



HAL
open science

BAYESIAN QUADRATURE AND ENERGY MINIMIZATION FOR SPACE-FILLING DESIGN

Luc Pronzato, Anatoly A. Zhigljavsky

► **To cite this version:**

Luc Pronzato, Anatoly A. Zhigljavsky. BAYESIAN QUADRATURE AND ENERGY MINIMIZATION FOR SPACE-FILLING DESIGN. 2018. hal-01864076v1

HAL Id: hal-01864076

<https://hal.science/hal-01864076v1>

Preprint submitted on 31 Aug 2018 (v1), last revised 20 May 2020 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

BAYESIAN QUADRATURE AND ENERGY MINIMIZATION FOR SPACE-FILLING DESIGN

LUC PRONZATO[†] AND ANATOLY ZHIGLJAVSKY*

Abstract. A standard objective in computer experiments is to approximate the behaviour of an unknown function on a compact domain from a few evaluations inside the domain. When little is known about the function, space-filling design is advisable: typically, points of evaluation spread out across the available space are obtained by minimizing a geometrical (for instance, covering radius) or a discrepancy criterion measuring distance to uniformity. The paper investigates connections between design for integration (quadrature design), construction of the (continuous) BLUE for the location model, space-filling design, and minimization of energy (kernel discrepancy) for signed measures. Integrally strictly positive definite kernels define strictly convex energy functionals, with an equivalence between the notions of potential and directional derivative, showing the strong relation between discrepancy minimization and more traditional design of optimal experiments. In particular, kernel herding algorithms, which are special instances of vertex-direction methods used in optimal design, can be applied to the construction of point sequences with suitable space-filling properties.

Keywords: Bayesian quadrature, BLUE, energy minimization, potential, discrepancy, space-filling design

AMS subject classifications: 62K99, 65D30, 65D99.

1. Introduction. The design of computer experiments, where observations of a real physical phenomenon are replaced by simulations of a complex mathematical model (e.g., based on PDEs), has emerged as a full discipline, central to uncertainty quantification. The final objective of the simulations is often goal-oriented, that is, precisely defined. It may correspond for example to the optimization of the response of a system with respect to its input factors, or to the estimation of the probability that the response will exceed a given threshold when input factors have a given probability distribution. Achieving this objective generally requires sequential learning of the behavior of the response in a particular domain of interest for input factors: the region where the response is close to its optimum, or is close to the given threshold; see, e.g., the references in [27]. When simulations are computationally expensive, sequential inference based on the direct use of the mathematical model is unfeasible due to the large number of simulations required and simplified prediction models, approximating the simulated response, have to be used. A most popular approach relies on Gaussian process modelling, where the response (unknown prior to simulation) is considered as the realization of a Gaussian Random Field (RF), with parameterized mean and covariance, and Bayesian inference gives access to the posterior distribution of the RF (after simulation). Typically, in a goal-oriented approach based on stepwise-uncertainty reduction [6, 7], the prediction model is used to select the input factors to be used for the next simulation, the selection being optimal in terms of predicted uncertainty on the target. The construction of a first, possibly crude, prediction model is necessary to initialize the procedure. This amounts at approximating the behaviour of an unknown function f (the model response) on a compact domain $\mathcal{X} \subset \mathbb{R}^d$ (the feasible set for d input factors) from a few evaluations inside the domain. That is the basic design objective we shall keep in mind throughout the paper, although we may use diverted paths where approximation/prediction will be shadowed by other objectives, integration in particular.

*Cardiff University, UK

[†]CNRS, Université Côte d'Azur, I3S, France

[‡]Luc.Pronzato@cnrs.fr (corresponding author)

In general, little is known about the function *a priori*, and it seems intuitively reasonable to spread out points of evaluation across the available space; see [9]. Such space-filling designs can be obtained by optimizing a geometrical measure of dispersion or a discrepancy criterion measuring distance to uniformity. When using a Gaussian RF model, minimizing the Integrated Mean-Squared Prediction Error (IMSPE) is also a popular approach, although not very much used due to its apparent complexity, see, e.g., [25, 30]. The paper promotes the use of designs optimized for integration with respect to the uniform measure for their good space-filling properties. It gives a survey of recent results on energy functionals that measure distance to uniformity and places recent approaches proposed for space-filling design, such as [41], in a general framework and perspective encompassing design for integration, construction of the (continuous) Best Linear Unbiased Estimator (BLUE) in a location model with correlated errors, and minimization of energy (kernel discrepancy) for signed measures.

We start by a quick introduction to Bayesian function approximation and integration (Section 2), where the function is considered as the realization of a Gaussian RF with covariance structure defined by some kernel K . Exploiting recent results on the minimization of energy functionals [16, 61, 62], we show in Section 3 that integrally strictly positive definite kernels define strictly convex energy functionals, with an equivalence between the notions of potential and directional derivative that reveals the strong relation between discrepancy minimization and more traditional design of optimal experiments. We show that Bayesian integration is equivalent to the construction of the BLUE in a model with modified correlation structure, so that the two associated design problems coincide. We also show that the posterior variance in Bayesian integration corresponds to the minimum of a squared kernel discrepancy for signed measures with total mass one and to the minimum of an energy functional for a reduced kernel. Since the posterior variance criterion in Bayesian integration takes a very simple form, its minimization constitutes an attractive alternative to the minimization of the IMSPE criterion for space-filling design. This is considered in Section 4. We consider in particular kernel herding algorithms from machine learning, which are special instances of vertex-direction methods used in optimal design and can be used for the construction of point sequences with suitable space-filling properties (any-time designs). Several auxiliary results are given in appendix. Appendix A provides convergence properties of algorithms presented in Section 4. Extension to design for the simultaneous estimation of several integrals is considered in Appendix B. A Karhunen-Loève expansion of the RF model is considered in Appendix C, that yields a Bayesian linear model for which minimization of the posterior variance in Bayesian integration corresponds to a c -optimal design problem.

2. Random-field models for function approximation and integration.

2.1. Space-filling design and kernel choice for function approximation.

Let $K(\cdot, \cdot)$ denote a symmetric positive definite kernel on $\mathcal{X} \times \mathcal{X}$, with associated Reproducing Kernel Hilbert Space (RKHS) \mathcal{H}_K . Denote $K_{\mathbf{x}}(\cdot) = K(\mathbf{x}, \cdot)$ and $\langle \cdot, \cdot \rangle_K$ the scalar product in \mathcal{H}_K , so that the reproducing property gives $\langle f, K_{\mathbf{x}} \rangle_K = f(\mathbf{x})$ for any $f \in \mathcal{H}_K$.

Consider first the common framework where the function f to be approximated is supposed to belong to \mathcal{H}_K . Let $\eta_n(\mathbf{x}) = \sum_{i=1}^n w_i f(\mathbf{x}_i) = \mathbf{w}_n^T \mathbf{y}_n$ be a linear predictor of $f(\mathbf{x})$ based on evaluations of f at the n -point design $\mathbf{X}_n = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$, with $\mathbf{x}_i \in \mathcal{X}$ for all i . Throughout the paper we denote $\mathbf{w}_n = (w_1, \dots, w_n)^T$, $\mathbf{y}_n = [f(\mathbf{x}_1), \dots, f(\mathbf{x}_n)]^T$, $\mathbf{k}_n(\cdot) = [K_{\mathbf{x}_1}(\cdot), \dots, K_{\mathbf{x}_n}(\cdot)]^T$ and $\{\mathbf{K}_n\}_{i,j} = K(\mathbf{x}_i, \mathbf{x}_j)$, $i, j =$

1, \dots, n. Cauchy-Schwarz inequality gives the classical result

$$|f(\mathbf{x}) - \eta_n(\mathbf{x})| = \left| f(\mathbf{x}) - \sum_{i=1}^n w_i f(\mathbf{x}_i) \right| = \left| \langle f, K_{\mathbf{x}} - \sum_{i=1}^n w_i K_{\mathbf{x}_i} \rangle_K \right| \\ \leq \|f\|_{\mathcal{H}_K} \left\| K_{\mathbf{x}} - \sum_{i=1}^n w_i K_{\mathbf{x}_i} \right\|_{\mathcal{H}_K},$$

where $\|f\|_{\mathcal{H}_K}$ depends on f but not on \mathbf{X}_n , and $\rho_n(\mathbf{x}, \mathbf{w}) = \|K_{\mathbf{x}} - \sum_{i=1}^n w_i K_{\mathbf{x}_i}\|_{\mathcal{H}_K}$ depends on \mathbf{X}_n (and \mathbf{w}_n) but not on f . Suppose that \mathbf{K}_n has full rank. For a given \mathbf{X}_n , the Best Linear Predictor (BLP) minimizes $\rho_n(\mathbf{x}, \mathbf{w})$ and corresponds to $\eta_n^*(\mathbf{x}) = (\mathbf{w}_n^*)^T \mathbf{y}_n$, with $\mathbf{w}_n^* = \mathbf{w}_n^*(\mathbf{x}) = \mathbf{K}_n^{-1} \mathbf{k}_n(\mathbf{x})$, which gives $\rho_n^{*2}(\mathbf{x}) = \rho_n^2(\mathbf{x}, \mathbf{w}_n^*) = K(\mathbf{x}, \mathbf{x}) - \mathbf{k}_n^T(\mathbf{x}) \mathbf{K}_n^{-1} \mathbf{k}_n(\mathbf{x})$.

A less restrictive assumption on f is to suppose that it corresponds to a realization of a RF Z_x , with zero mean ($\mathbb{E}\{Z_x\} = 0$) and covariance $\mathbb{E}\{Z_x Z_{x'}\} = \sigma^2 K(\mathbf{x}, \mathbf{x}')$ for all \mathbf{x}, \mathbf{x}' in \mathcal{X} , $\sigma > 0$. Then, straightforward calculation shows that $\eta_n^*(\mathbf{x})$ is still the BLP (the posterior mean if Z_x is Gaussian), and $\sigma^2 \rho_n^{*2}(\mathbf{x})$ is the Mean-Squared Prediction Error (MSPE) at \mathbf{x} . This construction corresponds to simple kriging; see, e.g., [3, 67]. IMSPE-optimal designs minimize the integrated squared error $\text{IMSPE}(\mathbf{X}_n) = \sigma^2 \int_{\mathcal{X}} \rho_n^{*2}(\mathbf{x}) d\mu(\mathbf{x})$, with μ generally taken as the uniform probability measure on \mathcal{X} , see, e.g., [25, 30, 56].

IMSPE-optimal designs \mathbf{X}_n^* depend on the chosen K . It is well known that the asymptotic rate of decrease of $\text{IMSPE}(\mathbf{X}_n^*)$ with n depends on the regularity properties of K (the same is true for the integration problem); see for instance [55]. It is rather usual to take K stationary (translation invariant), i.e., satisfying $K(\mathbf{x}, \mathbf{x}') = \Psi(\|\mathbf{x} - \mathbf{x}'\|)$ for all \mathbf{x} and \mathbf{x}' , with Ψ in some parametric class selected according to prior knowledge on the smoothness properties of f . A typical example is the Matérn class of covariances, see [63, Chap. 2]. On the other hand, for reasons explained in Section 1, computer experiments usually involve small values of n , and the asymptotic behavior of the approximation error is hardly observed. Its behavior on a short horizon is much more important and strongly depends on the correlation lengths in K , which are difficult to choose *a priori*. Robustness with respect to the choice of K favours space-filling designs, where the \mathbf{x}_i are suitably spread over \mathcal{X} . Noticeably, it is shown in [58] that for translation invariant and isotropic kernels (i.e., such that $K(\mathbf{x}, \mathbf{x}') = \Psi(\|\mathbf{x} - \mathbf{x}'\|)$, with $\|\cdot\|$ the Euclidean distance in \mathbb{R}^d), one has $\rho_n^2(\mathbf{x}) \leq S_K[h_r(\mathbf{x})]$ for some increasing function $S_K(\cdot)$. Here $h_r(\mathbf{x}) = \max_{\|\mathbf{x} - \mathbf{x}'\| \leq r} \min_{1 \leq i \leq n} \|\mathbf{x}' - \mathbf{x}_i\|$ measures the density of design points \mathbf{x}_i around \mathbf{x} , with r a fixed positive constant. It satisfies $\max_{\mathbf{x} \in \mathcal{X}} h_r(\mathbf{x}) \geq \max_{\mathbf{x} \in \mathcal{X}} h_0(\mathbf{x}) = \text{CR}(\mathbf{X}_n)$, with

$$\text{CR}(\mathbf{X}_n) = \max_{\mathbf{x} \in \mathcal{X}} \min_{1 \leq i \leq n} \|\mathbf{x} - \mathbf{x}_i\|, \quad (2.1)$$

the *covering radius* of \mathbf{X}_n : $\text{CR}(\mathbf{X}_n)$ defines the smallest r such that the n closed balls of radius r centred at the \mathbf{x}_i cover \mathcal{X} . $\text{CR}(\mathbf{X}_n)$ is also called the dispersion of \mathbf{X}_n [46, Chap. 6] and corresponds to the minimax-distance criterion [36] used in space-filling design. Loosely speaking, the property $\rho_n^2(\mathbf{x}) \leq S_K[h(\mathbf{x})]$ quantifies the intuition that designs with a small value of CR provide precise predictions over \mathcal{X} since for any \mathbf{x} in \mathcal{X} there always exists a design point \mathbf{x}_i at proximity where $f(\mathbf{x}_i)$ has been evaluated. Another standard geometrical criterion of spreadness is the *packing radius*

$$\text{PR}(\mathbf{X}_n) = \frac{1}{2} \min_{i \neq j} \|\mathbf{x}_i - \mathbf{x}_j\|. \quad (2.2)$$

It corresponds to the largest r such that the n open balls of radius r centred at the \mathbf{x}_i do not intersect; $2\text{PR}(\cdot)$ corresponds to the maximin-distance criterion [36] often used in computer experiments.

In this paper, we shall adopt the following point of view. We do not intend to construct designs adapted to a particular K chosen from *a priori* knowledge on f . Neither shall we estimate the parameters in K (such as correlation lengths) when K is taken from a parametric class. We shall rather consider the kernel K as a tool for constructing a space-filling design, the quality of which will be measured in particular through the value of CR. The motivation is twofold: (i) the construction will be much easier than the direct minimization of CR, (ii) it will facilitate the construction of *sequences of points* suitably spread over \mathcal{X} (*any-time* space-filling designs).

2.2. Bayesian quadrature. Denote by $\mathcal{M} = \mathcal{M}[\mathcal{X}]$ the set of finite signed Borel measures on a nonempty set \mathcal{X} , and by $\mathcal{M}(q)$, $q \in \mathbb{R}$, the set of signed measures with total mass q : $\mathcal{M}(q) = \{\mu \in \mathcal{M} : \mu(\mathcal{X}) = q\}$. The set of Borel probability measures on \mathcal{X} is denoted by $\mathcal{M}^+(1)$, \mathcal{M}^+ is the set of finite positive measures on \mathcal{X} . Typical applications correspond to \mathcal{X} being a compact subset of \mathbb{R}^d for some $d \in \mathbb{N}$.

Suppose we wish to integrate a real function defined on \mathcal{X} with respect to $\mu \in \mathcal{M}^+(1)$. Assume that $\mathbb{E}_\mu\{|f(X)|\} < +\infty$ and denote

$$I_\mu(f) = \mathbb{E}_\mu\{f(\mathbf{X})\} = \int_{\mathcal{X}} f(\mathbf{x}) d\mu(\mathbf{x}).$$

We set a prior on f , and assume that f is the realization of a Gaussian RF, with covariance $\sigma^2 K(\cdot, \cdot)$, $\sigma^2 > 0$, and unknown mean β_0 ; that is, we consider the location model with correlated errors

$$f(\mathbf{x}) = \beta_0 + Z_x, \quad (2.3)$$

where $\mathbb{E}\{Z_x\} = 0$ and $\mathbb{E}\{Z_x Z_{x'}\} = \sigma^2 K(\mathbf{x}, \mathbf{x}')$ for all $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$, $\sigma > 0$. Regression models more general than (2.3) are considered in Appendix B. Here K is a symmetric Positive Definite (PD) kernel; that is, $K(\mathbf{x}, \mathbf{x}') = K(\mathbf{x}', \mathbf{x})$, and for all $n \in \mathbb{N}$ and all pairwise different $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathcal{X}$, the matrix \mathbf{K}_n is non-negative definite; if \mathbf{K}_n is positive definite, then K is called Strictly Positive Definite (SPD). Note that $K^2(\mathbf{x}, \mathbf{x}') \leq K(\mathbf{x}, \mathbf{x})K(\mathbf{x}', \mathbf{x}') < +\infty$ for all $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$ since K corresponds to a covariance. We will call a general kernel K bounded when $K(\mathbf{x}, \mathbf{x}) < \infty$ for all $\mathbf{x} \in \mathcal{X}$, and uniformly bounded when there is a constant C such that $K(\mathbf{x}, \mathbf{x}) \leq C$ for all $\mathbf{x} \in \mathcal{X}$. Any PD kernel is bounded.

Similarly to Section 2.1, we denote by \mathcal{H}_K the associated RKHS and by $\langle \cdot, \cdot \rangle_K$ the scalar product in \mathcal{H}_K . The assumption that K is bounded will be relaxed in Section 3.2 where we shall also consider singular kernels, but throughout the paper we assume that K is symmetric, $K(\mathbf{x}, \mathbf{x}') = K(\mathbf{x}', \mathbf{x})$ for all $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$. Also, we always assume, as in [23, Sect. 2.1], that either K is non-negative on $\mathcal{X} \times \mathcal{X}$, or \mathcal{X} is compact.

We set a vague prior on β_0 and assume that $\beta_0 \sim \mathcal{N}(\hat{\beta}_0^0, \sigma^2 A)$ with $A \rightarrow +\infty$. This amounts to setting $1/A = 0$ in all Bayesian calculations; the choice of $\hat{\beta}_0^0$ is then irrelevant. Suppose that f has been evaluated at the n -point design $\mathbf{X}_n = \{\mathbf{x}_1, \dots, \mathbf{x}_n\} \in \mathcal{X}^n$. We assume that \mathbf{K}_n has full rank. For any $\mathbf{x} \in \mathcal{X}$, the posterior distribution of $f(\mathbf{x})$ (conditional on σ^2 and K) is normal, with mean

$$\hat{\eta}_n(\mathbf{x}) = \hat{\beta}_0^n + \mathbf{k}_n^T(\mathbf{x})\mathbf{K}_n^{-1}(\mathbf{y}_n - \hat{\beta}_0^n \mathbf{1}_n)$$

and variance (mean-squared error)

$$\sigma^2 \rho_n^2(\mathbf{x}) = \sigma^2 \left[K(\mathbf{x}, \mathbf{x}) - \mathbf{k}_n^T(\mathbf{x}) \mathbf{K}_n^{-1} \mathbf{k}_n(\mathbf{x}) + \frac{(1 - \mathbf{k}_n^T(\mathbf{x}) \mathbf{K}_n^{-1} \mathbf{1}_n)^2}{\mathbf{1}_n^T \mathbf{K}_n^{-1} \mathbf{1}_n} \right], \quad (2.4)$$

where

$$\hat{\beta}_0^n = \frac{\mathbf{1}_n^T \mathbf{K}_n^{-1} \mathbf{y}_n}{\mathbf{1}_n^T \mathbf{K}_n^{-1} \mathbf{1}_n} \quad (2.5)$$

and $\mathbf{1}_n$ is the n -dimensional vector $(1, \dots, 1)^T$, see for instance [57, Chap. 4]. The posterior mean of $I_\mu(f)$ is thus

$$\hat{I}_n = \int_{\mathcal{X}} \hat{\eta}_n(\mathbf{x}) \, d\mu(\mathbf{x}) = \mathbb{E}_\mu \{ \hat{\eta}_n(\mathbf{X}) \} = \hat{\beta}_0^n + \mathbf{p}_n(\mu)^T \mathbf{K}_n^{-1} (\mathbf{y}_n - \hat{\beta}_0^n \mathbf{1}_n), \quad (2.6)$$

with

$$\mathbf{p}_n(\mu) = (P_\mu(\mathbf{x}_1), \dots, P_\mu(\mathbf{x}_n))^T, \quad (2.7)$$

where, for any $\nu \in \mathcal{M}$ and $\mathbf{x} \in \mathcal{X}$, we denote

$$P_\nu(\mathbf{x}) = \int_{\mathcal{X}} K(\mathbf{x}, \mathbf{x}') \, d\nu(\mathbf{x}'). \quad (2.8)$$

$P_\nu(\cdot)$ is called the kernel imbedding of ν into \mathcal{H}_K , see [61, Def. 9]; $P_\nu(\mathbf{x})$ is well defined and finite for any $\nu \in \mathcal{M}$ and $\mathbf{x} \in \mathcal{X}$ when K is uniformly bounded. On the other hand, there always exists $\nu \in \mathcal{M}$ such that $P_\nu(\mathbf{x})$ is infinite for all $\mathbf{x} \in \mathcal{X}$ when K is not uniformly bounded on \mathcal{X} . The function $P_\nu(\cdot)$ is called potential in potential theory, see Section 3.2.

Similarly to (2.4), we obtain that the posterior variance of $I_\mu(f)$ becomes

$$\sigma^2 s_n^2 = \sigma^2 \left[\mathcal{E}_K(\mu) - \mathbf{p}_n^T(\mu) \mathbf{K}_n^{-1} \mathbf{p}_n(\mu) + \frac{(1 - \mathbf{p}_n^T(\mu) \mathbf{K}_n^{-1} \mathbf{1}_n)^2}{\mathbf{1}_n^T \mathbf{K}_n^{-1} \mathbf{1}_n} \right], \quad (2.9)$$

where, for any $\nu \in \mathcal{M}$, we denote

$$\mathcal{E}_K(\nu) = \int_{\mathcal{X}^2} K(\mathbf{x}, \mathbf{x}') \, d\nu(\mathbf{x}) d\nu(\mathbf{x}'). \quad (2.10)$$

This is one of the key notions in potential theory, called the energy of ν ; see Section 3.2. For μ in $\mathcal{M}^+(1)$, we have $\mathcal{E}_K(\mu) = \mathbb{E}_\mu \{ K(\mathbf{X}, \mathbf{X}') \}$ where \mathbf{X} and \mathbf{X}' are independently identically distributed (i.i.d.) with μ . The quantity $-\mathcal{E}_K(\mu)$ corresponds to the quadratic entropy introduced by C.R. Rao [53]; see also Remark 3.1. Define

$$\mathcal{M}_K^\alpha = \left\{ \nu \in \mathcal{M} : \int_{\mathcal{X}} K^\alpha(\mathbf{x}, \mathbf{x}) \, d|\nu|(\mathbf{x}) < +\infty \right\}, \quad \alpha > 0. \quad (2.11)$$

When $\mu \in \mathcal{M}_K^{1/2}$, the reproducing property and Cauchy-Schwarz inequality imply that

$$\begin{aligned} \mathcal{E}_K(\mu) &= \int_{\mathcal{X}^2} \langle K(\cdot, \mathbf{x}), K(\cdot, \mathbf{x}') \rangle_K \, d\mu(\mathbf{x}) d\mu(\mathbf{x}') \\ &\leq \left[\int_{\mathcal{X}} K^{1/2}(\mathbf{x}, \mathbf{x}) \, d|\mu|(\mathbf{x}) \right]^2 < +\infty. \end{aligned} \quad (2.12)$$

When β_0 is assumed to be known (equal to zero for instance), we simply substitute β_0 for $\hat{\beta}_0^n$ in (2.6) and the posterior variance is

$$\sigma^2 s_{n,0}^2 = \sigma^2 [\mathcal{E}_K(\mu) - \mathbf{p}_n^T(\mu) \mathbf{K}_n^{-1} \mathbf{p}_n(\mu)] . \quad (2.13)$$

Bayesian quadrature relies on the estimation of $I_\mu(f)$ by \hat{I}_n . An optimal design for estimating $I_\mu(f)$ should minimize s_n^2 given by (2.9). One may refer to [18] for a historical perspective and to [33] for a recent exposition on Bayesian numerical computation. The framework presented above corresponds to that considered in [47], restricted to the case (recommended in that paper) where the known trend function is simply the constant 1 (which corresponds to the presence of an unknown mean β_0 in the model (2.3)). In Section 4, we shall see that $s_{n,0}^2$ is equal to the minimum value of a (squared) kernel discrepancy between the measure μ and a signed measure supported on \mathbf{X}_n , and that s_n^2 corresponds to the minimum of a squared discrepancy for signed measures that are constrained to have total mass one, and also corresponds to the minimum of an energy functional for a modified kernel K_μ . Note that $\sigma^2 s_n^2 \leq \text{IMSPE}(\mathbf{X}_n) = \sigma^2 \int_{\mathcal{X}} \rho_n^2(\mathbf{x}) d\mu(\mathbf{x})$ (which requires $\mu \in \mathcal{M}_K^1 \subset \mathcal{M}_K^{1/2}$ to be well defined). One of the key ideas of the paper is that space-filling design may be based on the minimization of s_n^2 rather than the minimization of $\text{IMSPE}(\mathbf{X}_n)$.

3. Kernel discrepancy, energy and potentials.

3.1. Maximum mean discrepancy. Suppose now that K is bounded and f belongs to the RKHS \mathcal{H}_K . Let μ and ν be two probability measures in $\mathcal{M}^+(1) \cap \mathcal{M}_K^{1/2}$. Since $f \in \mathcal{H}_K$, using the reproducing property, we obtain $I_\mu(f) = \int_{\mathcal{X}} \langle f, K_{\mathbf{x}} \rangle_K d\mu(\mathbf{x})$, $I_\nu(f) = \int_{\mathcal{X}} \langle f, K_{\mathbf{x}} \rangle_K d\nu(\mathbf{x})$ and

$$|I_\mu(f) - I_\nu(f)| = \left| \int_{\mathcal{X}} \langle f, K_{\mathbf{x}} \rangle_K d(\mu - \nu)(\mathbf{x}) \right| = |\langle f, P_\mu - P_\nu \rangle_K| ,$$

with $P_\mu(\cdot)$ and $P_\nu(\cdot)$ the kernel imbeddings (2.8). Define

$$\gamma_K(\mu, \nu) = \|P_\mu - P_\nu\|_{\mathcal{H}_K} . \quad (3.1)$$

Cauchy-Schwarz inequality yields the Koksma-Hlawka type inequality [46, Chap. 2] $|I_\mu(f) - I_\nu(f)| \leq \|f\|_{\mathcal{H}_K} \gamma_K(\mu, \nu)$, and

$$\gamma_K(\mu, \nu) = \sup_{\|f\|_{\mathcal{H}_K}=1} |I_\mu(f) - I_\nu(f)| , \quad (3.2)$$

see, e.g., [62, Th. 1]. Also, the expansion of $\|P_\mu - P_\nu\|_{\mathcal{H}_K}^2$ gives

$$\begin{aligned} \gamma_K(\mu, \nu) &= (\|P_\mu\|_{\mathcal{H}_K}^2 + \|P_\nu\|_{\mathcal{H}_K}^2 - 2\langle P_\mu, P_\nu \rangle_K)^{1/2} \\ &= \left(\int_{\mathcal{X}^2} K(\mathbf{x}, \mathbf{x}') d(\nu - \mu)(\mathbf{x}) d(\nu - \mu)(\mathbf{x}') \right)^{1/2} . \end{aligned} \quad (3.3)$$

Therefore, $\gamma_K(\cdot, \cdot)$ is at the same time a pseudometric between kernel imbeddings (3.1) and an integral pseudometric on probability distributions (3.2). It defines a kernel discrepancy between distributions (3.3), $\gamma_K(\cdot, \cdot)$ is also called the *Maximum Mean Discrepancy* (MMD) between μ and ν in $\mathcal{M}^+(1) \cap \mathcal{M}_K^{1/2}$, see [61, Def. 10].

To define a metric on the whole $\mathcal{M}^+(1)$, we need P_μ to be well defined and so that $P_\mu = P_\nu$ for μ and ν in $\mathcal{M}^+(1)$ implies $\mu = \nu$. This corresponds to the notion of *characteristic kernel*, see [62, Def. 6], which is closely connected to the following definitions.

Definition 3.1. A kernel K is Integrally Strictly Positive Definite (ISPD) on \mathcal{M} when $\mathcal{E}_K(\nu) > 0$ for any nonzero measure $\nu \in \mathcal{M}$.

Definition 3.2. A kernel K is Conditionally Integrally Strictly Positive Definite (CISPD) on \mathcal{M} when it is ISPD on $\mathcal{M}(0)$; that is, when $\mathcal{E}_K(\nu) > 0$ for all nonzero signed measures $\nu \in \mathcal{M}$ such that $\nu(\mathcal{X}) = 0$.

An ISPD kernel is CISPD. A bounded ISPD kernel is SPD and defines an RKHS. In [62, Lemma 8], the authors show that a uniformly bounded kernel is characteristic if and only if it is CISPD. The proof is a direct consequence of the expression (3.3) for the MMD $\gamma_K(\mu, \nu)$. They also give (Corollary 4) a spectral interpretation of $\gamma_K(\mu, \nu)$ and show that a translation-invariant kernel such that $K(\mathbf{x}, \mathbf{x}') = \Psi(\mathbf{x} - \mathbf{x}')$, with Ψ a uniformly bounded continuous real-valued positive-definite function, satisfies, for any μ and ν in $\mathcal{M}^+(1)$,

$$\gamma_K(\mu, \nu) = \left[\int_{\mathbb{R}^d} |\phi_\mu(\boldsymbol{\omega}) - \phi_\nu(\boldsymbol{\omega})|^2 d\Lambda(\boldsymbol{\omega}) \right]^{1/2}.$$

Here, ϕ_μ and ϕ_ν denote the characteristic functions of μ and ν respectively and Λ is the spectral Borel measure on \mathbb{R}^d , defined by

$$\Psi(\mathbf{x}) = \int_{\mathbb{R}^d} e^{-i\mathbf{x}^T \boldsymbol{\omega}} d\Lambda(\boldsymbol{\omega}). \quad (3.4)$$

Using this spectral representation, they prove (Th. 9) that K is characteristic if and only if the support of the associated Λ coincides with \mathbb{R}^d . For example, the sinc squared kernel $K(x, x') = \sin^2[\theta(x - x')]/(x - x')^2$, $\theta > 0$, is SPD but is not characteristic (and therefore not CISPD) since the support of Λ equals $[-2\theta, 2\theta]$. When $\gamma_K(\mu, \delta_{\mathbf{x}})$ is well defined for all $\mathbf{x} \in \mathcal{X}$, with $\delta_{\mathbf{x}}$ the Dirac delta measure at \mathbf{x} (and thus in particular when K is characteristic), we may consider the empirical measure $\xi_{n,e} = \xi_{n,e}(\mathbf{X}_n) = (1/n) \sum_{i=1}^n \delta_{\mathbf{x}_i}$ associated with a given design $\mathbf{X}_n = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$, and $\gamma_K(\mu, \xi_{n,e})$ of (3.2) gives the worst-case integration error for $\xi_{n,e}$ when f has norm one in \mathcal{H}_K ; see Section 4.1.1.

Typical examples of uniformly bounded ISPD, and therefore characteristic, kernels are the squared exponential kernel $K_t(\mathbf{x}, \mathbf{x}') = \exp(-t \|\mathbf{x} - \mathbf{x}'\|^2)$, $t > 0$, and the isotropic Matérn kernels, in particular

$$K_{3/2,\theta}(\mathbf{x}, \mathbf{x}') = (1 + \sqrt{3}\theta \|\mathbf{x} - \mathbf{x}'\|) \exp(-\sqrt{3}\theta \|\mathbf{x} - \mathbf{x}'\|) \quad (\text{Matérn } 3/2), \quad (3.5)$$

and $K_{5/2,\theta}(\mathbf{x}, \mathbf{x}') = [1 + \sqrt{5}\theta \|\mathbf{x} - \mathbf{x}'\| + 5\theta^2 \|\mathbf{x} - \mathbf{x}'\|^2/3] \exp(-\sqrt{5}\theta \|\mathbf{x} - \mathbf{x}'\|)$ (Matérn 5/2), see, e.g., [63]. Two other important examples are given hereafter.

Example 3.1 (Generalized multiquadric kernel). The sum of ISPD kernels is ISPD. Since the squared exponential kernel $K_t(\mathbf{x}, \mathbf{x}')$ is ISPD for any $t > 0$, the integrated kernel obtained by setting a probability distribution on t is ISPD too. One may thus consider $K(\mathbf{x}, \mathbf{x}') = \int K_t(\mathbf{x}, \mathbf{x}') d\pi(t)$ for π bounded and non decreasing on $[0, +\infty)$, which generates the class of continuous isotropic autocovariance functions in arbitrary dimension, see [60] and [63, p. 44]. In particular, for any $\epsilon > 0$ and $s > 0$, we obtain

$$K(\mathbf{x}, \mathbf{x}') = \int_0^{+\infty} K_t(\mathbf{x}, \mathbf{x}') t^{s/2-1} \exp(-\epsilon t) dt = \frac{\Gamma(s/2)}{(\|\mathbf{x} - \mathbf{x}'\|^2 + \epsilon)^{s/2}},$$

showing that the generalized multiquadric kernel

$$K_{s,\epsilon}(\mathbf{x}, \mathbf{x}') = (\|\mathbf{x} - \mathbf{x}'\|^2 + \epsilon)^{-s/2}, \quad \epsilon > 0, \quad s > 0, \quad (3.6)$$

is ISPD, see also [62, Sect. 3.2]. ◁

Example 3.2 (distance-induced kernels). Consider the kernels defined by

$$K^{(s)}(\mathbf{x}, \mathbf{x}') = -\|\mathbf{x} - \mathbf{x}'\|^s, \quad s > 0, \quad (3.7)$$

which are CISPD for $s \in (0, 2)$ [66], and the related distance-induced kernels

$$K'^{(s)}(\mathbf{x}, \mathbf{x}') = \|\mathbf{x}\|^s + \|\mathbf{x}'\|^s - \|\mathbf{x} - \mathbf{x}'\|^s, \quad s > 0.$$

Note that $\mathcal{E}_{K'^{(s)}}(\mu) = \mathcal{E}_{K^{(s)}}(\mu)$ when $\mu(\mathcal{X}) = 0$; in [66] $\mathcal{E}_{K'^{(s)}}$ is called energy distance for $s = 1$ and generalized energy distance for general $s \in (0, 2]$. For $s > 0$, the set $\mathcal{M}_{K'^{(s)}}$ contains all signed measures μ such that $\int_{\mathcal{X}} \|\mathbf{x} - \mathbf{x}_0\|^s d|\mu|(\mathbf{x}) < +\infty$ for some $\mathbf{x}_0 \in \mathcal{X}$. This result is a direct consequence of the triangular inequality when $s \in (0, 1]$; for $s > 1$ it follows from considerations involving semimetrics generated by kernels, see [61, Remark 21]. $K'^{(s)}$ is CISPD for $s \in (0, 2)$ ($K'^{(s)}/2$ corresponds to the covariance function of the fractional Brownian motion), but is not SPD (one has in particular, $K'^{(s)}(\mathbf{0}, \mathbf{0}) = 0$); $K'^{(2)}$ is not CISPD since $\mathcal{E}_{K'^{(2)}}(\mu) = [\int_{\mathcal{X}} \mathbf{x}^T d\mu(\mathbf{x})][\int_{\mathcal{X}} \mathbf{x} d\mu(\mathbf{x})]$, $\mu \in \mathcal{M}$. $K(x, x') = 1 - K^{(1)}(x, x') = 1 - |x - x'|$ is ISPD for $\mathcal{X} = [0, 1]$. \triangleleft

3.2. Energy and potentials. In this section we extend the considerations of previous section to signed measures and kernels which may have singularity on the diagonal. Definitions 3.1 and 3.2 extend to singular kernels, with Riesz kernels as typical examples.

Example 3.3 (Riesz kernels). These fundamental kernels of potential theory are defined by

$$K_{(s)}(\mathbf{x}, \mathbf{x}') = \|\mathbf{x} - \mathbf{x}'\|^{-s}, \quad s > 0, \quad \text{and} \quad K_{(0)}(\mathbf{x}, \mathbf{x}') = -\log \|\mathbf{x} - \mathbf{x}'\|, \quad (3.8)$$

with $\mathbf{x}, \mathbf{x}' \in \mathcal{X} \subset \mathbb{R}^d$ and $\|\cdot\|$ the Euclidean norm. When $s \geq d$, $\mathcal{E}_{K_{(s)}}(\mu)$ is infinite for any nonzero signed measure, but for $s \in (0, d)$ $K_{(s)}$ is ISPD. Since the logarithmic kernel $K_{(0)}(\mathbf{x}, \mathbf{x}')$ has singularity at zero and tends to $-\infty$ when $\|\mathbf{x} - \mathbf{x}'\|$ tends to $+\infty$, it will only be considered for \mathcal{X} compact; $K_{(0)}$ is CISPD, see [38, p. 80]. \triangleleft

Consider again $\mathcal{E}_K(\mu)$ given by (2.10), with $\mu \in \mathcal{M}$. In potential theory, this quantity is called the energy of the signed measure μ for the kernel K . Denote

$$\mathcal{M}_K = \{\nu \in \mathcal{M} : |\mathcal{E}_K(\nu)| < +\infty\}.$$

In the following, we shall only consider kernels that are at least CISPD. When K is ISPD, $\mathcal{E}_K(\mu)$ is positive for any nonzero $\mu \in \mathcal{M}$, but when K is only CISPD, $\mathcal{E}_K(\mu)$ can be negative; this is the reason for the presence of absolute value in the definition of \mathcal{M}_K . Note that \mathcal{M}_K is the set of measures such that $\mathcal{E}_K(\mu^+)$, $\mathcal{E}_K(\mu^-)$ and $\mathcal{E}_K(\mu^+, \mu^-) = \int_{\mathcal{X}^2} K(\mathbf{x}, \mathbf{x}') d\mu^+(\mathbf{x})d\mu^-(\mathbf{x}')$ are all finite, with μ^+ and μ^- denoting the positive and negative parts of the Hahn-Jordan decomposition $\mu = \mu^+ - \mu^-$ of μ , see [23, Sect. 2.1]. Also note that when K is bounded and defines an RKHS, $\mathcal{M}_K^\alpha \subset \mathcal{M}_K$ for any $\alpha \geq 1/2$, see (2.11) and (2.12); when K is uniformly bounded, $\mathcal{M}_K = \mathcal{M}$.

For any $\mu \in \mathcal{M}_K$, $P_\mu(\mathbf{x})$ given by (2.8) is called the potential at \mathbf{x} associated with $\mathcal{E}_K(\mu)$. It is well-defined, with values in $\mathbb{R} \cup \{-\infty, +\infty\}$, when $P_{\mu^+}(\mathbf{x})$ and $P_{\mu^-}(\mathbf{x})$ are not both infinite. Also, $P_\mu(\mathbf{x})$ is finite for μ -almost any \mathbf{x} , even if K is singular, when $\mu \in \mathcal{M}^+(1) \cap \mathcal{M}_K^{1/2}$.

When K is ISPD, we can still define MMD through (3.3),

$$\gamma_K(\mu, \nu) = \mathcal{E}_K^{1/2}(\nu - \mu), \quad (3.9)$$

since $\mathcal{E}_K(\nu - \mu)$ is nonnegative whenever defined. The set \mathcal{M}_K forms a pre-Hilbert space, with scalar product the mutual energy $\mathcal{E}_K(\mu, \nu) = \int_{\mathcal{X}^2} K(\mathbf{x}, \mathbf{x}') d\mu(\mathbf{x})d\nu(\mathbf{x}')$ and norm $\mathcal{E}_K^{1/2}(\mu)$. Denote by \mathcal{P}_K the linear space of potential fields $P_\mu(\cdot)$, $\mu \in \mathcal{M}_K$; when K defines an RKHS \mathcal{H}_K , $\|P_\mu\|_{\mathcal{H}_K} = \mathcal{E}_K^{1/2}(\mu)$ so that $\mathcal{P}_K \subset \mathcal{H}_K$, and \mathcal{P}_K is dense in \mathcal{H}_K . For \mathcal{P}_K to contain all functions $K_{\mathbf{x}}(\cdot) = K(\cdot, \mathbf{x})$, $\mathbf{x} \in \mathcal{X}$, we need $\delta_{\mathbf{x}} \in \mathcal{M}_K$ for all \mathbf{x} , which requires $K(\mathbf{x}, \mathbf{x}) < \infty$ for all $\mathbf{x} \in \mathcal{X}$.

For $\mu, \nu \in \mathcal{M}_K$, $\mathcal{E}_K(\mu, \nu)$ defines a scalar product $\langle P_\mu, P_\nu \rangle_{\mathcal{P}_K}$ on \mathcal{P}_K , with $\gamma_K(\mu, \nu) = \|P_\mu - P_\nu\|_{\mathcal{P}_K}$. Similarly to Section 3.1, we obtain

$$\begin{aligned} \gamma_K(\mu, \nu) &= \sup_{\xi \in \mathcal{M}_K, \mathcal{E}_K(\xi)=1} \left| \int_{\mathcal{X}^2} K(\mathbf{x}, \mathbf{x}') d\xi(\mathbf{x})d(\mu - \nu)(\mathbf{x}') \right| \\ &= \sup_{\|h\|_{\mathcal{P}_K} \leq 1} |I_\mu(h) - I_\nu(h)|; \end{aligned} \quad (3.10)$$

that is, a result that extends (3.2) to general ISPD kernels. If K is only CISPD, we can also define $\gamma_K(\mu, \nu)$ in the same way when considering measures $\mu, \nu \in \mathcal{M}(1)$; we then define \mathcal{P}_K as the linear space of potentials fields $P_\mu(\cdot)$, $\mu \in \mathcal{M}_K \cap \mathcal{M}(0)$, and in (3.10) we restrict ξ to be in $\mathcal{M}(0)$.

When K is singular, there always exists ν in \mathcal{M}_K such that $P_\nu(\mathbf{x}_0) = +\infty$ for some \mathbf{x}_0 . Consider for example Riesz kernel $K_{(s)}(\mathbf{x}, \mathbf{x}')$ with $s \in (0, d)$; \mathcal{M}_K contains in particular all signed measures with compact support $\mathbb{S}(\mu)$ whose potential $P_\mu(\mathbf{x})$ is bounded on $\mathbb{S}(\mu)$, see [38, p. 81]. Take ν as the measure with density $c/\|\mathbf{x} - \mathbf{x}_0\|^{s-d}$ on \mathcal{X} , with $\mathbf{x}_0 \in \mathcal{X}$; we have $\mathcal{E}_{K_{(s)}}(\nu) < \infty$ for \mathcal{X} compact, but $P_\nu(\mathbf{x}_0) = +\infty$. As a consequence, as noted in [16], singular kernels have little interest for integration. Indeed, take $\mu, \nu \in \mathcal{M}_K$ and $h = P_\nu \in \mathcal{P}_K$, then $|I_\mu(h)| \leq \|h\|_{\mathcal{P}_K} \mathcal{E}_K^{1/2}(\mu) = \mathcal{E}_K^{1/2}(\nu) \mathcal{E}_K^{1/2}(\mu) < \infty$ whereas $|I_{\xi_n}(h)|$ may be infinite for some discrete approximation ξ_n of μ as h can be infinite at some points. Singular kernels may nevertheless be used for the construction of space-filling designs, see for instance the example in Section 4.3, and this is our motivation for considering them in the following.

The key difficulty with singular kernels is the fact that delta measures do not belong to \mathcal{M}_K . An expedient solution to circumvent the problem is replace a singular kernel with a bounded surrogate. For instance, in space-filling design we may replace Riesz kernel $K_{(s)}$, $s > 0$, by a generalized inverse multiquadric kernel $K_{s,\epsilon}$ given by (3.6), and consider the limiting behaviour of the designs obtained when $\epsilon \rightarrow 0$, see Section 4.1.1.

3.3. Minimum energy and equilibrium measures. In this section, we show that there exist strong connections between results in potential theory and optimal design theory, where one minimizes a convex functional of $\mu \in \mathcal{M}^+(1)$, with the particularity that here the functional is quadratic.

3.3.1. ISPD kernels and convexity of $\mathcal{E}_K(\cdot)$.

Lemma 3.1. *K is ISPD if and only if \mathcal{M}_K is convex and $\mathcal{E}_K(\cdot)$ is strictly convex on \mathcal{M}_K .*

Proof. For any K , any μ and ν in \mathcal{M}_K and any $\alpha \in [0, 1]$, direct calculation gives

$$(1 - \alpha) \mathcal{E}_K(\mu) + \alpha \mathcal{E}_K(\nu) - \mathcal{E}_K[(1 - \alpha)\mu + \alpha\nu] = \alpha(1 - \alpha) \mathcal{E}_K(\nu - \mu). \quad (3.11)$$

Assume that K is ISPD. For any μ and ν in \mathcal{M}_K , the mutual energy $\mathcal{E}_K(\mu, \nu)$ satisfies $|\mathcal{E}_K(\mu, \nu)| \leq \sqrt{\mathcal{E}_K(\mu)\mathcal{E}_K(\nu)} < +\infty$. Therefore, $\mathcal{E}_K(\mu - \nu) = \mathcal{E}_K(\mu) + \mathcal{E}_K(\nu) - 2\mathcal{E}_K(\mu, \nu)$ is finite and (3.11) implies that $\mathcal{E}_K[(1 - \alpha)\mu + \alpha\nu]$ is finite, showing that \mathcal{M}_K is convex. Since K is ISPD, $\mathcal{E}_K(\nu - \mu) > 0$ for $\mu, \nu \in \mathcal{M}$, $\nu \neq \mu$, and (3.11) implies that $\mathcal{E}_K(\cdot)$ is strictly convex on \mathcal{M}_K .

Conversely, assume that \mathcal{M}_K is convex and $\mathcal{E}_K(\cdot)$ is strictly convex on \mathcal{M}_K . Any $\xi \in \mathcal{M}_K$ can be written as $\xi = \nu - \mu$ with, for instance, $\nu = 2\xi$ and $\mu = \xi$, both in \mathcal{M}_K . If $\mathcal{E}_K(\cdot)$ is strictly convex on \mathcal{M}_K , (3.11) with $\alpha \in (0, 1)$ implies that $\mathcal{E}_K(\xi) > 0$ when $\nu \neq \mu$, that is, when $\xi \neq 0$. Therefore, K is ISPD. ■

Lemma 3.1 also applies to singular kernels. The lemma below concerns CISPD kernels, which are assumed to be uniformly bounded.

Lemma 3.2. *Assume that K is uniformly bounded. Then, K is CISPD if and only if $\mathcal{E}_K(\cdot)$ is strictly convex on $\mathcal{M}(1)$.*

Proof. Since K is uniformly bounded, $\mathcal{M}_K = \mathcal{M}$. Assume that K is CISPD. Then, $\mathcal{E}_K(\nu - \mu) > 0$ for any $\mu \neq \nu \in \mathcal{M}(1)$, and (3.11) implies that $\mathcal{E}_K(\cdot)$ is strictly convex on $\mathcal{M}(1)$.

Assume now that $\mathcal{E}_K(\cdot)$ is strictly convex on $\mathcal{M}(1)$. Take any non-zero signed measure ξ in $\mathcal{M}(0)$ and consider the Hahn-Jordan decomposition $\xi = \xi^+ - \xi^-$, with $\xi^+(\mathcal{X}) = \xi^-(\mathcal{X}) = c > 0$. Denote $\nu = \xi^+/c$, $\mu = \xi^-/c$, with ν and μ in $\mathcal{M}^+(1)$ (ν and μ are in \mathcal{M}_K since K is uniformly bounded). Then, for any $\alpha \in (0, 1)$, (3.11) and the strict convexity of $\mathcal{E}_K(\cdot)$ on $\mathcal{M}(1)$ gives $\mathcal{E}_K(\xi) = c^2 \mathcal{E}_K(\nu - \mu) > 0$. ■

Note that one may replace $\mathcal{M}(1)$ by $\mathcal{M}^+(1)$, or by any $\mathcal{M}(\gamma)$ with $\gamma \neq 0$, in Lemma 3.2.

3.3.2. Minimum-energy probability measures. In the remaining part of Section 3.3, we assume that K is such that $\mathcal{E}_K(\cdot)$ is strictly convex on $\mathcal{M}^+(1) \cap \mathcal{M}_K$ and $\mathcal{M}(1) \cap \mathcal{M}_K$, which is true under the conditions of Lemma 3.1 or Lemma 3.2.

For $\mu, \nu \in \mathcal{M}_K$, denote by $F_K(\mu; \nu)$ the directional derivative of $\mathcal{E}_K(\cdot)$ at μ in the direction ν ,

$$F_K(\mu; \nu) = \lim_{\alpha \rightarrow 0^+} \frac{\mathcal{E}_K[(1 - \alpha)\mu + \alpha\nu] - \mathcal{E}_K(\mu)}{\alpha}.$$

Straightforward calculation gives

$$F_K(\mu; \nu) = 2 \left[\int_{\mathcal{X}^2} K(\mathbf{x}, \mathbf{x}') d\nu(\mathbf{x}) d\mu(\mathbf{x}') - \mathcal{E}_K(\mu) \right]. \quad (3.12)$$

In particular, for any $\mathbf{x} \in \mathcal{X}$, the potential $P_\mu(\mathbf{x})$ associated with μ at \mathbf{x} satisfies

$$P_\mu(\mathbf{x}) = \frac{1}{2} F_K(\mu; \delta_{\mathbf{x}}) + \mathcal{E}_K(\mu).$$

Remark 3.1 (Bregman divergence and Jensen difference). The strict convexity of $\mathcal{E}_K(\cdot)$ implies that $\mathcal{E}_K(\nu) \geq \mathcal{E}_K(\mu) + F_K(\mu, \nu)$ for any $\mu, \nu \in \mathcal{M}_K$, with equality if and only if $\nu = \mu$. This can be used to define a Bregman divergence between measures in \mathcal{M}_K (and thus between probability measures in $\mathcal{M}^+(1) \cap \mathcal{M}_K$), as

$$B_K(\mu, \nu) = \mathcal{E}_K(\nu) - [\mathcal{E}_K(\mu) + F_K(\mu, \nu)];$$

see [54]. Direct calculation gives $B_K(\mu, \nu) = \mathcal{E}_K(\nu - \mu)$ (with therefore $B_K(\mu, \nu) = B_K(\nu, \mu)$), providing another interpretation for the MMD $\gamma_K(\mu, \nu)$, see (3.9).

The squared MMD is also proportional to dissimilarity coefficient, or Jensen difference, $\Delta_J(\mu, \nu) = (1/2)[\mathcal{E}_K(\mu) + \mathcal{E}_K(\nu)] - \mathcal{E}_K[\mu + \nu]/2$ of [53]; indeed, direct calculation gives $\gamma_K^2(\mu, \nu) = \mathcal{E}_K(\nu - \mu) = 4\Delta_J(\mu, \nu)$. \triangleleft

Since $\mathcal{E}_K(\cdot)$ is strictly convex on $\mathcal{M}^+(1)$, there exists a unique minimum-energy probability measure. The measure $\mu_K^+ \in \mathcal{M}^+(1)$ is the minimum-energy measure if and only if $F_K(\mu_K^+; \nu) \geq 0$ for all $\nu \in \mathcal{M}^+(1)$, or equivalently, since ν is a probability measure, if and only if $F_K(\mu_K^+; \delta_{\mathbf{x}}) \geq 0$ for all $\mathbf{x} \in \mathcal{X}$. We thus obtain the following property, called equivalence theorem in the optimal-design literature.

Theorem 3.1. *When $\mathcal{E}_K(\cdot)$ is strictly convex on $\mathcal{M}^+(1) \cap \mathcal{M}_K$, $\mu_K^+ \in \mathcal{M}^+(1)$ is the minimum-energy probability measure on \mathcal{X} if and only if*

$$\forall \mathbf{x} \in \mathcal{X}, P_{\mu_K^+}(\mathbf{x}) \geq \mathcal{E}_K(\mu_K^+).$$

Note that, by construction, $\int_{\mathcal{X}} P_{\mu_K^+}(\mathbf{x}) d\mu_K^+(\mathbf{x}) = \mathcal{E}_K(\mu_K^+)$, implying $P_{\mu_K^+}(\mathbf{x}) = \mathcal{E}_K(\mu_K^+)$ on the support of μ_K^+ . The quantity $C_K^+ = [\inf_{\mu \in \mathcal{M}^+(1)} \mathcal{E}_K(\mu)]^{-1}$, with K an ISPD kernel, is called the *capacity* of \mathcal{X} in potential theory; note that $C_K^+ \geq 0$. The minimizing measure $\mu_K^+ \in \mathcal{M}^+(1)$ is called the *equilibrium measure* of \mathcal{X} (μ_K^+ is sometimes renormalized into $C_K^+ \mu_K^+$, see [38, p. 138]). Theorem 3.1 thus gives a necessary and sufficient condition for a probability measure μ to be the equilibrium measure of \mathcal{X} .

Example 3.4 (Continuation of Example 3.2). Properties of minimum-energy probability measures $\mu^+ = \mu_{K^{(s)}}^+$ for $K^{(s)}$ given by (3.7) with \mathcal{X} a compact subset of \mathbb{R}^d , $d \geq 2$, are investigated in [10] and [52]. The mass of μ^+ is concentrated on the boundary of \mathcal{X} , and its support only comprises extreme points of the convex hull of \mathcal{X} when $s > 1$; for $0 < s < 2$, μ^+ is unique; it is supported on no more than $d + 1$ points when $s > 2$.

Take $\mathcal{X} = \mathcal{B}_d(\mathbf{0}, 1)$. For symmetry reasons, μ^+ for $0 < s < 2$ is uniform on the unit sphere $\mathcal{S}_d(\mathbf{0}, 1)$ and

$$\mathcal{E}_{K^{(s)}}(\mu^+) = - \int_{\mathcal{X}^2} \|\mathbf{x} - \mathbf{x}'\|^s d\mu^+(\mathbf{x})d\mu^+(\mathbf{x}') = - \int_{\mathcal{X}} \|\mathbf{x}_0 - \mathbf{x}'\|^s d\mu^+(\mathbf{x}'),$$

where $\mathbf{x}_0 = (1, 0, \dots, 0)^T$. Denote by $\psi_d(\cdot)$ the density of the first component $t = x'_1$ of $\mathbf{x}' = (x'_1, \dots, x'_d)^T$. We obtain $\psi_d(t) = (d-1)V_{d-1}(1-t^2)^{(d-3)/2}/(dV_d)$ and

$$\mathcal{E}_{K^{(s)}}(\mu^+) = - \int_{-1}^1 [(1-t)^2 + 1-t^2]^{s/2} \psi_d(t) dt = - \frac{2^{d-q-2} \Gamma(d/2) \Gamma[(d+s-1)/2]}{\sqrt{\pi} \Gamma(d+s/2-1)}.$$

In particular, $\mathcal{E}_{K^{(1)}}(\mu^+) = -4/\pi$ when $d = 2$ and is a decreasing function of d . When $s = 2$, the uniform distribution on the unit sphere is also optimal, and the minimum energy equals -2 for all $d \geq 1$, but μ^+ is not unique and the measure allocating equal weight $1/(d+1)$ at each of the $d+1$ vertices of a d regular simplex with vertices on the unit sphere is optimal too. \triangleleft

Example 3.5 (Continuation of Example 3.3). Consider Riesz kernels $K_{(s)}$, see (3.8), for $\mathcal{X} = \mathcal{B}_d(\mathbf{0}, 1)$, the closed unit ball in \mathbb{R}^d . When $s \geq d$, $\mathcal{E}_{K_{(s)}}(\nu)$ is infinite for any non-zero $\nu \in \mathcal{M}$, but for $0 < s < d$ there exists a minimum-energy probability measure $\mu^+ = \mu_{K_{(s)}}^+$. When $d > 2$ and $s \in (0, d-2]$, μ^+ is uniform on the unit sphere $\mathcal{S}_d(\mathbf{0}, 1)$ (the boundary of \mathcal{X}); the potential at all interior points satisfies

$P_{\mu^+}(\mathbf{x}) \geq \mathcal{E}_{K(s)}(\mu^+)$ with strict inequality when $s \in (0, d-2)$. When $s \in (d-2, d)$, μ^+ has a density $\varphi_s(\cdot)$ in $\mathcal{B}_d(\mathbf{0}, 1)$,

$$\varphi_s(\mathbf{x}) = \frac{\pi^{-d/2} \Gamma(1+s/2)}{\Gamma[1-(d-s)/2]} \frac{1}{(1-\|\mathbf{x}\|^2)^{(d-s)/2}},$$

and the potential $P_{\mu^+}(\cdot)$ is constant in $\mathcal{B}_d(\mathbf{0}, 1)$, see, e.g., [38, p. 163].

When $d \leq 2$ and $s = 0$, μ^+ has a density in $\mathcal{B}_2(\mathbf{0}, 1)$ and $P_{\mu^+}(\cdot) = \mathcal{E}_{K(0)}(\mu^+)$ in $\mathcal{B}_2(\mathbf{0}, 1)$. In particular, for $d = 1$, μ^+ has the arcsine density $1/(\pi\sqrt{1-x^2})$ in $[-1, 1]$ with potential $P_{\mu^+}(x) = \log(2) - \log(|x| + \sqrt{x^2 - 1})$, $x \in \mathbb{R}$ (and $P_{\mu^+}(x) = \log(2)$ for $x \in [-1, 1]$).

The energy $\mathcal{E}_{K(s)}$ is infinite for empirical measures associated with n -point designs \mathbf{X}_n . One may nevertheless consider the ‘‘physical’’ energy

$$\tilde{\mathcal{E}}_{K(s)}(\mathbf{X}_n) = [2/n(n-1)] \sum_{1 \leq i < j \leq n} \|\mathbf{x}_i - \mathbf{x}_j\|^{-s} \quad (3.13)$$

($\tilde{\mathcal{E}}_{K(s)}(\mathbf{X}_n) = -[2/n(n-1)] \sum_{1 \leq i < j \leq n} \log \|\mathbf{x}_i - \mathbf{x}_j\|$ when $s = 0$), which is finite provided that all \mathbf{x}_i are distinct, see [16]. An n -point set \mathbf{X}_n^* minimizing $\tilde{\mathcal{E}}_{K(s)}(\mathbf{X}_n)$ is called a set of Fekete points, and the limit $\lim_{n \rightarrow \infty} \tilde{\mathcal{E}}_{K(s)}^{-1}(\mathbf{X}_n^*)$ exists and is called the transfinite diameter of \mathcal{X} . A major result in potential theory, see, e.g., [32], is that the transfinite diameter coincides with the capacity $C_{K(s)}^+$ of \mathcal{X} . If $C_{K(s)}^+ > 0$, then $\mu_{K(s)}^+$ is the weak limit of a sequence of empirical probability measures associated with Fekete points in \mathbf{X}_n^* . In the example considered, $\tilde{\mathcal{E}}_{K(s)}(\mathbf{X}_n^*)$ tends to infinity when $s \geq d$, but any sequence of Fekete points is asymptotically uniformly distributed in \mathcal{X} ; $\tilde{\mathcal{E}}_{K(s)}(\mathbf{X}_n^*)$ grows like $n^{s/d-1}$ for $s > d$ (and like $\log n$ for $s = d$). \triangleleft

Remark 3.2 (Stein variational gradient descent and energy minimization). Variational inference using smooth transform based on kernelized Stein discrepancy provides a gradient descent method for the approximation of a target distribution; see [39] and the references therein. The fact that the construction does not require knowledge of the normalizing constant of the target distribution makes the method particularly attractive for approximating a posterior distribution in Bayesian inference. Direct calculation shows that when the kernel is translation invariant and the target distribution is uniform, then the method corresponds to steepest descent for the minimization of $\mathcal{E}_K(\xi_{n,e})$; that is, at iteration k each design point $\mathbf{x}_i^{(k)}$ is updated into

$$\mathbf{x}_i^{(k+1)} = \mathbf{x}_i^{(k)} + \gamma \sum_{i < j} \frac{\partial K(\mathbf{x}, \mathbf{x}_j^{(k)})}{\partial \mathbf{x}} \Big|_{\mathbf{x}=\mathbf{x}_i^{(k)}}$$

for some $\gamma > 0$. The construction of space-filling design through energy minimization has already been considered in the literature. For instance, it is suggested in [2] to construct designs in a compact subset \mathcal{X} of \mathbb{R}^d by minimizing $\tilde{\mathcal{E}}_{K(2)}(\mathbf{X}_n)$ given by (3.13) (note that for $d \geq 3$ design points constructed in this way are not asymptotically uniformly distributed in \mathcal{X}). This approach tends to push points to the border of \mathcal{X} , similarly to the maximization of the packing radius $\text{PR}(\mathbf{X}_n)$ defined by (2.2). This is generally not desirable, especially when d is large. \triangleleft

3.3.3. Minimum-energy signed measures. The situation is slightly different from that in previous section when we consider measures in $\mathcal{M}(1)$. In that case, μ_K^* is

the minimum-energy measure in $\mathcal{M}(1)$ if and only if $F_K(\mu_K^*; \nu) = 0$ for all $\nu \in \mathcal{M}(1)$, this condition being equivalent to $F_K(\mu_K^*; \delta_{\mathbf{x}}) = 0$ for all $\mathbf{x} \in \mathcal{X}$. We thus obtain the following property.

Theorem 3.2. *When $\mathcal{E}_K(\cdot)$ is strictly convex on $\mathcal{M}(1) \cap \mathcal{M}_K$, $\mu_K^* \in \mathcal{M}(1)$ is the minimum-energy signed measure with total mass one on \mathcal{X} if and only if*

$$\forall \mathbf{x} \in \mathcal{X}, P_{\mu_K^*}(\mathbf{x}) = \mathcal{E}_K(\mu_K^*). \quad (3.14)$$

If we define now a *signed equilibrium measure* on \mathcal{X} as a measure $\mu \in \mathcal{M}(1)$ such that $P_\mu(\mathbf{x})$ is constant on \mathcal{X} , from the definition of $P_\mu(\cdot)$, when such a measure exists it necessarily satisfies the condition of Theorem 3.2 and therefore coincides with μ_K^* . Similarly to the case where one considers probability measures in $\mathcal{M}^+(1)$, we can define the (generalized) capacity of \mathcal{X} for measures in $\mathcal{M}(1)$ as $C_K^* = [\inf_{\mu \in \mathcal{M}(1)} \mathcal{E}_K(\mu)]^{-1}$, with $C_K^* = 1/\mathcal{E}_K(\mu_K^*)$ when μ_K^* exists, see [16, p. 824] (note that C_K^* may be negative). However, μ_K^* may not exist. Notice in particular that $\mathcal{M}(1)$ is not vaguely compact, contrarily to $\mathcal{M}^+(1)$ (and for Riesz kernels (3.8) with $s < d - 1$, $\mathcal{M}_{K(s)}$ is not complete contrarily to $\mathcal{M}_{K(s)} \cap \mathcal{M}^+$ [38, Th. 1.19]).

Example 3.6 (Continuation of Examples 3.2 and 3.4). Take $K(x, x') = K^{(s)}(x, x') = -|x - x'|^s$ on $\mathcal{X} = [0, 1]$, $s \in (0, 2)$, see (3.7). K is CISPD, and there exists a unique minimum-energy probability measure $\mu^+ = \mu_{K(s)}^+$ in $\mathcal{M}^+(1)$. On the other hand, below we show that minimum-energy signed measures in $\mathcal{M}(1)$ do not belong to $\mathcal{M}^+(1)$ when $s \in (1, 2)$ and that there is no minimum-energy signed measure in $\mathcal{M}(1)$ when $s \geq 2$.

When $s \in (0, 1)$, μ^+ has a density $\varphi^{(s)}(\cdot)$ with respect to the Lebesgue measure on $[0, 1]$,

$$\varphi^{(s)}(x) = \frac{\Gamma[1 - s/2]}{2^s \sqrt{\pi} \Gamma[(1 - s)/2]} \frac{1}{[x(1 - x)]^{(1+s)/2}},$$

and $P_{\mu^+}(x) = \mathcal{E}(\mu^+) = -\sqrt{\pi} \Gamma(1 - s/2) / \{2^s \Gamma[(1 - s)/2] \cos(\pi s/2)\}$ for all $x \in \mathcal{X}$ (and $\mathcal{E}(\mu^+) \rightarrow -1/2$ as $s \rightarrow 1^-$). The fact that $P_{\mu^+}(x) = \mathcal{E}(\mu^+)$ for all $x \in \mathcal{X}$ indicates that μ^+ is the minimum-energy signed measure with total mass one when $s \in (0, 1]$.

When $s \in [1, 2)$, $\mu^+ = (\delta_0 + \delta_1)/2$; the associated potential is $P_{\mu^+}(x) = -(|x|^s + |1 - x|^s)/2 \geq \mathcal{E}(\mu^+) = -1/2$, $x \in \mathcal{X}$ (note that $P_{\mu^+}(x) = -1/2$ for all $x \in \mathcal{X}$ when $s = 1$).

Consider now the signed measure $\mu_w = [(1 + w)/2](\delta_0 + \delta_1) - w\delta_{1/2}$, $w > 0$, so that $\mu_w(\mathcal{X}) = 1$ (i.e., $\mu_w \in \mathcal{M}(1)$). Direct calculation gives $\mathcal{E}_{K(s)}(\mu_w) = -(1 + w)(1 + w - 2^{2-s}w)$, which is minimum for $w = w_*(s) = (1 - 2^{1-s})/(2^{2-s} - 1)$ when $s < 2$, with $\mathcal{E}_{K(s)}(\mu_{w_*(s)}) = 2(1 - 2^{2-s})/(4 - 2^s)^2$. For $s \in (1, 2)$ we get $\mathcal{E}_{K(s)}(\mu_{w_*(s)}) < \mathcal{E}(\mu^+) = -1/2$, and there exist signed measures in $\mathcal{M}(1)$ such that $\mathcal{E}_{K(s)}(\mu_w) < \mathcal{E}(\mu^+)$. Therefore, minimum-energy signed measures with total mass one are not probability measures. For $s \geq 2$, $\lim_{w \rightarrow +\infty} \mathcal{E}_{K(s)}(\mu_w) = -\infty$, and there is no minimum-energy signed measure; in particular, $\mathcal{E}_{K(s)}(\mu_w) = -(w + 1)/2$ for $s = 2$. \triangleleft

Example 3.7 (Continuation of Examples 3.3 and 3.5). Consider Riesz kernels $K_{(s)}$, see (3.8), for $\mathcal{X} = \mathcal{B}_d(\mathbf{0}, 1)$, $d > 2$ and $s \in (0, d - 2)$; the minimum-energy probability measure μ^+ is then uniform on the unit sphere $\mathcal{S}_d(\mathbf{0}, 1)$ and the potential at all interior points satisfies $P_{\mu^+}(\mathbf{x}) > \mathcal{E}_{K(s)}(\mu^+)$. Consider the signed measure $\mu_w = (1 + w)\mu^+ - w\mu^{(r)}$, with $\mu^{(r)}$ uniform on the sphere $\mathcal{S}_d(\mathbf{0}, r)$ with radius $r \in (0, 1)$. Calculations

similar to those in the proof of [38, Th. 1.32] show that $\mathcal{E}_{K(s)}(\mu_w) < \mathcal{E}_{K(s)}(\mu^+)$ for w small enough, indicating that μ^+ is not the minimum-energy signed measure with total mass one. \triangleleft

3.3.4. When minimum-energy signed measures are probability measures.

Theorem 3.3. *Assume that K is ISPD and translation invariant, with $K(\mathbf{x}, \mathbf{x}') = \Psi(\mathbf{x} - \mathbf{x}')$ and Ψ continuous, twice differentiable except at the origin, with $\Psi(\mathbf{0}) < \infty$ and Laplacian $\Delta_\Psi(\mathbf{x}) = \sum_{i=1}^d \partial^2 \Psi(\mathbf{x}) / \partial x_i^2 \geq 0$, $\mathbf{x} \neq \mathbf{0}$. Then there exists a unique minimum-energy signed measure μ_K^* in $\mathcal{M}(1)$, and μ_K^* is a probability measure.*

Proof. The conditions of Theorem 3.1 are satisfied, and there exists a unique minimum-energy probability measure μ^+ such that $P_{\mu^+}(\mathbf{x}) \geq \mathcal{E}_K(\mu^+)$ for all $\mathbf{x} \in \mathcal{X}$. It also satisfies $P_{\mu^+}(\mathbf{x}) = \mathcal{E}_K(\mu^+)$ on the support of μ^+ . On the other hand, the conditions on K imply that for any μ in $\mathcal{M}^+(1)$, $P_\mu(\cdot)$ is subharmonic outside the support of μ , see, e.g., [38, Sect. I.2]. The first maximum principle of potential theory thus holds [38, Th. 1.10]: $P_\mu(\mathbf{x}) \leq c$ on the support of μ implies $P_\mu(\mathbf{x}) \leq c$ everywhere. Applying this to μ^+ , we obtain that $P_{\mu^+}(\mathbf{x}) \leq \mathcal{E}_K(\mu^+)$ everywhere; therefore, $P_{\mu^+}(\mathbf{x}) = \mathcal{E}_K(\mu^+)$ for all $\mathbf{x} \in \mathcal{X}$. Theorem 3.2 implies that μ^+ is the minimum-energy signed measure with total mass one. \blacksquare

The central argument for the proof of the property above is that $P_\mu(\cdot)$ is subharmonic outside the support of μ for any probability measure μ with finite energy. Weaker conditions than those in the theorem may be sufficient in particular situations, such as $\Psi(\mathbf{x} - \mathbf{x}') = \psi(\|\mathbf{x} - \mathbf{x}'\|)$ with $\psi(\cdot)$ convex on $(0, \infty)$, which generalizes a result of Hájek (1956), see also Section 3.4.1.

Another generalization is to consider CISP kernels. For example, for the kernels $K^{(s)}$ of (3.7), we have $\Delta(-\|\mathbf{x}\|^s) = s(2 - s - d)/\|\mathbf{x}\|^{2-s}$, $\mathbf{x} \neq \mathbf{0}$. Potentials are superharmonic for $d \geq 2$. When $d = 1$, they are superharmonic for $s \in [1, 2)$; they are subharmonic and satisfy the maximum principle for $s \in (0, 1)$, see Example 3.6.

Extension of Theorem 3.3 to singular kernels requires advanced results from potential theory; see especially [24, 38]. In particular, for the Riesz kernels $K_{(s)}$ of (3.8), we have $\Delta(\|\mathbf{x}\|^{-s}) = s(s + 2 - d)/\|\mathbf{x}\|^{s+2}$, $\mathbf{x} \neq \mathbf{0}$. When $d > 2$ and $s \in (0, d - 2)$, P_μ can be proved to be superharmonic in \mathbb{R}^d , and when $s \in [d - 2, d)$, P_μ can be proved to be subharmonic outside the support of μ , μ^+ being then the minimum-energy signed measure. This is also true for the logarithmic kernel for $d \leq 2$, with $\Delta(-\log \|\mathbf{x}\|) = (2 - d)/\|\mathbf{x}\|^2$, $\mathbf{x} \neq \mathbf{0}$. Examples 3.5 and 3.7 give an illustration.

3.4. Best Linear Unbiased Estimator (BLUE) of β_0 .

3.4.1. Continuous BLUE. Consider again the situation of Section 2.2 where $\sigma^2 K$ corresponds to the covariance of a random field Z_x . Suppose that we may observe $f(\cdot)$ over \mathcal{X} in order to estimate β_0 in the regression (location) model with correlated errors (2.3). Any linear estimator of β_0 takes the general form

$$\hat{\beta}_0 = \hat{\beta}_0(\xi) = \int_{\mathcal{X}} f(\mathbf{x}) d\xi(\mathbf{x}) = I_\xi(f)$$

for some $\xi \in \mathcal{M}$, and $\hat{\beta}_0(\xi)$ is unbiased when $\xi \in \mathcal{M}(1)$. Its variance is

$$V_\xi = \mathbb{E}\{(\hat{\beta}_0(\xi) - \beta_0)^2\} = \sigma^2 \mathcal{E}_K(\xi);$$

see [45, Sect. 4.2]. The existence of a minimum-energy signed measure μ_K^* is then equivalent to the existence of the continuous BLUE $\hat{\beta}_0^*$ for β_0 , with $\hat{\beta}_0^* = \hat{\beta}_0(\mu_K^*)$;

the variance of $\hat{\beta}_0^*$ is proportional to the minimum energy $\mathcal{E}_K(\mu_K^*)$, and Theorem 3.2 corresponds to Grenander's theorem [31]. Also, from that theorem, the existence of μ_K^* is equivalent to the existence of an equilibrium measure that yields a constant potential on \mathcal{X} . It can be related to a property of the generalized capacity C_K^* , as shown in the following theorem.

Theorem 3.4. *When K is ISPD, the constant function $1_{\mathcal{X}}$ equal to 1 on \mathcal{X} belongs to the space \mathcal{P}_K of potential fields if and only if there exists a minimum-energy signed measure $\mu_K^* \in \mathcal{M}(1)$, with $\mathcal{E}_K(\mu_K^*) \neq 0$. Moreover, the generalized capacity C_K^* is finite and nonzero, and satisfies $\|1_{\mathcal{X}}\|_{\mathcal{P}_K}^2 = C_K^*$.*

Proof. Suppose that $1_{\mathcal{X}} \in \mathcal{P}_K$. There exists $\mu \in \mathcal{M}_K$ such that $P_\mu = 1_{\mathcal{X}}$; that is, $P_\mu(\mathbf{x}) = 1$ for all $\mathbf{x} \in \mathcal{X}$. The definition of P_μ yields $\mathcal{E}_K(\mu) = \mu(\mathcal{X})$, which is finite and strictly positive since K is ISPD and $\mu \neq 0$. Denote $\mu' = \mu/\mu(\mathcal{X}) \in \mathcal{M}(1)$. We obtain $P_{\mu'}(\mathbf{x}) = 1/\mu(\mathcal{X}) = \mathcal{E}_K(\mu') > 0$ for all $\mathbf{x} \in \mathcal{X}$. Theorem 3.2 implies that μ' is the minimum-energy measure μ_K^* . Also, $C_K^* = 1/\mathcal{E}_K(\mu') = \mu(\mathcal{X}) \neq 0$, with $\|1_{\mathcal{X}}\|_{\mathcal{P}_K}^2 = \mathcal{E}_K(\mu)$, see Section 3.2.

Suppose now that there exists a minimum-energy signed measure $\mu_K^* \in \mathcal{M}(1)$ with $\mathcal{E}_K(\mu_K^*) \neq 0$. Theorem 3.2 implies that $P_{\mu_K^*}(\mathbf{x}) = \mathcal{E}_K(\mu_K^*)$ for all $\mathbf{x} \in \mathcal{X}$. For $\mu = \mu_K^*/\mathcal{E}_K(\mu_K^*)$, we get $P_\mu(\mathbf{x}) = 1$ for all $\mathbf{x} \in \mathcal{X}$, and $\|1_{\mathcal{X}}\|_{\mathcal{P}_K}^2 = \mathcal{E}_K(\mu) = 1/\mathcal{E}_K(\mu_K^*)$. ■

Under the conditions of Theorem 3.3, the BLUE exists, $\hat{\beta}_0^* = \hat{\beta}_0(\mu_K^+)$, with μ_K^+ the minimum-energy probability measure, and its variance equals $\sigma^2 \mathcal{E}_K(\mu_K^+)$. This is also true when $K(\mathbf{x}, \mathbf{x}') = \psi(\|\mathbf{x} - \mathbf{x}'\|)$ with $\psi(\cdot)$ convex on $(0, \infty)$. For $d = 1$, this property was known to Hájek (1956), see [45, p. 56]. The existence of a minimum-energy signed measure is not guaranteed in other circumstances, in particular when $K(\mathbf{x}, \mathbf{x}') = \Psi(\mathbf{x} - \mathbf{x}')$ and Ψ is differentiable at 0; see Example 3.8 below.

3.4.2. Discrete BLUE. Consider the framework of Section 2.2, with the same notation, and suppose that the n design points \mathbf{x}_i in \mathbf{X}_n are fixed. Any linear estimator of β_0 in (2.3) has then the form $\tilde{\beta}_0^n = \mathbf{w}_n^T \mathbf{y}_n$, with $\mathbf{w}_n = (w_1, \dots, w_n)^T \in \mathbb{R}^n$. The unbiasedness constraint imposes $\mathbf{w}_n^T \mathbf{1}_n = 1$. The variance of $\tilde{\beta}_0^n$ equals $\sigma^2 \mathbf{w}_n^T \mathbf{K}_n \mathbf{w}_n$, and the BLUE corresponds to the estimator $\hat{\beta}_0^n$ given by (2.5) (we assume that \mathbf{K}_n is nonsingular). The minimum-energy signed measure in $\mathcal{M}(1)$ (here discrete) μ_K^* is defined by the weights $\mathbf{w}_n^* = \mathbf{K}_n^{-1} \mathbf{1}_n / (\mathbf{1}_n^T \mathbf{K}_n^{-1} \mathbf{1}_n)$ set on the points in \mathbf{X}_n ; its energy is $\mathcal{E}_K(\mu_K^*) = \mathbf{w}_n^{*T} \mathbf{K}_n \mathbf{w}_n^* = 1/(\mathbf{1}_n^T \mathbf{K}_n^{-1} \mathbf{1}_n)$ and the variance of the BLUE equals $\sigma^2 \mathcal{E}_K(\mu_K^*)$. Note that some components of \mathbf{w}_n^* may be negative and that the potential associated with the measure $\mu_K^*/\mathcal{E}_K(\mu_K^*)$ on $\mathcal{X} = \mathbf{X}_n$ gives the constant function $1_{\mathcal{X}} = \mathbf{1}_n$, see Theorem 3.4. The optimal design problem for the discrete BLUE thus corresponds to the determination of the n -point set maximizing $\mathbf{1}_n^T \mathbf{K}_n^{-1} \mathbf{1}_n$.

Example 3.8. Consider $K(x, x') = \exp(-\theta|x - x'|)$, $\theta > 0$, for $x, x' \in \mathcal{X} = [0, 1]$. K is ISPD and satisfies

$$1_{\mathcal{X}} = \frac{K(\cdot, 0) + K(\cdot, 1)}{2} + \frac{\theta}{2} \int_0^1 K(\cdot, x) dx,$$

so that $1_{\mathcal{X}} \in \mathcal{P}_K$, see [1]. The minimum-energy measure in $\mathcal{M}(1)$ is $\mu_K^* = (\delta_0 + \delta_1 + \theta \mu_L)/(\theta + 2)$, with μ_L the Lebesgue measure on \mathcal{X} , and $\mu_K^* \in \mathcal{M}^+(1)$. The BLUE of β_0 in (2.3) is $\hat{\beta}_0^* = \int_{\mathcal{X}} f(x) d\mu_K^*(x)$, its variance equals $\sigma^2 \mathcal{E}_K(\mu_K^*) = 2\sigma^2/(\theta + 2)$, see [45, p. 56]. Note that $K' = K - 2/(\theta + 2)$ is still positive definite, but $1_{\mathcal{X}} \notin \mathcal{H}_{K'}$ since $c^2 K' - 1$ is not positive definite for any $c \neq 0$, see, e.g., [8, p. 30], [48, p. 20].

Consider now the squared exponential kernel $K(x, x') = \exp(-\theta|x - x'|^2)$, $\theta > 0$. The constant $1_{\mathcal{X}}$ does not belong to \mathcal{H}_K [64] and the BLUE of β_0 in (2.3) is not defined for that kernel. On the other hand, the discrete BLUE (2.5) is well defined for any set of n distinct points x_i , $\hat{\beta}_0^n = \mathbf{w}_n^{*T} \mathbf{y}_n = \mathbf{1}_n^T \mathbf{K}_n^{-1} \mathbf{y}_n / (\mathbf{1}_n^T \mathbf{K}_n^{-1} \mathbf{1}_n)$. Suppose that the n points x_i are equally spaced in $\mathcal{X} = [0, 1]$. The process Z_x in (2.3) has mean square derivatives of all orders, and, roughly speaking, for large n the construction of the BLUE mimics the estimation of derivatives of f and the weights w_i^* strongly oscillate between large positive and negative values. Figure 3.1-Left shows the optimal weights $(w_i^*/|w_i^*|)(\log_{10}(\max\{|w_i^*|, 1\}))$, truncated to absolute values larger than 1 and in log scale, when $x_i = (i - 1)/(n - 1)$, $i = 1, \dots, n = 101$. In Figure 3.1-Right, the kernel is $K(x, x') = (1 + \sqrt{5}|x - x'| + 5|x - x'|^2/3) \exp(-\sqrt{5}|x - x'|)$ (Matérn 5/2), so that Z_x is twice mean-square differentiable; the construction of the BLUE mimics the estimation of the first and second order derivatives of f at 0 and 1, suggesting that $1_{\mathcal{X}} \notin \mathcal{H}_K$ in that case too; see [17] for more details. \triangleleft




Figure 3.1: BLUE weights $(w_i^*/|w_i^*|)(\log_{10}(\max\{|w_i^*|, 1\}))$ for $x_i = (i - 1)/(n - 1)$, $i = 1, \dots, n = 101$. Left: $K(x, x') = \exp(-|x - x'|^2)$, Right: $K(x, x') = (1 + \sqrt{5}|x - x'| + 5|x - x'|^2/3) \exp(-\sqrt{5}|x - x'|)$ (Matérn 5/2).

Although a minimum-energy signed measure may not exist, in the next section we shall see how, for any measure $\mu \in \mathcal{M}(1)$ and any CISPD kernel K , we can modify K in such a way that the minimum-energy signed measure for the modified kernel exists (and coincides with μ).

3.5. Equilibrium measure and kernel reduction. Minimum-energy signed measures, when they exist, satisfy the following property.

Lemma 3.3. *If K is CISPD and if a minimum-energy signed measure μ_K^* exists in $\mathcal{M}(1)$, we have $\mathcal{E}_K(\xi) = \mathcal{E}_K[\xi - \xi(\mathcal{X})\mu_K^*] + [\xi(\mathcal{X})]^2 \mathcal{E}_K(\mu_K^*)$, $\forall \xi \in \mathcal{M}_K$.*

Proof. For any $\xi \in \mathcal{M}_K$, direct calculation gives

$$\begin{aligned} \mathcal{E}_K[\xi - \xi(\mathcal{X})\mu_K^*] &= \mathcal{E}_K(\xi) + [\xi(\mathcal{X})]^2 \mathcal{E}_K(\mu_K^*) - 2\xi(\mathcal{X}) \int_{\mathcal{X}^2} K(\mathbf{x}, \mathbf{x}') d\mu_K^*(\mathbf{x}) d\xi(\mathbf{x}') \\ &= \mathcal{E}_K(\xi) - [\xi(\mathcal{X})]^2 \mathcal{E}_K(\mu_K^*), \end{aligned}$$

where the second equality follows from (3.14). \blacksquare

Under the conditions of Lemma 3.3, any $\xi \in \mathcal{M}(1)$ satisfies

$$\mathcal{E}_K(\xi) = \mathcal{E}_K(\xi - \mu_K^*) + \mathcal{E}_K(\mu_K^*),$$

where the first term on the right-hand side equals the squared MMD $\gamma_K^2(\xi, \mu_K^*)$, see (3.9), and the second term does not depend on ξ . Minimizing the energy $\mathcal{E}_K(\xi)$ is thus equivalent to minimizing the MMD $\gamma_K(\xi, \mu_K^*)$. However, (i) μ_K^* may not exist, (ii) in many situations we wish to select a measure ξ having small MMD $\gamma_K(\xi, \mu)$ for a given measure $\mu \in \mathcal{M}_K$. This is the case in particular when one aims at evaluating the integral of a function with respect to some $\mu \in \mathcal{M}^+(1)$ (Section 2.2), or when we want construct a space-filling design in \mathcal{X} , μ being then uniform.

3.5.1. Kernel reduction. Take any $\mu \in \mathcal{M}_K$ such that $\mu(\mathcal{X}) \neq 0$. Without any loss of generality, we assume $\mu \in \mathcal{M}(1)$. Following [16], we show how to modify the kernel K in such a way that minimizing the energy $\mathcal{E}_{K_\mu}(\xi)$, $\xi \in \mathcal{M}(1)$, for the new (reduced) kernel K_μ is equivalent to minimizing $\gamma_{K_\mu}(\xi, \mu)$.

Define

$$K_\mu(\mathbf{x}, \mathbf{x}') = K(\mathbf{x}, \mathbf{x}') - P_\mu(\mathbf{x}) - P_\mu(\mathbf{x}') + \mathcal{E}_K(\mu), \quad (3.15)$$

see [59]. One can readily check that the energy for this new reduced kernel K_μ satisfies $\mathcal{E}_{K_\mu}(\beta\mu) = 0$ for any real β and that the potential for μ associated with K_μ satisfies $\tilde{P}_\mu(\mathbf{x}) = \int_{\mathcal{X}} K_\mu(\mathbf{x}, \mathbf{x}') d\mu(\mathbf{x}') = 0$ for all \mathbf{x} .

Next theorem indicates that, for any given μ in $\mathcal{M}(1) \cap \mathcal{M}_K$, when considering signed measures ξ with total mass one, minimizing the energy $\mathcal{E}_{K_\mu}(\xi)$ is equivalent to minimizing the MMD $\gamma_K(\xi, \mu)$, provided that K is CISPD.

Theorem 3.5. *If K is CISPD, then for any $\mu \in \mathcal{M}(1) \cap \mathcal{M}_K$, we have*

- (i) *the reduced kernel K_μ defined by (3.15) is CISPD;*
- (ii) *μ is the minimum-energy measure in $\mathcal{M}(1)$ for K_μ , and*

$$\forall \xi \in \mathcal{M}_K, \quad \mathcal{E}_{K_\mu}(\xi) = \mathcal{E}_K[\xi - \xi(\mathcal{X})\mu] = \mathcal{E}_{K_\mu}[\xi - \xi(\mathcal{X})\mu].$$

Proof. For any nonzero $\xi \in \mathcal{M}_K$, direct calculation using (3.15) gives

$$\begin{aligned} \mathcal{E}_{K_\mu}(\xi) &= \mathcal{E}_K(\xi) - 2\xi(\mathcal{X}) \int_{\mathcal{X}^2} K(\mathbf{x}, \mathbf{x}') d\mu(\mathbf{x})d\xi(\mathbf{x}') + [\xi(\mathcal{X})]^2 \mathcal{E}_K(\mu) \\ &= \mathcal{E}_K[\xi - \xi(\mathcal{X})\mu]. \end{aligned} \quad (3.16)$$

(i) When $\xi(\mathcal{X}) = 0$ we get $\mathcal{E}_{K_\mu}(\xi) = \mathcal{E}_K(\xi)$ which is strictly positive when $\xi \neq 0$, showing that K_μ is CISPD. (ii) Since $[\xi - \xi(\mathcal{X})\mu](\mathcal{X}) = 0$ and K is CISPD, $\mathcal{E}_{K_\mu}(\xi) > 0$ for $\xi \neq \xi(\mathcal{X})\mu$, showing that μ is the (unique) minimum-energy signed measure in $\mathcal{M}(1)$ for K_μ . Since $\mathcal{E}_{K_\mu}(\mu) = 0$, Lemma 3.3 with K_μ substituted for K implies that $\mathcal{E}_{K_\mu}(\xi) = \mathcal{E}_{K_\mu}[\xi - \xi(\mathcal{X})\mu]$ for any $\xi \in \mathcal{M}_K$, which, together with (3.16), concludes the proof. ■

3.5.2. Kernel reduction, BLUE and Bayesian integration. Consider again the situation of Section 3.4, and define \mathcal{P}_1 as the orthogonal projection of $L^2(\mathcal{X}, \mu)$ onto the linear space spanned by the constant 1; see [26]. The model (2.3) can then be written as

$$f(\mathbf{x}) = \beta_0 + \mathcal{P}_1 Z_x + (\text{Id}_{L^2} - \mathcal{P}_1) Z_x = \beta'_0 + \tilde{Z}_x, \quad (3.17)$$

where $\beta'_0 = \beta_0 + \mathcal{P}_1 Z_x$ and $\tilde{Z}_x = (\text{Id}_{L^2} - \mathcal{P}_1) Z_x$, with \tilde{Z}_x having zero mean and covariance $\mathbb{E}\{\tilde{Z}_x \tilde{Z}_{x'}\} = \sigma^2 K_\mu(\mathbf{x}, \mathbf{x}')$. We have seen in Section 3.4 that the variance of the continuous BLUE of β_0 equals $\sigma^2 \mathcal{E}_K(\mu_K^*)$ provided that the minimum-energy signed measure μ_K^* exists. (Note that the prior on β'_0 remains non-informative when

the prior on β_0 is non-informative.) On the other hand, we obtain now that the continuous BLUE of β'_0 always exists: it coincides with $I_\mu(f)$ and its variance is $\sigma^2 \mathcal{E}_{K_\mu}(\mu) = 0$. Therefore, as mentioned in introduction, Bayesian integration for the model (2.3) with correlated errors is equivalent to parameter estimation in a location model with different correlation structure.

3.6. Tensor product kernels. From d kernels K_i respectively defined on $\mathcal{X}_i \times \mathcal{X}_i$, $i = 1, \dots, d$, we can construct a tensor product kernel as

$$K^\otimes(\mathbf{x}, \mathbf{x}') = \prod_{i=1}^d K_i(x_i, x'_i), \quad (3.18)$$

where $\mathbf{x} = (x_1, \dots, x_d)^T$ and $\mathbf{x}' = (x'_1, \dots, x'_d)^T$ belong to the product space $\mathcal{X} = \mathcal{X}_1 \times \dots \times \mathcal{X}_d$. The construction is particularly useful when considering product measures on \mathcal{X} , since, in some sense, it allows us to decompose an integration or space-filling design problem in a high dimensional space into its one-dimensional counterparts. Suppose that each K_i is uniformly bounded and CISPD on $\mathcal{M}^{(i)} = \mathcal{M}[\mathcal{X}_i]$; that is, K_i is ISPD on $\mathcal{M}^{(i)}(0)$, see Definitions 3.1 and 3.2. One can show that this is equivalent to K^\otimes being ISPD on $\otimes_{i=1}^d \mathcal{M}^{(i)}(0)$, see [65, Th. 2]. In the same paper, the authors prove (Th. 4) that if each K_i is moreover continuous and translation invariant, then K^\otimes is ISPD on $\mathcal{M}(0)$; that is, K^\otimes is CISPD on \mathcal{M} . Their proof relies on the equivalence between the CISPD and characteristic properties for uniformly bounded kernels, and on the characterization of characteristic continuous, uniformly bounded and translation invariant kernels through a property of the support of the measure Λ defined in (3.4); see Section 3.1.

An important property of tensor product kernels K^\otimes is that kernel reductions K_μ^\otimes , see (3.15), are easily obtained explicitly. Indeed, when $\mu = \otimes_{i=1}^d \mu^{(i)}$ is a product measure on \mathcal{X} , then, for all $\mathbf{x} \in \mathcal{X}$,

$$\mathcal{E}_{K^\otimes}(\mu) = \prod_{i=1}^d \mathcal{E}_{K_i}(\mu^{(i)}), \quad (3.19)$$

$$P_\mu(\mathbf{x}) = \prod_{i=1}^d \int_{\mathcal{X}_i} K_i(x_i, x'_i) d\mu^{(i)}(x'_i) = \prod_{i=1}^d P_{\mu^{(i)}}(x_i), \quad (3.20)$$

which facilitates the calculation of $\mathcal{E}_{K_\mu^\otimes}(\xi)$, in particular when ξ is a discrete measure as considered in Section 4. Table 3.1 gives the expressions of $\mathcal{E}_K(\mu)$ and $P_\mu(x)$ obtained for a few kernels, with μ uniform on $\mathcal{X} = [0, 1]$; the expressions for the squared exponential and Matérn kernels can be found in [28]. Note that in each case $\mathcal{E}_{K_\mu}(\xi) > 0$ for any $\xi \in \mathcal{M}(1)$, $\xi \neq \mu$.

Remark 3.3 (super- and subharmonicity for tensor product kernels). In complement of Theorem 3.3, we may notice that when each K_i in (3.18) is stationary and satisfies $K_i(x, x') = \Psi_i(x - x')$, then $K^\otimes(\mathbf{x}, \mathbf{x}') = \Psi(\mathbf{x} - \mathbf{x}')$ with $\Psi(\mathbf{z}) = \prod_{i=1}^d \Psi_i(z_i)$. Assume that all Ψ_i are twice continuously differentiable. Then, the Laplacian of Ψ is

$$\Delta_\Psi(\mathbf{x}) = \sum_{i=1}^d \left[\frac{\partial^2 \Psi_i(x_i)}{\partial x_i^2} \prod_{j \neq i} \Psi_j(x_j) \right]. \quad \triangleleft$$

4. Experimental design. Consider an n -point design $\mathbf{X}_n = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$, with $\mathbf{x}_i \in \mathcal{X}$ for all i . In this section, we shall restrict our attention to finite signed

Table 3.1: Energy $\mathcal{E}_K(\mu)$ and potential $P_\mu(x)$ for different kernels K with μ uniform on $\mathcal{X} = [0, 1]$; $P_\mu(x) = S_\mu(x) + S_\mu(1-x) + T_\mu(x)$; $S_\mu(\cdot)$ is continuously differentiable in $(0, 1]$, $T_\mu = 0$ when K is translation invariant.

$K(x, x')$	$\mathcal{E}_K(\mu)$	$S_\mu(x)$ [and $T_\mu(x)$]
$\exp(-\theta x-x')$	$2(\theta + e^{-\theta} - 1)/\theta^2$	$x(1 - e^{-\theta x })/(\theta x)$
$K_{3/2, \theta/\sqrt{3}}(x, x')$ in (3.5)	$2[\theta(2 + e^{-\theta}) + 3(e^{-\theta} - 1)]/\theta^2$	$x[2 - (2 + \theta x)e^{-\theta x }]/(\theta x)$
$[(x-x')^2 + \epsilon]^{-1}$ ($\epsilon \geq 0$)	$(2/\sqrt{\epsilon}) \arctan(1/\sqrt{\epsilon}) - \log(1 + 1/\epsilon)$	$(1/\sqrt{\epsilon}) \arctan(x/\sqrt{\epsilon})$
$(x-x' + \epsilon)^{-1}$ ($\epsilon > 0$)	$2[(1 + \epsilon) \log(1 + 1/\epsilon) - 1]$	$\text{sign}(x) \log(1 + x /\epsilon)$
$(x-x' + \epsilon)^{-1/2}$ ($\epsilon > 0$)	$4\epsilon^{3/2} [2(1 + 1/\epsilon)^{3/2} - 2 - 3/\epsilon]/3$	$2\sqrt{\epsilon} \text{sign}(x)(\sqrt{1 + x /\epsilon} - 1)$
$1 - \theta x-x' $ ($0 < \theta \leq 1$)	$1 - \theta/3$	$1/2 - \theta x /2$
$ x-x' ^{-s}$ ($0 < s < 1$)	$2/(s^2 - 3s + 2)$	$x/[(1-s) x ^s]$
$-\log x-x' $	$3/2$	$1/2 - x \log x $
$ x + x' - x-x' $	$2/3$	$1/4 - x x /2$ [$T_\mu(x) = x $]
$\sqrt{ x + x' } - \sqrt{ x-x' }$	$4/5$	$1/3 - 2x\sqrt{ x }/3$ [$T_\mu(x) = \sqrt{ x }$]

measures $\xi_n = \sum_{i=1}^n w_i \delta_{\mathbf{x}_i}$ supported on \mathbf{X}_n , and denote $\mathbf{w}_n = (w_1, \dots, w_n)^T$. As in Section 2.2, we consider a measure $\mu \in \mathcal{M}^+(1)$, with special attention to space-filling design for which μ is uniform on a compact subset \mathcal{X} of \mathbb{R}^d . We assume that K is SPD and μ has finite energy $\mathcal{E}_K(\mu)$, see (2.10). Direct calculation gives

$$\begin{aligned} \gamma_K^2(\xi_n, \mu) &= \mathcal{E}_K(\xi_n - \mu) = \mathbf{w}_n^T \mathbf{K}_n \mathbf{w}_n - 2\mathbf{w}_n^T \mathbf{p}_n(\mu) + \mathcal{E}_K(\mu), \\ &= \sum_{i,j} w_i w_j K(\mathbf{x}_i, \mathbf{x}_j) - 2 \sum_{i=1}^n w_i P_\mu(\mathbf{x}_i) + \mathcal{E}_K(\mu), \end{aligned} \quad (4.1)$$

where $\{\mathbf{K}_n\}_{i,j} = K(\mathbf{x}_i, \mathbf{x}_j)$, $i, j = 1, \dots, n$, and $\mathbf{p}_n(\mu)$ is given by (2.7). Note that $\mathcal{E}_K(\mu)$ and the $P_\mu(\mathbf{x}_i)$ have simple expressions when K is a tensor product kernel and $\mu = \otimes_{i=1}^d \mu^{(i)}$ is a product measure on $\mathcal{X} = \mathcal{X}_1 \times \dots \times \mathcal{X}_d$, see (3.19, 3.20). Monte-Carlo approximation, based on a large i.i.d. sample from μ , or a low-discrepancy sequence, can always be used instead.

4.1. One-shot designs.

4.1.1. Support of empirical measures. Denote by $\xi_{n,e} = \xi_{n,e}(\mathbf{X}_n)$ the empirical measure associated with a given design $\mathbf{X}_n = \{\mathbf{x}_1, \dots, \mathbf{x}_n\} \subset \mathbb{R}^{nd}$, $\xi_{n,e} = (1/n) \sum_{i=1}^n \delta_{\mathbf{x}_i}$. As indicated hereafter, the literature on space-filling design provides several examples of construction of n -point designs through the minimization of the squared MMD $\mathcal{E}_K(\xi_{n,e} - \mu)$ with respect to \mathbf{X}_n .

For $\mathcal{X} = [0, 1]^d$, tensorised kernels based on variants of Brownian motion covariance yield L_2 discrepancies (symmetric, centred, wrap-around and so on); see, e.g., [34], [20, Chap. 3]. For instance, for $\mathcal{X} = [0, 1]$ and $K(x, x') = 1 - |x - x'|$ (for which the expressions of $\mathcal{E}_K(\mu)$ and $P_\mu(x)$ are given in Table 3.1), $\mathcal{E}_{K_\mu}(\xi_{n,e})$ is twice the squared L_2 star discrepancy for $d = 1$.

The ISPD kernel $K_{s,\epsilon}^\otimes(\mathbf{x}, \mathbf{x}') = \prod_{i=1}^d K_{s,\epsilon}(x_i, x'_i)$, with $K_{s,\epsilon}$ given by (3.6) with $s > 0$ and $\epsilon > 0$, is called *projection kernel* in [40]. For very small ϵ , the minimization of $\mathcal{E}_{K_{1,\epsilon}^\otimes}(\xi_{n,e})$ corresponds to the construction of a maximum-projection design, as defined in [37]. Note that minimizing $\mathcal{E}_{K_{s,\epsilon}^\otimes}(\xi_{n,e})$ is not equivalent to minimizing $\mathcal{E}_{K_{s,\epsilon}^\otimes}(\xi_{n,e} - \mu)$: in particular, when μ is uniform on \mathcal{X} , which is assumed to be compact and convex, the former tends to push design points to the boundary of \mathcal{X} whereas the latter keeps all points in the interior of \mathcal{X} ; see [40].

In [41], space-filling designs in a compact set $\mathcal{X} \subset \mathbb{R}^d$ are constructed by minimizing $\mathcal{E}_{K^{(1)}}(\xi_{n,e} - \mu)$ for μ uniform on \mathcal{X} , see (3.7). They call *support points* the optimal support \mathbf{X}_n^* , which they determine via a majorization-minimization algorithm using the property that the problem can be formulated as a difference-of-convex optimization problem. Values of $\mathcal{E}_{K^{(1)}}(\mu)$ and $P_\mu(\mathbf{x})$ are not available even for $\mathcal{X} = [0, 1]^d$ and Monte-Carlo approximation is used.

4.1.2. Space-filling design through Bayesian quadrature. Since s_n^2 given by (2.9) does not depend on the function f considered, a design \mathbf{X}_n for Bayesian integration can in principle be chosen beforehand, by direct minimization of s_n^2 . This corresponds to the approach followed in [47] where several quadrature rules are tabulated (for several values of n). Next theorem shows the connection between the minimum of $\mathcal{E}_K(\xi_n - \mu)$ with respect to weights \mathbf{w}_n and the posterior variances s_n^2 and $s_{n,0}^2$, see (2.9) and (2.13). We assume that all points in \mathbf{X}_n are pairwise different and μ is not fully supported on \mathbf{X}_n .

Theorem 4.1. *Let K be an SPD kernel and let $\mu \in \mathcal{M}^+(1) \cap \mathcal{M}_K$.*

(i) *The optimal unconstrained weights \mathbf{w}_n^* that minimize $\mathcal{E}_K(\xi_n - \mu)$ are $\mathbf{w}_n^* = \mathbf{K}_n^{-1} \mathbf{p}_n(\mu)$ and the corresponding measure ξ_n^* , with weights \mathbf{w}_n^* , satisfies*

$$\mathcal{E}_K(\xi_n^* - \mu) = s_{n,0}^2, \quad (4.2)$$

with $s_{n,0}^2$ given by (2.13).

(ii) *The optimal weights $\hat{\mathbf{w}}_n$ that minimize $\mathcal{E}_K(\xi_n - \mu)$ under the constraint $\mathbf{w}_n^T \mathbf{1}_n = \sum_{i=1}^n w_i = 1$ are*

$$\hat{\mathbf{w}}_n = \left(\mathbf{K}_n^{-1} - \frac{\mathbf{K}_n^{-1} \mathbf{1}_n \mathbf{1}_n^T \mathbf{K}_n^{-1}}{\mathbf{1}_n^T \mathbf{K}_n^{-1} \mathbf{1}_n} \right) \mathbf{p}_n(\mu) + \frac{\mathbf{K}_n^{-1} \mathbf{1}_n}{\mathbf{1}_n^T \mathbf{K}_n^{-1} \mathbf{1}_n}, \quad (4.3)$$

and the corresponding measure $\hat{\xi}_n$, with weights $\hat{\mathbf{w}}_n$, satisfies

$$\mathcal{E}_K(\hat{\xi}_n - \mu) = s_n^2, \quad (4.4)$$

with s_n^2 given by (2.9); the estimator (2.6) of the integral $I_\mu(f)$ is $\hat{I}_n = \hat{\mathbf{w}}_n^T \mathbf{y}_n$.

(iii) *For any bounded signed measure $\xi_n = \sum_{i=1}^n w_i \delta_{\mathbf{x}_i}$ we can write*

$$\mathcal{E}_K(\xi_n - \mu) = (\mathbf{w}_n - \mathbf{w}_n^*)^T \mathbf{K}_n (\mathbf{w}_n - \mathbf{w}_n^*) + \mathcal{E}_K(\xi_n^* - \mu), \quad (4.5)$$

and when the weights w_i sum to one, we have

$$\mathcal{E}_K(\xi_n - \mu) = (\mathbf{w}_n - \hat{\mathbf{w}}_n)^T \mathbf{K}_n (\mathbf{w}_n - \hat{\mathbf{w}}_n) + \mathcal{E}_K(\hat{\xi}_n - \mu). \quad (4.6)$$

Proof. The expression for \mathbf{w}_n^* , (4.2) and (4.5) directly follow from the fact that $\mathcal{E}_K(\xi_n - \mu)$ is quadratic in \mathbf{w}_n , see (4.1). Since K is SPD, straightforward calculation using Lagrangian theory indicates that the minimization of $\mathcal{E}_K(\xi_n - \mu)$ under the constraint $\mathbf{w}_n^T \mathbf{1}_n = 1$ gives (4.3) and (4.4). Suppose that $\mathbf{w}_n^T \mathbf{1}_n = 1$, then $\mathcal{E}_K(\xi_n - \mu) = (\mathbf{w}_n - \hat{\mathbf{w}}_n + \hat{\mathbf{w}}_n - \mathbf{w}_n^*)^T \mathbf{K}_n (\mathbf{w}_n - \hat{\mathbf{w}}_n + \hat{\mathbf{w}}_n - \mathbf{w}_n^*) + \mathcal{E}_K(\xi_n^* - \mu)$ gives (4.6) since $\mathbf{K}_n(\hat{\mathbf{w}}_n - \mathbf{w}_n^*)$ is proportional to $\mathbf{1}_n$ and $(\mathbf{w}_n - \hat{\mathbf{w}}_n)^T \mathbf{1}_n = 0$. ■

In the discrete case considered here, the minimum-energy signed measure $\hat{\xi}_n$ with total mass one always exists, but note that it is not necessarily a probability measure; that is, some weights \hat{w}_i may be negative. Theorem 4.1 can be extended to the case where K is only conditionally SPD, but the computation of optimal weights $\hat{\mathbf{w}}_n$ is more involved when \mathbf{K}_n is singular; see Remark 4.1.

Denote by $\tilde{\mathbf{K}}_n$ the $n \times n$ matrix with elements $\{\tilde{\mathbf{K}}_n\}_{i,j} = K_\mu(\mathbf{x}_i, \mathbf{x}_j)$, where K_μ is the reduced kernel (3.15); the corresponding vector of potential values at the \mathbf{x}_i is then $\tilde{\mathbf{p}}_n = (\tilde{P}_\mu(\mathbf{x}_1), \dots, \tilde{P}_\mu(\mathbf{x}_n))^T = \mathbf{0}$. For measures ξ_n in $\mathcal{M}(1)$, in complement of (ii) of Theorem 4.1, we also have the following property.

Theorem 4.2. For K an SPD kernel, $\mu \in \mathcal{M}^+(1) \cap \mathcal{M}_K$ and $\xi_n \in \mathcal{M}(1)$, we have

$$\mathcal{E}_K(\xi_n - \mu) = \mathcal{E}_{K_\mu}(\xi_n) = \mathbf{w}_n^T \tilde{\mathbf{K}}_n \mathbf{w}_n. \quad (4.7)$$

The posterior mean (2.6) and variance (2.9) of $I_\mu(f)$ are respectively given by

$$\hat{I}_n = \frac{\mathbf{1}_n^T \tilde{\mathbf{K}}_n^{-1} \mathbf{y}_n}{\mathbf{1}_n^T \tilde{\mathbf{K}}_n^{-1} \mathbf{1}_n}, \quad (4.8)$$

$$\sigma^2 s_n^2 = \sigma^2 (\mathbf{1}_n^T \tilde{\mathbf{K}}_n^{-1} \mathbf{1}_n)^{-1}. \quad (4.9)$$

Proof. Equation (4.7) follows from Proposition 3.5. Since we assumed that μ is not fully supported on \mathbf{X}_n and K is SPD, (4.7) gives $\inf_{\|\mathbf{w}_n\|=1} \mathbf{w}_n^T \tilde{\mathbf{K}}_n \mathbf{w}_n > 0$, which implies that $\tilde{\mathbf{K}}_n$ has full rank. Direct calculation using (3.15) gives $\tilde{\mathbf{K}}_n = \mathbf{K}_n - \mathbf{p}_n(\mu) \mathbf{1}_n^T - \mathbf{1}_n \mathbf{p}_n^T(\mu) + \mathcal{E}_K(\mu) \mathbf{1}_n \mathbf{1}_n^T$. The expression for $\tilde{\mathbf{K}}_n^{-1}$ then yields $\mathbf{1}_n^T \tilde{\mathbf{K}}_n^{-1} \mathbf{1}_n = 1/s_n^2$, with s_n^2 given by (2.9), proving (4.9). The expansion of $(\mathbf{1}_n^T \tilde{\mathbf{K}}_n^{-1} \mathbf{y}_n)/(\mathbf{1}_n^T \tilde{\mathbf{K}}_n^{-1} \mathbf{1}_n)$ gives (2.6), which proves (4.8). ■

Remark 4.1 (Optimal weights for CISPD kernels). Lagrangian theory indicates that the solution $\hat{\mathbf{w}}_n$ is obtained by solving the linear equation $\mathbf{R}_n(\hat{\mathbf{w}}_n^T \lambda)^T = (\mathbf{0}^T \ 1)^T$, where

$$\mathbf{R}_n = \begin{pmatrix} \tilde{\mathbf{K}}_n & \mathbf{1}_n \\ \mathbf{1}_n^T & 0 \end{pmatrix}.$$

When K is conditionally SPD, K_μ is conditionally SPD too, and the matrix \mathbf{R}_n has full rank $n + 1$. Indeed, $\mathbf{R}_n(\mathbf{z}_n^T \ z)^T = \mathbf{0}$ implies $\mathbf{1}_n^T \mathbf{z}_n = 0$ and $\tilde{\mathbf{K}}_n \mathbf{z}_n + z \mathbf{1}_n = 0$. Multiplying the second equation by \mathbf{z}_n^T , we get $\mathbf{z}_n^T \tilde{\mathbf{K}}_n \mathbf{z}_n = 0$. Since K_μ is conditionally SPD, this is incompatible with $\mathbf{1}_n^T \mathbf{z}_n = 0$ unless $\mathbf{z}_n = \mathbf{0}$ and $z = 0$. We obtain

$$\hat{\mathbf{w}}_n = \frac{(\tilde{\mathbf{K}}_n + \mathbf{1}_n \mathbf{1}_n^T)^{-1} \mathbf{1}_n}{\mathbf{1}_n^T (\tilde{\mathbf{K}}_n + \mathbf{1}_n \mathbf{1}_n^T)^{-1} \mathbf{1}_n},$$

and $s_n^2 = \hat{\mathbf{w}}_n^T \tilde{\mathbf{K}}_n \hat{\mathbf{w}}_n = (\mathbf{1}_n^T (\tilde{\mathbf{K}}_n + \mathbf{1}_n \mathbf{1}_n^T)^{-1} \mathbf{1}_n)^{-1} - 1$. When K is SPD and $\tilde{\mathbf{K}}_n$ has full rank (Theorem 4.2), we recover $\hat{\mathbf{w}}_n = \tilde{\mathbf{K}}_n^{-1} \mathbf{1}_n / (\mathbf{1}_n^T \tilde{\mathbf{K}}_n^{-1} \mathbf{1}_n)$ and $\hat{I}_n = \hat{\mathbf{w}}_n^T \mathbf{y}_n$ given by (4.8). ◁

Remark 4.2 (BLUE and kernel reduction). Equations (4.8) and (4.9) indicate that \hat{I}_n is the BLUE of β'_0 and $\sigma^2 s_n^2$ is its variance in the model (3.17), $f(\mathbf{x}) = \beta'_0 + \tilde{Z}_x$, see Sections 3.4.2 and 3.5.2. A possible interpretation is as follows. Predictions are not modified when using the reduced kernel K_μ instead of K , that is, when considering model $f(\mathbf{x}) = \beta'_0 + \tilde{Z}_x$ instead of (2.3), see [26, Sect. 5.4]. It implies that the expressions (2.6) and (2.9) of \hat{I}_n and s_n^2 are unchanged when replacing K by K_μ . Since, by construction, $\tilde{\mathbf{p}}_n(\mu) = \mathbf{0}$ and $\mathcal{E}_{K_\mu}(\mu) = 0$ ($(\text{Id}_{L^2} - \mathcal{P}_1)Z_x$ has no contribution to the integral of f), we directly obtain (4.8) and (4.9). ◁

Remark 4.3 (IMSPE for tensor product kernels). The use of a tensor product kernel (3.18) and a product measure $\mu = \otimes_{i=1}^d \mu^{(i)}$ on $\mathcal{X} = \mathcal{X}_1 \times \dots \times \mathcal{X}_d$ facilitates the calculations of $\tilde{\mathbf{K}}_n$ and $\mathcal{E}_K(\xi_n - \mu)$, see (4.1), since $\mathcal{E}_K(\mu)$ and $P_\mu(\mathbf{x}_i)$ have the simple

expressions (3.19, 3.20). The calculation of the IMSPE is facilitated too, but to a lesser extent. Indeed, we have

$$\int_{\mathcal{X}} \rho_n^2(\mathbf{x}) d\mu(\mathbf{x}) = \mathcal{E}_K(\mu) + \frac{1}{\mathbf{1}_n^T \mathbf{K}_n^{-1} \mathbf{1}_n} - 2 \frac{\mathbf{p}_n^T(\mu) \mathbf{K}_n^{-1} \mathbf{1}_n}{\mathbf{1}_n^T \mathbf{K}_n^{-1} \mathbf{1}_n} - \text{trace} [\mathbf{K}_n^{-1} \mathbf{Q}_n^\perp \mathbf{H}_n(\mu)],$$

see (2.4), where $\mathbf{Q}_n^\perp = \mathbf{I}_n - \mathbf{1}_n \mathbf{1}_n^T \mathbf{K}_n^{-1} / (\mathbf{1}_n^T \mathbf{K}_n^{-1} \mathbf{1}_n)$, with \mathbf{I}_n the n -dimensional identity matrix, is a projector onto the linear space orthogonal to $\mathbf{1}_n$, and where $\mathbf{H}_n(\mu)$ is the symmetric non-negative definite $n \times n$ matrix with elements

$$\{\mathbf{H}_n(\mu)\}_{j,k} = \int_{\mathcal{X}} K(\mathbf{x}, \mathbf{x}_j) K(\mathbf{x}, \mathbf{x}_k) d\mu(\mathbf{x}) = \prod_{i=1}^d \int_{\mathcal{X}_i} K_i(x_i, x_{j_i}) K_i(x_i, x_{k_i}) d\mu^{(i)}(x_i). \triangleleft$$

Theorems 4.1 and 4.2 indicate that, if K is SPD, $(\mathbf{1}_n^T \tilde{\mathbf{K}}_n^{-1} \mathbf{1}_n)^{-1}$ is the minimum value of $\mathcal{E}_K(\xi_n - \mu)$ for measures $\xi_n \in \mathcal{M}(1)$. Hence, we can construct space-filling designs on a compact and convex subset \mathcal{X} of \mathbb{R}^d by maximizing $\mathbf{1}_n^T \tilde{\mathbf{K}}_n^{-1} \mathbf{1}_n$ with respect to $\mathbf{X}_n \in \mathbb{R}^{nd}$, taking μ uniform on \mathcal{X} . This can be performed using any unconstrained nonlinear programming algorithm, as Example 4.3 will illustrate. Note that, from (4.7) and Cauchy-Schwarz inequality, $(\mathbf{1}_n^T \tilde{\mathbf{K}}_n^{-1} \mathbf{1}_n)^{-1} \leq \mathcal{E}_K(\xi_{n,e} - \mu) = (\mathbf{1}_n^T \tilde{\mathbf{K}}_n \mathbf{1}_n) / n^2$, the minimization of which was considered in Section 4.1.1.

4.2. Any-time designs. There exist situations where the number n of design points ultimately used (for integration, or function approximation) differs from that initially planned, say N . It is the case in particular when function evaluations are computationally more expensive than expected, and numerical experimentation is stopped after $n < N$ simulations, or when simulations fail at some design points and testing at more than N points is required to obtain N valid evaluations in total. In such circumstances, it is convenient to have sequences of nested designs at one's disposal. The objective is then to construct any-time designs; that is, ordered sequences $\mathbf{x}_1, \mathbf{x}_2, \dots$ of design points such that any design $\mathbf{X}_n = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ made of the first n points of the sequence has good space-filling properties. A typical example is given by Low Discrepancy Sequences (LDS) in $[0, 1]^d$, see [46].

When K is SPD, we may exploit expression (4.9) of the conditional variance of $I_\mu(f)$ in a greedy sequential construction: at step n we choose \mathbf{x}_{n+1} that minimizes s_{n+1}^2 . This sequential construction, called Sequential Bayesian Quadrature in [13], is straightforward to implement compared with global minimization of s_n^2 , see (2.9). Direct calculation, using formulae for the inversion of the block matrix

$$\tilde{\mathbf{K}}_{n+1} = \begin{pmatrix} \tilde{\mathbf{K}}_n & \tilde{\mathbf{k}}_n(\mathbf{x}_{n+1}) \\ \tilde{\mathbf{k}}_n^T(\mathbf{x}_{n+1}) & K_\mu(\mathbf{x}_{n+1}, \mathbf{x}_{n+1}) \end{pmatrix},$$

where $\{\tilde{\mathbf{K}}_n\}_{i,j} = K_\mu(\mathbf{x}_i, \mathbf{x}_j)$ and $\{\tilde{\mathbf{k}}_n(\mathbf{x})\}_i = K_\mu(\mathbf{x}, \mathbf{x}_i)$, $i, j = 1, \dots, n$, $\mathbf{x} \in \mathcal{X}$, gives

$$s_{n+1}^2 = \left[\mathbf{1}_n^T \tilde{\mathbf{K}}_n^{-1} \mathbf{1}_n + \frac{(1 - \tilde{\mathbf{k}}_n(\mathbf{x}_{n+1})^T \tilde{\mathbf{K}}_n^{-1} \mathbf{1}_n)^2}{K_\mu(\mathbf{x}_{n+1}, \mathbf{x}_{n+1}) - \tilde{\mathbf{k}}_n^T(\mathbf{x}_{n+1}) \tilde{\mathbf{K}}_n^{-1} \tilde{\mathbf{k}}_n(\mathbf{x}_{n+1})} \right]^{-1}. \quad (4.10)$$

The sequential construction is thus

$$\mathbf{x}_{n+1} \in \text{Arg max}_{\mathbf{x} \in \mathcal{X}} \frac{(1 - \tilde{\mathbf{k}}_n(\mathbf{x})^T \tilde{\mathbf{K}}_n^{-1} \mathbf{1}_n)^2}{K_\mu(\mathbf{x}, \mathbf{x}) - \tilde{\mathbf{k}}_n^T(\mathbf{x}) \tilde{\mathbf{K}}_n^{-1} \tilde{\mathbf{k}}_n(\mathbf{x})}.$$

The conditional gradient algorithm of [22] yields a simpler construction, particularly well adapted to the situation and also applicable when K is unbounded. It relies on the sequential selection of points that minimize the current directional derivative of $\mathcal{E}_K(\xi - \mu) = \mathcal{E}_{K_\mu}(\xi)$, with ξ supported on design points previously selected. The algorithm is initialized at a measure $\xi^{(n_0)}$ supported on $\mathbf{X}_{n_0} \in \mathcal{X}^{n_0}$ (with for instance $n_0 = 1$ and $\xi^{(1)} = \delta_{\mathbf{x}_1}$ for some $\mathbf{x}_1 \in \mathcal{X}$). Let $\xi^{(n)}$ denote the measure associated with the current design \mathbf{X}_n of iteration n , with weights $w_i^{(n)}$, i.e., $\xi^{(n)} = \sum_{i=1}^n w_i^{(n)} \delta_{\mathbf{x}_i}$. Next design point is chosen in $\text{Arg min}_{\mathbf{x} \in \mathcal{X}} F_{K_\mu}(\xi^{(n)}, \delta_{\mathbf{x}})$ (any minimizer can be selected in case there are several). Straightforward calculation using (3.12) gives $\mathbf{x}_{n+1} \in \text{Arg min}_{\mathbf{s} \in \mathcal{X}} [P_{\xi^{(n)}}(\mathbf{x}) - P_\mu(\mathbf{x})]$, that is,

$$\mathbf{x}_{n+1} \in \text{Arg min}_{\mathbf{x} \in \mathcal{X}} \left[\sum_{i=1}^n w_i^{(n)} K(\mathbf{x}, \mathbf{x}_i) - P_\mu(\mathbf{x}) \right]. \quad (4.11)$$

Note that this construction is well defined even if K is singular: in that case, it ensures that all design points are different ($\mathbf{x}_i \neq \mathbf{x}_j$ for all i, j); the same is true for all one-dimensional canonical projections when K is the tensorised product of singular kernels.

After choosing \mathbf{x}_{n+1} , the measure $\xi^{(n)}$ is updated into

$$\xi^{(n+1)} = (1 - \alpha_n) \xi^{(n)} + \alpha_n \delta_{\mathbf{x}_{n+1}} \quad (4.12)$$

for some $\alpha_n \in [0, 1]$, so that $\xi^{(n+1)} \in \mathcal{M}^+(1)$ when $\xi^{(n)} \in \mathcal{M}^+(1)$. When $\xi^{(n_0)}$ is the empirical (uniform) measure on \mathbf{X}_{n_0} , the choice $\alpha_n = 1/(n+1)$ implies that $\xi^{(n)}$ remains uniform on its support \mathbf{X}_n for all n , see [68] for an early contribution in the design context. The method is called *kernel herding* in the machine-learning literature, see [4, 14, 35]. It is shown in [14] that $\mathcal{E}_K(\xi^{(n)} - \mu) = \mathcal{O}(1/n^2)$ when \mathcal{H}_K is finite dimensional, but we only have the weaker result $\mathcal{E}_K(\xi^{(n)} - \mu) = \mathcal{O}(1/n)$ when \mathcal{H}_K is infinite dimensional, see [4].

Remark 4.4. Denote $\xi^{(n+)}(\mathbf{x}) = [n/(n+1)]\xi^{(n)} + [1/(n+1)]\delta_{\mathbf{x}}$. The direct minimization of $\mathcal{E}_K(\xi^{(n+)}(\mathbf{x}) - \mu)$ with respect to \mathbf{x} yields

$$\mathbf{x}^{(n+1)} \in \text{Arg min}_{\mathbf{x} \in \mathcal{X}} \left[\frac{1}{n+1} \sum_{i=1}^k K(\mathbf{x}, \mathbf{x}_i) - P_\mu(\mathbf{x}) + \frac{1}{2(k+1)} K(\mathbf{x}, \mathbf{x}) \right],$$

that is, a selection very close to (4.11) when $K(\mathbf{x}, \mathbf{x})$ is constant (Matérn kernel for instance). Note that this construction requires $K(\mathbf{x}, \mathbf{x}) < \infty$ for all $\mathbf{x} \in \mathcal{X}$, contrary to (4.11). \triangleleft

In practice n is always smaller than some given n_{\max} , and to facilitate the construction we can restrict the choice of the \mathbf{x}_i to a finite subset $\mathcal{X}_\Omega = \{\mathbf{s}_1, \dots, \mathbf{s}_\Omega\}$ of \mathcal{X} , with $\Omega \gg n_{\max}$ (when $\mathcal{X} = [0, 1]^d$, \mathcal{X}_Ω can be given by the first Ω points of a LDS). For any $n \leq n_{\max}$, we can write $\mathbf{X}_n = \{\mathbf{x}_1, \dots, \mathbf{x}_n\} = \{\mathbf{s}_{i_1}, \dots, \mathbf{s}_{i_n}\}$, the construction being initialized at some n_0 -point design $\mathbf{X}_{n_0} \subset \mathcal{X}_\Omega$. A measure ξ supported on \mathbf{X}_n can thus be written as $\xi = \sum_{i=1}^\Omega \omega_i \delta_{\mathbf{s}_i}$, with $\omega_i = 0$ when $\mathbf{s}_i \notin \mathbf{X}_n$. Therefore, for all n , $\xi^{(n)}$ is fully characterized by a Ω -dimensional vector $\boldsymbol{\omega}^{(n)} = (\omega_1^{(n)}, \dots, \omega_\Omega^{(n)})^T$, with $\boldsymbol{\omega}^{(n)}$ in the probability simplex \mathbb{P}_Ω when $\xi^{(n)} \in \mathcal{M}^+(1)$. The updating equations (4.11, 4.12) then imply that $\boldsymbol{\omega}^{(n+1)}$ is obtained by moving $\boldsymbol{\omega}^{(n)}$ in the direction of a vertex of \mathbb{P}_Ω , hence the name *vertex-direction* given to methods based on (4.12) in the literature on optimal design, see, e.g., [51, Chap. 9] and the references therein. A

summary of results on the rate of decrease of $\mathcal{E}_K(\xi^{(n)} - \mu)$ in this situation is given in Appendix A. The cost of the determination of \mathbf{x}_{n+1} in (4.11) is $\mathcal{O}(\Omega)$ (we need to compute $K(\mathbf{x}, \mathbf{x}_n)$ for all $\mathbf{x} \in \mathcal{X}_\Omega$), and the cost for n iterations scales as $\mathcal{O}(n\Omega)$ (including the initial cost for the computation of $P_\mu(\mathbf{x})$ for all $\mathbf{x} \in \mathcal{X}_\Omega$). An n -point design constructed in this way can be used as initialization for the (unconstrained) minimization of $\mathcal{E}_K(\xi_{n,e} - \mu) = \mathbf{1}_n^T \tilde{\mathbf{K}}_n \mathbf{1}_n / n^2$ (Section 4.1.1), or the maximization of $\mathbf{1}_n^T \tilde{\mathbf{K}}_n^{-1} \mathbf{1}_n$ (Section 4.1.2), with respect to \mathbf{X}_n . The resulting design \mathbf{X}_n^* can in turn be used as candidate set for the greedy construction of [29] (also called coffee-house design in [44]), yielding a sequence of nested designs $\mathbf{X}_1, \dots, \mathbf{X}_n = \mathbf{X}_n^*$.

4.3. Illustrative examples. We take $\mathcal{X} = [0, 1]^2$, $n_{\max} = 100$, $n_0 = 1$ with $\mathbf{X}_1 = \{(0.5, 0.5)\}$; μ is uniform on \mathcal{X} and \mathcal{X}_Ω is given by the first 2^{12} points of Sobol' LDS. The kernel K is the tensor product of uni-dimensional Matérn 3/2 covariance functions $K_{3/2, \theta}$, see (3.5).

Figure 4.1-Left shows the design \mathbf{X}_{100} obtained after 99 iterations of (4.11, 4.12) with $\alpha_n = 1/(n+1)$ (\mathbf{X}_{100} is the support of $\xi^{(100)}$), for $\theta = 10$ in $K_{3/2, \theta}$. This design has visually better space-filling properties than the first 100 points \mathbf{S}_{100} of Sobol' sequence presented on the right part of the figure. This is confirmed by the numerical values of the covering and packing radii, respectively given by (2.1) and (2.2): $\text{CR}(\mathbf{X}_{100}) \simeq 0.0925 < \text{CR}(\mathbf{S}_{100}) \simeq 0.1377$, and $\text{PR}(\mathbf{X}_{100}) \simeq 0.0262 > \text{PR}(\mathbf{S}_{100}) \simeq 0.0204$. Local maximization of $\mathbf{1}_n^T \tilde{\mathbf{K}}_n^{-1} \mathbf{1}_n$ (Section 4.1.2) with respect to \mathbf{X}_n , initialized at \mathbf{X}_{100} , yields a design \mathbf{X}_{100}^* with better space-filling properties: $\text{CR}(\mathbf{X}_{100}^*) \simeq 0.0869$ and $\text{PR}(\mathbf{X}_{100}^*) \simeq 0.0364$. When minimizing $\mathcal{E}_K(\xi_{n,e} - \mu)$ with respect to \mathbf{X}_n (Section 4.1.1) we obtain $\text{CR}(\mathbf{X}_{100}^*) \simeq 0.0955$ and $\text{PR}(\mathbf{X}_{100}^*) \simeq 0.0359$. The performance is significantly worse, both in terms of CR and PR, when the optimization is initialized at \mathbf{S}_{100} .




Figure 4.1: Left: $\mathbf{X}_{100} = \text{Supp}(\xi^{(100)})$ generated by (4.11, 4.12) with $\alpha_n = 1/(n+1)$ and $\theta = 10$ in $K_{3/2, \theta}$. Right: first 100 points \mathbf{S}_{100} of Sobol' LDS.

Figure 4.2-Left presents the efficiencies, in terms of covering radius (2.1) (solid line) and packing radius (2.2) (dashed line) of \mathbf{S}_n corresponding to the first n points of Sobol' sequence relatively to $\mathbf{X}_n = \text{Supp}(\xi^{(n)})$, when $\xi^{(n)}$ is generated by (4.11, 4.12) with $\alpha_n = 1/(n+1)$ and $\theta = 10$ in $K_{3/2, \theta}$. Values smaller than one (shown by an horizontal line) indicate that \mathbf{X}_n has better space-filling properties than \mathbf{S}_n . A greedy coffee-house construction [29, 44] ensures $\text{CR}(\mathbf{X}_n) \leq 2 \min_{\mathbf{X}'_n} \text{CR}(\mathbf{X}'_n)$ and $\text{PR}(\mathbf{X}_n) \geq (1/2) \max_{\mathbf{X}'_n} \text{PR}(\mathbf{X}'_n)$, but the method is computationally expensive and

tends to choose points on the border of \mathcal{X} . On the other hand, the construction is straightforward (and efficient) when we use the design $\mathbf{X}_{n_{\max}}^*$ obtained by maximizing $\mathbf{1}_n^T \tilde{\mathbf{K}}_n^{-1} \mathbf{1}_n$ as candidate set. The first point $\mathbf{x}^{(1)}$ is chosen as the point in $\mathbf{X}_{n_{\max}}^*$ closest to the center of \mathcal{X} (here $(1/2, 1/2)$); then $\mathbf{x}^{(n+1)}$ is the point in $\mathbf{X}_{n_{\max}}^*$ furthest away from $\mathbf{X}_n = \{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(n)}\}$, for $n = 1, \dots, n_{\max} - 1$. The efficiencies of \mathbf{S}_n , in terms of covering radius (solid line) and packing radius (dashed line), relatively to this new sequence, are plotted in Figure 4.2-Right.




Figure 4.2: Efficiencies $\text{CR}(\mathbf{X}_n)/\text{CR}(\mathbf{S}_n)$ (solid line) and $\text{PR}(\mathbf{S}_n)/\text{PR}(\mathbf{X}_n)$ (dashed line), with \mathbf{S}_n corresponding to the first n points of Sobol' sequence. Left: $\mathbf{X}_n = \text{Supp}(\xi^{(n)})$ is generated by (4.11, 4.12) with $\alpha_n = 1/(n+1)$ and $\theta = 10$ in $K_{3/2, \theta}$. Right: \mathbf{X}_n is generated by coffee-house design with candidate set \mathbf{X}_{100}^* obtained by local maximization of $\mathbf{1}_n^T \tilde{\mathbf{K}}_n^{-1} \mathbf{1}_n$.

Not surprisingly, the performance of $\xi^{(n)}$ in terms of MMD are better than those of the empirical measure $\xi_{n,e}(\mathbf{S}_n)$ associated with \mathbf{S}_n , see Figure 4.3-Left. Figure 4.3-Right illustrates the fact that a larger correlation length in K yields a faster decrease of $\mathcal{E}_K[\xi_{n,e}(\mathbf{S}_n) - \mu]$ (compare with the figure on the left). On the other hand, this faster decrease does not mean that design points are better distributed, compare Figure 4.4-Left with Figures 4.1-Left.

Finally, Figure 4.4-Right illustrates an application of Algorithm (4.11, 4.12) in higher dimension with a singular kernel: $\mathcal{X} = [0, 1]^{10}$ and K is the tensor product of the one-dimensional logarithmic kernel $K_{(0)}$ in (3.8). Similarly to Figure 4.2, the figure presents the efficiencies $\text{CR}(\mathbf{X}_n)/\text{CR}(\mathbf{S}_n)$ and $\text{PR}(\mathbf{S}_n)/\text{PR}(\mathbf{X}_n)$ for $n = 1, \dots, 100$. Due to the large value of d , for a given design \mathbf{Z}_n , $\text{CR}(\mathbf{Z}_n)$ defined by (2.1) is under-approximated by $\text{CR}'(\mathbf{Z}_n) = \max_{\mathbf{x} \in \mathcal{X}'_{\Omega'}} \min_{\mathbf{z} \in \mathbf{Z}_n} \|\mathbf{x} - \mathbf{z}\|$, with $\mathcal{X}'_{\Omega'}$ formed by the first 2^{19} points of Sobol' sequence complemented with a 3^d full factorial design.

Appendix A. Some convergence properties of conditional gradient algorithms.

We consider a conditional gradient algorithm with iterations given by (4.12) when \mathcal{X} is replaced by a finite set $\mathcal{X}_{\Omega} = \{\mathbf{s}_1, \dots, \mathbf{s}_{\Omega}\}$. We do not assume that \mathcal{H}_K is finite dimensional. A measure ξ on \mathcal{X}_{Ω} is characterized by a vector of weights $\boldsymbol{\omega}$ in the probability simplex \mathbb{P}_{Ω} , and we shall write $J_K(\boldsymbol{\omega}) = \mathcal{E}_K(\xi - \mu)$. When $\boldsymbol{\omega}^T \mathbf{1}_{\Omega} = 1$, equation (4.6) gives

$$J_K(\boldsymbol{\omega}) = \|\boldsymbol{\omega} - \hat{\boldsymbol{\omega}}\|_{\tilde{\mathbf{K}}}^2 + J_K(\hat{\boldsymbol{\omega}}),$$




Figure 4.3: $\mathcal{E}_K(\xi^{(n)} - \mu)$ (solid line) and $\mathcal{E}_K[\xi_{n,e}(\mathbf{S}_n) - \mu]$ (dashed line) as functions of n (log scale). $\mathbf{X}_n = \text{Supp}(\xi^{(n)})$ is generated by (4.11, 4.12) with $\alpha_n = 1/(n+1)$ and the kernel $K_{3/2,\theta}$; \mathbf{S}_n corresponds to the first n points of Sobol' sequence. Left: $\theta = 10$ in $K_{3/2,\theta}$; Right: $\theta = 1$.




Figure 4.4: Left: $\mathbf{X}_{100} = \text{Supp}(\xi^{(100)})$; $\mathcal{X} = [0, 1]^2$, \mathbf{X}_n is generated by (4.11, 4.12) with $\alpha_n = 1/(n+1)$ and the kernel $K_{3/2,\theta}$ with $\theta = 1$. Right: Efficiencies $\text{CR}(\mathbf{X}_n)/\text{CR}(\mathbf{S}_n)$ (solid line) and $\text{PR}(\mathbf{S}_n)/\text{PR}(\mathbf{X}_n)$ (dashed line); $\mathcal{X} = [0, 1]^{10}$, $\mathbf{X}_n = \text{Supp}(\xi^{(n)})$ generated by (4.11, 4.12) with $\alpha_n = 1/(n+1)$, and $K(\mathbf{x}, \mathbf{x}') = \prod_{i=1}^{10} \log(1/|x_i - x'_i|)$, \mathbf{S}_n corresponds to the first n points of Sobol' sequence.

where we write $\mathbf{K} = \mathbf{K}_\Omega$, $\|\omega - \hat{\omega}\|_{\mathbf{K}}^2 = (\omega - \hat{\omega})^T \mathbf{K} (\omega - \hat{\omega})$ and where $\hat{\omega} = \hat{\omega}_\Omega$ is given by (4.3). We denote by $\lambda_{\max}(\mathbf{K})$ the largest eigenvalue of \mathbf{K} .

For $i = 1, \dots, \Omega$, we denote by \mathbf{e}_i the i -th basis vector, with component number i equal to one. Iteration (4.12) has the form

$$\omega^{(n+1)} = \omega^{(n)} + \alpha_n \Delta_n$$

for some step-size α_n and direction $\Delta_n = \mathbf{e}_{i_n^+} - \omega^{(n)}$, with the index i_n^+ taken in $\text{Arg min}_{i=1, \dots, \Omega} \mathbf{e}_i^T \nabla J_K(\omega^{(n)})$, where the gradient $\nabla J_K(\omega)$ is given by

$$\nabla J_K(\omega) = 2\mathbf{K}(\omega - \hat{\omega}).$$

This is equivalent to $\mathbf{s}_{i_n^+} \in \text{Arg min}_{\mathbf{s} \in \mathcal{X}_\Omega} [P_{\xi^{(n)}}(\mathbf{s}) - P_\mu(\mathbf{s})]$, see (4.11).

A.1. Vertex-direction, predefined step-size. Take $\alpha_n = 1/(n+1)$ in (4.12). We first mention a simple result indicating that $\|\boldsymbol{\omega}^{(n)} - \mathbf{1}_\Omega/\Omega\|_{\mathbf{K}}^2 = \mathcal{O}(1/n)$ during the initial $n_1 \leq \Omega$ iterations when all i_n^+ are distinct for $n \leq n_1$.

Lemma A.1. *Algorithm (4.12) with $\alpha_n = 1/(n+1)$, initialized at $\boldsymbol{\omega}^{(1)} = \mathbf{e}_{i_0}$ for some $i_0 \in \{1, \dots, \Omega\}$, satisfies*

$$\|\boldsymbol{\omega}^{(n)} - \mathbf{1}_\Omega/\Omega\|_{\mathbf{K}}^2 \leq \frac{\lambda_{\max}(\mathbf{K})}{n}, \quad 1 \leq n \leq n_1 \leq \Omega,$$

where n_1 is such that all i_n^+ are distinct for $n \leq n_1$.

Proof. For $n \leq n_1$, after a suitable reordering of indices we have $\boldsymbol{\omega}^{(n)} = (1/n, \dots, 1/n, 0, \dots, 0)^T$. Therefore, $\|\boldsymbol{\omega}^{(n)} - \mathbf{1}_\Omega/\Omega\|_{\mathbf{K}}^2 \leq \lambda_{\max}(\mathbf{K})\|\boldsymbol{\omega}^{(n)} - \mathbf{1}_\Omega/\Omega\|^2 = \lambda_{\max}(\mathbf{K})(\Omega - n)/(n\Omega) \leq \lambda_{\max}(\mathbf{K})/n$. ■

Note that this property is independent of the order in which the vertices of \mathbb{P}_Ω (the \mathbf{e}_{i_n}) are selected. It is therefore also valid for MC sampling without replacement within \mathcal{X}_Ω . Also note that the optimal step-size $\hat{\alpha}_n$ at iteration n for the minimization of $\|\boldsymbol{\omega}^{(n)} - \mathbf{1}_\Omega/\Omega\|^2$ equals $\alpha_n = 1/(n+1)$.

The following lemma shows that (4.12) with $\alpha_n = 1/(n+1)$ ensures that $\|\boldsymbol{\omega}^{(n)} - \hat{\boldsymbol{\omega}}\|_{\mathbf{K}}^2 = \mathcal{O}(\log n/n)$, independently of Ω and of the positions of $\boldsymbol{\omega}^{(1)}$ and $\hat{\boldsymbol{\omega}}$ in \mathbb{P}_Ω .

Lemma A.2. *Algorithm (4.12) with $\alpha_n = 1/(n+1)$, initialized at any $\boldsymbol{\omega}^{(1)}$ in \mathbb{P}_Ω , satisfies*

$$\|\boldsymbol{\omega}^{(n)} - \hat{\boldsymbol{\omega}}\|_{\mathbf{K}}^2 \leq 2\lambda_{\max}(\mathbf{K}) \frac{1 + 2\log(n+1)}{n}, \quad n \geq 1. \quad (\text{A.1})$$

Proof. The proof follows the same lines as in [15, Sect. 3]. Denote $g(\boldsymbol{\omega}) = \|\boldsymbol{\omega} - \hat{\boldsymbol{\omega}}\|_{\mathbf{K}}^2$ and $\boldsymbol{\omega}^{(n+)}(\alpha) = \boldsymbol{\omega}^{(n)} + \alpha\Delta_n$. Notice that $\boldsymbol{\omega}^{(n)} \in \mathbb{P}_\Omega$ for all $n \geq 1$. We have

$$\begin{aligned} g[\boldsymbol{\omega}^{(n+)}(\alpha)] &= g(\boldsymbol{\omega}^{(n)}) + 2\alpha\Delta_n^T \mathbf{K}(\boldsymbol{\omega}^{(n)} - \hat{\boldsymbol{\omega}}) + \alpha^2\|\Delta_n\|_{\mathbf{K}}^2 \\ &\leq g(\boldsymbol{\omega}^{(n)}) + 2\alpha\Delta_n^T \mathbf{K}(\boldsymbol{\omega}^{(n)} - \hat{\boldsymbol{\omega}}) + \alpha^2\lambda_{\max}(\mathbf{K})\|\Delta_n\|^2 \\ &\leq g(\boldsymbol{\omega}^{(n)}) + 2\alpha\Delta_n^T \mathbf{K}(\boldsymbol{\omega}^{(n)} - \hat{\boldsymbol{\omega}}) + 2\alpha^2\lambda_{\max}(\mathbf{K}). \end{aligned}$$

The convexity of $g(\cdot)$ and the definition of Δ_n imply that

$$g(\boldsymbol{\omega}^{(n)}) \geq g(\hat{\boldsymbol{\omega}}) = 0 \geq g(\boldsymbol{\omega}^{(n)}) + (\hat{\boldsymbol{\omega}} - \boldsymbol{\omega}^{(n)})^T \nabla J_K(\boldsymbol{\omega}^{(n)}) \geq g(\boldsymbol{\omega}^{(n)}) + \Delta_n^T \nabla J_K(\boldsymbol{\omega}^{(n)}).$$

Therefore, $\Delta_n^T \nabla J_K(\boldsymbol{\omega}^{(n)}) = 2\Delta_n^T \mathbf{K}(\boldsymbol{\omega}^{(n)} - \hat{\boldsymbol{\omega}}) \leq -g(\boldsymbol{\omega}^{(n)})$ and

$$g[\boldsymbol{\omega}^{(n+)}(\alpha)] \leq (1 - \alpha)g(\boldsymbol{\omega}^{(n)}) + 2\alpha^2\lambda_{\max}(\mathbf{K}). \quad (\text{A.2})$$

The rest of the proof is by induction on n . The bound (A.1) is valid for $n = 1$ since $\|\boldsymbol{\omega}^{(1)} - \hat{\boldsymbol{\omega}}\|_{\mathbf{K}}^2 \leq 2\lambda_{\max}(\mathbf{K})$. Suppose that it is satisfied by $\boldsymbol{\omega}^{(n)}$; (A.2) gives

$$\begin{aligned} \|\boldsymbol{\omega}^{(n+1)} - \hat{\boldsymbol{\omega}}\|_{\mathbf{K}}^2 &\leq 2\lambda_{\max}(\mathbf{K}) \left\{ \frac{1 + 2\log(n+2)}{n+1} - \frac{2(n+1)\log[1 + 1/(n+1)] - 1}{(n+1)^2} \right\} \\ &\leq 2\lambda_{\max}(\mathbf{K}) \frac{1 + 2\log(n+2)}{n+1} \end{aligned}$$

since $\log(1+t) \geq t/2$ for $t \in [0, 1]$. ■

Using $\alpha = 2/(n+3)$ in (A.2), one can easily prove by induction that $g(\boldsymbol{\omega}^{(n)}) \leq 8\lambda_{\max}(\mathbf{K})/(n+3)$ for all n , see [15], which means that (4.12) with $\alpha_n = 2/(n+3)$

instead of $1/(n+1)$ satisfies $\|\boldsymbol{\omega}^{(n)} - \hat{\boldsymbol{\omega}}\|_{\mathbf{K}}^2 \leq 8 \lambda_{\max}(\mathbf{K})/(n+3)$, $n \geq 1$, with thus a much faster decrease than (A.1). Using a different approach, it is shown in [19] that a rate of decrease of $\mathcal{O}(1/n)$ is also obtained when α_n corresponds to the sequence $\alpha_{n+1} = \alpha_n - \alpha_n^2/2$ with $\alpha_0 = 1$.

Next lemma, based on [14], shows that $\|\boldsymbol{\omega}^{(n)} - \hat{\boldsymbol{\omega}}\|_{\mathbf{K}}^2$ decreases in C/n^2 when $\hat{\boldsymbol{\omega}}$ lies in the interior of \mathbb{P}_Ω . Here, contrary to [14], we do not assume that \mathcal{H} is finite dimensional and use instead the finite dimensionality of $\boldsymbol{\omega}$.

Lemma A.3. *When $\hat{\boldsymbol{\omega}}$ is in the interior of \mathbb{P}_Ω , (4.12) with $\alpha_n = 1/(n+1)$, initialized at any $\boldsymbol{\omega}^{(1)}$ in \mathbb{P}_Ω , satisfies*

$$\|\boldsymbol{\omega}^{(n)} - \hat{\boldsymbol{\omega}}\|_{\mathbf{K}}^2 \leq 4R_*^2 \left(1 + \frac{R_*^2}{\alpha_*^2}\right) \frac{1}{n^2}, \quad n \geq 1,$$

where $R_* = [\lambda_{\max}(\mathbf{K}) (1 - 1/\Omega)]^{1/2}$ and $\alpha_* = w_*/L$, with $w_* = \min_{i=1, \dots, \Omega} \{\hat{\mathbf{w}}\}_i$ (so that $w_* \leq 1/\Omega$) and $L = (\max_{i=1, \dots, \Omega} \{\mathbf{K}^{-1}\}_{ii})^{1/2}$.

Proof. Denote $\mathbf{v}(\alpha) = \hat{\boldsymbol{\omega}} - \alpha(\boldsymbol{\omega}^{(n)} - \hat{\boldsymbol{\omega}})/\|\boldsymbol{\omega}^{(n)} - \hat{\boldsymbol{\omega}}\|_{\mathbf{K}}$, $\alpha > 0$. Then, for any $i = 1, \dots, \Omega$,

$$\begin{aligned} \frac{|\mathbf{e}_i^T(\boldsymbol{\omega}^{(n)} - \hat{\boldsymbol{\omega}})|}{\|\boldsymbol{\omega}^{(n)} - \hat{\boldsymbol{\omega}}\|_{\mathbf{K}}} &\leq \max_{\mathbf{u}: \mathbf{u}^T \mathbf{1}_\Omega = 0, \mathbf{u}^T \mathbf{K} \mathbf{u} = 1} |\mathbf{e}_i^T \mathbf{u}| \\ &\leq \max_{\mathbf{u}: \mathbf{u}^T \mathbf{K} \mathbf{u} = 1} |\mathbf{e}_i^T \mathbf{u}| = \sqrt{\mathbf{e}_i^T \mathbf{K}^{-1} \mathbf{e}_i} \leq L = w_*/\alpha_*, \end{aligned}$$

so that $\{\mathbf{v}(\alpha)\}_i \geq \{\hat{\boldsymbol{\omega}}\}_i - w_* \geq 0$, and $\mathbf{v}(\alpha) \in \mathbb{P}_\Omega$, for any $\alpha \leq \alpha_*$. The definition of $\mathbf{e}_{i_n}^+$ then implies that

$$(\mathbf{e}_{i_n}^+ - \hat{\boldsymbol{\omega}})^T \mathbf{K}(\boldsymbol{\omega}^{(n)} - \hat{\boldsymbol{\omega}}) \leq [\mathbf{v}(\alpha_*) - \hat{\boldsymbol{\omega}}]^T \mathbf{K}(\boldsymbol{\omega}^{(n)} - \hat{\boldsymbol{\omega}}) = -\alpha_* \|\boldsymbol{\omega}^{(n)} - \hat{\boldsymbol{\omega}}\|_{\mathbf{K}}. \quad (\text{A.3})$$

The rest of the proof is based on [14]. Denote $\mathbf{e}_{i_0}^+ = \boldsymbol{\omega}^{(1)}$ and $\mathbf{z}_n = \sum_{i=1}^n (\mathbf{e}_{i_{n-1}}^+ - \hat{\boldsymbol{\omega}})$. We can write $\boldsymbol{\omega}^{(n)} = (1/n) \sum_{i=1}^n \mathbf{e}_{i_{n-1}}^+$, so that $\mathbf{z}_n = n(\boldsymbol{\omega}^{(n)} - \hat{\boldsymbol{\omega}})$, $n^2 \|\boldsymbol{\omega}^{(n)} - \hat{\boldsymbol{\omega}}\|_{\mathbf{K}}^2 = \|\mathbf{z}_n\|_{\mathbf{K}}^2$, and we only need to bound $\|\mathbf{z}_n\|_{\mathbf{K}}^2$. We have

$$\|\mathbf{z}_n\|_{\mathbf{K}}^2 - \|\mathbf{z}_{n+1}\|_{\mathbf{K}}^2 = -2(\mathbf{e}_{i_n}^+ - \hat{\boldsymbol{\omega}})^T \mathbf{K} \mathbf{z}_n - \|\mathbf{e}_{i_n}^+ - \hat{\boldsymbol{\omega}}\|_{\mathbf{K}}^2,$$

where $\|\mathbf{e}_{i_n}^+ - \hat{\boldsymbol{\omega}}\|_{\mathbf{K}} \leq 2R_*$ and $(\mathbf{e}_{i_n}^+ - \hat{\boldsymbol{\omega}})^T \mathbf{K} \mathbf{z}_n \leq -\alpha_* \|\mathbf{z}_n\|_{\mathbf{K}}$ from (A.3). Therefore,

$$\|\mathbf{z}_{n+1}\|_{\mathbf{K}}^2 \leq \|\mathbf{z}_n\|_{\mathbf{K}}^2 - 2\alpha_* (\|\mathbf{z}_n\|_{\mathbf{K}} - 2R_*^2/\alpha_*).$$

Suppose that $\|\mathbf{z}_n\|_{\mathbf{K}} > 2R_*^2/\alpha_*$. Then, $\|\mathbf{z}_{n+1}\|_{\mathbf{K}}^2 \leq \|\mathbf{z}_n\|_{\mathbf{K}}^2$, and $\|\mathbf{z}_n\|_{\mathbf{K}}^2$ decreases until some n_0 when $\|\mathbf{z}_{n_0}\|_{\mathbf{K}} \leq 2R_*^2/\alpha_*$. But then,

$$\|\mathbf{z}_{n_0+1}\|_{\mathbf{K}}^2 \leq \|\mathbf{z}_{n_0}\|_{\mathbf{K}}^2 - 2\alpha_* (\|\mathbf{z}_{n_0}\|_{\mathbf{K}} - 2R_*^2/\alpha_*) \leq 4R_*^2(1 + R_*^2/\alpha_*^2),$$

so that $\|\mathbf{z}_n\|_{\mathbf{K}}^2 \leq 4R_*^2(1 + R_*^2/\alpha_*^2)$ for all $n > n_0$. \blacksquare

Lemma A.3 indicates that $\|\boldsymbol{\omega}^{(n)} - \hat{\boldsymbol{\omega}}\|_{\mathbf{K}}^2 \leq C/n^2$. However, for large Ω the constant C grows like $\mathcal{O}(\Omega^2)$ (since $\alpha_* \leq 1/(L\Omega)$) which makes this result of theoretical interest only. Note that kernel herding typically concerns situations where Ω is very large (or even infinite when \mathcal{X} is not discretized).

A.2. Vertex-direction, optimal step-size. The choice of a predefined step-size $\alpha_n = 1/(n+1)$ in (4.12) does not ensure a monotonic decrease of $\mathcal{E}_K(\xi^{(n)} - \mu)$. An alternative option is to choose α_n that minimizes $\mathcal{E}_K[\xi^{(n+)}(\alpha) - \mu]$ with respect to $\alpha \in [0, 1]$, with $\xi^{(n+)}(\alpha) = (1-\alpha)\xi^{(n)} + \alpha\delta_{\mathbf{x}_{n+1}}$ and \mathbf{x}_{n+1} given by (4.11). Straightforward calculation gives $\alpha_n = \min\{1, \hat{\alpha}_n\}$, with

$$\begin{aligned}\hat{\alpha}_n &= \frac{\langle P_{\xi^{(n)}} - P_\mu, P_{\xi^{(n)}} - P_{\delta_{\mathbf{x}_{n+1}}} \rangle_K}{\|P_{\xi^{(n)}} - P_{\delta_{\mathbf{x}_{n+1}}}\|_{\mathcal{H}_K}^2} \\ &= \frac{\mathcal{E}_K(\xi^{(n)}) - P_{\xi^{(n)}}(\mathbf{x}_{n+1}) - \sum_{i=1}^n w_i^{(n)} P_\mu(\mathbf{x}_i) + P_\mu(\mathbf{x}_{n+1})}{\mathcal{E}_K(\xi^{(n)}) - 2P_{\xi^{(n)}}(\mathbf{x}_{n+1}) + K(\mathbf{x}_{n+1}, \mathbf{x}_{n+1})}\end{aligned}\quad (\text{A.4})$$

(which requires that $\mathcal{E}_K(\xi^{(n)}) < \infty$). Direct calculation shows that $\hat{\alpha}_n$ given by (A.4) satisfies

$$\hat{\alpha}_n = \frac{(\mathbf{e}_{i_n^+} - \boldsymbol{\omega}^{(n)})^T \mathbf{K}(\hat{\boldsymbol{\omega}} - \boldsymbol{\omega}^{(n)})}{\|\mathbf{e}_{i_n^+} - \boldsymbol{\omega}^{(n)}\|_{\mathbf{K}}^2}.\quad (\text{A.5})$$

Next Lemma indicates that $\hat{\alpha}_n \leq 1$ when $\hat{\boldsymbol{\omega}} \in \mathbb{P}_\Omega$, so that setting $\alpha_n = \hat{\alpha}_n$ in (4.12) ensures that $\boldsymbol{\omega}^{(n)}$ remains in \mathbb{P}_Ω for all n . It should be noticed that the global decrease of $\|\boldsymbol{\omega}^{(n)} - \hat{\boldsymbol{\omega}}\|_{\mathbf{K}}^2$ over *many* iterations with this optimal α_n is not necessarily better than with the predefined step-size $\alpha_n = 1/(n+1)$ of Section A.1; one may refer in particular to [19] for such considerations; see also [4].

Lemma A.4. *When $\hat{\boldsymbol{\omega}} \in \mathbb{P}_\Omega$, $\hat{\alpha}_n$ given by (A.5) is less than one.*

Proof. We can write $(\mathbf{e}_{i_n^+} - \boldsymbol{\omega}^{(n)})^T \mathbf{K}(\hat{\boldsymbol{\omega}} - \boldsymbol{\omega}^{(n)}) = \|\mathbf{e}_{i_n^+} - \boldsymbol{\omega}^{(n)}\|_{\mathbf{K}}^2 + (\mathbf{e}_{i_n^+} - \hat{\boldsymbol{\omega}})^T \mathbf{K}(\boldsymbol{\omega}^{(n)} - \hat{\boldsymbol{\omega}}) - \|\mathbf{e}_{i_n^+} - \hat{\boldsymbol{\omega}}\|_{\mathbf{K}}^2$. When $\hat{\boldsymbol{\omega}} \in \mathbb{P}_\Omega$, $\hat{\omega}_i \geq 0$ for all i , and $\sum_{i=1}^\Omega \hat{\omega}_i (\mathbf{e}_i - \hat{\boldsymbol{\omega}}) = \mathbf{0}$ implies that $\sum_{i=1}^\Omega \hat{\omega}_i (\mathbf{e}_i - \hat{\boldsymbol{\omega}})^T \mathbf{K}(\boldsymbol{\omega}^{(n)} - \hat{\boldsymbol{\omega}}) = 0$. Therefore $\min_{i=1, \dots, \Omega} (\mathbf{e}_i - \hat{\boldsymbol{\omega}})^T \mathbf{K}(\boldsymbol{\omega}^{(n)} - \hat{\boldsymbol{\omega}}) = (\mathbf{e}_{i_n^+} - \hat{\boldsymbol{\omega}})^T \mathbf{K}(\boldsymbol{\omega}^{(n)} - \hat{\boldsymbol{\omega}}) \leq 0$, which gives $\hat{\alpha}_n \leq 1$. ■

Lemma A.5. *Algorithm (4.12) with $\alpha_n = \hat{\alpha}_n$ given by (A.5), initialized at any $\boldsymbol{\omega}^{(1)} \in \mathbb{P}_\Omega$, satisfies*

$$\|\boldsymbol{\omega}^{(n)} - \hat{\boldsymbol{\omega}}\|_{\mathbf{K}}^2 \leq 8 \lambda_{\max}(\mathbf{K}) \frac{1}{n+3}, \quad n \geq 1.\quad (\text{A.6})$$

Proof. The proof follows [15, Sect. 2] and uses the same notation as in the proof of Lemma A.2. The right-hand side of (A.2) is minimum for $\hat{\alpha} = g(\boldsymbol{\omega}^{(n)})/[4\lambda_{\max}(\mathbf{K})] \leq 1/2$. Therefore,

$$\begin{aligned}g(\boldsymbol{\omega}^{(n+1)}) &= \min_{\alpha \in [0,1]} g[\boldsymbol{\omega}^{(n+)}(\alpha)] \leq (1 - \hat{\alpha})g(\boldsymbol{\omega}^{(n)}) + 2\hat{\alpha}^2 \lambda_{\max}(\mathbf{K}) \\ &= g(\boldsymbol{\omega}^{(n)}) \left[1 - \frac{g(\boldsymbol{\omega}^{(n)})}{8\lambda_{\max}(\mathbf{K})} \right].\end{aligned}$$

Since $1 - t \leq 1/(1+t)$ for all $t > -1$, we obtain

$$g(\boldsymbol{\omega}^{(n+1)})/[8\lambda_{\max}(\mathbf{K})] \leq \frac{1}{1 + \{g(\boldsymbol{\omega}^{(n)})/[8\lambda_{\max}(\mathbf{K})]\}^{-1}}$$

which, by induction, implies that $g(\boldsymbol{\omega}^{(n)}) \leq 8\lambda_{\max}(\mathbf{K})/(n+3)$; that is, (A.6). ■

Lemma A.6. *When $\hat{\boldsymbol{\omega}}$ is in the interior of \mathbb{P}_Ω , (4.12) with $\alpha_n = \hat{\alpha}_n$ given by (A.5), initialized at any $\boldsymbol{\omega}^{(1)} \in \mathbb{P}_\Omega$, satisfies*

$$\|\boldsymbol{\omega}^{(n+1)} - \hat{\boldsymbol{\omega}}\|_{\mathbf{K}}^2 \leq \|\boldsymbol{\omega}^{(1)} - \hat{\boldsymbol{\omega}}\|_{\mathbf{K}}^2 \exp\left(-\frac{\alpha_*^2 k}{4R_*^2}\right), \quad n \geq 1,\quad (\text{A.7})$$

where $R_* = [\lambda_{\max}(\mathbf{K})(1 - 1/\Omega)]^{1/2}$ and $\alpha_* = w_*/L$, with $w_* = \min_{i=1, \dots, \Omega} \{\hat{\mathbf{w}}\}_i$ (so that $w_* \leq 1/\Omega$) and $L = (\max_{i=1, \dots, \Omega} \{\mathbf{K}^{-1}\}_{ii})^{1/2}$.

Proof. We use the same approach as in [5] and use the same notation as in the proof of Lemma A.2. We can write $g(\boldsymbol{\omega}^{(n+1)}) = g[\boldsymbol{\omega}^{(n)}(\hat{\alpha}_n)]$, with $\hat{\alpha}_n$ given by (A.5). Therefore,

$$\begin{aligned} g(\boldsymbol{\omega}^{(n+1)}) &= g(\boldsymbol{\omega}^{(n)}) + 2\hat{\alpha}_n \Delta_n^T \mathbf{K}(\boldsymbol{\omega}^{(n)} - \hat{\boldsymbol{\omega}}) + \alpha_n^2 \|\Delta_n\|_{\mathbf{K}}^2 \\ &= g(\boldsymbol{\omega}^{(n)}) - \frac{[\Delta_n^T \mathbf{K}(\boldsymbol{\omega}^{(n)} - \hat{\boldsymbol{\omega}})]^2}{\|\Delta_n\|_{\mathbf{K}}^2}. \end{aligned}$$

Equation (A.3) implies that $[\Delta_n^T \mathbf{K}(\boldsymbol{\omega}^{(n)} - \hat{\boldsymbol{\omega}})]^2 \geq \alpha_*^2 g(\boldsymbol{\omega}^{(n)})$, and thus

$$g(\boldsymbol{\omega}^{(n+1)}) \leq g(\boldsymbol{\omega}^{(n)}) \left[1 - \frac{\alpha_*^2}{\|\Delta_n\|_{\mathbf{K}}^2} \right] \leq g(\boldsymbol{\omega}^{(n)}) \left[1 - \frac{\alpha_*^2}{4R_*^2} \right].$$

This implies $g(\boldsymbol{\omega}^{(n+1)}) \leq g(\boldsymbol{\omega}^{(1)}) \exp[-\alpha_*^2 k / (4R_*^2)]$, that is, (A.7). \blacksquare

Similarly to Lemma A.3, the small value of the constant α_* makes the linear convergence rate in (A.7) of theoretical interest only.

A.3. Vertex-exchange. Following [42, 43], one may also use a vertex-exchange method based on the true steepest-descent direction, see also [11, 12]. The iterations are then

$$\boldsymbol{\xi}^{(n+1)} = \boldsymbol{\xi}^{(n)} + \alpha_n (\delta_{\mathbf{x}_{n+1}} - \delta_{\mathbf{x}_n^-}), \quad (\text{A.8})$$

where \mathbf{x}_{n+1} is given by (4.11) and

$$\mathbf{x}_n^- \in \text{Arg} \max_{\mathbf{x} \in \text{Supp}(\boldsymbol{\xi}^{(n)})} [P_{\boldsymbol{\xi}^{(n)}}(\mathbf{x}) - P_{\mu}(\mathbf{x})], \quad (\text{A.9})$$

with $\text{Supp}(\boldsymbol{\xi}^{(n)}) = \mathbf{X}_n$ the support of $\boldsymbol{\xi}^{(n)}$. The step-size α_n is then given by $\min\{\hat{\alpha}_n, \boldsymbol{\xi}^{(n)}(\mathbf{x}_n^-)\}$, where $\hat{\alpha}_n$ minimizes $\mathcal{E}_K[\{\boldsymbol{\xi}^{(n)} + \alpha(\delta_{\mathbf{x}_{n+1}} - \delta_{\mathbf{x}_n^-})\} - \mu]$ with respect to α (the constraint $\alpha_n \leq \boldsymbol{\xi}^{(n)}(\mathbf{x}_n^-)$ ensures that $\boldsymbol{\xi}^{(n+1)} \in \mathcal{M}^+(1)$ when $\boldsymbol{\xi}^{(n)} \in \mathcal{M}^+(1)$). Direct calculation gives

$$\begin{aligned} \hat{\alpha}_n &= \frac{\langle P_{\boldsymbol{\xi}^{(n)}} - P_{\mu}, P_{\delta_{\mathbf{x}_n^-}} - P_{\delta_{\mathbf{x}_{n+1}}} \rangle_K}{\|P_{\delta_{\mathbf{x}_n^-}} - P_{\delta_{\mathbf{x}_{n+1}}}\|_{\mathcal{H}_K}^2} \\ &= \frac{[P_{\boldsymbol{\xi}^{(n)}}(\mathbf{x}_n^-) - P_{\mu}(\mathbf{x}_n^-)] - [P_{\boldsymbol{\xi}^{(n)}}(\mathbf{x}_{n+1}) - P_{\mu}(\mathbf{x}_{n+1})]}{K(\mathbf{x}_n^-, \mathbf{x}_n^-) + K(\mathbf{x}_{n+1}, \mathbf{x}_{n+1}) - 2K(\mathbf{x}_n^-, \mathbf{x}_{n+1})}. \end{aligned} \quad (\text{A.10})$$

For the algorithm defined by (A.8, A.9), we have $\boldsymbol{\omega}^{(n+1)} = \boldsymbol{\omega}^{(n)} + \alpha_n \Delta_n$ with now $\Delta_n = \mathbf{e}_{i_n^+} - \mathbf{e}_{i_n^-}$, where we take $i_n^+ \in \text{Arg} \min_{i=1, \dots, \Omega} \mathbf{e}_i^T \mathbf{K}(\boldsymbol{\omega}^{(n)} - \hat{\boldsymbol{\omega}})$ and $i_n^- \in \text{Arg} \max_{i: \mathbf{e}_i^T \boldsymbol{\omega}^{(n)} > 0} \mathbf{e}_i^T \mathbf{K}(\boldsymbol{\omega}^{(n)} - \hat{\boldsymbol{\omega}})$. The step size (A.10) equals

$$\hat{\alpha}_n = \frac{\Delta_n^T \mathbf{K}(\hat{\boldsymbol{\omega}} - \boldsymbol{\omega}^{(n)})}{\|\Delta_n\|_{\mathbf{K}}^2}. \quad (\text{A.11})$$

Take $\alpha_n = \min\{\boldsymbol{\omega}_{i_n^-}^{(n)}, \hat{\alpha}_n\}$ in (A.8), so that $\boldsymbol{\omega}^{(n)}$ remains in \mathbb{P}_{Ω} for all n . Using the same notation as in the proof of Lemma A.2, we have

$$g[\boldsymbol{\omega}^{(n+1)}(\alpha)] \leq g(\boldsymbol{\omega}^{(n)}) + 2\alpha \Delta_n^T \mathbf{K}(\boldsymbol{\omega}^{(n)} - \hat{\boldsymbol{\omega}}) + 2\alpha^2 \lambda_{\max}(\mathbf{K}),$$

and, since $\boldsymbol{\omega}^{(n)} \in \mathbb{P}_\Omega$, the convexity of $g(\cdot)$ and the definition of Δ_n imply that

$$g(\hat{\boldsymbol{\omega}}) = 0 \geq g(\boldsymbol{\omega}^{(n)}) + 2(\hat{\boldsymbol{\omega}} - \boldsymbol{\omega}^{(n)})^T \mathbf{K}(\boldsymbol{\omega}^{(n)} - \hat{\boldsymbol{\omega}}) \geq g(\boldsymbol{\omega}^{(n)}) + 2\Delta_n^T \mathbf{K}(\boldsymbol{\omega}^{(n)} - \hat{\boldsymbol{\omega}}).$$

We obtain the following property; the proof is identical to that of Lemma A.5.

Lemma A.7. *Suppose that $\hat{\boldsymbol{\omega}}$ and \mathbf{K} are such that $\hat{\alpha}_n \leq \omega_{i_n}^{(n)}$ for any $\boldsymbol{\omega}^{(n)} \in \mathbb{P}_\Omega$. Then, algorithm (A.8, A.9) with $\alpha_n = \hat{\alpha}_n$ given by (A.11), initialized at any $\boldsymbol{\omega}^{(1)} \in \mathbb{P}_\Omega$, satisfies*

$$\|\boldsymbol{\omega}^{(n)} - \hat{\boldsymbol{\omega}}\|_{\mathbf{K}}^2 \leq \frac{8 \lambda_{\max}(\mathbf{K})}{n+3}, \quad n \geq 1.$$

There exist situations where the condition $\hat{\alpha}_n \leq \omega_{i_n}^{(n)}$ is not satisfied. Take for instance $\Omega = 3$, \mathbf{K} the identity matrix and $\hat{\boldsymbol{\omega}} = (0, 0, 1)^T$, $\boldsymbol{\omega}^{(n)} = (1/3, 1/3, 1/3)^T$; then $\hat{\alpha}_n = 1/2 > \omega_{i_n}^{(n)} = 1/3$. On the other hand, the condition is satisfied for instance for $\hat{\boldsymbol{\omega}} = \mathbf{1}_\Omega/\Omega$ and \mathbf{K} the identity matrix (we have $i_{n-} = \text{Arg max}_{i: \omega_i^{(n)} > 0} (\omega_i^{(n)} - \hat{\omega}_i)$, and $\sum_{i=1}^\Omega (\omega_i^{(n)} - \hat{\omega}_i) = 0$ implies that $\omega_{i_{n-}}^{(n)} > \hat{\omega}_{i_{n-}}$ and similarly $\omega_{i_{n+}}^{(n)} < \hat{\omega}_{i_{n+}}$; we get $\hat{\alpha}_n = (\omega_{i_{n-}}^{(n)} - \omega_{i_{n+}}^{(n)})/2 \leq \omega_{i_{n-}}^{(n)}/2 < \omega_{i_{n-}}^{(n)}$, and numerical experiments indicate that it holds true in most situations.

Appendix B. Bayesian quadrature: several integrals.

Following [47], consider a generalization of the situation considered in Section 2.2 where one wishes to estimate

$$\mathbf{I}_\mu(f) = \mathbb{E}_\mu\{f(\mathbf{X})\mathbf{r}(\mathbf{X})\} = \int_{\mathcal{X}} f(\mathbf{x})\mathbf{r}(\mathbf{x}) \, d\mu(\mathbf{x}),$$

with $\mathbf{r}(\mathbf{x}) = (r_0(\mathbf{x}), \dots, r_p(\mathbf{x}))^T$ a vector of $p+1$ known functions of \mathbf{x} , such that the $(p+1) \times (p+1)$ matrix

$$\mathbf{M}_r = \mathbb{E}_\mu\{\mathbf{r}(\mathbf{X})\mathbf{r}^T(\mathbf{X})\}$$

exists and is nonsingular. Without any loss of generality, we may assume that $r_0(\mathbf{x}) \equiv 1$.

We also slightly generalize model (2.3) by introducing a linear trend $\mathbf{h}^T(\mathbf{x})\boldsymbol{\beta}$; that is, we consider

$$f(\mathbf{x}) = \mathbf{h}^T(\mathbf{x})\boldsymbol{\beta} + Z_x, \quad (\text{B.1})$$

where $\mathbf{h}(\mathbf{x}) = (h_0(\mathbf{x}), \dots, h_{p'}(\mathbf{x}))^T$ is a vector of $p'+1$ known functions of \mathbf{x} and $\boldsymbol{\beta} \in \mathbb{R}^{p'+1}$ has the normal prior $\mathcal{N}(\hat{\boldsymbol{\beta}}^0, \sigma^2 \mathbf{A})$, non-informative so that we can replace \mathbf{A}^{-1} by the null matrix $\mathbf{0}$ in all calculations (the choice of $\hat{\boldsymbol{\beta}}^0$ being then irrelevant). We assume that the matrix $\mathbb{E}_\mu\{\mathbf{h}(\mathbf{X})\mathbf{h}^T(\mathbf{X})\}$ is well defined. For reasons that will become clear below, we shall consider in particular the case where $\mathbf{h} = \mathbf{r}$.

The posterior mean and variance of $f(\mathbf{x})$, conditional on σ^2 and K , are now, respectively,

$$\begin{aligned} \hat{\eta}_n(\mathbf{x}) &= \mathbf{h}^T(\mathbf{x})\hat{\boldsymbol{\beta}}^n + \mathbf{k}_n^T(\mathbf{x})\mathbf{K}_n^{-1}(\mathbf{y}_n - \mathbf{H}_n\hat{\boldsymbol{\beta}}^n), \\ \sigma^2 \rho_n^2(\mathbf{x}) &= \sigma^2 \{K(\mathbf{x}, \mathbf{x}) - \mathbf{k}_n^T(\mathbf{x})\mathbf{K}_n^{-1}\mathbf{k}_n(\mathbf{x}) \\ &\quad + [\mathbf{h}(\mathbf{x}) - \mathbf{H}_n^T \mathbf{K}_n^{-1} \mathbf{k}_n(\mathbf{x})]^T (\mathbf{H}_n^T \mathbf{K}_n^{-1} \mathbf{H}_n)^{-1} [\mathbf{h}(\mathbf{x}) - \mathbf{H}_n^T \mathbf{K}_n^{-1} \mathbf{k}_n(\mathbf{x})]\}, \end{aligned}$$

where $\{\mathbf{H}_n(\mathbf{x})\}_{i,j} = h_j(\mathbf{x}_i)$, $i = 1, \dots, n$, $j = 0, \dots, p'$, and

$$\hat{\boldsymbol{\beta}}^n = (\mathbf{H}_n^T \mathbf{K}_n^{-1} \mathbf{H}_n)^{-1} \mathbf{H}_n^T \mathbf{K}_n^{-1} \mathbf{y}_n.$$

The posterior mean and covariance matrix of $\mathbf{I}_\mu(f)$ are

$$\begin{aligned} \hat{\mathbf{I}}_n &= \mathbf{B}(\mu) \hat{\boldsymbol{\beta}}^n + \mathbf{P}_n(\mu) \mathbf{K}_n^{-1} (\mathbf{y}_n - \mathbf{H}_n \hat{\boldsymbol{\beta}}^n), \\ \sigma^2 \mathbf{V}_n &= \sigma^2 \left\{ \mathbf{U}(\mu) - \mathbf{P}_n(\mu) \mathbf{K}_n^{-1} \mathbf{P}_n^T(\mu) \right. \\ &\quad \left. + [\mathbf{B}(\mu) - \mathbf{P}_n(\mu) \mathbf{K}_n^{-1} \mathbf{H}_n] (\mathbf{H}_n^T \mathbf{K}_n^{-1} \mathbf{H}_n)^{-1} [\mathbf{B}(\mu) - \mathbf{P}_n(\mu) \mathbf{K}_n^{-1} \mathbf{H}_n]^T \right\}, \end{aligned} \quad (\text{B.2})$$

$$+ [\mathbf{B}(\mu) - \mathbf{P}_n(\mu) \mathbf{K}_n^{-1} \mathbf{H}_n] (\mathbf{H}_n^T \mathbf{K}_n^{-1} \mathbf{H}_n)^{-1} [\mathbf{B}(\mu) - \mathbf{P}_n(\mu) \mathbf{K}_n^{-1} \mathbf{H}_n]^T \Big\}, \quad (\text{B.3})$$

where $\mathbf{B}(\mu) = \mathbb{E}_\mu\{\mathbf{r}(\mathbf{X})\mathbf{h}^T(\mathbf{X})\}$, $\mathbf{P}_n(\mu) = \mathbb{E}_\mu\{\mathbf{r}(\mathbf{X})\mathbf{k}_n^T(\mathbf{X})\}$ and

$$\mathbf{U}(\mu) = \mathbb{E}_\mu\{\mathbf{r}(\mathbf{X})\mathbf{r}^T(\mathbf{X}')K(\mathbf{X}, \mathbf{X}')\}, \quad (\text{B.4})$$

with \mathbf{X} and \mathbf{X}' i.i.d. $\sim \mu$.

Consider now the special choice $\mathbf{h} = \mathbf{r}$ in (B.1). Following Section 3.5.2, we can write $f(\mathbf{x}) = \mathbf{r}^T(\mathbf{x})\boldsymbol{\beta} + \mathcal{P}_r Z_x + (\text{Id}_{L^2} - \mathcal{P}_r)Z_x$, where \mathcal{P}_r denotes the orthogonal projection of $L^2(\mathcal{X}, \mu)$ onto the linear space spanned by $\mathbf{r}(\cdot)$; that is, $\mathcal{P}_r g(\mathbf{x}) = \mathbf{r}^T(\mathbf{x})\mathbf{M}_r^{-1} \int_{\mathcal{X}} \mathbf{r}(\mathbf{x}')g(\mathbf{x}')d\mu(\mathbf{x}')$ for all $g \in L^2(\mathcal{X}, \mu)$. This gives

$$\mathcal{P}_r Z_x = \mathbf{r}^T(\mathbf{x})\mathbf{M}_r^{-1} \int_{\mathcal{X}} \mathbf{r}(\mathbf{x}')Z_{x'}d\mu(\mathbf{x}').$$

In absence of prior information on $\boldsymbol{\beta}$ ($\mathbf{A}^{-1} = \mathbf{0}$), the prior on the parameters $\boldsymbol{\beta}' = \boldsymbol{\beta} + \mathbf{M}_r^{-1} \int_{\mathcal{X}} \mathbf{r}(\mathbf{x}')Z_{x'}d\mu(\mathbf{x}')$ remains non-informative, and the covariance kernel of $\tilde{Z}_x = (\text{Id}_{L^2} - \mathcal{P}_r)Z_x$ is

$$\begin{aligned} K_\mu(\mathbf{x}, \mathbf{x}') &= K(\mathbf{x}, \mathbf{x}') - \mathbf{u}_\mu^T(\mathbf{x})\mathbf{M}_r^{-1}\mathbf{r}(\mathbf{x}') - \mathbf{r}^T(\mathbf{x})\mathbf{M}_r^{-1}\mathbf{u}_\mu(\mathbf{x}') \\ &\quad + \mathbf{r}^T(\mathbf{x})\mathbf{M}_r^{-1}\mathbf{U}(\mu)\mathbf{M}_r^{-1}\mathbf{r}(\mathbf{x}'), \end{aligned}$$

where $\mathbf{U}(\mu)$ is given by (B.4) and $\mathbf{u}_\mu(\mathbf{x}) = \mathbb{E}_\mu\{\mathbf{r}(\mathbf{X})K(\mathbf{X}, \mathbf{x})\}$, $\mathbf{x} \in \mathcal{X}$.

Similarly to Remark 4.2 (see [26, Sect. 5.4]), this kernel reduction does not modify predictions, and direct calculation shows that $\mathbb{E}_\mu\{\mathbf{r}(\mathbf{X})\mathbf{r}^T(\mathbf{X}')K_\mu(\mathbf{X}, \mathbf{X}')\} = \mathbf{0}$ and $\mathbb{E}_\mu\{\mathbf{r}(\mathbf{X})\tilde{\mathbf{k}}_n^T(\mathbf{X})\} = \mathbf{0}$. We thus obtain the following property, where we denote by \mathbf{R}_n the $n \times (p+1)$ matrix $\{\mathbf{R}_n(\mathbf{x})\}_{i,j} = r_j(\mathbf{x}_i)$, $i = 1, \dots, n$, $j = 0, \dots, p$.

Lemma B.1. *When K is SPD, $\hat{\mathbf{I}}_n$ given by (B.2) satisfies*

$$\hat{\mathbf{I}}_n = \mathbf{M}_r(\mathbf{R}_n^T \tilde{\mathbf{K}}_n^{-1} \mathbf{R}_n)^{-1} (\mathbf{R}_n^T \tilde{\mathbf{K}}_n^{-1} \mathbf{y}_n),$$

and the posterior covariance matrix (B.3) satisfies

$$\mathbf{V}_n = \mathbf{M}_r(\mathbf{R}_n^T \tilde{\mathbf{K}}_n^{-1} \mathbf{R}_n)^{-1} \mathbf{M}_r. \quad (\text{B.5})$$

To ensure a precise estimation of $\mathbf{I}_\mu(f)$, we may select a design \mathbf{X}_n that minimizes $\mathcal{J}(\mathbf{V}_n)$, with $\mathcal{J}(\cdot)$ a Loewner increasing function defined on the set of symmetric non-negative definite matrices. Typical choices are $\mathcal{J}(\mathbf{V}_n) = \det(\mathbf{V}_n)$ (D-optimality) and $\mathcal{J}(\mathbf{V}_n) = \text{trace}(\mathbf{V}_n)$ (A-optimality). Greedy minimization of $\mathcal{J}(\mathbf{V}_n)$ corresponds to Sequential Bayesian Quadrature, see Section 4.2. Using (B.5) and formulae for the

inversion of a block matrix, we obtain the following expressions for $\det(\mathbf{V}_{n+1})$ and $\text{trace}(\mathbf{V}_{n+1})$:

$$\det(\mathbf{V}_{n+1}) = \det(\mathbf{V}_n) \frac{K_\mu(\mathbf{x}, \mathbf{x}) - \tilde{\mathbf{k}}_n^T(\mathbf{x}) \tilde{\mathbf{K}}_n^{-1} \tilde{\mathbf{k}}_n(\mathbf{x})}{\tilde{\rho}_n^2(\mathbf{x})},$$

$$\text{trace}(\mathbf{V}_{n+1}) = \text{trace}(\mathbf{V}_n) - \frac{[\mathbf{r}(\mathbf{x}) - \mathbf{R}_n^T \tilde{\mathbf{K}}_n^{-1} \tilde{\mathbf{k}}_n(\mathbf{x})]^T (\mathbf{R}_n^T \tilde{\mathbf{K}}_n^{-1} \mathbf{R}_n)^{-1} \mathbf{M}_r^2 (\mathbf{R}_n^T \tilde{\mathbf{K}}_n^{-1} \mathbf{R}_n)^{-1} [\mathbf{r}(\mathbf{x}) - \mathbf{R}_n^T \tilde{\mathbf{K}}_n^{-1} \tilde{\mathbf{k}}_n(\mathbf{x})]}{\tilde{\rho}_n^2(\mathbf{x})},$$

with

$$\tilde{\rho}_n^2(\mathbf{x}) = \left[K_\mu(\mathbf{x}, \mathbf{x}) - \tilde{\mathbf{k}}_n^T(\mathbf{x}) \tilde{\mathbf{K}}_n^{-1} \tilde{\mathbf{k}}_n(\mathbf{x}) + \frac{(1 - \tilde{\mathbf{k}}_n^T(\mathbf{x}) \tilde{\mathbf{K}}_n^{-1} \mathbf{1}_n)^2}{\mathbf{1}_n^T \tilde{\mathbf{K}}_n^{-1} \mathbf{1}_n} \right].$$

When $p = 0$ ($\mathbf{r}(\mathbf{x}) \equiv 1$), $\mathbf{V}_n = s_n^2$ in (4.9) and $\det(\mathbf{V}_{n+1}) = \text{trace}(\mathbf{V}_{n+1}) = s_{n+1}^2$ given by (4.10).

Appendix C. Karhunen-Loève decomposition and Bayesian c-optimal design.

As shown in Section 3.5.2, estimation of $I_\mu(f)$ is equivalent to estimation of β'_0 in the model (3.17). Following [26], we may consider the Karhunen-Loève decomposition of the Gaussian RF \tilde{Z}_x , which has zero mean and covariance $\mathbb{E}\{\tilde{Z}_x \tilde{Z}_{x'}\} = \sigma^2 K_\mu(\mathbf{x}, \mathbf{x}')$ and is such that $\int_\mu Z_x d\mu(\mathbf{x}) = 0$ a.s.; see also [21]. We thus write $f(\mathbf{x}) = \beta_0 + \sum_{k \geq 1} \beta_k \varphi_k(\mathbf{x}) + \tilde{Z}_{x,0}$, where the φ_k are orthonormal (in $L^2(\mathcal{X}, \mu)$) eigenfunctions of the spectral decomposition of the integral operator T_μ defined on $L^2(\mathcal{X}, \mu)$ by $T_\mu[f](\mathbf{x}) = \int_{\mathcal{X}} K_\mu(\mathbf{x}, \mathbf{x}') f(\mathbf{x}') d\mu(\mathbf{x}')$, $f \in L^2(\mathcal{X}, \mu)$, $\mathbf{x} \in \mathcal{X}$, with associated eigenvalues Λ_k (strictly positive and assumed to be ordered by decreasing values), and where the β_k , $k \geq 1$, are independent random variables distributed $\mathcal{N}(0, \sigma^2 \Lambda_k)$. Here $\tilde{Z}_{x,0}$ is a zero-mean Gaussian RF with covariance $\mathbb{E}\{\tilde{Z}_{x,0} \tilde{Z}_{x',0}\} = \sigma^2 \left[K_\mu(\mathbf{x}, \mathbf{x}') - \sum_{k \geq 1} \Lambda_k \varphi_k(\mathbf{x}) \varphi_k(\mathbf{x}') \right]$. Spectral truncation at $k = M - 1$ yields the approximated model

$$f(\mathbf{x}) \approx \beta'_0 + \sum_{k=1}^{M-1} \beta_k \varphi_k(\mathbf{x}) + \tilde{Z}_{x,1}, \quad (\text{C.1})$$

where $\tilde{Z}_{x,1}$ has zero mean and satisfies

$$\mathbb{E}\{\tilde{Z}_{x,1} \tilde{Z}_{x',1}\} = \sigma^2 \left[K_\mu(\mathbf{x}, \mathbf{x}') - \sum_{k=1}^{M-1} \Lambda_k \varphi_k(\mathbf{x}) \varphi_k(\mathbf{x}') \right].$$

The eigenfunctions φ_k satisfy $\int_{\mathcal{X}} \varphi_k(\mathbf{x}) d\mu(\mathbf{x}) = 0$ for all k . When evaluating f at \mathbf{x} outside the support of μ we replace all φ_k in (C.1) by their canonical extensions φ'_k defined by $\varphi'_k(\mathbf{x}) = (1/\Lambda_k) T_\mu[\varphi_k](\mathbf{x})$, for all $k \geq 1$ and all \mathbf{x} in \mathcal{X} ; they also satisfy $\int_{\mathcal{X}} \varphi'_k(\mathbf{x}) d\mu(\mathbf{x}) = 0$ for all k . Note that the φ'_k are defined over the whole \mathcal{X} , with $\varphi'_k = \varphi_k$ μ -almost everywhere, which is important when μ is only supported on a subset of \mathcal{X} .

The errors $\tilde{Z}_{x,1}$ in (C.1) are correlated. To facilitate the construction of experimental designs ensuring a precise estimation of β'_0 , an additional approximation is introduced which yields a Bayesian Linear Model (BLM) with M parameters

$\boldsymbol{\beta} = (\beta'_0, \beta_1, \dots, \beta_{M-1})^T$ and uncorrelated errors

$$f(\mathbf{x}) \approx \beta'_0 + \sum_{k=1}^{M-1} \beta_k \varphi_k(\mathbf{x}) + \tilde{Z}'_{x,1}. \quad (\text{C.2})$$

Here $\tilde{Z}'_{x,1}$ has zero mean and satisfies $\mathbb{E}\{\tilde{Z}_{x,1}\tilde{Z}'_{x',1}\} = \sigma^2 v(\mathbf{x})$ if $\mathbf{x}' = \mathbf{x}$ and is zero otherwise, with

$$v(\mathbf{x}) = \left[K_\mu(\mathbf{x}, \mathbf{x}) - \sum_{k=1}^{M-1} \Lambda_k \varphi_k(\mathbf{x}) \varphi_k(\mathbf{x}) \right].$$

After observation of $\mathbf{y}_n = [f(\mathbf{x}_1), \dots, f(\mathbf{x}_n)]^T$, $\boldsymbol{\beta}$ has the posterior normal distribution $\mathcal{N}(\hat{\boldsymbol{\beta}}^n, \sigma^2 \mathbf{M}_B^{-1}(\mathbf{X}_n))$, with

$$\mathbf{M}_B(\mathbf{X}_n) = \boldsymbol{\Psi}_n^T \boldsymbol{\Sigma}_n^{-1} \boldsymbol{\Psi}_n + \boldsymbol{\Lambda}^{-1},$$

where $\{\boldsymbol{\Psi}_n\}_{i,0} = 1$ and $\{\boldsymbol{\Psi}_n\}_{i,k} = \varphi_k(\mathbf{x}_i)$, $k = 1, \dots, M-1$, $i = 1, \dots, n$, $\boldsymbol{\Sigma}_n = \text{diag}\{v(\mathbf{x}_i), i = 1, \dots, n\}$, $\boldsymbol{\Lambda} = \text{diag}\{\infty, \Lambda_1, \dots, \Lambda_{M-1}\}$, $\hat{\boldsymbol{\beta}}^n = \mathbf{M}_B^{-1}(\mathbf{X}_n) \boldsymbol{\Psi}_n^T \boldsymbol{\Sigma}_n^{-1} \mathbf{y}_n$. According to the BLM (C.2), an optimal design for the estimation of $I_\mu(f)$, or equivalently for the estimation of β'_0 , is a Bayesian c-optimal design minimizing $\{\mathbf{M}_B^{-1}(\mathbf{X}_n)\}_{1,1}$; see [49]. The design criterion is different when one is interested in approximation rather than integration: as shown in [26], a design minimizing the integrated MSPE in this model is an A-optimal design minimizing $\text{trace}[\mathbf{M}_B^{-1}(\mathbf{X}_n)]$.

Approximate design theory can also be used. There, an optimal design ξ^* is a probability measure on \mathcal{X} that minimizes $\{\mathbf{M}_{B,m}^{-1}(\xi)\}_{1,1}$ with respect to $\xi \in \mathcal{M}^+(1)$, where

$$\mathbf{M}_{B,m}(\xi) = \int_{\mathcal{X}} \frac{1}{v(\mathbf{x})} \boldsymbol{\psi}(\mathbf{x}) \boldsymbol{\psi}^T(\mathbf{x}) d\xi(\mathbf{x}) + \boldsymbol{\Lambda}^{-1}/m, \quad m > 0, \quad (\text{C.3})$$

with $\boldsymbol{\psi}(\mathbf{x}) = (1, \varphi_1(\mathbf{x}), \dots, \varphi_M(\mathbf{x}))^T$. The scalar m defines a projected number of observations, in the sense that $\mathbf{M}_{B,n}(\xi_{n,e}) = (1/n) \mathbf{M}_B(\mathbf{X}_n)$ when $\xi_{n,e}$ is the empirical measure associated with \mathbf{X}_n . The determination of ξ^* forms a convex optimization problem [49] for which many algorithms are available; see, e.g., [51, Chap. 9]. Note that ξ^* depends both on m and M . When M is small enough, ξ^* has generally few support points; those points can be used for the construction of an exact design \mathbf{X}_n and an extraction procedure is described in [26, 50]. Choosing m and M of the same order of magnitude as the projected number n of observations is recommended. The construction of ξ^* can also be sequential, using a vertex-direction algorithm similar to those in Section 4.2. In general, the eigenfunctions φ_k and eigenvalues Λ_k corresponding to a given kernel K (here K_μ) are unknown and a numerical construction is required. It may rely on the substitution of a quadrature approximation for μ in the definition of the integral operator T_μ ; the construction is much facilitated when a tensor-product kernel can be used: the φ_k and Λ_k are then direct products of one-dimensional eigenfunctions and eigenvalues. In that case, the BLM (C.2) also allows for straightforward estimation of the so-called Sobol' indices used in sensitivity analysis; see [50].

Example. We present a simple illustrative example where $\mathcal{X} = [0, 1]^2$, μ is uniform on \mathcal{X} and K is the tensor product of Matérn 3/2 kernels (3.5) with $\theta = 2$. We consider designs supported on \mathcal{X}_Ω formed by a 32×32 regular grid in \mathcal{X} ; $m = M = 15$ in (C.3).

Thanks to the the tensor structure, numerical approximations of eigenfunctions and eigenvalues can be obtained through one-dimensional quadrature approximations; we use here the uniform distribution on the 100 points $0, 1/99, 2/99, \dots, 1$; see [50]. The optimal design measure ξ^* is constructed with a vertex-exchange algorithm [12] with optimal step size, initialized at the uniform design on \mathcal{X}_Ω . Iterations are stopped at iteration k when

$$\min_{\mathbf{x} \in \mathcal{X}_\Omega} \frac{\partial \left(\{ \mathbf{M}_{B,m}^{-1} [(1-\alpha)\xi^{(k)} + \alpha\delta_{\mathbf{x}}] \}_{1,1} \right)}{\partial \alpha} \Big|_{\alpha=0} > -10^{-6}.$$

Figure C.1-Left shows the corresponding measure $\xi^{(k)}$: there are 44 support points, their weights are proportional to the disk areas shown on the figure. Figure C.1-right presents the exact design \mathbf{X}_n extracted, with $n = 25$ points; the circles have radius equal to $\text{CR}(\mathbf{X}_n)$ and illustrate the good covering of \mathcal{X} by \mathbf{X}_{25} . In general, the adjustment of m and M to obtain a design \mathbf{X}_n with good space-filling properties for a prescribed value of n remains a difficult task and requires further investigations. \triangleleft




Figure C.1: Left: design measure ξ^* minimizing $\{ \mathbf{M}_{B,m}^{-1}(\xi) \}_{1,1}$ (44 support points with weights proportional to disk areas); Right: exact design \mathbf{X}_{25} extracted from ξ^* , the circles have radius $\text{CR}(\mathbf{X}_n)$.

Acknowledgements. The first author thanks the Isaac Newton Institute for Mathematical Sciences, Cambridge, for support and hospitality during the programme UQ for inverse problems in complex systems where work on this paper was partly undertaken. This work was supported by EPSRC grant no EP/K032208/1. We gratefully acknowledge Sébastien Da Veiga (Safran, Paris) who draw our attention to the machine learning literature on kernel herding and to connection with minimization of L_2 discrepancy.

REFERENCES

- [1] A. Antoniadis. Analysis of variance on function spaces. *Math. Operationsforsch. u. Statist.*, 15(1):59–71, 1984.
- [2] P. Audze and V. Eglais. New approach for planning out of experiments. *Problems of Dynamics and Strengths*, 35:104–107, 1977.
- [3] Y. Auffray, P. Barbillon, and J.-M. Marin. Maximin design on non hypercube domains and kernel interpolation. *Statistics and Computing*, 22(3):703–712, 2012.

- [4] F. Bach, S. Lacoste-Julien, and G. Obozinski. On the equivalence between herding and conditional gradient algorithms. *arXiv preprint arXiv:1203.4523*, 2012.
- [5] A. Beck and M. Teboulle. A conditional gradient method with linear rate of convergence for solving convex linear systems. *Mathematical Methods of Operations Research*, 59(2):235–247, 2004.
- [6] J. Bect, F. Bachoc, and D. Ginsbourger. A supermartingale approach to Gaussian process based sequential design of experiments. *arXiv preprint arXiv:1608.01118*, 2016.
- [7] J. Bect, D. Ginsbourger, L. Li, V. Picheny, and E. Vazquez. Sequential design of computer experiments for the estimation of a probability of failure. *Statistics and Computing*, 22(3):773–793, 2012.
- [8] A. Berlinet and C. Thomas-Agnan. *Reproducing Kernel Hilbert Spaces in Probability and Statistics*. Kluwer, Boston, 2004.
- [9] S. Biedermann and H. Dette. Minimax optimal designs for nonparametric regression — a further optimality property of the uniform distribution. In P. Hackl A.C. Atkinson and W.G. Müller, editors, *mODa’6 – Advances in Model-Oriented Design and Analysis, Proceedings of the 76th Int. Workshop, Puchberg/Schneeberg (Austria)*, pages 13–20, Heidelberg, June 2001. Physica Verlag.
- [10] G. Björck. Distributions of positive mass, which maximize a certain generalized energy integral. *Arkiv för Matematik*, 3(21):255–269, 1956.
- [11] D. Böhning. Numerical estimation of a probability measure. *Journal of Statistical Planning and Inference*, 11:57–69, 1985.
- [12] D. Böhning. A vertex-exchange-method in D -optimal design theory. *Metrika*, 33:337–347, 1986.
- [13] F.-X. Briol, C. Oates, M. Girolami, and M.A. Osborne. Frank-Wolfe Bayesian quadrature: Probabilistic integration with theoretical guarantees. In *Advances in Neural Information Processing Systems*, pages 1162–1170, 2015.
- [14] Y. Chen, M. Welling, and A. Smola. Super-samples from kernel herding. *arXiv preprint arXiv:1203.3472*, 2012.
- [15] K.M. Clarkson. Coresets, sparse greedy approximation, and the Frank-Wolfe algorithm. *ACM Transactions on Algorithms (TALG)*, 6(4):63, 2010.
- [16] S.B. Damelin, F.J. Hickernell, D.L. Ragozin, and X. Zeng. On energy, discrepancy and group invariant measures on measurable subsets of Euclidean space. *J. Fourier Anal. Appl.*, 16:813–839, 2010.
- [17] H. Dette, A. Pepelyshev, and A. Zhigljavsky. Best linear unbiased estimators in continuous time regression models. *arXiv preprint arXiv:1611.09804*, 2016.
- [18] P. Diaconis. Bayesian numerical analysis. *Statistical Decision Theory and Related Topics IV*, 1:163–175, 1988.
- [19] J.C. Dunn and S. Harshbarger. Conditional gradient algorithms with open loop step size rules. *Journal of Mathematical Analysis and Applications*, 62:432–444, 1978.
- [20] K.-T. Fang, R. Li, and A. Sudjianto. *Design and Modeling for Computer Experiments*. Chapman & Hall/CRC, Boca Raton, 2006.
- [21] V.V. Fedorov. Design of spatial experiments: model fitting and prediction. In S. Gosh and C.R. Rao, editors, *Handbook of Statistics, vol. 13*, chapter 16, pages 515–553. Elsevier, Amsterdam, 1996.
- [22] M. Frank and P. Wolfe. An algorithm for quadratic programming. *Naval Res. Logist. Quart.*, 3:95–110, 1956.
- [23] B. Fuglede. On the theory of potentials in locally compact spaces. *Acta mathematica*, 103:139–215, 1960.
- [24] B. Fuglede and N. Zorii. Green kernels associated with Riesz kernels. *Annales Academiae Scientiarum Fennicae, Mathematica*, 43(to appear), 2018. arXiv preprint arXiv:1610.00268.
- [25] B. Gauthier and L. Pronzato. Spectral approximation of the IMSE criterion for optimal designs in kernel-based interpolation models. *SIAM/ASA J. Uncertainty Quantification*, 2:805–825, 2014. DOI 10.1137/130928534.
- [26] B. Gauthier and L. Pronzato. Convex relaxation for IMSE optimal design in random field models. *Computational Statistics and Data Analysis*, 113:375–394, 2017.
- [27] D. Ginsbourger. Sequential design of computer experiments. *Wiley StatsRef*, 99:1–11, 2017.
- [28] D. Ginsbourger, O. Roustant, D. Schuhmacher, N. Durrande, and N. Lenz. On ANOVA decompositions of kernels and Gaussian random field paths. *preprint arXiv:1409.6008*, 2014.
- [29] T.F. Gonzalez. Clustering to minimize the maximum intercluster distance. *Theoretical Computer Science*, 38:293–306, 1985.
- [30] A. Gorodetsky and Y. Marzouk. Mercer kernels and integrated variance experimental design: connections between Gaussian process regression and polynomial approximation. *SIAM/ASA J. Uncertainty Quantification*, 4(1):796–828, 2016.

- [31] U. Grenander. Stochastic processes and statistical inference. *Arkiv för Matematik*, 1(3):195–277, 1950.
- [32] D.P. Hardin and E.B. Saff. Discretizing manifolds via minimum energy points. *Notices of the AMS*, 51(10):1186–1194, 2004.
- [33] P. Hennig, M.A. Osborne, and M. Girolami. Probabilistic numerics and uncertainty in computations. *Proc. Royal Soc. A*, 471(2179):20150142, 2015.
- [34] F.J. Hickernell. A generalized discrepancy and quadrature error bound. *Mathematics of Computation*, 67(221):299–322, 1998.
- [35] F. Huszár and D. Duvenaud. Optimally-weighted herding is Bayesian quadrature. *arXiv preprint arXiv:1204.1664*, 2012.
- [36] M.E. Johnson, L.M. Moore, and D. Ylvisaker. Minimax and maximin distance designs. *Journal of Statistical Planning and Inference*, 26:131–148, 1990.
- [37] V.R. Joseph, E. Gul, and S. Ba. Maximum projection designs for computer experiments. *Biometrika*, 102(2):371–380, 2015.
- [38] N.S. Landkof. *Foundations of Modern Potential Theory*. Springer, Berlin, 1972.
- [39] Q. Liu and D. Wang. Stein variational gradient descent: a general purpose Bayesian inference algorithm. *arXiv preprint arXiv:1608.04471v2*, 2016.
- [40] S. Mak and V.R. Joseph. Projected support points, with application to optimal MCMC reduction. *arXiv preprint arXiv:1708.06897*, 2017.
- [41] S. Mak and V.R. Joseph. Support points. *arXiv preprint arXiv:1609.01811*, 2017. To appear in the *Annals of Statistics*.
- [42] I. Molchanov and S. Zuyev. Variational calculus in the space of measures and optimal design. In A. Atkinson, B. Bogacka, and A. Zhigljavsky, editors, *Optimum Design 2000*, chapter 8, pages 79–90. Kluwer, Dordrecht, 2001.
- [43] I. Molchanov and S. Zuyev. Steepest descent algorithm in a space of measures. *Statistics and Computing*, 12:115–123, 2002.
- [44] W.G. Müller. Coffee-house designs. In A. Atkinson, B. Bogacka, and A. Zhigljavsky, editors, *Optimum Design 2000*, chapter 21, pages 241–248. Kluwer, Dordrecht, 2001.
- [45] W. Näther. *Effective Observation of Random Fields*, volume 72. Teubner-Texte zur Mathematik, Leipzig, 1985.
- [46] H. Niederreiter. *Random Number Generation and Quasi-Monte Carlo Methods*. SIAM, Philadelphia, 1992.
- [47] A. O’Hagan. Bayes-Hermite quadrature. *Journal of Statistical Planning and Inference*, 29(3):245–260, 1991.
- [48] V.I. Paulsen. An introduction to the theory of reproducing kernel Hilbert spaces, 2009. <https://www.math.uh.edu/~vern/rkhs.pdf>.
- [49] J. Pilz. *Bayesian Estimation and Experimental Design in Linear Regression Models*, volume 55. Teubner-Texte zur Mathematik, Leipzig, 1983. (also Wiley, New York, 1991).
- [50] 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*, 2018. to appear, hal-01545604v2.
- [51] L. Pronzato and A. Pázman. *Design of Experiments in Nonlinear Models. Asymptotic Normality, Optimality Criteria and Small-Sample Properties*. Springer, LNS 212, New York, 2013.
- [52] L. Pronzato, H.P. Wynn, and A. Zhigljavsky. Extremal measures maximizing functionals based on simplicial volumes. *Statistical Papers*, 57(4):1059–1075, 2016. hal-01308116.
- [53] C.R. Rao. Diversity and dissimilarity coefficients: a unified approach. *Theoret. Popn Biol.*, 21(1):24–43, 1982.
- [54] C.R. Rao and T.K. Nayak. Cross entropy, dissimilarity measures and characterizations of quadratic entropy. *IEEE Transactions on Information Theory*, 31(5):589–593, 1985.
- [55] K. Ritter, G.W. Wasilkowski, and H. Woźniakowski. Multivariate integration and approximation for random fields satisfying Sacks-Ylvisaker conditions. *The Annals of Applied Probability*, 5(2):518–540, 1995.
- [56] J. Sacks, W.J. Welch, T.J. Mitchell, and H.P. Wynn. Design and analysis of computer experiments. *Statistical Science*, 4(4):409–435, 1989.
- [57] T.J. Santner, B.J. Williams, and W.I. Notz. *The Design and Analysis of Computer Experiments*. Springer, Heidelberg, 2003.
- [58] R. Schaback. Error estimates and condition numbers for radial basis function interpolation. *Advances in Computational Mathematics*, 3(3):251–264, 1995.
- [59] R. Schaback. Native Hilbert spaces for radial basis functions I. In *New Developments in Approximation Theory*, pages 255–282. Springer, 1999.
- [60] I.J. Schoenberg. Metric spaces and positive definite functions. *Transactions of the American*

- Mathematical Society*, 44(3):522–536, 1938.
- [61] S. Sejdinovic, B. Sriperumbudur, A. Gretton, and K. Fukumizu. Equivalence of distance-based and RKHS-based statistics in hypothesis testing. *The Annals of Statistics*, 41(5):2263–2291, 2013.
 - [62] B.K. Sriperumbudur, A. Gretton, K. Fukumizu, B. Schölkopf, and G.R.G. Lanckriet. Hilbert space embeddings and metrics on probability measures. *Journal of Machine Learning Research*, 11(Apr):1517–1561, 2010.
 - [63] M.L. Stein. *Interpolation of Spatial Data. Some Theory for Kriging*. Springer, Heidelberg, 1999.
 - [64] I. Steinwart, D. Hush, and C. Scovel. An explicit description of the reproducing kernel Hilbert spaces of Gaussian RBF kernels. *IEEE Transactions on Information Theory*, 52(10):4635–4643, 2006.
 - [65] Z. Szabó and B. Sriperumbudur. Characteristic and universal tensor product kernels, 2017. Preprint hal-01585727.
 - [66] G.J. Székely and M.L. Rizzo. Energy statistics: A class of statistics based on distances. *Journal of Statistical Planning and Inference*, 143(8):1249–1272, 2013.
 - [67] E. Vazquez and J. Bect. Sequential search based on kriging: convergence analysis of some algorithms. Proc. 58th World Statistics Congress of the ISI, August 21-26, Dublin, Ireland, arXiv preprint arXiv:1111.3866v1, 2011.
 - [68] H.P. Wynn. The sequential generation of D -optimum experimental designs. *Annals of Math. Stat.*, 41:1655–1664, 1970.