{x})] \leq \mathcal{L}_\lambda(\mathbf{x}) - \mathcal{D}_\lambda[\mathbf{y}_1(\mathbf{x})]$. Direct calculation gives

$$\mathcal{L}_\lambda(\mathbf{x}) - \mathcal{D}_\lambda[\mathbf{y}_1(\mathbf{x})] = \frac{\|\mathbf{v}\|_\infty^2}{\lambda} + \lambda \|\mathbf{x}\|_1^2 - 2\mathbf{x}^\top \mathbf{v},$$

which gives the expression (3.22) for D_1 . \blacksquare

Corollary 3.8. *Let $\mathbf{x} \in \mathbb{R}^p \setminus \{\mathbf{0}\}$, and define*

$$\begin{aligned} D_2(\mathbf{a}_i; \mathbf{x}) &= D_0\left(\mathbf{a}_i; \mathbf{y}_2, \lambda \mathbf{c}^\top \mathbf{M}^{-1}(\widehat{\mathbf{w}}(\mathbf{x}))\mathbf{c} - \mathcal{D}_\lambda(\mathbf{y}_2)\right) \\ &= \|\mathbf{v}'\|_\infty - |v'_i| - \left(\frac{\|\mathbf{v}'\|_\infty^2}{\lambda} + \mathbf{y}_2^\top [\mathbf{y}_2 - \mathbf{c}]\right)^{1/2} \sqrt{\|\mathbf{a}_i\|^2 + \lambda}, \end{aligned} \quad (3.23)$$

where $\mathbf{y}_2 = \mathbf{y}_2(\widehat{\mathbf{w}}(\mathbf{x})) = \lambda \mathbf{M}^{-1}(\widehat{\mathbf{w}}(\mathbf{x}))\mathbf{c}$ and $\mathbf{v}' = \mathbf{v}'(\mathbf{x}) = \mathbf{A}^\top \mathbf{y}_2(\widehat{\mathbf{w}}(\mathbf{x}))$. Then, any \mathbf{a}_i satisfying $D_2(\mathbf{a}_i; \mathbf{x}) > 0$ cannot support an optimal design.

Proof This time, we use $\lambda\phi_c(\widehat{\mathbf{w}}(\mathbf{x})) - \mathcal{D}_\lambda[\mathbf{y}_2(\widehat{\mathbf{w}}(\mathbf{x}))]$ as an upper bound on the duality gap at \mathbf{x} , which follows from $\mathcal{D}_\lambda(\mathbf{y}^*) - \mathcal{D}_\lambda[\mathbf{y}_2(\widehat{\mathbf{w}}(\mathbf{x}))] = \lambda\phi_c(\mathbf{w}^*) - \mathcal{D}_\lambda[\mathbf{y}_2(\widehat{\mathbf{w}}(\mathbf{x}))] \leq \lambda\phi_c(\widehat{\mathbf{w}}(\mathbf{x})) - \mathcal{D}_\lambda[\mathbf{y}_2(\widehat{\mathbf{w}}(\mathbf{x}))]$. Direct calculation gives $\lambda \mathbf{c}^\top \mathbf{M}^{-1}(\widehat{\mathbf{w}}(\mathbf{x}))\mathbf{c} - \mathcal{D}_\lambda[\mathbf{y}_2(\widehat{\mathbf{w}}(\mathbf{x}))] = \|\mathbf{v}'\|_\infty^2/\lambda + \mathbf{y}_2^\top(\widehat{\mathbf{w}}(\mathbf{x}))[\mathbf{y}_2(\widehat{\mathbf{w}}(\mathbf{x})) - \mathbf{c}]$ and therefore (3.23). \blacksquare

More generally, as $\lambda\phi_c(\widehat{\mathbf{w}}(\mathbf{x})) \leq \mathcal{L}_\lambda(\mathbf{x})$, see (3.8), the substitution of $\lambda\phi_c(\widehat{\mathbf{w}}(\mathbf{x}))$ for $\mathcal{L}_\lambda(\mathbf{x})$ in the construction of $D_1(\mathbf{a}_i; \mathbf{x})$ yields a better bound

$$\widetilde{D}_1(\mathbf{a}_i; \mathbf{x}) = D_0(\mathbf{a}_i; \mathbf{y}_1(\mathbf{x}), \lambda\phi_c(\widehat{\mathbf{w}}(\mathbf{x})) - \mathcal{D}_\lambda(\mathbf{y}_1(\mathbf{x}))) \geq D_1(\mathbf{a}_i; \mathbf{x}).$$

The reason for using this refined bound on the duality gap in D_2 (see Corollary 3.8) but not in D_1 is computational efficiency: the calculation of D_1 is dominated by the computation of \mathbf{v} , which requires $O(mp)$ operations, while the calculation of D_2 is dominated by the computation of $\mathbf{y}_2(\widehat{\mathbf{w}}(\mathbf{x}))$ that takes $O(m^2p)$ operations. However, if we are willing to pay this higher computational cost then we might as well take advantage of the improved bound on the duality gap, hence its use in D_2 .

In practice, the exact complexity of a screening operation depends on the algorithm used to solve the optimization problem (3.1). The above discussion assumes that this algorithm works directly on the primal variables \mathbf{x} , such as FISTA or the block coordinate descent proposed in (Sagnol and Pauwels, 2019). An example comparing the efficiencies of screening with D_1 and D_2 is given in Section 5.1. However, the situation is quite different when the algorithm operates on the design weights \mathbf{w} , such as the multiplicative update algorithm (3.9): the vector $\mathbf{y}_2(\mathbf{w})$ is then readily obtained from intermediate calculations when computing the gradient $\nabla\phi'_c(\mathbf{w})$, and there is no computational incentive to use \widetilde{D}_1 rather than \widetilde{D}_1 anymore. Moreover, in this case, starting from \mathbf{w} , we use $\mathbf{x} = \widehat{\mathbf{x}}(\mathbf{w})$ in \widetilde{D}_1 and $\widehat{\mathbf{w}}(\mathbf{x}) = \mathbf{w}$ in D_2 , see (3.2) and (3.3). Lemma 3.4 then implies that $\mathbf{y}_2(\mathbf{w}) = \mathbf{y}_1(\mathbf{x}) = \lambda \mathbf{M}^{-1}(\mathbf{w})\mathbf{c}$, so that \widetilde{D}_1 and D_2 coincide. Rewriting this bound as a function of \mathbf{w} , we obtain

$$D'_2(\mathbf{a}_i; \mathbf{w}) = \|\mathbf{v}\|_\infty - |v_i| - \left(\frac{\|\mathbf{v}\|_\infty^2}{\lambda} + \mathbf{y}^\top(\mathbf{y} - \mathbf{c})\right)^{1/2} \sqrt{\|\mathbf{a}_i\|^2 + \lambda},$$

where $\mathbf{y} = \mathbf{y}(\mathbf{w}) = \lambda \mathbf{M}^{-1}(\mathbf{w})\mathbf{c}$ and $\mathbf{v} = \mathbf{A}^\top \mathbf{y}$. We thus have the following property.

Corollary 3.9. *Let \mathbf{M} be any matrix in the convex hull of \mathbb{H} with the \mathbf{H}_i satisfying (2.5). Then, any matrix \mathbf{H}_i in \mathbb{H} such that $B(\mathbf{M}, \mathbf{H}_i) > 0$, where*

$$B(\mathbf{M}, \mathbf{H}_i) = \left\{ (1 + \delta) \Phi_c(\mathbf{M}) - \lambda \mathbf{c}^\top \mathbf{M}^{-2} \mathbf{c} \right\}^{1/2} - \left[\delta \Phi_c(\mathbf{M}) \left(1 + \frac{\|\mathbf{a}_i\|^2}{\lambda} \right) \right]^{1/2} - |\mathbf{c}^\top \mathbf{M}^{-1} \mathbf{a}_i|$$

and $\delta = \max_{\mathbf{H}_i \in \mathbb{H}} \mathbf{c}^\top \mathbf{M}^{-1} \mathbf{H}_i \mathbf{M}^{-1} \mathbf{c} / (\mathbf{c}^\top \mathbf{M}^{-1} \mathbf{c}) - 1$, does not support a c -optimal design.

Proof We have $\mathbf{M} = \mathbf{M}(\mathbf{w})$ for some $\mathbf{w} \in \mathcal{P}_p$. Direct calculation gives $D'_2(\mathbf{a}_i; \mathbf{w}) = \lambda B(\mathbf{M}, \mathbf{H}_i)$. \blacksquare

Note that $\delta \geq 0$ and $\Phi_c(\mathbf{M}) \leq (1 + \delta)\phi_c(\mathbf{w}^*)$ with \mathbf{w}^* a c -optimal design, see Pronzato and Sagnol (2021).

4. A homotopy algorithm for the quadratic Lasso

In the standard Lasso problem (3.20) (with non-squared penalty $\alpha\|\mathbf{x}\|_1$), the dual problem geometrically corresponds to the projection of the point \mathbf{c}/α onto the polytope $\mathcal{Q} = \{\mathbf{z} : \|\mathbf{A}^\top \mathbf{z}\|_\infty \leq 1\}$. This was used by Osborne et al. (2000) to design a homotopy algorithm which starts from a large value of α (such that $\mathbf{c}/\alpha \in \mathcal{Q}$ and the projection is trivial) and stores the required data to maintain the projection of \mathbf{c}/α onto \mathcal{Q} as α decreases. One advantage of this technique is that it can be used to compute the full *regularization path*, i.e., a set of optimal solutions $\mathbf{x}^*(\alpha)$ to Problem (3.20) for all $\alpha > 0$. This is particularly interesting in applications of the Lasso algorithms for which the value of the regularization parameter α must be tuned using cross-validation, as in the least angle regression (LARS) method for model selection (Efron et al., 2004). In this section, we use Theorem 3.5 to adapt the homotopy method to the case of a quadratic Lasso penalty, and hence to the case of c -optimality.

In c -optimal design, λ is a fixed constant related to the variance of the observations and the variance-covariance matrix of the prior (see Section 2), and computing the full regularization path $(\mathbf{w}^*(\lambda))_{\lambda \geq 0}$ is not essential. The example in Section 5.1 nevertheless illustrates that the proposed homotopy method can yield a faster design construction than standard first-order algorithms, by several orders of magnitude, although the objective is to solve Problem (2.2) for a unique $\lambda > 0$. On the negative side, one should notice that pathological cases are known for the standard lasso problem (3.20) for which the regularization path can have up to 3^n breakpoints (Mairal and Yu, 2012).

The main idea of the homotopy algorithm is that the projection $\mathbf{z}^*(\alpha)$ of \mathbf{c}/α over \mathcal{Q} has a closed form solution if we know on which face of \mathcal{Q} the projection lies. By tracking the subset of equalities of the form $\mathbf{a}_i^\top \mathbf{z}^*(\alpha) = \pm 1$ that hold as α is decreased, the algorithm can detect the next “breakpoint” in the regularization path, that is, the largest value of α below the current value such that $\mathbf{z}^*(\alpha)$ changes of face of \mathcal{Q} . Moreover, it can be seen that $\mathbf{z}^*(\alpha)$ is linear between two successive breakpoints, and we also obtain a corresponding primal optimal solution $\mathbf{x}^*(\alpha)$ that is continuous and piecewise linear with respect to α . In the end, the algorithm produces a decreasing sequence of breakpoints $\|\mathbf{A}^\top \mathbf{c}\|_\infty = \alpha_1 > \alpha_2 > \dots > \alpha_N = 0$ and of corresponding solutions $\mathbf{x}_1^*, \dots, \mathbf{x}_N^*$, such that $\mathbf{x}^*(\alpha) = \mathbf{x}_1^* = \mathbf{0}$ is optimal for all $\alpha \geq \alpha_1$, and for all k and α of the form $\alpha = (1 - \theta)\alpha_{k+1} + \theta\alpha_k \in [\alpha_{k+1}, \alpha_k]$ for some $\theta \in [0, 1]$, an optimal solution is given by

$$\mathbf{x}^*(\alpha) = (1 - \theta)\mathbf{x}_{k+1}^* + \theta\mathbf{x}_k^*. \tag{4.1}$$

The situation is in fact similar for the quadratic lasso. Recall the geometric interpretation of the dual problem (3.16): $\bar{\mathbf{y}}^* = (\mathbf{y}^{*\top}, u^*)^\top$ is the projection of the vector $\bar{\mathbf{c}} = (\mathbf{c}^\top, 0)^\top$ onto the polyhedral cone $\mathcal{P}_\lambda(\mathbf{A})$ given by (3.17). The projection $\bar{\mathbf{y}}^*$ lies on a face of $\mathcal{P}_\lambda(\mathbf{A})$, which

can be identified with the subset of equalities of the form $\mathbf{a}_i^\top \mathbf{y}^* = \pm\sqrt{\lambda}u^*$ satisfied at $\bar{\mathbf{y}}^*$. The sequence of faces on which $\bar{\mathbf{y}}^*$ lies as λ decreases is in one-to-one correspondence with the faces visited by the homotopy algorithm for the standard lasso problem. Rather than re-inventing the wheel by following the path of $\bar{\mathbf{y}}^*$ on the faces of $\mathcal{P}_\lambda(\mathbf{A})$ as λ decreases, we are going to use Theorem 3.5 to reparametrize the regularization path between two successive breakpoints as a function of λ .

Theorem 4.1. *Let $(\alpha_1, \mathbf{x}_1^*), \dots, (\alpha_N, \mathbf{x}_N^*)$ with $\|\mathbf{A}^\top \mathbf{c}\|_\infty = \alpha_1 > \dots > \alpha_N = 0$ denote the sequence of breakpoints and corresponding solutions to the standard lasso problem returned by the homotopy algorithm. Define $\lambda_1 = \infty$ and $\lambda_k = \frac{\alpha_k}{\|\mathbf{x}_k^*\|_1}$ for all $k \in \{2, \dots, N\}$. Then, we have $\lambda_1 > \lambda_2 > \dots > \lambda_N = 0$, and for all $\lambda \in [\lambda_{k+1}, \lambda_k)$, an optimal solution to the quadratic lasso (3.1) is given by*

$$\mathbf{x}^* = \frac{(\alpha_k - \lambda \|\mathbf{x}_k^*\|_1) \mathbf{x}_{k+1}^* + (\lambda \|\mathbf{x}_{k+1}^*\|_1 - \alpha_{k+1}) \mathbf{x}_k^*}{\alpha_k - \alpha_{k+1} + \lambda(\|\mathbf{x}_{k+1}^*\|_1 - \|\mathbf{x}_k^*\|_1)}. \quad (4.2)$$

and a c -optimal design is

$$\mathbf{w}^* = \widehat{\mathbf{w}}(\mathbf{x}^*) = \frac{(\alpha_k - \lambda \|\mathbf{x}_k^*\|_1) \cdot |\mathbf{x}_{k+1}^*| + (\lambda \|\mathbf{x}_{k+1}^*\|_1 - \alpha_{k+1}) \cdot |\mathbf{x}_k^*|}{\alpha_k \|\mathbf{x}_{k+1}^*\|_1 - \alpha_{k+1} \|\mathbf{x}_k^*\|_1}, \quad (4.3)$$

where $|\mathbf{x}_k^*| \in \mathbb{R}^p$ represents the coordinate-wise absolute values of \mathbf{x}_k^* , that is, $(|\mathbf{x}_k^*|)_i = |\mathbf{x}_{k,i}^*|, \forall i \in [p]$. In particular, $\lambda \mapsto \mathbf{w}^*$ is continuous and piecewise linear.

Proof We first observe that $0 = \|\mathbf{x}_1^*\| \leq \|\mathbf{x}_2^*\| \leq \dots \leq \|\mathbf{x}_N^*\|$, which follows from the convexity of the Pareto-front for the minimization of the two objectives $\|\mathbf{A}\mathbf{x} - \mathbf{c}\|^2$ and $\|\mathbf{x}\|_1$; see Boyd and Vandenberghe (2004, §4.7.5). Thus, we have $\lambda_1 > \dots > \lambda_N$, and for all $\lambda > 0$ there exists $k \in \{1, \dots, N-1\}$ such that $\lambda_{k+1} \leq \lambda < \lambda_k$. By Theorem 3.5 and (4.1), for all $\theta \in [0, 1]$ the point $\mathbf{x}^*((1-\theta)\alpha_{k+1}^* + \theta\alpha_k^*) = (1-\theta)\mathbf{x}_{k+1}^* + \theta\mathbf{x}_k^*$ solves the quadratic lasso with regularizer

$$\lambda = \frac{(1-\theta)\alpha_{k+1}^* + \theta\alpha_k^*}{\|(1-\theta)\mathbf{x}_{k+1}^* + \theta\mathbf{x}_k^*\|_1} = \frac{(1-\theta)\alpha_{k+1}^* + \theta\alpha_k^*}{(1-\theta)\|\mathbf{x}_{k+1}^*\|_1 + \theta\|\mathbf{x}_k^*\|_1},$$

where the second equality comes from the fact that no coordinate of $\mathbf{x}^*(\alpha)$ changes sign between two consecutive breakpoints, that is, $x_{k,i}^* \cdot x_{k+1,i}^* \geq 0$ holds for all $i \in [p]$. Solving for θ yields $\theta = \frac{\lambda \|\mathbf{x}_{k+1}^*\|_1 - \alpha_{k+1}}{\alpha_k - \alpha_{k+1} + \lambda(\|\mathbf{x}_{k+1}^*\|_1 - \|\mathbf{x}_k^*\|_1)}$. It is easy to verify that θ lies in $[0, 1]$, as desired, by using the inequalities $\alpha_{k+1} \leq \lambda \|\mathbf{x}_{k+1}^*\|_1$, $\lambda \|\mathbf{x}_k^*\|_1 \leq \alpha_k$, $\alpha_k > \alpha_{k+1}$ and $\|\mathbf{x}_{k+1}^*\|_1 \geq \|\mathbf{x}_k^*\|_1$. Substituting the value of θ in $\mathbf{x}^* = (1-\theta)\mathbf{x}_{k+1}^* + \theta\mathbf{x}_k^*$ yields the formula (4.2). Finally, we replace \mathbf{x}^* by its value in $\mathbf{w}^* = \widehat{\mathbf{w}}(\mathbf{x}^*)$ to obtain (4.3). \blacksquare

The steps of the homotopy algorithm for the standard lasso problem are summarized in lines 1–13 of Algorithm 1, following the description of the algorithm given by Mairal and Yu (2012). Here, the notation $\mathbf{A}_J = [\mathbf{a}_j]_{j \in J} \in \mathbb{R}^{m \times |J|}$ stands for the submatrix of \mathbf{A} formed by the columns indexed in J , and $\bar{J} := [p] \setminus J$ denotes the complement of J in $[p]$. Note that practical enhancements can easily be added. In particular, one can replace lines 8–11 by a

Algorithm 1 Homotopy algorithm for the quadratic lasso/ c -optimal design.

```

1: procedure HOMOTOPY( $\mathbf{A} = [\mathbf{a}_1, \dots, \mathbf{a}_n] \in \mathbb{R}^{m \times n}, \mathbf{c} \in \mathbb{R}^m, \lambda > 0$ )
2:   Initialization:  $\alpha_1 \leftarrow \|\mathbf{A}^\top \mathbf{c}\|_\infty; \mathbf{x}_1^* \leftarrow \mathbf{0}; \lambda_1 \leftarrow \infty$ 
3:    $j_0 \leftarrow \arg \max_{j \in [p]} |\mathbf{a}_j^\top \mathbf{c}|; J \leftarrow \{j_0\}; k \leftarrow 0$ 
4:   while  $\lambda < \lambda_{k+1}$  do
5:      $k \leftarrow k + 1$ 
6:      $\boldsymbol{\varepsilon} \leftarrow \text{sign}(\mathbf{A}^\top (\mathbf{c} - \mathbf{A} \mathbf{x}_k^*)) \in \{-1, 0, 1\}^p$ 
7:     For all  $\alpha \in \mathbb{R}$  define  $\boldsymbol{\xi}(\alpha) \in \mathbb{R}^p$  by  $\begin{cases} \boldsymbol{\xi}_J(\alpha) = (\mathbf{A}_J^\top \mathbf{A}_J)^{-1} (\mathbf{A}_J^\top \mathbf{c} - \alpha \cdot \boldsymbol{\varepsilon}_J); \\ \boldsymbol{\xi}_{\bar{J}}(\alpha) = \mathbf{0}. \end{cases}$ 
8:     Find the largest  $\alpha \in [0, \alpha_k]$  such that either:
9:     •  $|\mathbf{a}_j^\top (\mathbf{c} - \mathbf{A}_J \boldsymbol{\xi}_J(\alpha))| = \alpha$  for some  $j \in \bar{J}$ ; In that case set  $J \leftarrow J \cup \{j\}$ 
10:    •  $\boldsymbol{\xi}_j^*(\alpha) = 0$  for some  $j \in J$ ; In that case, set  $J \leftarrow J \setminus \{j\}$ 
11:    If no such  $\alpha$  exists, set  $\alpha \leftarrow 0$ 
12:     $\alpha_{k+1} \leftarrow \alpha; \mathbf{x}_{k+1}^* \leftarrow \boldsymbol{\xi}^*(\alpha_{k+1}); \lambda_{k+1} \leftarrow \frac{\alpha_{k+1}}{\|\mathbf{x}_{k+1}^*\|_1}$ 
13:   end while
14:   return  $\mathbf{x}^*$  given by (4.2) and/or the  $c$ -optimal design  $\mathbf{w}^*$  given by (4.3).
15: end procedure

```

closed-form expression to find the next break point α . One can also store the inverse matrix $(\mathbf{A}_J^\top \mathbf{A}_J)^{-1}$ to accelerate the computation of the vectors $(\mathbf{A}_J^\top \mathbf{A}_J)^{-1} \mathbf{A}_J^\top \mathbf{c}$ and $(\mathbf{A}_J^\top \mathbf{A}_J)^{-1} \boldsymbol{\varepsilon}$, and use rank-one update formulas at the end of each iteration; see Osborne et al. (2000) for more details.

Since by Theorem 4.1 we have $\infty = \lambda_1 > \lambda_2 > \dots > \lambda_N = 0$, the algorithm exits the while loop with $\lambda \in [\lambda_{k+1}, \lambda_k)$, so the vector \mathbf{x}^* (or \mathbf{w}^*) returned by Algorithm 1 is an optimal solution to the quadratic lasso (respectively, a c -optimal design).

Remark 2. *Degeneracy can occur if the polytope \mathcal{Q} has a particular geometrical structure, which can prevent Algorithm 1 from terminating (this might happen if the largest α we are looking for at line 8 is realized for several distinct indices $j \in J$ or $j \in \bar{J}$). To avoid this, a typical approach consists in adding random noise to the data, so that the problem becomes non-degenerate with probability 1. Alternatively, we can choose the index entering or leaving J arbitrarily among the maximizers, and implement an anti-cycling rule to ensure that the same face of \mathcal{Q} is not visited twice; see, e.g., Gill et al. (1989). One should note that, even for non-degenerate problems, numerical issues can occur when the decrease of α at a given iteration is less than machine precision. \triangleleft*

Remark 3. *The generalization of this approach to the quadratic group lasso (and thus to L -optimality) is not straightforward, as the polyhedral structure of the dual feasible region is lost when $r > 1$; see (A.7). Nevertheless we point out that there have been attempts to design homotopy algorithms for more general penalty functions (Zhou and Wu, 2014), such as group-lasso (Yau and Hui, 2017), relying on the fact that the optimal solution $\mathbf{X}^*(\alpha)$ solves an ordinary differential equation in the interior of the segment (α_1, α_2) between two breakpoints of the regularization path.*

5. Examples

The example of Section 5.1 is directly taken from the literature on Lasso regression. A second example, concerning L -optimal design, is presented in Section 5.2.

5.1 Optimal design for image classification

We use the training set of the well known MNIST database (LeCun and Cortes, 2010), which contains 60 000 28×28 images in grey levels of handwritten digits, each associated with a label for the digit in $\{0, \dots, 9\}$ it represents. For our purpose, we build a smaller training set of 600 images of each digit, which we arrange in a matrix $\mathbf{A} \in \mathbb{R}^{784 \times 6000}$. The i th column \mathbf{a}_i is a vector representing the levels of grey of the $m = 28^2 = 784$ pixels of one image, and has been normalized so that $\|\mathbf{a}_i\| = 1$. As there are 600 images for each digit in $\{0, \dots, 9\}$, $p = 6000$. We choose an arbitrary image that does not belong to the training set, and transform it as above into a vector $\mathbf{c} \in \mathbb{R}^{784}$ of unit norm. The c -optimal design problem thus attempts to find which images should be labeled to best predict the label of the image represented by the vector \mathbf{c} , and one can legitimately assume that these images will represent the same digit.

This problem illustrates well the equivalence between a design problem (where we try to select samples that “span” the subspace $\mathbf{c} \cdot \mathbb{R}$ well) and a Lasso problem (where the data matrix has been transposed, so samples become features and vice-versa, and we want to select features that “explain” the target vector \mathbf{c} well). For a given value of λ penalizing non-sparse solutions, the Lasso problem (3.1) tries to express the vector \mathbf{c} as a linear combination of a few images \mathbf{a}_i . If the label of \mathbf{c} is unknown, this can be interpreted as a task of supervised learning: for well chosen values of the penalty parameter λ , we expect that the vast majority of \mathbf{a}_i ’s supporting the solution \mathbf{x}^* will have the same label as \mathbf{c} , which can be used to classify the unknown image. Alternatively, if no labels were provided, we could use this approach repeatedly in a leave-one-out manner (i.e., select one sample to be the vector \mathbf{c} and use the other $p - 1$ samples to build the matrix \mathbf{A}) to perform an unsupervised learning task and construct clusters of images that are likely to share the same label.

The 10 support points of the optimal solution for $\lambda = 0.4$ are displayed in Figure 3 (left), together with the image corresponding to the vector \mathbf{c} : as expected, all images represent the same digit (here, a six). Figure 2 shows the evolution of the duality gap $\mathcal{L}_\lambda(\mathbf{x}) - \mathcal{D}_\lambda(\mathbf{y}_1(\mathbf{x}))$ for $\lambda = 0.4$, when the Coordinate Descent (CD) algorithm² described in (Sagnol and Pauwels, 2019) is used (left), as well as the percentage of support points $\rho(k)$ that have been eliminated after k iterations (right). As m is large and $r = 1$ (this is a c -optimal design problem), one iteration of the CD algorithm is much faster than one iteration of the multiplicative algorithm (3.9): the former has complexity $\mathcal{O}(m \cdot r)$ to update a single row of the matrix \mathbf{X} ; when $r = 1$, this means that a single coordinate of \mathbf{x} is updated in $\mathcal{O}(m)$; an iteration cycling over $p_i = (1 - \rho(i)) \cdot p$ remaining coordinates has thus complexity $\mathcal{O}(p_i \cdot m)$. In contrast, the multiplicative algorithm is dominated by the computation of the information matrix $\mathbf{M}(\mathbf{w})$, which takes $\mathcal{O}(p_i \cdot m^2)$ when p_i coordinates remain. See Table 1 and its presentation below for an empirical comparison of the running times³ of the different algorithms.

2. Sagnol and Pauwels (2019) actually present a *block* CD algorithm for Bayes L -optimality. For c -optimal designs, the blocks are of size $r = 1$ so this is indeed a CD algorithm.

3. Calculations are in python on a PC at 2.7 GHz and 16 GB RAM.

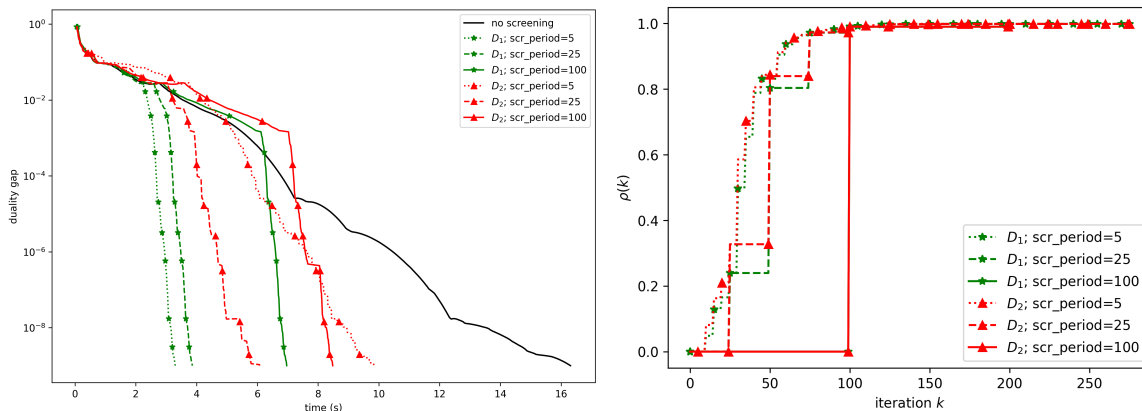


Figure 2: Evolution of the duality gap $\mathcal{L}_\lambda(\mathbf{x}) - \mathcal{D}_\lambda(\mathbf{y}_1(\mathbf{x}))$ (left) and of the proportion $\rho(k)$ of points eliminated after k iterations (right) for the screening bounds D_1 and D_2 applied with different periodicities.

We only compare the bounds D_1 and D_2 , as for large m the overhead of using B_1 , B_2 or B_3 is too high and B_4 cannot be applied for $r < m$. Several observations can be made. First, when screening is performed every 5 or 25 iterations, on the right panel $\rho(k)$ is larger for D_2 than for D_1 , with the consequence that D_2 eliminates points slightly earlier than D_1 (we observe no difference when screening is performed every 100 iterations, the two curves superimpose perfectly). Second, the curves on the left panel show that the computational cost of the screening test applied to one \mathbf{a}_i is generally higher for D_2 than for D_1 . Indeed, the rate of convergence is related to the slope of the curve showing the evolution of the duality gap. Ignore the initial part of the plot, where the number of remaining points to be tested varies much according to the method and periodicity. Once enough points have been eliminated, the slope for D_2 decreases when the period between two successive screenings decreases. On the opposite, the slope stays roughly constant for D_1 : screening with D_1 every 5, 25 or 100 iterations yields the same convergence rate; the more frequent the screening, the earlier the elimination, and thus the induced acceleration.

The effect of the regularization parameter λ is shown on Figure 3 (right), together with the convergence of the homotopy algorithm. The problem is badly conditioned for small λ , which slows down convergence and delays elimination of non-supporting points (we use D_1 every 10 iterations). Note, however, that the problem of c -optimal design obtained in the limit $\lambda \rightarrow 0$ can be solved efficiently by linear programming (Harman and Jurík, 2008). Interestingly, we see that the homotopy algorithm usually requires less time to get an *exact* solution than the CD algorithm needs to get a *reasonable approximate* solution, for most values of λ . The speed-up can attain several orders of magnitude: for $\lambda = 10^{-2}$, the homotopy algorithm computes an *exact* solution to the problem in less than 4 s, while the duality gap for the CD algorithm with D_1 is still 10^{-3} after 67 s; for $\lambda = 10^{-4}$, the homotopy algorithm finds an exact solution within 45 s, while the CD algorithm with D_1 still has a duality gap or more than 10^{-4} after 2000 s.

Table 1: Time (in s) to reach an exact solution with the Homotopy algorithm and a near-optimal solution (tolerance= 10^{-4}) with different algorithms, for several values of λ .

λ	10^0	10^{-1}	10^{-2}	10^{-3}	10^{-4}
Homotopy	0.22	0.73	3.16	13.05	44.66
CD	1.68	7.15	29.00	201.23	2012.15
MWU	53.92	67.39	204.67	432.39	619.09
FISTA	141.86	674.51	2011.01	–	–
FW	14.86	–	–	–	–
SOCP-1	62.57	66.41	72.59	72.16	64.86
SOCP-2	13.54	13.71	13.53	13.81	13.91

For the above experiment, the CD algorithm was used as a reference because it performed better than first-order algorithms. Table 1 compares the running time of the homotopy algorithm and CD with several first order algorithms (Multiplicative weight updates (MWU), FISTA, Frank-Wolfe (FW)), as well the time required by a commercial solver (Gurobi (Gurobi Optimization, LLC, 2023) used with the PICOS interface (Sagnol and Stahlberg, 2022) in Python) to solve two different Second-Order Cone Programming (SOCP) formulations of the problem. The first formulation (SOCP-1) comes from (Sagnol, 2011) and read $\min\{\mathbf{c}^\top \mathbf{x} : \|\mathbf{x}\|^2 \leq v, (\mathbf{a}_i^\top \mathbf{x})^2 \leq u_i, u_i + \lambda v \leq 1, \forall i \in [p]\}$. The second one (SOCP-2) is a direct reformulation of (3.1): $\min\{u + \lambda t : \|\mathbf{A}\mathbf{x} - \mathbf{c}\|^2 \leq u, \|\mathbf{x}\|_1 \leq v, v^2 \leq t\}$, and performs significantly better. Note that the reported times are for the computation of an exact solution (up to machine precision) of the problem by using the homotopy method, while we have used a tolerance of 10^{-4} for all other algorithms. We used the screening rule D_1 with $\tau = 10$ for CD, MWU, FISTA and FW. Observe that although the CD iterations are faster than the MWU iterations, this algorithm has a slow convergence for small values of λ . Thus, more iterations are required and MWU is faster than CD for $\lambda = 10^{-4}$. Also, due to the very slow convergence of FW and FISTA for small values of λ , we have not been able to reach the desired tolerance within 1 hour in some of the experiments. In contrast, the SOCP solver appears to be insensitive to the value of λ .

Finally, in Figure 4 we compare the effect of the screening rule D_1 on three algorithms for an L -optimal design problem with $r = 50$. The training set contains $p = 1\,200$ images (there are 120 samples for each digit), the images have been down-sampled to $m = 20^2 = 400$ pixels, and the matrix $\mathbf{K} \in \mathbb{R}^{m \times r}$ is formed by a collection of 5 images of each digit, randomly chosen outside the training set. The top-left panel of the figure shows the support of the corresponding L -optimal design (for $\lambda = 0.4$): it contains at least one representative of each digit, and thus fulfills the objective of selecting a subset of samples that spans the subspace of all handwritten digits well. The other three panels show that application of the screening rule speeds up computations for the three algorithms considered, with, however, different gains. In this example, the multiplicative weight update algorithm benefits the most from the screening rule, but block CD yields the fastest convergence, with an acceleration factor of about two when D_1 is used every 10 iterations.

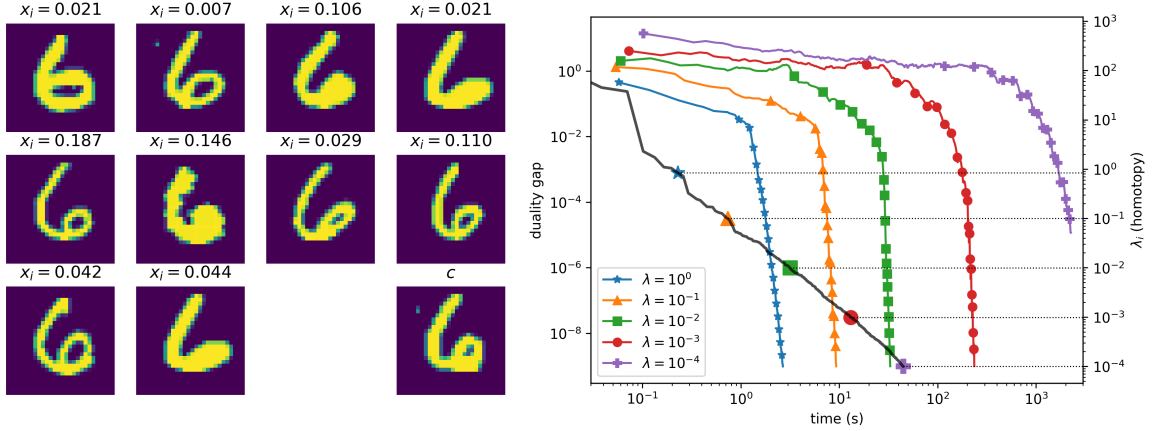


Figure 3: Left: Image corresponding to \mathbf{c} and support points of the optimal solution for $\lambda = 0.4$. Right: The colored curves indicate the convergence of the CD algorithm for different values of the regularization parameter λ (duality gap shown on leftmost y -axis), when D_1 is used every 10 iterations. The black curve shows the progression of λ_i (rightmost y -axis) along time for the homotopy algorithm. The markers on this curve indicate the points in time where the homotopy algorithm has reached an *exact* solution for the corresponding value of λ .

5.2 An example of L -optimal design: IMSE optimal design in random-field models

We consider the framework of (Gauthier and Pronzato, 2017) for the minimization of the Integrated Mean-Squared Error (IMSE) in a random-field interpolation model $Z_{\mathbf{s}}$, indexed by a space variable $\mathbf{s} \in \mathcal{S} \subset \mathbb{R}^d$. We assume that $Z_{\mathbf{s}}$ is Gaussian, centered, with covariance $\mathbb{E}\{Z_{\mathbf{s}_1} Z_{\mathbf{s}_2}\} = \gamma(\mathbf{s}_1, \mathbf{s}_2)$. After observation at n design points $\mathbf{S}_n = \{\mathbf{s}_1, \dots, \mathbf{s}_n\}$ in \mathcal{S} , the IMSE for a given measure of interest μ on \mathcal{S} is

$$\text{IMSE}(\mathbf{S}_n) = \int_{\mathcal{S}} \left[\gamma(\mathbf{s}, \mathbf{s}) - \boldsymbol{\gamma}_n^\top(\mathbf{s}) \boldsymbol{\Gamma}_n^{-1} \boldsymbol{\gamma}_n(\mathbf{s}) \right] d\mu(\mathbf{s}),$$

with $\boldsymbol{\gamma}_n(\mathbf{s}) = (\gamma(\mathbf{s}_1, \mathbf{s}), \dots, \gamma(\mathbf{s}_n, \mathbf{s}))^\top$ and $(\boldsymbol{\Gamma}_n)_{i,j} = \gamma(\mathbf{s}_i, \mathbf{s}_j)$, $i, j = 1, \dots, n$.

The truncation of a particular Karhunen-Loève expansion of $Z_{\mathbf{s}}$ yields a Bayesian linear model, with the eigenfunctions as regression functions, and the construction of an IMSE-optimal design can be cast as an L -optimal design problem. The solution is much facilitated when μ is replaced by a discrete measure μ_p supported on a finite set $\mathcal{S}_p \subset \mathcal{S}$ and design points are selected within \mathcal{S}_p . Denote $\mathbf{D}_p = \text{diag}\{\mu_p(\mathbf{s}_1), \dots, \mu_p(\mathbf{s}_p)\}$. We diagonalize the $p \times p$ matrix $\mathbf{D}_p^{1/2} \boldsymbol{\Gamma}_p \mathbf{D}_p^{1/2}$ into $\mathbf{U}_p \boldsymbol{\Lambda}_p \mathbf{U}_p^\top$, with $\boldsymbol{\Lambda}_p = \text{diag}\{\lambda_1, \dots, \lambda_p\}$ and $\mathbf{U}_p \mathbf{U}_p^\top = \mathbf{I}_p$, and define $\mathbf{V}_p = \mathbf{D}_p^{-1/2} \mathbf{U}_p$. The i th column \mathbf{v}_i of \mathbf{V}_p corresponds to an eigenvector of $\boldsymbol{\Gamma}_p \mathbf{D}_p$ for the eigenvalue λ_i and satisfies $\mathbf{v}_i^\top \mathbf{D}_p \mathbf{v}_j = \delta_{i,j}$ (the Kronecker symbol) for all i, j . We assume that the λ_i are ordered by decreasing values, and introduce a spectral truncation by using the m largest eigenvalues only. Denoting by $\boldsymbol{\Phi}_m$ the $p \times m$ matrix with i th column equal to \mathbf{v}_i , $i = 1, \dots, m$, we obtain the Bayesian linear model $Z_{\mathbf{s}_i} = \boldsymbol{\phi}_{m,i}^\top \boldsymbol{\beta} + \varepsilon_{\mathbf{s}_i}$, where $\boldsymbol{\beta}$ has the normal prior $\mathcal{N}(\mathbf{0}_m, \boldsymbol{\Lambda}_m)$, $\boldsymbol{\phi}_{m,i}^\top$ is the i th row of $\boldsymbol{\Phi}_m$, and where the $\varepsilon_{\mathbf{s}_i}$ are normal random variables, independent of $\boldsymbol{\beta}$, with zero mean and covariance $\mathbb{E}\{\varepsilon_{\mathbf{s}_i}, \varepsilon_{\mathbf{s}_j}\} = \gamma(\mathbf{s}_i, \mathbf{s}_j) - \boldsymbol{\phi}_{m,i}^\top \boldsymbol{\Lambda}_m \boldsymbol{\phi}_{m,j}$ for

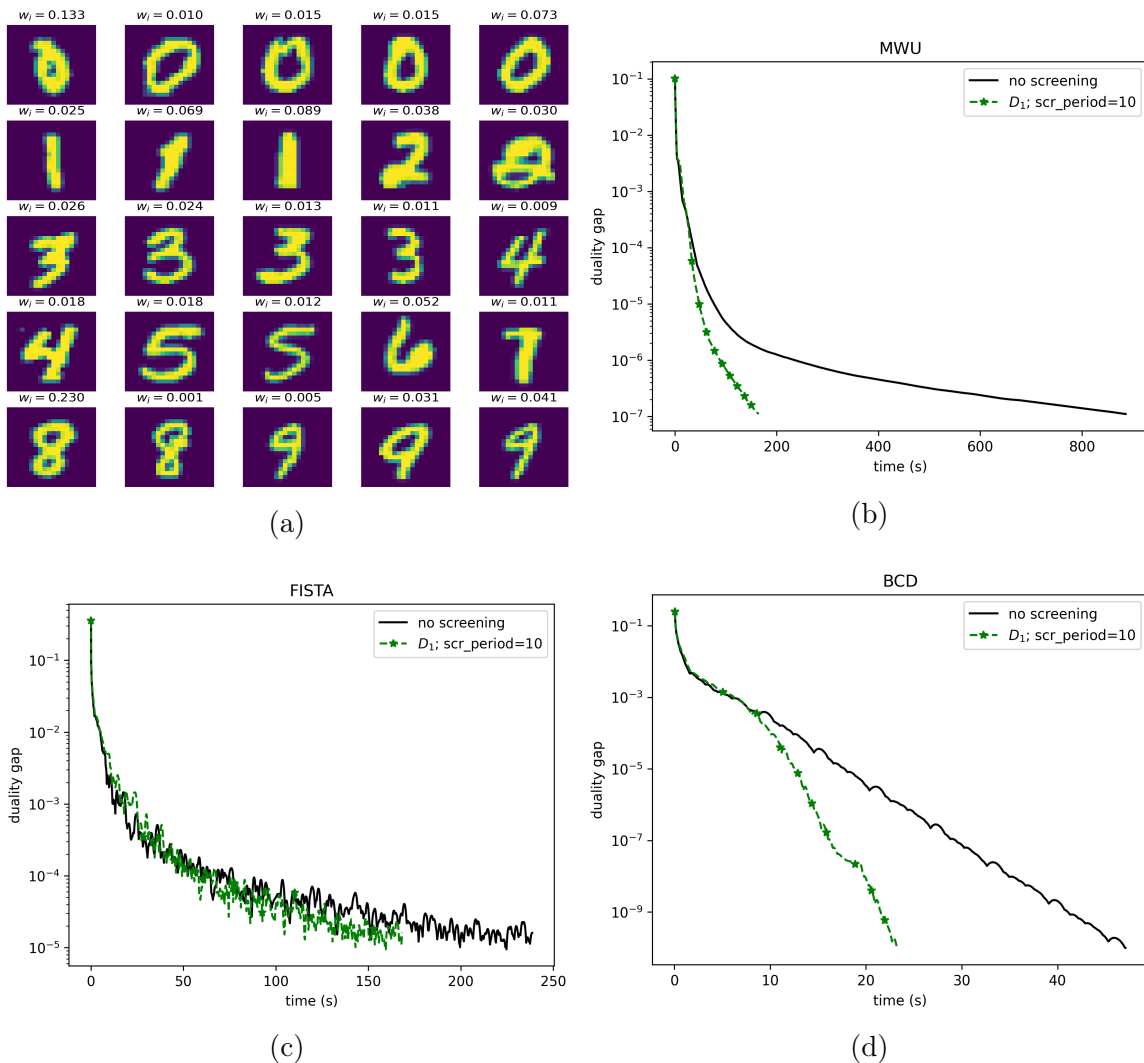


Figure 4: (a) L -optimal design for $\lambda = 0.4$ over the MNIST database, with a training set of $p = 1200$ images (120 per digit) of size $m = 400 = 20 \times 20$, with \mathbf{K} containing $r = 50$ images, i.e., 5 images of each digit (top left). (b-d) Convergence of three algorithms, with and without the screening test D_1 run every 10 iterations: Multiplicative Weight Update (b); FISTA (c); Block CD (d).

all i, j . Finally, we neglect correlations and approximate the IMSE via the estimation of β by its posterior mean in

$$Z'_{\mathbf{s}_i} = \phi_{m,i}^\top \beta + \varepsilon'_{\mathbf{s}_i}, \quad (5.1)$$

where the errors $\varepsilon'_{\mathbf{s}_i}$ are uncorrelated but heteroscedastic, with variance $\sigma_i^2 = \mathbf{E}\{\varepsilon_{\mathbf{s}_i}, \varepsilon_{\mathbf{s}_i}\} = \gamma(\mathbf{s}_i, \mathbf{s}_i) - \phi_{m,i}^\top \mathbf{\Lambda}_m \phi_{m,i}$; see Gauthier and Pronzato (2017). After observations at n points

\mathbf{s}_i with indices in $J_n \subseteq \{1, \dots, p\}$, the approximated IMSE is then

$$\widehat{\text{IMSE}}(J_n) = \text{trace} \left[\left(\sum_{i \in J_n} \frac{1}{\sigma_i^2} \phi_{m,i} \phi_{m,i}^\top + \Lambda_m^{-1} \right)^{-1} \right] = \frac{1}{N} \text{trace} [\mathbf{C}\mathbf{M}^{-1}(\mathbf{w}(J_n))] , \quad (5.2)$$

where $\mathbf{C} = \Lambda_m$ and $\mathbf{M}(\mathbf{w})$ is defined by (2.1), with \mathbf{H}_i given by (2.5), $\lambda = 1/n$, $\mathbf{a}_i = \Lambda_m^{1/2} \phi_{m,i}/\sigma_i$ for all i , and $w_i(J_n)$ equal to $1/n$ for $i \in J_n$ and zero otherwise. We thus consider the minimization of $\text{trace} [\mathbf{C}\mathbf{M}^{-1}(\mathbf{w})]$ for the construction of IMSE optimal designs.

We take μ uniform and $\mu_p(\mathbf{s}_i) = 1/p$ for all i . The covariance of the random field is given by Matérn 3/2 kernel,

$$\gamma(\mathbf{s}_1, \mathbf{s}_2) = (1 + \theta \|\mathbf{s}_1 - \mathbf{s}_2\|) \exp(-\theta \|\mathbf{s}_1 - \mathbf{s}_2\|) ,$$

where we take $\theta = 10$. The left panel of Figure 5 shows the type of designs that are obtained. Here \mathcal{S}_p corresponds to a 33×33 regular grid in $\mathcal{S} = [0, 1]^2$ and the truncation level m equals 10. The construction relies on approximate design theory and we still need to extract an exact design (a finite set of support points) once optimal weights have been determined. Several approaches can be used, see Gauthier and Pronzato (2017, 2016), which are beyond the scope of this paper. In particular, the method does not allow full control of the size of resulting design (which is of the order of magnitude of m). A straightforward approach is to simply remove points with negligible weights: on Figure 5-left, the 8 design points marked by a blue \bullet have total weight less than $7.14 \cdot 10^{-5}$, the 12 other points are well spread over \mathcal{S} , with associated weights $w_i \in (0.0688, 0.0955)$. An important limitation of this approach comes from the need to diagonalize a $p \times p$ matrix, with p being necessarily large when d is large. A possible solution when μ is uniform on $\mathcal{X} = [0, 1]^d$ is to use a separable covariance so that eigenvalues (respectively, eigenfunctions) are products (respectively, tensor products) of their one-dimensional counterparts; see Pronzato (2019b). Alternatively, one can use a sparse kernel induced by $k \ll p$ points, which reduces the size of the matrix to be diagonalized ($k \times k$ instead of $p \times p$); see (Sagnol et al., 2016) where this approach is used for the sequential design of computer experiments.

In the rest of the example, \mathcal{S} is the hypercube $[0, 1]^d$ with $d = 5$, \mathcal{S}_p corresponds to the first p points of Sobol' low-discrepancy sequence in \mathcal{S} , with $p = 8192$. The truncation level is set at $m = 50$. We consider the elimination of inessential \mathbf{H}_i , or equivalently of inessential $\mathbf{s}_i \in \mathcal{S}_p$, during the construction of an optimal design with the algorithm

$$w_i^{k+1} = \frac{w_i^k \mathbf{a}_i^\top \mathbf{M}^{-1}(\mathbf{w}^k) \mathbf{C}\mathbf{M}^{-1}(\mathbf{w}^k) \mathbf{a}_i}{\sum_{j=1}^p w_j^k \mathbf{a}_j^\top \mathbf{M}^{-1}(\mathbf{w}^k) \mathbf{C}\mathbf{M}^{-1}(\mathbf{w}^k) \mathbf{a}_j} , \quad (5.3)$$

initialized at $w_i^0 = 1/p$ for all i . As shown in (Pronzato and Sagnol, 2021), screening is computationally more efficient when applied periodically, every τ iterations, rather than at each k , and we use $\tau = 100$ in the example. When inessential \mathbf{s}_i are eliminated by a screening test performed with $\mathbf{M}(\mathbf{w}^k)$, the weights of the remaining points, with indices in I_k , are renormalized into $w_i^k / \sum_{i \in I_k} w_i^k$ before application of (5.3). The algorithm is stopped when $\delta < 10^{-6}$, with δ , given by (A.9), measuring the (sub-)optimality of \mathbf{w}^k .

We compare the bound D_2 of Corollary A.5 (i.e., the bound B of Corollary A.6) with the bounds B_1 to B_4 given in (Pronzato and Sagnol, 2021) for elimination of inessential

points, both in terms of efficiency and computational time. Consider iteration k where a screening test is performed, using the matrix $\mathbf{M}_k = \mathbf{M}(\mathbf{w}^k)$. B_1 and B_2 are consequences of the Equivalence Theorem (see Theorem 3.3 for the case of c -optimality) and require the calculation of the minimum and maximum eigenvalues of $\mathbf{M}_k^{-1/2} \mathbf{H}_i \mathbf{M}_k^{-1/2}$ for each \mathbf{H}_i tested. B_3 is derived from a second-order-cone-programming formulation of the design problem (Sagnol, 2011), its calculation requires the solution of one-dimensional convex minimization problem for each \mathbf{H}_i ; we use the dichotomy line-search algorithm of Pronzato and Sagnol (2021) with the precision parameter ϵ fixed at 0.01. B_4 corresponds to the method proposed in (Pronzato, 2013)⁴ and requires the calculation of the maximum eigenvalue of \mathbf{M}_k .

The right panel of Figure 5 shows the proportion $\rho(k)$ of points eliminated along iterations for the five methods considered (note the staircase growth due to periodic screening every $\tau = 100$ iterations only). There is a rather clear global hierarchy in terms of elimination efficiency in this example, with $B_3 \succ D_2 \succ B_2 \succ B_1 \succ B_4$ (B_1 eliminates more points than B_2 and D_2 in the early iterations, although this is not visible on the plot). Since the cost of one iteration of (5.3) is roughly proportional to the size of \mathbf{w} , that is, $p(1 - \rho(k))$ at iteration k , we rescale the iteration counter k into $C(k) = \sum_{i=1}^k (1 - \rho(i))$ to count pseudo iterations that consider the decreasing size of \mathbf{w}^k . The left panel of Figure 6 presents δ (log scale) as a function of $C(k)$: it shows the acceleration provided by elimination of inessential points by each one of the method considered (compare with the curve with black \square for which $\rho(k) \equiv 0$) if one neglects the computational time of the screening tests. The hierarchy observed on Figure 5-right is confirmed. The right panel of Figure 6 shows δ as a function of the true computational time⁵. The most efficient method in terms of elimination, B_3 (magenta +), yields the slowest convergence due to its high computational cost⁶. Screening by D_2 (red \diamond), which does not require any eigenvalue calculation or numerical optimization, provides a speed-up factor of about 5 compared to the direct application of (5.3) without screening (black \square). Due to its reasonable computational cost, B_4 (black \times) ensures a significant acceleration too despite its low screening efficiency. The poor performance of screening with B_1 (red \star) and B_2 (blue ∇) is explained by the dimension 50×50 of the matrices \mathbf{M}_k . Notice that elimination of inessential points does not provide any speedup in terms of number of iterations required for a given accuracy: 18 738 iterations are needed to reach $\delta \leq 10^{-6}$ with or without screening (217 points have then a positive weight, among which 15 have a total mass less than $5 \cdot 10^{-5}$).

The hierarchy between methods in terms of computational time depends on m and may vary with the choice of the kernel $\gamma(\cdot, \cdot)$ (in particular, with its correlation length), but in this example the ranking was fairly stable when varying γ (we also used the Matérn 1/2 and 5/2 kernels and the Gaussian, i.e., squared exponential, kernel) and the number p of Sobol' points in \mathcal{S}_p .

-
4. The method proposed there is for A -optimality, but since the matrix \mathbf{C} in (5.2) has full rank, the problem can be straightforwardly transformed into an A -optimal design problem.
 5. Calculations are in Matlab, on a PC with a clock speed of 2.5 GHz and 32 GB RAM — only the comparison between curves is of interest.
 6. The situation is reversed when $\lambda = 0$ in (2.5) and all \mathbf{H}_i have rank 1 (so that neither D_2 nor B_2 can be used): B_3 has an explicit expression and due to its high efficiency can provide an important acceleration factor; see Pronzato and Sagnol (2021).




Figure 5: Left: Design obtained when \mathcal{S}_p is a 33×33 regular grid in $\mathcal{S} = [0, 1]^2$, $m = 10$ (the 8 points marked with a blue \bullet have negligible weight). Right: evolution of the proportion $\rho(k)$ of points eliminated by the five methods considered along iterations ($d = 5$, $p = 8192$, $m = 50$).

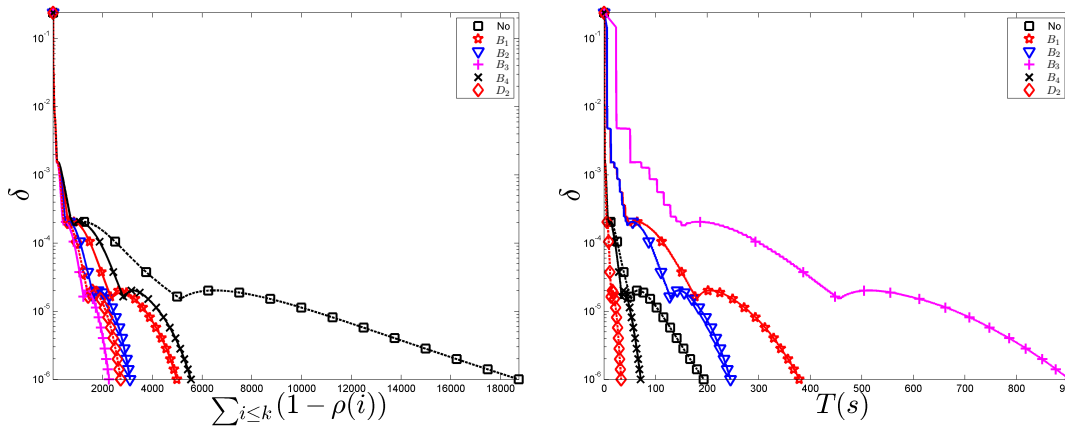


Figure 6: Left: Evolution of δ given by (A.9) as a function of the pseudo iteration number (the computational cost of each screening test is ignored). Right: Evolution of δ as a function of computational time (in seconds).

6. Conclusion

We have shown that the reformulation of c - and L -optimal design problems as a linear regression problem with squared sparsity-inducing penalty can be used to obtain safe screening rules. These rules are much faster to compute than those previously proposed in the optimal design community, especially for problems with a high dimensional parameter space. Their

efficiency has been demonstrated on several experiments using real data sets. In addition, we have leveraged the polyhedral conic geometry of the dual quadratic-lasso problem to develop a new homotopy algorithm for Bayesian c -optimal design. This algorithm terminates after a finite number of steps with an optimal design and can be several orders of magnitude faster than standard first-order algorithms.

Acknowledgements

The authors thank three anonymous referees whose comments contributed to improve the presentation of these results.

The work of Luc Pronzato was partly supported by project INDEX (INcremental Design of EXperiments) ANR-18-CE91-0007 of the French National Research Agency (ANR). The work of Guillaume Sagnol is supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy — The Berlin Mathematics Research Center MATH+ (EXC-2046/1, project ID: 390685689).

References

- K.M. Anstreicher. Efficient solution of maximum-entropy sampling problems. *Operations Research*, 68(6):1826–1835, 2020.
- A. Bonnefoy, V. Emiya, L. Ralaivola, and R. Gribonval. Dynamic screening: accelerating first-order algorithms for the Lasso and Group-Lasso. *IEEE Transactions on Signal Processing*, 63(19):5121–5132, 2015.
- S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, Cambridge, 2004.
- M. Černý and M. Hladík. Two complexity results on c -optimality in experimental design. *Computational Optimization and Applications*, 51(3):1397–1408, 2012.
- A. Civril and M. Magdon-Ismail. On selecting a maximum volume sub-matrix of a matrix and related problems. *Theoretical Computer Science*, 410(47-49):4801–4811, 2009.
- D.A. Cohn, Z. Ghahramani, and M.I. Jordan. Active learning with statistical models. *Journal of Artificial Intelligence Research*, 4:129–145, 1996.
- B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani. Least angle regression. *The Annals of Statistics*, 32(2):407–499, 2004.
- L. El Ghaoui, V. Viallon, and T. Rabbani. Safe feature elimination in sparse supervised learning. *Pacific Journal of Optimization*, 8(4):667–698, 2012.
- V.V. Fedorov. *Theory of Optimal Experiments*. New York : Academic Press, 1972.
- J. Fellman. On the allocation of linear observations. *Comment. Phys. Math.*, 44:27–78, 1974.
- O. Fercoq, A. Gramfort, and J. Salmon. Mind the duality gap: safer rules for the Lasso. In *International Conference on Machine Learning*, pages 333–342, 2015.
- B. Gauthier and L. Pronzato. Optimal design for prediction in random field models via covariance kernel expansions. In J. Kunert, Ch.H. Müller, and A.C. Atkinson, editors, *mODa’11 – Advances in Model-Oriented Design and Analysis, Proceedings of the 11th Int. Workshop, Hamminkeln-Dingden (Germany)*, pages 103–111, Heidelberg, 2016. Springer.

- B. Gauthier and L. Pronzato. Convex relaxation for IMSE optimal design in random field models. *Computational Statistics and Data Analysis*, 113:375–394, 2017.
- P.E. Gill, W. Murray, M.A. Saunders, and M.H. Wright. A practical anti-cycling procedure for linearly constrained optimization. *Mathematical Programming*, 45(1):437–474, 1989.
- Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2023. URL <https://www.gurobi.com>.
- R. Harman and T. Jurík. Computing c -optimal experimental designs using the simplex method of linear programming. *Computational Statistics & Data Analysis*, 53(2):247–254, 2008.
- R. Harman and L. Pronzato. Improvements on removing non-optimal support points in D -optimum design algorithms. *Statistics & Probability Letters*, 77:90–94, 2007.
- R. Harman and S. Rosa. Removal of points that do not support an E -optimal experimental design. *Statistics & Probability Letters*, 147:83–89, 2019.
- R. Harman, L. Filová, and P. Richtárik. A randomized exchange algorithm for computing optimal approximate designs of experiments. *Journal of the American Statistical Association*, 115(529):348–361, 2020.
- J. Kiefer. General equivalence theory for optimum designs (approximate theory). *The Annals of Statistics*, 2(5):849–879, 1974.
- Y. LeCun and C. Cortes. MNIST handwritten digit database, 2010. <http://yann.lecun.com/exdb/mnist/>.
- Y. Li, M. Fampa, J. Lee, F. Qiu, W. Xie, and R. Yao. D-optimal data fusion: Exact and approximation algorithms, 2022. arXiv:2208.03589.
- J. Mairal and B. Yu. Complexity analysis of the lasso regularization path. In *Proceedings of the 29th International Conference on Machine Learning, ICML 2012, Edinburgh, Scotland, UK, June 26 - July 1, 2012*. icml.cc / Omnipress, 2012. URL <http://icml.cc/2012/papers/202.pdf>.
- E. Ndiaye, O. Fercoq, A. Gramfort, and J. Salmon. Gap safe screening rules for sparsity enforcing penalties. *Journal of Machine Learning Research*, 18(1):4671–4703, 2017.
- M.R. Osborne, B. Presnell, and B.A. Turlach. A new approach to variable selection in least squares problems. *IMA Journal of Numerical Analysis*, 20(3):389–403, 2000.
- J. Pilz. *Bayesian Estimation and Experimental Design in Linear Regression Models*, volume 55. Teubner-Texte zur Mathematik, Leipzig, 1983. (also Wiley, New York, 1991).
- L. Pronzato. Removing non-optimal support points in D -optimum design algorithms. *Statistics & Probability Letters*, 63:223–228, 2003.
- L. Pronzato. A delimitation of the support of optimal designs for Kiefer’s ϕ_p -class of criteria. *Statistics & Probability Letters*, 83:2721–2728, 2013.
- L. Pronzato. On the elimination of inessential points in the smallest enclosing ball problem. *Optimization Methods and Software*, 34(2):225–247, 2019a.
- L. Pronzato. Sensitivity analysis via Karhunen-Loève expansion of a random field model: estimation of Sobol’ indices and experimental design. *Reliability Engineering and System Safety*, 187:93–109, 2019b.

- L. Pronzato and G. Sagnol. Removing inessential points in c - and A -optimal design. *Journal of Statistical Planning and Inference*, 213:233–252, 2021.
- L. Pronzato and A.A. Zhigljavsky. Measures minimizing regularized dispersion. *Journal of Scientific Computing*, 78(3):1550–1570, 2019.
- F. Pukelsheim. *Optimal Experimental Design*. Wiley, New York, 1993. (also SIAM, Philadelphia, 2006).
- F. Pukelsheim and S. Rieder. Efficient rounding of approximate designs. *Biometrika*, 79(4):763–770, 1992.
- G. Sagnol. Computing optimal designs of multiresponse experiments reduces to second-order cone programming. *Journal of Statistical Planning and Inference*, 141:1684–1708, 2011.
- G. Sagnol and E. Pauwels. An unexpected connection between Bayes A -optimal design criteria and the group Lasso. *Statistical Papers*, 60(2):215–234, 2019.
- G. Sagnol, H.-C. Hege, and M. Weiser. Using sparse kernels to design computer experiments with tunable precision. In *Proceedings of the 22nd International Conference on Computational Statistics*, pages 397–408, 2016.
- Guillaume Sagnol and Maximilian Stahlberg. PICOS: A Python interface to conic optimization solvers. *Journal of Open Source Software*, 7(70):3915, February 2022. ISSN 2475-9066. doi: 10.21105/joss.03915.
- S.D. Silvey. *Optimal Design*. Chapman & Hall, London, 1980.
- M. Singh and W. Xie. Approximation algorithms for D -optimal design. *Mathematics of Operations Research*, 45(4):1512–1534, 2020.
- J. Wang, P. Wonka, and J. Ye. Lasso screening rules via dual polytope projection. *Journal of Machine Learning Research*, 16(1):1063–1101, 2015.
- W.J. Welch. Algorithmic complexity: three NP-hard problems in computational statistics. *Journal of Statistical Computation and Simulation*, 15(1):17–25, 1982.
- H.P. Wynn. The sequential generation of D -optimum experimental designs. *Annals of Mathematical Statistics*, 41:1655–1664, 1970.
- Z.J. Xiang and P.J. Ramadge. Fast Lasso screening tests based on correlations. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2137–2140, 2012.
- Z.J. Xiang, Y. Wang, and P.J. Ramadge. Screening tests for Lasso problems. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 39(5):1008–1027, 2016.
- C.Y. Yau and T.S. Hui. Lars-type algorithm for group lasso. *Statistics and Computing*, 27:1041–1048, 2017.
- Y. Yu. Monotonic convergence of a general algorithm for computing optimal designs. *The Annals of Statistics*, 38(3):1593–1606, 2010.
- H. Zhou and Y. Wu. A generic path algorithm for regularized statistical estimation. *Journal of the American Statistical Association*, 109(506):686–699, 2014.

Appendix A. L -optimality and quadratic group Lasso

In order to generalize the results of Section 3 to the case of L -optimality, we start with the counterpart of Theorem 3.1 for matrices \mathbf{H}_i of the form $\mathbf{H}_i = \mathbf{A}_i \mathbf{A}_i^\top + \lambda \mathbf{I}$ for some $\mathbf{A}_i \in \mathbb{R}^{m \times q}$, for all $i \in [p]$. This result is interesting in its own right, as this situation arises naturally for the linear model with multiresponse experiments, where (2.4) is replaced by a q -dimensional observation $Y_i = \mathbf{A}_i^\top \boldsymbol{\theta} + \varepsilon_i$.

Theorem A.1. *Let $\mathbf{H}_i = \mathbf{A}_i \mathbf{A}_i^\top + \lambda \mathbf{I}$, with $\mathbf{A}_i \in \mathbb{R}^{m \times q}$. Then, the c -optimal design problem (2.2) is equivalent to the following problem:*

$$\min_{\mathbf{x}_1, \dots, \mathbf{x}_p \in \mathbb{R}^q} \mathcal{L}_\lambda^q(\mathbf{x}) = \left\| \sum_{i=1}^p \mathbf{A}_i \mathbf{x}_i - \mathbf{c} \right\|^2 + \lambda \left(\sum_{i=1}^p \|\mathbf{x}_i\| \right)^2, \quad (\text{A.1})$$

in the following sense:

- (i) the optimal value of (A.1) is equal to $\lambda \phi_c(\mathbf{w}^*)$, where \mathbf{w}^* is a c -optimal design;
- (ii) If $\mathbf{x}^* = (\mathbf{x}_1^*, \dots, \mathbf{x}_p^*)$ solves (A.1) and $\mathbf{A}_i^\top \mathbf{c} \neq \mathbf{0}$ for some $i \in [p]$, then $\mathbf{x}^* \neq \mathbf{0}$ and $\widehat{\mathbf{w}}^* = \widehat{\mathbf{w}}(\mathbf{x}^*)$ is c -optimal, where

$$\widehat{w}_i(\mathbf{x}) = \frac{\|\mathbf{x}_i\|}{\sum_{j=1}^p \|\mathbf{x}_j\|}, \quad \forall i \in [p], \quad \mathbf{x} \neq \mathbf{0}; \quad (\text{A.2})$$

- (iii) If $\mathbf{w}^* \in \mathcal{P}_p$ is c -optimal, then $\widehat{\mathbf{x}}^* = (\widehat{\mathbf{x}}_1(\mathbf{w}^*), \dots, \widehat{\mathbf{x}}_p(\mathbf{w}^*))$ is optimal for (A.1), where

$$\widehat{\mathbf{x}}_i(\mathbf{w}) = w_i \mathbf{A}_i^\top \mathbf{M}^{-1}(\mathbf{w}) \mathbf{c}, \quad \forall i \in [p]. \quad (\text{A.3})$$

- (iv) In the pathological case $\mathbf{A}_i^\top \mathbf{c} = \mathbf{0}$, $\forall i \in [p]$, the unique optimal solution to the quadratic lasso is $\mathbf{x}^* = \mathbf{0}$, while every design $\mathbf{w} \in \mathcal{P}_p$ is c -optimal.

Proof The proof follows almost the same lines as the proof of Theorem 3.1, with the difference that we define $\mathbf{A} = [\mathbf{A}_1, \dots, \mathbf{A}_p] \in \mathbb{R}^{m \times (pq)}$ and $\mathbf{D}(\mathbf{w}) = \text{diag}(\mathbf{w}) \otimes \mathbf{I}_q \in \mathbb{R}^{pq \times pq}$, i.e., the diagonal matrix with p diagonal blocks of the form $w_i \mathbf{I}_q$. Then, we still have $\mathbf{M}(\mathbf{w}) = \mathbf{A} \mathbf{D}(\mathbf{w}) \mathbf{A}^\top + \lambda \mathbf{I}_m$. We also change the definition of the function v to

$$v(\mathbf{x}, \mathbf{w}) = \left\| \sum_{i=1}^p \mathbf{A}_i \mathbf{x}_i - \mathbf{c} \right\|^2 + \lambda \sum_{i=1}^p \frac{\|\mathbf{x}_i\|^2}{w_i} = \|\mathbf{A} \mathbf{x} - \mathbf{c}\|^2 + \lambda \mathbf{x}^\top \mathbf{D}^{-1}(\mathbf{w}) \mathbf{x}.$$

Here, the function $(\mathbf{x}_i, w_i) \mapsto \frac{\|\mathbf{x}_i\|^2}{w_i}$ is the perspective function of $\mathbf{x}_i \mapsto \|\mathbf{x}_i\|^2$, hence it is convex, and defined by continuity in $w_i = 0$ (i.e., $\|\mathbf{x}_i\|^2/0 = 0$ if $\mathbf{x}_i = \mathbf{0}$ and $\|\mathbf{x}_i\|^2/0 = \infty$ otherwise).

The rest of the proof is identical to that of Theorem 3.1. We obtain that (3.5) and (3.6) hold, which imply (i), (ii) and (iii). That $\mathbf{x} = \mathbf{0}$ cannot be optimal whenever $\mathbf{A}^\top \mathbf{c} \neq \mathbf{0}$ follows from

$$\mathcal{L}_\lambda \left(\frac{\mathbf{a}_k^\top \mathbf{c}}{\|\mathbf{a}_k\|^2 + \lambda} \mathbf{e}_k \right) = \|\mathbf{c}\|^2 - \frac{(\mathbf{a}_k^\top \mathbf{c})^2}{\|\mathbf{a}_k\|^2 + \lambda} < \|\mathbf{c}\|^2 = \mathcal{L}_\lambda(\mathbf{0}),$$

if k is such that the k th column \mathbf{a}_k of \mathbf{A} satisfy $\mathbf{a}_k^\top \mathbf{c} \neq 0$. The proof of point (iv) for the case $\mathbf{A}^\top \mathbf{c} = \mathbf{0}$ is also the same as for Theorem 3.1. \blacksquare

Let us return to the case where $\mathbf{H}_i = \mathbf{a}_i \mathbf{a}_i^\top + \lambda \mathbf{I}_m$. The generalization of Theorem 3.1 to the case of L -optimality is now straightforward following (Sagnol, 2011). Let $\mathbf{K} \in \mathbb{R}^{m \times r}$ and define $\bar{\mathbf{c}} = \text{vec}(\mathbf{K}) \in \mathbb{R}^{mr}$. Then, for all $\mathbf{w} \in \mathcal{P}_p$ we have

$$\phi_L(\mathbf{w}) = \Phi_L[\mathbf{M}(\mathbf{w})] = \text{trace}[\mathbf{K}^\top \mathbf{M}^{-1}(\mathbf{w}) \mathbf{K}] = \bar{\mathbf{c}}^\top \widetilde{\mathbf{M}}^{-1}(\mathbf{w}) \bar{\mathbf{c}},$$

where

$$\widetilde{\mathbf{M}}(\mathbf{w}) = \mathbf{I}_r \otimes \mathbf{M}(\mathbf{w}) = \sum_{i=1}^p w_i (\mathbf{I}_r \otimes (\mathbf{a}_i \mathbf{a}_i^\top)) + \lambda \mathbf{I}_{rm} = \sum_{i=1}^p w_i \widetilde{\mathbf{H}}_i,$$

and $\widetilde{\mathbf{H}}_i = (\mathbf{I}_r \otimes \mathbf{a}_i)(\mathbf{I}_r \otimes \mathbf{a}_i)^\top + \lambda \mathbf{I}_{rm}$. Note that $\widetilde{\mathbf{H}}_i$ has the form $\widetilde{\mathbf{A}}_i \widetilde{\mathbf{A}}_i^\top + \lambda \mathbf{I}$ with $\widetilde{\mathbf{A}}_i = \mathbf{I}_r \otimes \mathbf{a}_i \in \mathbb{R}^{rm \times r}$. Theorem A.1 can thus be applied to L -optimality. Straightforward manipulations allows one to get rid off Kronecker products by reshaping the vector $\mathbf{x} \in \mathbb{R}^{pr}$ into a matrix $\mathbf{X} \in \mathbb{R}^{p \times r}$, which yields the following result.

Theorem A.2. *Let \mathbf{H}_i be given by (2.5) and denote by \mathbf{A} the $m \times p$ -matrix with columns $(\mathbf{a}_i)_{i \in [p]}$. Consider the L -optimal design problem where we minimize $\phi_L(\mathbf{w})$ given by (2.3) with respect to $\mathbf{w} \in \mathcal{P}_p$ with \mathbf{K} an $m \times r$ matrix. This problem is equivalent to the following one, which we call quadratic group Lasso,*

$$\min_{\mathbf{X} \in \mathbb{R}^{p \times r}} \mathcal{L}'_\lambda(\mathbf{X}) = \|\mathbf{A}\mathbf{X} - \mathbf{K}\|_F^2 + \lambda (\|\mathbf{X}\|_{1,2})^2, \quad (\text{A.4})$$

in the following sense:

- (i) the optimal value of (A.4) is equal to $\lambda \phi_L(\mathbf{w}^*)$, where \mathbf{w}^* is an L -optimal design;
- (ii) If \mathbf{X}^* solves the quadratic group Lasso problem and $\mathbf{A}^\top \mathbf{K} \neq \mathbf{O}$, then $\widehat{\mathbf{w}}^* = \widehat{\mathbf{w}}(\mathbf{X}^*)$ is L -optimal, where

$$\widehat{w}_i(\mathbf{X}) = \frac{\|\mathbf{X}_{i,\cdot}\|}{\|\mathbf{X}\|_{1,2}} \quad \forall i \in [p], \quad \mathbf{X} \neq \mathbf{O}; \quad (\text{A.5})$$

- (iii) If $\mathbf{w}^* \in \mathcal{P}_p$ is L -optimal, then $\widehat{\mathbf{X}}^* = \widehat{\mathbf{X}}(\mathbf{w}^*)$ is optimal for (A.4), where

$$\widehat{\mathbf{X}}_{i,\cdot}(\mathbf{w}) = w_i \mathbf{a}_i^\top \mathbf{M}^{-1}(\mathbf{w}) \mathbf{K}. \quad (\text{A.6})$$

- (iv) In the pathological case $\mathbf{A}^\top \mathbf{K} = \mathbf{O}$, the unique optimal solution to the quadratic lasso is $\mathbf{x}^* = \mathbf{O}$, while every design $\mathbf{w} \in \mathcal{P}_p$ is c -optimal.

Proceeding as above, we can also show that the dual of the quadratic group Lasso problem can be written as

$$\max_{\mathbf{Y} \in \mathbb{R}^m} \mathcal{D}'_\lambda(\mathbf{Y}) = \|\mathbf{K}\|_F^2 - \|\mathbf{Y} - \mathbf{K}\|_F^2 - \frac{\max_{i \in [p]} \|\mathbf{Y}^\top \mathbf{a}_i\|^2}{\lambda}, \quad (\text{A.7})$$

and that the following relations between optimal primal and dual variables hold:

$$\mathbf{Y}^* = \mathbf{K} - \mathbf{A}\mathbf{X}^* = \lambda \mathbf{M}^{-1}(\mathbf{w}^*) \mathbf{K},$$

with \mathbf{w}^* an L -optimal design. The generalization of Theorem 3.6 is not completely straightforward, though. We first prove the following lemma:

Lemma A.3. *Let $\mathbf{Y} \in \mathbb{R}^{m \times r}$ and $\mathbf{a} \in \mathbb{R}^m$. Then,*

$$\sup_{\mathbf{Z} \in \mathbb{R}^{m \times r}} \{ \|\mathbf{Z}^\top \mathbf{a}\| : \|\mathbf{Y} - \mathbf{Z}\|_F \leq R \} = \|\mathbf{Y}^\top \mathbf{a}\| + R \|\mathbf{a}\|,$$

the supremum being reached for $\mathbf{Z} = \mathbf{Z}_* = \mathbf{Y} + R \mathbf{a} \mathbf{a}^\top \mathbf{Y} / (\|\mathbf{a}\| \|\mathbf{Y}^\top \mathbf{a}\|)$.

Proof Denote by \mathbf{y}_i and \mathbf{z}_i the i -th column of \mathbf{Y} and \mathbf{Z} , respectively, and define $u_i = \|\mathbf{y}_i - \mathbf{z}_i\|$, $i \in [r]$, and $\mathbf{u} = (u_1, \dots, u_r)^\top$. Then, $|\mathbf{z}_i^\top \mathbf{a}| \leq |\mathbf{y}_i^\top \mathbf{a}| + u_i \|\mathbf{a}\|$, $i \in [r]$. Therefore,

$$\begin{aligned} \|\mathbf{Z}^\top \mathbf{a}\|^2 &= \sum_{i=1}^r (\mathbf{z}_i^\top \mathbf{a})^2 \leq \sum_i (|\mathbf{y}_i^\top \mathbf{a}| + u_i \|\mathbf{a}\|)^2 = \|\mathbf{Y}^\top \mathbf{a}\|^2 + \|\mathbf{u}\|^2 \|\mathbf{a}\|^2 + 2\|\mathbf{a}\| \sum_{i=1}^r u_i |\mathbf{y}_i^\top \mathbf{a}| \\ &\leq \|\mathbf{Y}^\top \mathbf{a}\|^2 + \|\mathbf{u}\|^2 \|\mathbf{a}\|^2 + 2\|\mathbf{a}\| \|\mathbf{u}\| \|\mathbf{Y}^\top \mathbf{a}\| \\ &\leq \|\mathbf{Y}^\top \mathbf{a}\|^2 + R^2 \|\mathbf{a}\|^2 + 2R \|\mathbf{a}\| \|\mathbf{Y}^\top \mathbf{a}\| = (\|\mathbf{Y}^\top \mathbf{a}\| + R\|\mathbf{a}\|)^2, \end{aligned}$$

where we used the Cauchy-Schwarz inequality and $\|\mathbf{u}\|^2 = \|\mathbf{Y} - \mathbf{Z}\|_F^2 \leq R$. The supremum is reached when $u_i = \alpha |\mathbf{y}_i^\top \mathbf{a}|$, $i \in [r]$ for some $\alpha > 0$ and $\mathbf{z}_i - \mathbf{y}_i = \beta_i \mathbf{a}$, i.e., $\mathbf{Z} = \mathbf{Y} + \mathbf{a}\boldsymbol{\beta}^\top$ for some $\boldsymbol{\beta} = (\beta_1, \dots, \beta_r)^\top \in \mathbb{R}^r$. Direct calculation shows that it implies $\mathbf{Z} = \mathbf{Z}_*$. \blacksquare

Theorem A.4. Let $\mathbf{Y} \in \mathbb{R}^{m \times r}$ be an ϵ -suboptimal dual solution, i.e., $\mathcal{D}'_\lambda(\mathbf{Y}) \geq \mathcal{D}'_\lambda(\mathbf{Y}^*) - \epsilon$. Then, any \mathbf{a}_i satisfying

$$D_0(\mathbf{a}_i; \mathbf{Y}, \epsilon) = \max_{j \in [p]} \|\mathbf{Y}^\top \mathbf{a}_j\| - \|\mathbf{Y}^\top \mathbf{a}_i\| - \sqrt{\epsilon(\|\mathbf{a}_i\|^2 + \lambda)} > 0$$

cannot support an optimal design.

Proof Let $u = \max_{i \in [p]} \|\mathbf{Y}^\top \mathbf{a}_i\|/\sqrt{\lambda}$ and $u^* = \max_{i \in [p]} \|\mathbf{Y}^{*\top} \mathbf{a}_i\|/\sqrt{\lambda}$. Following the same lines as in the proof of Theorem 3.6, we get $\|\mathbf{Y} - \mathbf{Y}^*\|_F^2 + (u - u^*)^2 \leq \epsilon$.

Now, the complementary slackness condition of optimality of (A.4) and (A.7) indicates that the i -th row of \mathbf{X}^* vanishes whenever $\|\mathbf{Y}^{*\top} \mathbf{a}_i\| < \sqrt{\lambda}u^*$, and thus \mathbf{a}_i is inessential. Using $\|\mathbf{Y} - \mathbf{Y}^*\|_F \leq \sqrt{\epsilon - (u - u^*)^2}$ and Lemma A.3, we get

$$\begin{aligned} \|\mathbf{Y}^{*\top} \mathbf{a}_i\| - \sqrt{\lambda}u^* &\leq \sup_{u'} -\sqrt{\lambda}u' + \sup_{\mathbf{Y}' \in \mathbb{R}^{m \times r}} \{\|\mathbf{Y}'^\top \mathbf{a}_i\| : \|\mathbf{Y} - \mathbf{Y}'\|_F \leq \sqrt{\epsilon - (u - u')^2}\} \\ &= \sup_{u'} -\sqrt{\lambda}u' + \|\mathbf{Y}^\top \mathbf{a}_i\| + \sqrt{\epsilon - (u - u')^2} \|\mathbf{a}_i\|. \end{aligned}$$

Simple calculations show that the supremum is attained at $u' = u \pm \sqrt{\lambda\epsilon}/\sqrt{\|\mathbf{a}_i\|^2 + \lambda}$, and substitution yields

$$\|\mathbf{Y}^{*\top} \mathbf{a}_i\| - \sqrt{\lambda}u^* \leq \|\mathbf{Y}^\top \mathbf{a}_i\| - \sqrt{\lambda}u + \sqrt{\epsilon(\|\mathbf{a}_i\|^2 + \lambda)}.$$

Therefore, \mathbf{a}_i is inessential if $\|\mathbf{Y}^\top \mathbf{a}_i\| - \max_{j \in [p]} \|\mathbf{Y}^\top \mathbf{a}_j\| + \sqrt{\epsilon(\|\mathbf{a}_i\|^2 + \lambda)} < 0$. \blacksquare

As in Section 3.4, we obtain two screening rules for L -optimality by considering the dual points $\mathbf{Y}_1(\mathbf{X}) = \mathbf{K} - \mathbf{A}\mathbf{X}$ and $\mathbf{Y}_2(\widehat{\mathbf{w}}(\mathbf{X})) = \lambda\mathbf{M}^{-1}(\widehat{\mathbf{w}}(\mathbf{X}))\mathbf{K}$, where

$$\mathbf{M}(\widehat{\mathbf{w}}(\mathbf{X})) = \frac{1}{\|\mathbf{X}\|_{1,2}} \sum_{i=1}^p \|\mathbf{X}_{i,\cdot}\| \mathbf{a}_i \mathbf{a}_i^\top + \lambda \mathbf{I}_m.$$

Corollary A.5. Let $\mathbf{X} \in \mathbb{R}^{p \times r}$. If

$$D_1(\mathbf{a}_i; \mathbf{X}) = D_0(\mathbf{a}_i; \mathbf{Y}_1(\mathbf{X}), \mathcal{L}'_\lambda(\mathbf{X}) - \mathcal{D}'_\lambda[\mathbf{Y}_1(\mathbf{X})]) > 0,$$

or $\mathbf{X} \neq \mathbf{O}$ and

$$D_2(\mathbf{a}_i; \mathbf{X}) = D_0(\mathbf{a}_i; \mathbf{Y}_2(\widehat{\mathbf{w}}(\mathbf{X})), \lambda \text{trace}[\mathbf{K}^\top \mathbf{M}^{-1}(\widehat{\mathbf{w}}(\mathbf{X}))\mathbf{K}] - \mathcal{D}'_\lambda[\mathbf{Y}_2(\widehat{\mathbf{w}}(\mathbf{X}))]) > 0,$$

then \mathbf{a}_i cannot support an L -optimal design.

As in Section 3.4, in a design context one may start with a given \mathbf{w} and use $\mathbf{X} = \widehat{\mathbf{X}}(\mathbf{w})$ in D_1 and $\widehat{\mathbf{w}}(\mathbf{X}) = \mathbf{w}$ in D_2 . Similarly to Lemma 3.4, we have $\mathbf{Y}_1(\widehat{\mathbf{X}}(\mathbf{w})) = \mathbf{Y}_2(\mathbf{w})$ and we obtain the following generalization of Corollary 3.9.

Corollary A.6. *Let \mathbf{M} be any matrix in the convex hull of \mathbb{H} with the \mathbf{H}_i satisfying (2.5). Then, any matrix \mathbf{H}_i in \mathbb{H} such that $B(\mathbf{M}, \mathbf{H}_i) > 0$, where*

$$B(\mathbf{M}, \mathbf{H}_i) = \left\{ (1 + \delta) \Phi_L(\mathbf{M}) - \lambda \operatorname{trace}(\mathbf{K}^\top \mathbf{M}^{-2} \mathbf{K}) \right\}^{1/2} - \left[\delta \Phi_L(\mathbf{M}) \left(1 + \frac{\|\mathbf{a}_i\|^2}{\lambda} \right) \right]^{1/2} - |\mathbf{K}^\top \mathbf{M}^{-1} \mathbf{a}_i| \quad (\text{A.8})$$

and

$$\delta = \max_{\mathbf{H}_i \in \mathbb{H}} \frac{\operatorname{trace}(\mathbf{K}^\top \mathbf{M}^{-1} \mathbf{H}_i \mathbf{M}^{-1} \mathbf{K})}{\operatorname{trace}(\mathbf{K}^\top \mathbf{M}^{-1} \mathbf{K})} - 1, \quad (\text{A.9})$$

does not support a c -optimal design.

We have $\delta \geq 0$ and $\Phi_L(\mathbf{M}) \leq (1 + \delta) \phi_L(\mathbf{w}^*)$ with \mathbf{w}^* an L -optimal design.