

KL-optimum designs: theoretical properties and practical computation*

Giacomo Aletti Caterina May Chiara Tommasi

August 22, 2014

Abstract

In this paper some new properties and computational tools for finding KL-optimum designs are provided. KL-optimality is a general criterion useful to select the best experimental conditions to discriminate between statistical models. A KL-optimum design is obtained from a minimax optimization problem, which is defined on an infinite-dimensional space. In particular, continuity of the KL-optimality criterion is proved under mild conditions; as a consequence, the first-order algorithm converges to the set of KL-optimum designs for a large class of models. It is also shown that KL-optimum designs are invariant to any scale-position transformation. Some examples are given and discussed, together with some practical implications for numerical computation purposes.

1 Introduction

In presence of families of competing models, one of the main tasks of the optimum experimental design is the choice of the best experimental conditions to discriminate between rival models.

This problem was studied by many researchers in the field. See for instance, [2], [7, 8] and [16], among many others. In [3, 4], the authors provided the T-optimality criterion to select between two, or more, competing homoscedastic Gaussian models. This criterion was extended in [27] to heteroscedastic Gaussian models. More recently, the KL-optimality, which is based on the Kullback-Leibler divergence, has been proposed in [19]. The idea of using the Kullback-Leibler divergence to construct an optimality criterion is not new. For instance, [6] proposed a Bayesian sequential method to discriminate among several models; at each stage of their sequential scheme, the experimental conditions were chosen by maximizing a weighted sum of Kullback-Leibler divergences between any couple of models. In this way, the maximum change in entropy expected from the observations to be taken at that stage was maximized. See also [18], who applied the Box and Hill criterion to the problem of augmenting a multifactor design. Finally, we refer the reader to [30] for another optimality criterion based on the Kullback-Leibler divergence.

*The final publication is available at Springer via <http://dx.doi.org/10.1007/s11222-014-9515-8>

KL-optimum designs proposed by [19] are obtained from a minimax optimization problem. The KL-criterion is very general since it can be applied when the rival models are nested or not, homoscedastic or heteroscedastic and with any error distribution; moreover, it may be seen as an extension of the T-criterion and its generalizations. It was also applied to discriminate among several models in [25] and used in compound criteria for the double goal of discrimination and estimation in [24, 20]. T- and KL-optimum designs are usually computed numerically. Differently, [9] provide analytical results for T-optimal designs when the interest is in discriminating between two polynomial models which differ in the degree of two.

In this paper, some interesting properties of KL-optimum designs are proved. The notation and framework of the problem are given in Section 2. In Section 3 it is shown that KL-optimum designs are invariant to a scale-position transformation. Moreover, it is proved the continuity of the KL-optimality criterion with respect to the experimental design; mild conditions are assumed. This result is very important also for the numerical computation of a KL-optimum design. In fact, the continuity property is crucial to prove that the first-order algorithm, recalled in Section 4, converges to the set of KL-optimum designs, whenever it moves on regular designs (for a detailed proof see [1]). In Section 4.1 it is shown that, at least for the case of Generalized Linear Models, any design with non-singular Fisher information matrix is KL-regular and viceversa. As a consequence, the algorithm always converges at least in this class of models. In Section 5 some examples are provided and discussed. In addition, some computational hints for the application of the algorithm are described. A final discussion concludes the paper.

2 Framework and notations

Let $f_1(y|\mathbf{x};\boldsymbol{\beta}_1)$ and $f_2(y|\mathbf{x};\boldsymbol{\beta}_2)$ be two rival statistical models, i.e. two parametric families of conditional probability densities of an experimental response Y under the experimental condition \mathbf{x} , where \mathbf{x} belongs to a compact experimental domain $\mathcal{X} \subset \mathbb{R}^q$, $q \geq 1$, the parameters $\boldsymbol{\beta}_i \in \Theta_i$, $i = 1, 2$ and Θ_i is an open set of \mathbb{R}^{d_i} , $i = 1, 2$.¹

A design ξ is a probability distribution having support on \mathcal{X} . In a discrimination problem, the choice of ξ should be done in order to maximize the “separation” between the competing models.

The KL-optimality criterion is based on the Kullback-Leibler divergence between the two conditional distributions $f_1(y|\mathbf{x};\boldsymbol{\beta}_1)$ and $f_2(y|\mathbf{x};\boldsymbol{\beta}_2)$:

$$\mathcal{I}(\mathbf{x}, \boldsymbol{\beta}_1, \boldsymbol{\beta}_2) = \int_{\mathcal{Y}} \log \frac{f_1(y|\mathbf{x};\boldsymbol{\beta}_1)}{f_2(y|\mathbf{x};\boldsymbol{\beta}_2)} f_1(y|\mathbf{x};\boldsymbol{\beta}_1) dy. \quad (1)$$

Note that the statistical models $f_1(y|\mathbf{x};\boldsymbol{\beta}_1)$ and $f_2(y|\mathbf{x};\boldsymbol{\beta}_2)$ may be conditional densities with respect to a common general measure μ ; to include discrete models it is enough to replace dy with $\mu(dy)$ in Equation (1) and in the rest of the paper. The notation used in (1) is maintained for simplicity.

¹Note that the parameter spaces Θ_1 and Θ_2 can be required instead to be compact.

The quantity in Equation (1) is known to be non-negative, and it is zero if and only if the two responses are equal almost surely. The Kullback-Leibler divergence is often called distance, although it is not symmetric and does not satisfy the triangular inequality. In this context, the Kullback-Leibler divergence in Equation (1) measures the dissimilarity between the two different distributions with parameters β_1 and β_2 , when the experimental condition is \mathbf{x} .

If a design ξ is chosen to maximize the power function in the worst case in a hypothesis test where $f_1(y|\mathbf{x};\beta_1)$ is the true completely known model under the alternative $f_2(y|\mathbf{x};\beta_2)$, then this design is the maximum of the KL-optimality criterion proposed in [19]:

$$I_{2,1}(\xi; \beta_1) = \inf_{\beta_2 \in \Theta_2} \int_{\mathcal{X}} \mathcal{I}(\mathbf{x}, \beta_1, \beta_2) \xi(d\mathbf{x}). \quad (2)$$

The criterion (2) is the minimum Kullback-Leibler distance between the joint distribution $f_1(y|\mathbf{x};\beta_1)\xi(\mathbf{x})$ and the joint statistical model $f_2(y|\mathbf{x};\beta_2)\xi(\mathbf{x})$.

For a given value $\beta_1 \in \Theta_1$ of the first model, a design ξ is *KL-regular* if the following set

$$\Omega_2(\xi; \beta_1) = \left\{ \hat{\beta}_2 : \hat{\beta}_2(\xi) = \arg \inf_{\beta_2 \in \Theta_2} \int_{\mathcal{X}} \mathcal{I}(\mathbf{x}, \beta_1, \beta_2) \xi(d\mathbf{x}) \right\} \quad (3)$$

is a singleton. Otherwise ξ is called *KL-singular*.

A KL-optimum design

$$\xi^* \in \arg \max_{\xi} I_{2,1}(\xi; \beta_1) \quad (4)$$

exists since the KL-criterion function

$$I_{2,1}(\xi; \beta_1) : (\Xi, d_w) \rightarrow [0, +\infty) \quad (5)$$

is concave (as proved in [19]) and upper semi-continuous (as proved in [20]), where Ξ is the set of probability distributions ξ with support $\mathcal{X} \subset \mathbb{R}^q$ endowed with a metric d_w which metrizes the weak convergence on \mathcal{X} . Since \mathcal{X} is compact, the metric space (Ξ, d_w) , which is an infinite-dimensional space, is complete and compact. In fact, any sequence of probability distribution on \mathcal{X} is tight, and hence, by Prokhorov's Theorem, it admits a converging subsequence (in (Ξ, d_w)). In what follows, we take the Kantorovich-Wasserstein metric (see [15]):

$$d_w(\xi_1, \xi_2) = \inf \{ E(|X_1 - X_2|) : X_1 \sim \xi_1, X_2 \sim \xi_2 \}.$$

Since β_1 is assumed to be known, for sake of simplicity, $I_{2,1}(\xi; \beta_1)$ will be denoted by $I_{2,1}(\xi)$.

3 Invariance property and continuity of the KL-criterion

In this section two important theoretical properties are presented. First, we prove that KL-optimum designs are invariant to a scale-position transformation of the design region, whenever this transformation on the experimental condition results in a new parametrization of the rival models. Then, we prove the continuity of the KL-optimality criterion (5),

with respect to the design ξ .

Theorem 1. *Let*

$$\mathcal{Z} = \{\mathbf{z} = \mathbf{a} + \mathbf{B}\mathbf{x} | \mathbf{x} \in \mathcal{X}\}$$

be a rescaled experimental domain, where \mathbf{a} and \mathbf{B} are, respectively, a column vector and a non singular matrix of known constants.

If

1. $f_i(y|\mathbf{x}; \beta_i) = f_i(y|\mathbf{z}; \gamma_i)$, where $\gamma_i = g_i(\beta_i)$ and $g_i(\cdot)$ is a one to one function, and $i = 1, 2$;
2. the following mean log-likelihood

$$\mathcal{ML}_2(\beta_2) = \int_{\mathcal{X}} \int_{\mathcal{Y}} \log f_2(y|\mathbf{x}; \beta_2) f_1(y|\mathbf{x}; \beta_1) dy \xi(d\mathbf{x}) \quad (6)$$

has at least a maximizer $\hat{\beta}_2 \in \arg \sup_{\beta_2} \mathcal{ML}_2(\beta_2)$ for any fixed design $\xi \in \Xi$;

then a KL-optimum design on \mathcal{Z} is

$$\eta^* = \xi^* \circ \mathbf{z}^{-1}, \quad \mathbf{x} \in \mathcal{X} \quad \text{and} \quad \mathbf{z}(\mathbf{x}) = \mathbf{a} + \mathbf{B}\mathbf{x}.$$

Proof. From hypothesis 1 the rival models can be expressed as $f_i(y|\mathbf{z}; \gamma_i)$ where $\mathbf{z} = \mathbf{a} + \mathbf{B}\mathbf{x}$, $i = 1, 2$. The KL-criterion applied to models $f_1(y|\mathbf{z}; \gamma_1)$ and $f_2(y|\mathbf{z}; \gamma_2)$ is

$$\begin{aligned} I_{2,1}^{\mathcal{Z}}(\eta) &= \inf_{\gamma_2} \int_{\mathcal{Z}} \int_{\mathcal{Y}} \log \frac{f_1(y|\mathbf{z}; \gamma_1)}{f_2(y|\mathbf{z}; \gamma_2)} f_1(y|\mathbf{z}; \gamma_1) dy \eta(d\mathbf{z}) \\ &= \int_{\mathcal{Z}} \int_{\mathcal{Y}} \log f_1(y|\mathbf{z}; \gamma_1) f_1(y|\mathbf{z}; \gamma_1) dy \eta(d\mathbf{z}) \\ &\quad - \sup_{\gamma_2} \int_{\mathcal{Z}} \int_{\mathcal{Y}} \log f_2(y|\mathbf{z}; \gamma_2) f_1(y|\mathbf{z}; \gamma_1) dy \eta(d\mathbf{z}). \end{aligned}$$

Let

$$\widetilde{\mathcal{ML}}_2(\gamma_2) = \int_{\mathcal{Z}} \int_{\mathcal{Y}} \log f_2(y|\mathbf{z}; \gamma_2) f_1(y|\mathbf{z}; \gamma_1) dy \eta(d\mathbf{z});$$

from hypothesis 1 and setting $\xi(d\mathbf{x}) = \eta(d\mathbf{z})$

$$\widetilde{\mathcal{ML}}_2(\gamma_2) = \mathcal{ML}_2(g_2^{-1}(\gamma_2))$$

therefore

$$\hat{\gamma}_2 \in \arg \sup \widetilde{\mathcal{ML}}_2(\gamma_2) \quad \text{if and only if} \quad \hat{\beta}_2 \in \arg \sup \mathcal{ML}_2(\beta_2) \quad (7)$$

where $\hat{\gamma}_2 = g_2(\hat{\beta}_2)$.

Using hypothesis 1 and Equation (7),

$$\begin{aligned}
I_{2,1}^{\mathcal{Z}}(\eta) &= \int_{\mathcal{Z}} \int_{\mathcal{Y}} \log f_1(y|\mathbf{z}; \gamma_1) f_1(y|\mathbf{z}; \gamma_1) dy \eta(d\mathbf{z}) \\
&\quad - \int_{\mathcal{Z}} \int_{\mathcal{Y}} \log f_2(y|\mathbf{z}; \hat{\gamma}_2) f_1(y|\mathbf{z}; \gamma_1) dy \eta(d\mathbf{z}) \\
&= \int_{\mathcal{X}} \int_{\mathcal{Y}} \log f_1(y|\mathbf{x}; \beta_1) f_1(y|\mathbf{x}; \beta_1) dy \xi(d\mathbf{x}) \\
&\quad - \int_{\mathcal{X}} \int_{\mathcal{Y}} \log f_2(y|\mathbf{x}; \hat{\beta}_2) f_1(y|\mathbf{x}; \beta_1) dy \xi(d\mathbf{x}) = I_{2,1}^{\mathcal{X}}(\xi),
\end{aligned}$$

which proves the theorem. \square

Continuity is a very nice property which is necessary to prove the convergence of the first order algorithm; see [1], also recalled in Section 4. We require that the Kullback-Leibler divergence between the two conditional distributions is continuous and also Lipschitz with respect to the experimental condition.

Assumption 1. *The Kullback-Leibler divergence $\mathcal{I}(\mathbf{x}, \beta_1, \beta_2)$ given in Equation (1) is a Lipschitz function with respect to \mathbf{x} .*

Note that the Kullback-Leibler divergence between different models is the difference between the cross entropy of the two models and the information entropy of the first one. The information entropy is a continuously differentiable function of the parameters in almost all the classes of parametric models. This smoothness is also observed for the cross entropy between different classes of models with the same support (for instance, Lognormal, Weibull, Gamma, ...). Therefore, since Ξ is compact, if the parameters of the rival models (which are assumed to depend on \mathbf{x} through some coefficients) are continuously differentiable functions of the experimental condition \mathbf{x} , then the Kullback-Leibler divergence is a continuously differentiable function of the experimental conditions, and hence a Lipschitz function.

Denote by $\mathcal{J}(\xi, \beta_1, \beta_2)$ the average of the function $\mathcal{I}(\mathbf{x}, \beta_1, \beta_2)$ in (1) with respect to the probability measure ξ , namely

$$\mathcal{J}(\xi, \beta_1, \beta_2) = \int_{\mathcal{X}} \mathcal{I}(\mathbf{x}, \beta_1, \beta_2) d\xi(\mathbf{x}) = \int_{\mathcal{X}} \int_{\mathcal{Y}} \log \frac{f_1(y|\mathbf{x}; \beta_1)}{f_2(y|\mathbf{x}; \beta_2)} f_1(y|\mathbf{x}; \beta_1) dy \xi(d\mathbf{x}).$$

Again, since \mathcal{X} is compact, Assumption 1 implies that the function \mathcal{J} is continuous with respect to ξ . Moreover, this function is linear in ξ . The problem of finding a KL-optimal design as in equation (4) is an infinite dimension minmax problem. Our goal is to prove that $I_{2,1}(\xi) = \inf_{\beta_2} \mathcal{J}(\xi, \beta_1, \beta_2)$ is continuous as an extension of classical results for semi-infinite problem (see for instance [21]) to our context. We start with a counterexample, which shows that Assumption 1 is not sufficient for $I_{2,1}(\xi)$ to be continuous.

Example 1 ($I_{2,1}(\xi; \beta_1)$ is not continuous). *Take $\mathcal{X} = [0, 1]$, $\Theta_2 = (0, \infty)$, and define*

$$\mathcal{I}(\mathbf{x}, \beta_1, \beta_2) = \begin{cases} 2((2\beta_2 - 1)\mathbf{x} + (1 - \beta_2)) & \text{if } 0 < \beta_2 \leq 1 \\ (\beta_2 + 1)\mathbf{x}^{\beta_2} & \text{if } 1 < \beta_2 \end{cases}$$

We have:

- $\mathcal{I}(\mathbf{x}, \boldsymbol{\beta}_1, \boldsymbol{\beta}_2)$ is a continuous function on $\mathcal{X} \times \Theta_2$;
- $\mathcal{I}(\mathbf{x}, \boldsymbol{\beta}_1, \boldsymbol{\beta}_2)$ is a convex and Lipschitz function of \mathbf{x} , for any $\boldsymbol{\beta}_2 \in \Theta_2$;
- $I_{2,1}(\delta_{\mathbf{x}}; \boldsymbol{\beta}_1) = 0$ for any $\mathbf{x} \in \mathcal{X}$.

Take ξ_n be the uniform distribution on $[0, 1 - 1/n]$; it can be easily proved that $d_w(\xi_n, \xi) \rightarrow 0$, where ξ is the uniform distribution on $[0, 1]$. We have

$$\begin{aligned} \int_{\mathcal{X}} \mathcal{I}(\mathbf{x}, \boldsymbol{\beta}_1, \boldsymbol{\beta}_2) d\xi_n(\mathbf{x}) &= \int_0^{1-1/n} \frac{\mathcal{I}(\mathbf{x}, \boldsymbol{\beta}_1, \boldsymbol{\beta}_2)}{1-1/n} d\mathbf{x} \\ &= \begin{cases} 1 - \frac{2\boldsymbol{\beta}_2 - 1}{n} & \text{if } 0 < \boldsymbol{\beta}_2 \leq 1 \\ (1 - 1/n)^{\boldsymbol{\beta}_2} & \text{if } 1 < \boldsymbol{\beta}_2 \end{cases} \end{aligned}$$

while

$$\int_{\mathcal{X}} \mathcal{I}(\mathbf{x}, \boldsymbol{\beta}_1, \boldsymbol{\beta}_2) d\xi(\mathbf{x}) = \int_0^1 \mathcal{I}(\mathbf{x}, \boldsymbol{\beta}_1, \boldsymbol{\beta}_2) d\mathbf{x} \equiv 1.$$

Hence,

$$\begin{aligned} I_{2,1}(\xi_n) &= \inf_{\boldsymbol{\beta}_2 \in \Theta_2} \int_{\mathcal{X}} \mathcal{I}(\mathbf{x}, \boldsymbol{\beta}_1, \boldsymbol{\beta}_2) d\xi_n(\mathbf{x}) = 0 \\ &\neq 1 = \inf_{\boldsymbol{\beta}_2 \in \Theta_2} \int_{\mathcal{X}} \mathcal{I}(\mathbf{x}, \boldsymbol{\beta}_1, \boldsymbol{\beta}_2) d\xi(\mathbf{x}) = I_{2,1}(\xi). \end{aligned}$$

We give here a mild assumption which is satisfied in many situations. In fact, when we fix $\boldsymbol{\beta}_1$, we can expect that that the Kullback-Leibler divergence $\mathcal{I}(\mathbf{x}, \boldsymbol{\beta}_1, \boldsymbol{\beta}_2)$ is “dominated” in $\boldsymbol{\beta}_2$, in that if $\mathcal{I}(\mathbf{x}, \boldsymbol{\beta}_1, \boldsymbol{\beta}_2)$ is too big for some \mathbf{x} , there is another model $f_2(y|\mathbf{x}; \tilde{\boldsymbol{\beta}}_2)$ that is always closer to $f_1(y|\mathbf{x}; \boldsymbol{\beta}_1)$ and dominated by a constant $M(\boldsymbol{\beta}_1)$.

Assumption 2. For any fixed $\boldsymbol{\beta}_1$, there exists $M = M(\boldsymbol{\beta}_1) > 0$ such that if $\mathcal{I}(\mathbf{x}, \boldsymbol{\beta}_1, \boldsymbol{\beta}_2) > M$ for some $\mathbf{x} \in \mathcal{X}$, then there will exists $\tilde{\boldsymbol{\beta}}_2$ such that

$$\mathcal{I}(\mathbf{x}, \boldsymbol{\beta}_1, \boldsymbol{\beta}_2) \geq \mathcal{I}(\mathbf{x}, \boldsymbol{\beta}_1, \tilde{\boldsymbol{\beta}}_2), \quad \forall \mathbf{x} \in \mathcal{X},$$

and

$$\sup_{\mathbf{x} \in \mathcal{X}} \mathcal{I}(\mathbf{x}, \boldsymbol{\beta}_1, \tilde{\boldsymbol{\beta}}_2) \leq M.$$

Theorem 2. Assume 1 and 2. The KL-criterion (2) is a locally Lipschitz function and hence it is a continuous function of ξ .

Proof. Let $\boldsymbol{\beta}_1$ be fixed and M be as in Assumption 2. Define

$$\Theta_2(\boldsymbol{\beta}_1) = \{\tilde{\boldsymbol{\beta}}_2 \in \Theta_2 : \sup_{\mathbf{x} \in \mathcal{X}} \mathcal{I}(\mathbf{x}, \boldsymbol{\beta}_1, \tilde{\boldsymbol{\beta}}_2) \leq M\}. \quad (8)$$

The KL-criterion (2), for any $\xi \in \Xi$, may be rewritten as

$$I_{2,1}(\xi) = \inf_{\beta_2 \in \Theta_2(\beta_1)} \int_{\mathcal{X}} \mathcal{I}(\mathbf{x}, \beta_1, \beta_2) \xi(d\mathbf{x}), \quad (9)$$

for Assumption 2.

Let \mathcal{V} be the real vector space of all signed finite measures on \mathcal{X} (equipped with the usual Borel σ -algebra \mathcal{B}), which contains Ξ as proper, closed, convex subset. Denoting by $\|h\|_L$ the operator norm of the Banach space of Lipschitz function on \mathcal{X} , the vector space \mathcal{V} , equipped with the norm

$$\|\xi\|_{\mathcal{V}} = \sup \left\{ \left| \int_{\mathcal{X}} h(\mathbf{x}) \xi(d\mathbf{x}) \right|, \|h\|_L \leq 1 \right\},$$

is a Banach space. (This is a consequence of the results in [28] and [12]).

The map $\xi \mapsto \int_{\mathcal{X}} \mathcal{I}(\mathbf{x}, \beta_1, \beta_2) \xi(d\mathbf{x})$, is a linear functional on \mathcal{V} . Moreover, assumption 1 guarantees its boundness on the unit ball $\|\xi\|_{\mathcal{V}} \leq 1$:

$$\begin{aligned} \sup_{\|\xi\|_{\mathcal{V}} \leq 1} \left| \int_{\mathcal{X}} \mathcal{I}(\mathbf{x}, \beta_1, \beta_2) \xi(d\mathbf{x}) \right| &\leq \sup_{\substack{\|\xi\|_{\mathcal{V}} \leq 1 \\ \|h\|_L \leq \|\mathcal{I}(\cdot, \beta_1, \beta_2)\|_L}} \left| \int_{\mathcal{X}} h(\mathbf{x}) \xi(d\mathbf{x}) \right| \\ &= \|\mathcal{I}(\cdot, \beta_1, \beta_2)\|_L \sup_{\substack{\|\xi\|_{\mathcal{V}} \leq 1 \\ \|g\|_L \leq 1}} \left| \int_{\mathcal{X}} g(\mathbf{x}) \xi(d\mathbf{x}) \right| \\ &= \|\mathcal{I}(\cdot, \beta_1, \beta_2)\|_L, \end{aligned}$$

and hence $\xi \mapsto \int_{\mathcal{X}} \mathcal{I}(\mathbf{x}, \beta_1, \beta_2) \xi(d\mathbf{x})$ is a continuous function (see, e.g., [5]). The function

$$\xi \mapsto \inf_{\beta_2 \in \Theta_2(\beta_1)} \int_{\mathcal{X}} \mathcal{I}(\mathbf{x}, \beta_1, \beta_2) \xi(d\mathbf{x})$$

is concave and upper semi-continuous function since it is the point-wise infimum of linear continuous functions. Moreover, $-\infty < \inf_{\beta_2 \in \Theta_2(\beta_1)} \int_{\mathcal{X}} \mathcal{I}(\mathbf{x}, \beta_1, \beta_2) \xi(d\mathbf{x}) < \infty$ since, for any $\beta_2 \in \Theta_2(\beta_1)$,

$$\left| \int_{\mathcal{X}} \mathcal{I}(\mathbf{x}, \beta_1, \beta_2) \xi(d\mathbf{x}) \right| \leq M \|\xi\|_{\mathcal{V}},$$

as a consequence of (8). Therefore (see [5]) the function

$$\xi \mapsto \inf_{\beta_2 \in \Theta_2(\beta_1)} \int_{\mathcal{X}} \mathcal{I}(\mathbf{x}, \beta_1, \beta_2) \xi(d\mathbf{x})$$

is locally Lipschitz and hence continuous on the Banach space $(\mathcal{V}, \|\cdot\|_{\mathcal{V}})$.

Recall that the set Ξ is the set of the possible experimental designs ξ . When $\xi_1, \xi_2 \in \Xi$, the Kantorovich-Wasserstein distance d_w can be rewritten also as (see [15])

$$d_w(\xi_1, \xi_2) = \sup \left\{ \left| \int_{\mathcal{X}} h(\mathbf{x}) d\xi_1(\mathbf{x}) - \int_{\mathcal{X}} h(\mathbf{x}) d\xi_2(\mathbf{x}) \right|, \|h\|_L \leq 1 \right\} = \|\xi_1 - \xi_2\|_{\mathcal{V}}.$$

Hence the KL-criterion (2) is a locally Lipschitz and continuous function on Ξ . \square

4 First order algorithm to obtain KL-optimum designs

To construct a KL-optimum design ξ^* , [19] propose the use of the first order algorithm:

a) Given ξ_n , find

a1)

$$\beta_{2,n} = \arg \min_{\beta_2 \in \Omega_2} \int_{\mathcal{X}} \int_{\mathcal{Y}} \log \frac{f_1(y|\mathbf{x}; \beta_1)}{f_2(y|\mathbf{x}; \beta_2)} f_1(y|\mathbf{x}; \beta_1) dy \xi_n(d\mathbf{x})$$

a2)

$$\mathbf{x}_n = \arg \max_{\mathbf{x} \in \mathcal{X}} \int_{\mathcal{Y}} \log \frac{f_1(y|\mathbf{x}; \beta_1)}{f_2(y|\mathbf{x}; \beta_{2,n})} f_1(y|\mathbf{x}; \beta_1) dy$$

b) Properly choose $0 \leq \alpha_n \leq 1$ and construct $\xi_{n+1} = (1 - \alpha_n)\xi_n + \alpha_n\delta_{\mathbf{x}_n}$.

Moreover, [19] prove that, for any ξ , $U(\xi) \leq I_{2,1}(\xi)/I_{2,1}(\xi^*) \leq 1$ where

$$U(\xi) = \left[1 + \frac{\max_{\mathbf{x} \in \mathcal{X}} \psi(\mathbf{x}; \xi)}{I_{2,1}(\xi)} \right]^{-1}, \quad (10)$$

and

$$\psi(\mathbf{x}; \xi) = \mathcal{I}(\mathbf{x}, \beta_1, \hat{\beta}_2) - \int_{\mathcal{X}} \mathcal{I}(\mathbf{s}, \beta_1, \hat{\beta}_2) \xi(d\mathbf{s}) \quad (11)$$

is the directional derivative of $I_{2,1}(\xi)$ at a regular design ξ in the direction of $\delta_{\mathbf{x}} - \xi$ and $\hat{\beta}_2$ is the unique element of (3). According to [19], the iterative procedure should stop at step N if the upper bound of the efficiency $U(\xi_N) > \delta$, where $0 < \delta < 1$ is a value chosen by the experimenter, for instance $\delta = 0.99$.

However, [19] give no specific rule to choose α_n , neither they provide any proof of the convergence of the algorithm. Differently, authors [1] suggest to use

$$\alpha_n = \arg \max_{\alpha \in [0,1]} I_{2,1}[(1 - \alpha)\xi_n + \alpha\delta_{\mathbf{x}}] \quad (12)$$

and using the theory of point-to-set mappings they provide useful convergence results. As a consequence of Theorem 1 in [1] and of the continuity of $I_{2,1}(\xi)$ proved in Section 3, the following theorem states that if the algorithm explores regular designs (see (3) in Section 2), then it converges to the set of KL-optimum designs.

Theorem 3. *Let ξ_0 be an initial design. For any $n \geq 0$, let ξ_{n+1} be one of the designs obtained by the first order algorithm at step $(n + 1)$, with α_n as in (12).*

If ξ_n is a sequence of KL-regular designs then, as $n \rightarrow \infty$,

$$\inf_{\xi^* \in \arg \max I_{2,1}(\xi)} d_w(\xi_n, \xi^*) \rightarrow 0$$

and

$$|I_{2,1}(\xi_n) - \max_{\xi} I_{2,1}(\xi)| \rightarrow 0.$$

In particular, if the set $\{\xi^ \in \Xi : \xi^* = \arg \max_{\xi} I_{2,1}(\xi)\}$ is a singleton, then $\xi_n \rightarrow \xi^*$.*

Proof. We prove the first part of the theorem by way of contradiction. Suppose that for every $\epsilon > 0$ and for every $N \geq 1$ there exist $n_N \geq N$ such that

$$d_w(\xi_{n_N}, \xi^*) \geq \epsilon \quad (13)$$

for all $\xi^* \in \arg \max I_{2,1}(\xi)$.

In [1, Theorem 1] it is stated that if $(\xi_n)_n$ is a sequence of regular designs generated by the above described algorithm, then the limit of any converging subsequence of $(\xi_n)_n$ is a KL-optimum design. In [1, Theorem 1] it is required that the sub-level $\{\xi \in \Xi: I_{2,1}(\xi) \geq I_{2,1}(\xi_0)\}$ is compact. This is always verified as a consequence of Theorem 2: the sub-level is a closed subset of the compact set Ξ .

The sequence $(\xi_{n_N})_N$ in (13) has at least one convergent subsequence since Ξ is compact; then, from [1, Theorem 1] this converges to a KL-optimum design, which contradicts (13). It follows that for every $\epsilon > 0$ there exists N such that for every $n \geq N$

$$\inf_{\xi^* \in \arg \max I_{2,1}(\xi)} d_w(\xi_n, \xi^*) < \epsilon, \quad (14)$$

and hence the first part of the theorem is proved. Since $I_{2,1}(\xi)$ is continuous and $\{\xi^* \in \arg \max I_{2,1}(\xi)\}$ is compact, the second part of the result follows from (14). \square

A drawback of the application of the first order algorithm is that there is no guarantee that ξ_{n+1} is a regular design in the sense of (3), even if ξ_n is. In the next section we prove that, for generalized linear models (GLM), the algorithm always moves at regular designs; thus it converges, provided that the initial design ξ_0 has a non singular information matrix.

4.1 Regular designs for GLM

As recalled in Section 2, if the KL-criterion is used to find an optimum design to discriminate between two models, a regular design is such that there is a unique estimate of the parameters β_2 as in (3). In many situations, however, the aim of an experiment is to estimate as precisely as possible the parameters (or a function of the parameters) of the assumed known model. In this context a design is called regular if its Fisher information matrix is non singular. In this section it is proved that these two different definitions of “regular design” are equivalent for a large class of useful statistical models.

In what follows, we specialize the notation for a GLM. \mathbf{X} is the design matrix whose i -th row is \mathbf{x}_i^T (i.e. \mathbf{x}_i is the $d_2 \times 1$ vector of experimental conditions for the i -th unit), $\beta_2 \in \Theta_2$ is the $d_2 \times 1$ vector of regression parameters of $f_2(y|\mathbf{x}, \beta_2)$, $\mu_i = E_2(Y|\mathbf{x}_i)$ and $\text{Var}_i(Y) = \text{Var}_2(Y|\mathbf{x}_i)$ are the response mean and the response variance under model $f_2(y|\mathbf{x}_i, \beta_2)$, and $\eta_i = \mathbf{x}_i^T \beta_2$ is the linear predictor. Both μ_i and $\text{Var}_i(Y)$ depend on \mathbf{x}_i and β_2 through η_i , in particular $g(\mu_i) = \eta_i$ where $g(\cdot)$ is a link function. It is well known that for the generalized linear model the Fisher information matrix is

$$\mathbf{J} = \mathbf{X}^T \mathbf{W} \mathbf{X}, \quad (15)$$

where

$$\mathbf{W} = \text{diag} \left[\frac{1}{\text{Var}_i(Y)} \left(\frac{\partial \mu_i}{\partial \eta_i} \right)^2 \right]$$

is a diagonal matrix which depends on both the design matrix \mathbf{X} and the parameter vector $\boldsymbol{\beta}_2$.

From (15), $\mathbf{J} = \tilde{\mathbf{X}}^T \tilde{\mathbf{X}}$, where $\tilde{\mathbf{X}} = \mathbf{W}^{1/2} \mathbf{X}$; thus

$$\text{Range } \mathbf{J} = \text{Range } \tilde{\mathbf{X}}^T = \text{Range } \mathbf{X}^T,$$

since \mathbf{W} is a diagonal non-singular matrix. Therefore, the Fisher information matrix is non-singular (i.e. $\text{Rank}(\mathbf{J}) = d_2$) if and only if $\text{Rank}(\mathbf{X}) = d_2$.

As recalled in Section 2, in the context of discrimination between rival models, an exact design $\xi = \frac{1}{n} \sum_i \delta_{\mathbf{x}_i}$ is regular if the set

$$\Omega_2 \left(\frac{1}{n} \sum_i \delta_{\mathbf{x}_i} \right) = \left\{ \hat{\boldsymbol{\beta}}_2 : \hat{\boldsymbol{\beta}}_2 = \arg \min_{\boldsymbol{\beta}_2} \sum_{i=1}^n \int_{\mathcal{Y}} \log \frac{f_1(y|\mathbf{x}_i, \boldsymbol{\beta}_1)}{f_2(y|\mathbf{x}_i, \boldsymbol{\beta}_2)} f_1(y|\mathbf{x}_i, \boldsymbol{\beta}_1) dy \right\}$$

is singleton. Since

$$\begin{aligned} & \sum_{i=1}^n \int_{\mathcal{Y}} \log \frac{f_1(y|\mathbf{x}_i, \boldsymbol{\beta}_1)}{f_2(y|\mathbf{x}_i, \boldsymbol{\beta}_2)} f_1(y|\mathbf{x}_i, \boldsymbol{\beta}_1) dy \\ &= \sum_{i=1}^n \int_{\mathcal{Y}} \log f_1(y|\mathbf{x}_i, \boldsymbol{\beta}_1) f_1(y|\mathbf{x}_i, \boldsymbol{\beta}_1) dy - \sum_{i=1}^n \int_{\mathcal{Y}} \log f_2(y|\mathbf{x}_i, \boldsymbol{\beta}_2) f_1(y|\mathbf{x}_i, \boldsymbol{\beta}_1) dy, \end{aligned}$$

finding

$$\arg \min_{\boldsymbol{\beta}_2} \sum_{i=1}^n \int_{\mathcal{Y}} \log \frac{f_1(y|\mathbf{x}_i, \boldsymbol{\beta}_1)}{f_2(y|\mathbf{x}_i, \boldsymbol{\beta}_2)} f_1(y|\mathbf{x}_i, \boldsymbol{\beta}_1) dy$$

is equivalent to compute $\arg \max_{\boldsymbol{\beta}_2} \mathcal{ML}_2(\boldsymbol{\beta}_2)$ where $\xi = \frac{1}{n} \sum_i \delta_{\mathbf{x}_i}$ and thus

$$\mathcal{ML}_2(\boldsymbol{\beta}_2) = \sum_{i=1}^n \int_{\mathcal{Y}} \frac{1}{n} \log f_2(y|\mathbf{x}_i, \boldsymbol{\beta}_2) f_1(y|\mathbf{x}_i, \boldsymbol{\beta}_1) dy.$$

Assuming that $\mathcal{ML}_2(\boldsymbol{\beta}_2)$ is differentiable with respect to $\boldsymbol{\beta}_2$ for any fixed \mathbf{x}_i , then the maximization of $\mathcal{ML}_2(\boldsymbol{\beta}_2)$ is performed by setting the partial derivatives of $\mathcal{ML}_2(\boldsymbol{\beta}_2)$ equal to zero. Let us assume that

$$\begin{aligned} \frac{\partial \mathcal{ML}_2(\boldsymbol{\beta}_2)}{\partial \beta_{2j}} &\propto \sum_{i=1}^n \frac{\partial}{\partial \beta_{2j}} \int_{\mathcal{Y}} \log f_2(y|\mathbf{x}_i, \boldsymbol{\beta}_2) f_1(y|\mathbf{x}_i, \boldsymbol{\beta}_1) dy \\ &= \sum_{i=1}^n \int_{\mathcal{Y}} \frac{\partial}{\partial \beta_{2j}} \log f_2(y|\mathbf{x}_i, \boldsymbol{\beta}_2) f_1(y|\mathbf{x}_i, \boldsymbol{\beta}_1) dy. \end{aligned} \quad (16)$$

It is well known that if $f_2(y|\mathbf{x}, \boldsymbol{\beta}_2)$ is a GLM, then at \mathbf{x}_i , the i -th experimental condition,

$$\frac{\partial \log f_2(y|\mathbf{x}_i, \boldsymbol{\beta}_2)}{\partial \beta_{2j}} = \frac{(y - \mu_i)x_{ij}}{\text{Var}_i(Y)} \left(\frac{\partial \mu_i}{\partial \eta_i} \right), \quad j = 1, \dots, d_2.$$

Therefore, from equation (16)

$$\begin{aligned} \frac{\partial \mathcal{ML}_2(\boldsymbol{\beta}_2)}{\partial \beta_{2j}} &\propto \sum_{i=1}^n \int_{\mathcal{Y}} \frac{(y - \mu_i)x_{ij}}{\text{Var}_i(Y)} \left(\frac{\partial \mu_i}{\partial \eta_i} \right) f_1(y|\mathbf{x}_i, \boldsymbol{\beta}_1) dy \\ &= \sum_{i=1}^n \frac{[\mathbf{E}_1(Y|\mathbf{x}_i) - \mu_i]x_{ij}}{\text{Var}_i(Y)} \left(\frac{\partial \mu_i}{\partial \eta_i} \right) \\ &= \mathbf{v}^T(\mathbf{X}; \boldsymbol{\beta}_2) \mathbf{x}_{.j}, \quad j = 1 \dots, d_2, \end{aligned} \quad (17)$$

where $\mathbf{x}_{.j}$ denotes the j -th column of the design matrix \mathbf{X} and $\mathbf{v}(\mathbf{X}; \boldsymbol{\beta}_2)$ is a $n \times 1$ vector whose i -th item is

$$v_i(\mathbf{x}_i; \boldsymbol{\beta}_2) = \frac{\mathbf{E}_1(Y|\mathbf{x}_i) - \mathbf{E}_2(Y|\mathbf{x}_i)}{\text{Var}_i(Y)} \left(\frac{\partial \mu_i}{\partial \eta_i} \right)$$

and $\mathbf{E}_1(Y|\mathbf{x}_i)$ is the response mean under the known model $f_1(y|\mathbf{x}_i, \boldsymbol{\beta}_1)$.

Setting equal to zero the right-hand side of equations (17), we have a system of d_2 non-linear equations whose solution is the maximum likelihood estimator. For the most commonly used GLMs (for instance when the link function is canonical and for the link functions described in [29]) this solution exists and is unique if and only if $\text{Rank}(\mathbf{X}) = d_2$ i.e. if and only if the Fisher information matrix is non-singular.

Let us denote by ‘‘regular GLM’’ a GLM for which there exists a unique solution of the system $\mathbf{v}^T(\mathbf{X}; \boldsymbol{\beta}_2) \mathbf{x}_{.j} = 0$, $j = 1, \dots, d_2$ if and only if $\text{Rank}(\mathbf{X}) = d_2$. Theorem 3 has the following corollary at least for regular GLMs.

Corollary 1. *Let $f_1(y|\mathbf{x}, \boldsymbol{\beta}_1)$ and $f_2(y|\mathbf{x}, \boldsymbol{\beta}_2)$ be regular GLMs and let $\xi_0 \in \Xi$ have a non-singular Fisher information matrix. For any $n \geq 0$, let ξ_{n+1} be one of the designs obtained by the first order algorithm at step $(n + 1)$, with α_n as in (12). Then*

$$|I_{2,1}(\xi_n) - \max_{\xi} I_{2,1}(\xi)| \rightarrow 0$$

as $n \rightarrow \infty$. In particular, if the optimum ξ^* is unique, $\xi_n \rightarrow \xi^*$.

Proof. In the classes of the GLMs, any KL-regular design has non-singular Fisher information matrix, and any design with non-singular Fisher information matrix is KL-regular; therefore if ξ_n is regular then ξ_{n+1} is also regular. For these models, it is then guaranteed that the algorithm moves on regular designs if it starts from a regular design ξ_0 . \square

5 Examples and computational aspects

In this section we show some practical implications of the theoretical results proved in Sections 3 and 4.

For what concerns the invariance property, the practitioner can observe that the algorithm performs differently in different experimental regions. If it is numerically difficult to compute the KL-optimum design in a specified \mathcal{X} , the experimenter can take advantage of the invariance property. He can compute the KL-optimum design in a suitable transformed space \mathcal{Z} (where less computational effort is required) and then find the KL-optimum design in \mathcal{X} through the invariance property. In the example provided in Subsection 5.1 the algorithm needs more iterations than in a transformed space.

Continuity of KL-criterion is necessary to assure that the approximated solution obtained with the stopping rule is close to the KL-optimum design. Furthermore, it also guarantees that the solution obtained through the regularization procedure is a “nearly” KL-optimum design. Subsection 5.2 gives an example where the problem is regularized since the KL-optimum design is singular.

In Subsection 5.3 some practical hints are proposed, in comparison with the suggestions provided in [13].

5.1 A benchmark test

In this example we develop a benchmark test for the convergence of the algorithm, where the analytical result for the KL-optimum design is obtained through the Chebyshev polynomials.

We take $x \in \mathcal{X} = [-1; 1]$ and we test a polynomial of degree three versus a polynomial of degree two. In other words, $f_1(y|\mathbf{x}, \boldsymbol{\beta}_1)$ and $f_2(y|\mathbf{x}, \boldsymbol{\beta}_2)$ are the probability densities of two Gaussian random variables with means $\beta_{01} + \beta_{11}x + \beta_{21}x^2 + \beta_{31}x^3$ and $\beta_{02} + \beta_{12}x + \beta_{22}x^2$, respectively and common variance σ^2 . We assume the polynomial of degree 3 is completely known. Except for the constant of proportionality β_{31} , given a design ξ , the Kullback-Leibler divergence between these two Gaussian models is $\mathcal{I}(x, \boldsymbol{\theta}) = (\theta_0 + \theta_1x + \theta_2x^2 + x^3)^2$, where $\theta_i = \frac{\beta_{i1} - \beta_{i2}}{\beta_{31}}$ and $i = 0, 1, 2$. Let

$$\mathcal{J}(\xi, \boldsymbol{\theta}) = \int_{\mathcal{X}} (\theta_0 + \theta_1x + \theta_2x^2 + x^3)^2 \xi(dx),$$

then, except for a constant of proportionality, the KL-criterion is

$$I_{2,1}(\xi) = \inf_{\boldsymbol{\theta}} \mathcal{J}(\xi, \boldsymbol{\theta}).$$

In order to solve the previous optimization problem we set the partial derivatives of $\mathcal{J}(\xi, \boldsymbol{\theta})$ equal to zero, i.e.

$$\frac{\partial \mathcal{J}(\xi, \boldsymbol{\theta})}{\partial \theta_i} = 2 \int_{\mathcal{X}} x^i (\theta_0 + \theta_1x + \theta_2x^2 + x^3) \xi(dx) = 0, \quad i = 0, 1, 2. \quad (18)$$

The Chebyshev polynomial of degree three, $T_3(x) = 4x^3 - 3x$, fulfills the following orthogonality conditions,

$$\int_{\mathcal{X}} x^i T_3(x) \tilde{\xi}(dx) = 0, \quad i = 0, 1, 2,$$

where $\tilde{\xi} = \begin{Bmatrix} -1 & -1/2 & 1/2 & 1 \\ 1/6 & 1/3 & 1/3 & 1/6 \end{Bmatrix}$. Let us note that the design points of $\tilde{\xi}$ are the singular points of $T_3(x)$ in $\mathcal{X} = [-1, 1]$.

Therefore, if we denote the solutions of Equation (18) by $\hat{\theta} = (\hat{\theta}_0, \hat{\theta}_1, \hat{\theta}_2)^T$, we have that $\hat{\theta}_0 = \hat{\theta}_2 = 0$ and $\hat{\theta}_1 = -3/4$.

Let us prove that $\tilde{\xi}$ is actually a KL-optimum design by checking the equivalence theorem inequality (see [19], for more details). Since $T_3(x)$ is a Shabat polynomial (i.e. it assumes the same minimum and maximum values except for the sign) we have that the directional derivative of $I_{2,1}(\xi)$ at $\tilde{\xi}$ in the direction of $\xi_x - \tilde{\xi}$ is

$$\psi(x; \tilde{\xi}) = \mathcal{I}(x, \hat{\theta}) - \int_{\mathcal{X}} \mathcal{I}(x, \hat{\theta}) \tilde{\xi}(dx) = \left(-\frac{3}{4}x + x^3\right)^2 - \max_x \left(-\frac{3}{4}x + x^3\right)^2 \leq 0$$

with equality at the singular points -1, -1/2, 1/2 and 1. This proves that $\tilde{\xi}$ is a KL-optimum design. The same result has been shown in a different way by [23] and extended by [8].

We have applied the first order algorithm described in Section 4 using Matlab 7.6 optimization toolbox, obtaining the following results. The iterative procedure converges to the KL-optimum design $\tilde{\xi}$ after 107 iterations with $\delta = 0.99$, and after 384 iterations with $\delta = 0.995$.

We have then tested the invariance property in this benchmark setting. If the experimental condition is $z = 2 + 4x \in \mathcal{Z} = [-2; 6]$ instead of $x \in \mathcal{X}$, then after 430 iteration with $\delta = 0.95$ the KL-optimum design $\eta^* = \begin{Bmatrix} -2 & 0 & 4 & 6 \\ 1/6 & 1/3 & 1/3 & 1/6 \end{Bmatrix}$ is reached. This solution is consistent with the invariance property given in Theorem 1.

Note that in the transformed experimental region the optimum is reached with a larger number of iterations.

5.2 Discrimination between two logistic regression models

The purpose of this Subsection is to handle an example where the optimal design is singular. The experimental conditions are assumed to vary in the interval $\mathcal{X} = [0, 1]$ and logistic regression models are considered instead of linear regression models. In other words, Y is a binary response variable such that

$$P(Y = 1|x; \beta_i) = F(\eta_i) = \frac{e^{\eta_i}}{1 + e^{\eta_i}}, \quad i = 1, 2$$

where $\eta_1 = \beta_{01} + \beta_{11}x + \beta_{21}x^2$ and $\eta_2 = \beta_{12}x + \beta_{22}x^2$ are two rival models for the expected response. For these models the Kullback-Leibler divergence is

$$\mathcal{I}(x, \beta_1, \beta_2) = (\eta_1 - \eta_2) \frac{\exp(\eta_1)}{1 + \exp(\eta_1)} + \log \frac{1 + \exp(\eta_2)}{1 + \exp(\eta_1)}.$$

The algorithm reaches in few steps the one point design which concentrates the whole mass at zero, i.e. $\xi^* = \delta_0$. Once the algorithm has reached this design, it fails in capturing the whole $\Omega_2(\xi^*)$, and it enters in a loop.

It is well known that a design ξ^* is optimum if and only if $\psi(\mathbf{x}, \xi^*) \leq 0$ for any \mathbf{x} . However, $\psi(\mathbf{x}, \xi)$ can be computed only for regular designs. When the optimum design is singular, in order to check its optimality we can regularize the problem (see [1]), that is to use the function

$$I_\gamma(\xi) = I_{2,1}[(1 - \gamma)\xi + \gamma\tilde{\xi}] \quad (19)$$

instead of $I_{2,1}(\xi)$, where $0 < \gamma < 1$ and $\tilde{\xi}$ is a regular design, that is a design with a non singular information matrix. Let $\xi_1 = (1 - \gamma)\xi + \gamma\tilde{\xi}$; then, $I_\gamma(\xi) = I_{21}(\xi_1)$. The first order algorithm may be specialized for $I_\gamma(\xi)$ using $\xi_{1n} = (1 - \gamma)\xi_n + \gamma\tilde{\xi}$ instead of ξ_n at the step (a1) of the algorithm described in Section 4. The stopping rule recalled in Section 4 may also be specialized for $I_\gamma(\xi)$.

Let $\bar{\xi}_1 = (1 - \gamma)\bar{\xi} + \gamma\tilde{\xi}$; then, the directional derivative of $I_\gamma(\xi)$ in the direction of $\bar{\xi} - \xi$ is

$$\partial I_\gamma(\xi; \bar{\xi}) = \partial I_{21}(\xi_1; \bar{\xi}_1) = \int_{\mathcal{X}} \psi(\mathbf{x}; \xi_1) \bar{\xi}_1(d\mathbf{x}), \quad (20)$$

where the last equality is proved by [19]. From (20) and (11), it follows that

$$\begin{aligned} \partial I_\gamma(\xi; \bar{\xi}) &= \int_{\mathcal{X}} \left[\mathcal{I}(f_1, f_2, \mathbf{x}, \hat{\beta}_2) - \int_{\mathcal{X}} \mathcal{I}(f_1, f_2, \mathbf{s}, \hat{\beta}_2) \xi_1(ds) \right] \bar{\xi}_1(d\mathbf{x}) \\ &= \int_{\mathcal{X}} (1 - \gamma) \psi(\mathbf{x}; \xi) \bar{\xi}(d\mathbf{x}), \end{aligned} \quad (21)$$

where $\psi(\mathbf{x}; \xi) = \left[\mathcal{I}(f_1, f_2, \mathbf{x}, \hat{\beta}_2) - \int_{\mathcal{X}} \mathcal{I}(f_1, f_2, \mathbf{s}, \hat{\beta}_2) \xi(ds) \right]$ with

$$\hat{\beta}_2 = \beta_2(\xi_1) = \arg \min_{\beta_2 \in \Omega_2} \int_{\mathcal{X}} \mathcal{I}(f_1, f_2, \mathbf{x}, \beta_2) \xi_1(d\mathbf{x}). \quad (22)$$

instead of $\hat{\beta}_2 = \beta_2(\xi)$.

Since $I_\gamma(\xi)$ is a concave function, the iterative procedure based on $I_\gamma(\xi)$ stops at the step n if

$$\left[1 + \frac{\max_{\mathbf{x} \in \mathcal{X}} (1 - \gamma) \psi(\mathbf{x}; \xi_n)}{I_\gamma(\xi_n)} \right]^{-1} > \delta, \quad (23)$$

where $0 < \delta < 1$ is a suitable value chosen by the experimenter, e.g. $\delta = 0.995$.

It is well known that in maxmin problems it may be difficult to check the optimality of a design because the expression of the directional derivative is not easy to be found when a design is not regular (see, for instance, [14, Theorem 2.6.1]). More specifically, in order to check the directional derivative it is necessary introduce a measure on β_2 , which is not easy to be found. Differently, since $I_\gamma(\xi) = I_{21}(\xi_1)$ and ξ_1 is a regular design by definition, equation (21) provides the directional derivative of $I_\gamma(\xi)$ at any design ξ , without assuming the regularity of ξ . Therefore, if the above described algorithm stops at the n -th step, it is always possible to check the optimality of ξ_n through the following directional derivative of $I_\gamma(\xi)$ at ξ_n in the direction of $\delta_{\mathbf{x}} - \xi_n$,

$$\psi_\gamma(\mathbf{x}; \xi_n) = (1 - \gamma) \psi(\mathbf{x}; \xi_n). \quad (24)$$

If $\psi_\gamma(\mathbf{x}; \xi_n) \leq 0$ for any \mathbf{x} and the equality is reached at the design points of ξ_n , then ξ_n is an optimal design.

Applying the first order algorithm to $I_\gamma(\xi)$, after only two iterations we have that $\xi_n^* = \delta_0$, as shown by Figure 5.2 which displays the function $\psi_\gamma(x; \xi_n^*)$. As proved in




Figure 1: Directional derivative $\psi_\gamma(x; \delta_0)$

[1], from the continuity of the KL-criterion ξ_n^* is a “nearly” KL-optimal design and, in particular, the smaller γ the better this design. Note that different regular designs $\tilde{\xi}$ lead to different expressions of $\psi_\gamma(\mathbf{x}; \xi_n^*)$ but to the same “nearly” KL-optimum design $\xi_n^* = \delta_0$.

5.3 Computational aspects

In this section we suggest some tricks to get numerically a KL-optimum design through the first order algorithm. We also compare these tricks to [13, pag. 109ff.], underlying possible improvements.

1. When ξ_n is singular, the directional derivative of the KL-criterion function at ξ_n in the direction of $\delta_{\mathbf{x}} - \xi_n$ cannot be expressed as in (11) with $\xi = \xi_n$ (for more details see [14, pag. 41]). Therefore, it cannot be used in the algorithm. On the other hand, the directional derivative of the regularized criterion (19) at ξ_n , in the direction of $\delta_{\mathbf{x}} - \xi_n$, is given by (24), whether or not ξ_n is regular. Therefore, before each step (a1), we suggest to decide between the original KL-criterion or the regularized one according to the following rule: use the regularized criterion if ξ_n has less than d_2 support points (see the discussion in Section 4.1); otherwise, use the original KL-criterion.
2. At step (a1) of the first order algorithm described in Section 4, in order to find $\beta_{2,n}$ a starting point $\beta_{2,n}^{(0)}$ is necessary. We suggest to use a point close to $\beta_{2,n-1}$, for instance $\beta_{2,n}^{(0)} = \beta_{2,n-1} + \varepsilon$, where ε is a random error with the magnitude of $\beta_{2,n-1}$. Sometimes, the reasonable choice of $\beta_{2,n}^{(0)} = \beta_{2,n-1}$ led the algorithm to a situation of stagnation around a local maximum, for this reason a random error is also considered.

3. As in [13], at step (a2), in order to choose a starting point $\mathbf{x}_n^{(0)}$ we suggest to take the maximum point in a regular grid or in a set of randomly selected starting points of \mathcal{X} , as \mathcal{X} is a compact space. In our benchmarks, we did not find significant differences between these two choices.
4. We suggest to check the terminal condition $U(\xi) > \delta$ after step (a2), where the point \mathbf{x}_n is computed, which gives the direction of the maximum increasing $\delta_{\mathbf{x}_n} - \xi_{n-1}$. Note that if ξ_n is a singular design, even though it has at least d_2 support points, then it is not computationally guaranteed that $\delta_{\mathbf{x}_n} - \xi_{n-1}$ is an increasing direction. In that case, the value α_n computed at step (b) is zero, and thus we suggest to regularize the criterion function, as described in Example 5.2.
5. At step (b) we construct a new design ξ_{n+1} adding the new point \mathbf{x}_n with weight α_n . Before going back to step (a1), we suggest
 - (i) as in [13], to collapse the points $\mathbf{x}_s \in B(\mathbf{x}_n, r_n)$ (where r_n is a radius tending to zero) in a new point $\bar{\mathbf{x}}_n$ with weight given by the sum of the weights of \mathbf{x}_s 's and \mathbf{x}_n . To speed up the algorithm we suggest to increase the weight of \mathbf{x}_n in the computation of the barycenter $\bar{\mathbf{x}}_n$, as n becomes larger (we found that a good choice is to increase the weight of \mathbf{x}_n by a factor of magnitude $n^{0.8}$, while the radius decreases as $n^{-0.65}$);
 - (ii) as in [13], to remove the support points with low weights (in absolute value as in [13], and also when the weight is low compared with the mean weight of the others).

6 Conclusions

In this paper we have provided some nice properties for the KL-optimality criterion, which are useful from a practical point of view. For instance, since the KL-optimum design is invariant to a scale-position transformation of the design region, if the experimental domain changes from \mathcal{X} to $\mathcal{Z} = \{\mathbf{z} : \mathbf{z} = \mathbf{a} + \mathbf{B}\mathbf{x} | \mathbf{x} \in \mathcal{X}\}$, in order to compute a new KL-optimum design on \mathcal{Z} it is enough to change in the same way the support points of the KL-optimum design on \mathcal{X} .

We have proved a crucial theoretical property, that is, the KL-criterion is continuous with respect to the design ξ . Continuity also guarantees the computational stability, since when designs are obtained in practice they are approximated at each step. Moreover it is a key property to prove the convergence of the first order algorithm.

We have showed that, at least in a large class of models, if the first order algorithm starts from an initial design with non-singular information matrix, then it moves on regular designs and thus it converges to the set of KL-optimum designs. If the KL-optimum design is singular then, in order to check its optimality, a regularized KL-criterion can be used: from the continuity of the KL-criterion, an optimum design for the regularized criterion is also nearly optimum for the KL-criterion.

Let us note that Kullback (see [17]) has already investigated some invariance properties of the Kullback-Leibler divergence and some relations with the Fisher information matrix. The KL-criterion, however, is the minimum Kullback-Leibler divergence between two parametric families. Therefore, the invariance property and the connection with the Fisher information matrix herein proved, are original results.

In this work, we also provide a list of useful advices to implement the first order algorithm, in order to get numerically a KL-optimum design. Some interesting examples show the application of the algorithm.

The KL-criterion depends on the unknown model parameters. For this reason, KL-optimum designs computed in this work are only locally optimal. In [10] it is showed that T-optimal designs (which are also KL-optimal) are sensitive with respect to parameter misspecification. To construct efficient discriminating designs which are instead robust to parameter misspecification, [26] propose a Bayesian version of the KL-criterion, while [20] follow a sequential approach. Another possibility to solve the problem of parameter dependence could be to use the following standardized maxmin KL-optimality criterion:

$$I_{2,1}^M(\xi) = \inf_{\beta_1 \in \Theta_1} \frac{I_{2,1}(\xi; \beta_1)}{I_{2,1}(\xi^*; \beta_1)}.$$

Another drawback of the KL-criterion is that the Kullback-Leibler divergence is not symmetric and thus, when the rival models are not nested, two different KL-criteria may be defined. This problem is solved in [26] by using a prior distribution for the models. Another solution could be to use the Jeffreys divergence instead of the Kullback-Leibler one, as done in [11].

As a matter of future work we intend to develop a new optimality criterion based on the Jeffreys divergence as well as to apply the standardized maxmin approach to compute robust discriminating design.

Acknowledgements. We are grateful to Professor Valerii Fedorov who suggested us some relevant references and to two anonymous referees for their useful remarks. We thank also Professor Giancarlo Manzi for helping us in improving the original draft of the paper.

References

- [1] G. Aletti, C. May and C. Tommasi. A convergent algorithm for finding KL-optimum designs and related properties. In: mODa 10—Advances in model-oriented design and analysis, Contrib. Statist., pp. . Physica-Verlag/Springer, Heidelberg, 2013.
- [2] A. C. Atkinson and D. R. Cox. Planning experiments for discriminating between models. *Journal of the Royal Statistical Society. Series B*, 36(3):321–348, 1974.
- [3] A. C. Atkinson and V. V. Fedorov. The design of experiments for discriminating between two rival models. *Biometrika*, 62:57–70, 1975.
- [4] A. C. Atkinson and V. V. Fedorov. Optimal design: experiments for discriminating between several models. *Biometrika*, 62(2):289–303, 1975.

- [5] J. M. Borwein and A. S. Lewis. *Convex analysis and nonlinear optimization*. CMS Books in Mathematics/Ouvrages de Mathématiques de la SMC, 3. Springer, New York, second edition, 2006. Theory and examples.
- [6] G. E. P. Box and W. J. Hill. Discrimination among mechanistic models. *Technometrics*, 9(1):57–71, 1967.
- [7] H. Dette. Discrimination designs for polynomial regression on a compact interval. *The Annals of Statistics*, 22:890–904, 1994.
- [8] H. Dette. Optimal designs for identifying the degree of a polynomial regression. *The Annals of Statistics*, 23(4):1248–1266, 1995.
- [9] H. Dette, V. B. Melas and P. Shpilev. T -optimal designs for discrimination between two polynomial models. *The Annals of Statistics*, 40:188–205, 2012.
- [10] H. Dette, V.B. Melas and P. Shpilev. Robust T -optimal discriminating designs. *The Annals of Statistics*, 41(4):1693–1715, 2013.
- [11] V. Dragalin, V. Fedorov, S. Patterson and B. Jones. Kullback-Leibler divergence for evaluating bioequivalence. *Statistics in Medicine*, 22:913-930, 2003.
- [12] R. M. Dudley. Convergence of Baire measures. *Studia Math.*, 27:251–268, 1966.
- [13] V. V. Fedorov. *Theory of Optimal Experiments*. Academic Press, New York and London, 1972. Probability and mathematical statistics.
- [14] V. V. Fedorov and P. Hackl. *Model-Oriented Design of Experiments*. Springer-Verlag, New York, 1997. Lecture Notes in Statistics.
- [15] A. L. Gibbs and F. E. Su. On choosing and bounding probability metrics. *International Statistical Review*, 70(3):419–435, 2002.
- [16] P. D. Hill. A review of experimental design procedures for regression model discrimination. *Technometrics*, 20(1):15–21, 1978.
- [17] S. Kullback. *Information theory and statistics*. Reprint of the second (1968) edition, Dover Publications Inc., 1997.
- [18] R. D. Meyer, D. M. Steinberg and G. Box. Follow-up designs to resolve confounding in multifactor experiments. *Technometrics*, 38(4):303–332, 1967.
- [19] J. López-Fidalgo, C. Tommasi, and P. C. Trandafir. An optimal experimental design criterion for discriminating between non-normal models. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 69(2):231–242, 2007.
- [20] C. May and C. Tommasi. Model selection and parameter estimation in non-linear nested models: a sequential generalized DKL-optimum design. *Statistica Sinica*, 2012.

- [21] E. Polak. *Optimization*, volume 124 of *Applied Mathematical Sciences*. Springer-Verlag, New York, 1997. Algorithms and consistent approximations.
- [22] S. D. Silvey. *Optimal design*. Chapman & Hall, London, 1980. An introduction to the theory for parameter estimation.
- [23] M. C. Spruill. Good designs for testing the degree of a polynomial mean. *Sankhyā: the indian journal of statistics, series B*, 52(1): 67–74, 1990.
- [24] C. Tommasi. Optimal designs for both model discrimination and parameter estimation. *Journal of Statistical Planning and Inference*, 139(12):4123–4132, 2009.
- [25] C. Tommasi. Optimal designs for discriminating among several non-normal models. In *mODa 8—Advances in model-oriented design and analysis*, Contrib. Statist., pages 213–220. Physica-Verlag/Springer, Heidelberg, 2007.
- [26] C. Tommasi and J. López-Fidalgo. Bayesian optimum designs for discriminating between models with any distribution. *Computational statistics and data analysis*, 54, 143–150, 2010.
- [27] D. Uciński and B. Bogacka. T -optimum designs for discrimination between two multiresponse dynamic models. *J. R. Stat. Soc. B* **67**, 3–18, 2005
- [28] V. S. Varadarajan. Measures on topological spaces. *Mat. Sb. (N.S.)*, 55 (97):35–100, 1961.
- [29] R. W. M. Wedderburn. On the existence and uniqueness of the maximum likelihood estimates for certain generalized linear models. *Biometrika*, 63: 27–32, 1976.
- [30] D. P. Wiens. Robust discrimination designs. *Journal of the Royal Statistical Society. Series B*, 71:805–829, 2009.