Incremental space-filling design based on coverings and spacings: improving upon low discrepancy sequences

A. Nogales Gómez, L. Pronzato, M.-J. Rendas

UCA, CNRS, Laboratoire I3S, UMR 7271,
Bât. Euclide, Les Algorithmes, 2000 route des lucioles,
06900 Sophia Antipolis cedex, France
{amaya.nogales-gomez,pronzato,rendas}@i3s.unice.fr

May 26, 2021

Abstract

The paper addresses the problem of defining families of ordered sequences $\{\mathbf{x}_i\}_{i \in \mathbb{N}}$ of elements of a compact subset $\mathcal{X}$ of $\mathbb{R}^d$ whose prefixes $\mathbf{X}_n = \{\mathbf{x}_i\}_{i=1}^n$, for all orders n , have good space-filling properties as measured by the dispersion (covering radius) criterion. Our ultimate aim is the definition of incremental algorithms that generate sequences $\mathbf{X}_n$ with small optimality gap, i.e., with a small increase in the maximum distance between points of $\mathcal{X}$ and the elements of $\mathbf{X}_n$ with respect to the optimal solution $\mathbf{X}_n^*$. The paper is a first step in this direction, presenting incremental design algorithms with proven optimality bound for one-parameter families of criteria based on coverings and spacings that both converge to dispersion for large values of their parameter. The examples presented show that the covering-based method outperforms state-of-the-art competitors, including coffee-house, suggesting that it inherits from its guaranteed 50% optimality gap.

keywords Covering; Spacing; Submodularity; Greedy algorithm; Computer experiments; Space-filling design

MSC 62K99, 65D99

1 Introduction and motivation

The paper discusses algorithmic constructions to define space-filling designs. Given a compact subset $\mathcal{X} \subset \mathbb{R}^d$, we say that a finite subset $\mathbf{Z}_n \subset \mathcal{X}$ is a space-filling design if $\mathbf{Z}_n$ fills $\mathcal{X}$ evenly. Several mathematical definitions of this intuitive notion have been proposed in the literature, and we refer the interested reader to the reviews [28, 30] and the books [6], [35, Chap. 5] for a comprehensive presentation and discussion. Generically, a space-filling criterion is a set function, that maps (possibly finite) subsets of $\mathcal{X}$ to $\mathbb{R}$.

In this work, the space-filling quality of a design $\mathbf{Z}_n$ is measured through the distances from the domain points to their closest point in $\mathbf{Z}_n$, i.e., their distance to the design. Motivated by interpolation problems, the covering radius, equal to the maximum of these distances, is often considered to be the ultimate measure of the space-filling quality of $\mathbf{Z}_n$: good space-filling sequences should have small covering radius. In colloquial terms, “good designs leave no large holes in $\mathcal{X}$ ”.

The packing radius is in some sense a dual space-filling metric, assessing how well $\mathbf{Z}_n$ is spread inside $\mathcal{X}$ through the distance between its closest points. In terms of this criterion, a design is space-filling if it has a large packing radius. It is easy to see that the designs optimal for the covering radius are not optimal according to the packing radius: the latter must have some points in the boundary of $\mathcal{X}$ (otherwise a larger packing radius would be obtained by expanding the design uniformly), while this clearly does not lead to minimal covering radius. Nevertheless, the packing radius is widely used as a space-filling criterion, given the much smaller numerical complexity involved in its computation when compared to the covering radius, and in the paper we also consider the optimization of the packing radius.

Both the covering and packing radii have important limitations as space-filling criteria, since they only reflect local characteristics. Indeed, they are generally determined by a few points of $\mathbf{Z}_n$, being insensitive to perturbations of the other points, with the consequence that many extensions of a given design may keep these indicators constant. Moreover, when d is large the numerical estimation of the covering radius is difficult and, as we shall see, very sensitive to the presence of vertices in $\mathcal{X}$. As already discussed in [26], the α -quantile of the distribution of the distance of a random point of $\mathcal{X}$ to $\mathbf{Z}_n$, which we shall call the covering α -quantile of the design, with α close to 1, is a more stable space-filling indicator. For completeness, we will also compare designs according to this performance measure.

Most literature on space-filling design considers the construction of designs of a given target size. We are instead interested in the more complex situation where the number of design points that will effectively be used, $n \leq n_{\max}$, is to be decided in a later step, and our goal is to specify an ordered design $\mathbf{X}_{n_{\max}}$ such that all prefix designs $\mathbf{X}_n = \{\mathbf{x}_i\}_{i=1}^n$, $n \leq n_{\max}$, have good space-filling properties. Such a situation typically occurs in the construction of the initial design on which the metamodel subsequently used in a computer experiment will be identified. In this context, performance is no longer assessed through a single scalar, but rather by the trajectory of criterion values for increasing design sizes.

Background. Optimization of the covering radius is a highly non-linear (the cost function involves a maximum) and inherently multi-dimensional (the search space has dimension $n \times d$) combinatorial problem. Actually, the problem is known to be NP-hard, see [16, p. 414], meaning that except for toy-problems we can only hope to find tractable algorithms that produce reasonably good solutions. In these circumstances, it is important to know how far the solutions found by a given algorithm can be from optimality, commonly designated in algorithmics by optimality gap. The definition of algorithms with guaranteed optimality gap for NP-hard problems is an active research topic which has produced important achievements for many combinatorial optimization problems, and our work is a contribution in this sense.

Algorithms which define the solution by incrementally adding a point at a time are specially interesting given their more manageable complexity: at each iteration an optimization problem in $\mathcal{X}$, i.e., in only d variables, needs to be solved. It is a well known — and easily verified — fact that besides additive cost functions, for which the optimal solution can be found by greedily appending the point in $\mathcal{X}$ that produces the largest improvement, the optimality gap of greedy algorithms can be bounded if the criterion optimized is a submodular set function (this notion will be made precise below). Unfortunately, neither the covering radius nor the packing radius are submodular set functions.

Claims and hint of the contents. In this paper, we present a parameterized submodular design criterion, based of the cumulative distribution function (c.d.f.) of the distance of a random point to the design, which is asymptotically (for large values of its parameter) related to the covering radius, submodularity implying that its greedy solutions have a bounded optimality gap. We show numerically that the performance of the designs produced is robust with respect to the precise tuning parameters of the criterion. Interestingly, our construction is free from knowledge of the detailed geometry of the domain $\mathcal{X}$, requiring only the specification of two (possibly coincident) sets: a set of eligible design points and a set of points used to compute the c.d.f., which must finely cover $\mathcal{X}$.

Despite its simplicity, the incremental construction of designs by greedy maximization of the packing radius, sometimes called coffee-house design in the literature [22, 23], ideally ensures an optimality gap of 50% for the covering radius [10]. Practical constructions use a finite set of candidates $\mathcal{X}_C$ at each iteration, and the 50% guarantee above holds for $\mathcal{X}_C$ only. As explained in the paper, the quality of the resulting design heavily depends on the imbedding of $\mathcal{X}_C$ in $\mathcal{X}$. We show that a slight modification of the original coffee-house algorithm based on the notion of spacings, which keeps the design points away from the boundary of $\mathcal{X}$ with the aim of reducing the covering radius, improves robustness with respect to the choice of $\mathcal{X}_C$.

Our numerical studies indicate that the greedy optimization of our c.d.f.-based criterion yields design with better covering properties than all alternative techniques tested: low-discrepancy sequences (LDS), variants of coffee-house design, and incremental constructions based on other relaxed versions of the covering radius¹. In particular, when the domain is a d -dimensional hypercube and performance is measured by the covering radius, it always outperforms the other methods.

Paper organization. This paper extends preliminary work presented at the SIAM Conference on Uncertainty Quantification in 2016 [31]. It is organized as follows. In Section 2 we recall the definition of submodular set functions and the fundamental theorem that establishes a bound on the optimality gap of their greedy solutions. Section 3 presents the novel c.d.f.-based design criterion, addressing in detail its numerical evaluation and illustrating its behavior for several parametrization choices. Other greedy constructions, based on packing radius and spacings, are presented in Section 4. Finally, Section 5 presents a comparative performance study of the proposed methods against a set of state-of-the-art space-filling design methods.

Basic definitions and notation. Throughout the paper, $\mathcal{X}$ is a compact subset of $\mathbb{R}^d$ with nonempty interior, with the hypercube $\mathcal{C}_d = [0, 1]^d$ as a typical example. $\mathbf{1}_d$ denotes the d -dimensional vector with all components equal to one, and the center of $\mathcal{C}_d$ is thus $\mathbf{1}_d/2$. Let $\mathbf{Z}_n = \{\mathbf{z}_1, \dots, \mathbf{z}_n\}$ be any n -point design in $\mathcal{X}$. For any point $\mathbf{x} \in \mathcal{X}$, we denote $d(\mathbf{x}, \mathbf{Z}_n) = \min_{i=1, \dots, n} \|\mathbf{x} - \mathbf{z}_i\|$, with $\|\cdot\|$ the ℓ_2 norm. The covering radius $\text{CR}(\mathbf{Z}_n)$ of $\mathbf{Z}_n$ is defined by

$$\text{CR}(\mathbf{Z}_n) = \text{CR}_{[\mathcal{X}]}(\mathbf{Z}_n) = \max_{\mathbf{x} \in \mathcal{X}} d(\mathbf{x}, \mathbf{Z}_n).$$

¹There exist other incremental constructions that generate space-filling designs with good covering performance, in particular those based on maximization of mutual information [2, 17] or on minimization of a kernel discrepancy by kernel herding [33]. They rely on the choice of a suitable positive definite kernel and are especially adapted to interpolation based on Gaussian process models. We do not follow this approach here, and the methods we consider are based only on geometrical considerations. A thorough comparison would certainly be of interest, but is beyond the scope of this paper.

It is called dispersion in the theory of quasi-Monte Carlo methods [25, Chap. 6] and coincides with the minimax-distance, or fill criterion, used in computer experiments and function interpolation; see [5, 14, 28, 30]. When the design objective is to ensure an accurate prediction of the values of an unknown function f over all $\mathcal{X}$ based on evaluations at $\mathbf{Z}_n \in \mathcal{X}$, it is important to ensure that for any $\mathbf{x}$ in $\mathcal{X}$ there always exists a $\mathbf{z}_i$ at proximity of $\mathbf{x}$ where $f(\mathbf{z}_i)$ has been evaluated (see, e.g., [36] for a precise formulation and error bounds), making designs with a small value of $\text{CR}(\mathbf{Z}_n)$ particularly desirable.

Let $\mathcal{X}_N = \{\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(N)}\}$ be a finite subset of $\mathcal{X}$, with $N \gg n$, that is well spread over $\mathcal{X}$; it may be a regular grid when $\mathcal{X} = \mathcal{C}_d$ and d is small, or (possibly after a suitable rescaling of $\mathcal{X}$) the first N points of a LDS in $\mathcal{C}_d$ that fall in $\mathcal{X}$. With a slight abuse of notation, we assume that $\mathcal{X}_N$ and $\mathcal{X}_{N'}$ do not necessarily share any elements when $N \neq N'$. Note that the value

$$\text{CR}_{[\mathcal{X}_N]}(\mathbf{Z}_n) = \max_{\mathbf{x} \in \mathcal{X}_N} d(\mathbf{x}, \mathbf{Z}_n) \quad (1)$$

underestimates $\text{CR}(\mathbf{Z}_n)$ by an amount depending on how well $\mathcal{X}_N$ is itself spread in $\mathcal{X}$. In particular, we have

$$\text{CR}_{[\mathcal{X}_N]}(\mathbf{Z}_n) \leq \text{CR}(\mathbf{Z}_n) \leq \max_{\mathbf{x} \in \mathcal{X}} \min_{\mathbf{z}_i \in \mathbf{Z}_n, \mathbf{x}^{(j)} \in \mathcal{X}_N} \left(\|\mathbf{x} - \mathbf{x}^{(j)}\| + \|\mathbf{x}^{(j)} - \mathbf{z}_i\| \right) \leq \text{CR}_{[\mathcal{X}_N]}(\mathbf{Z}_n) + \text{CR}(\mathcal{X}_N). \quad (2)$$

In the appendix we give an upper bound on $\text{CR}(\mathcal{X}_N)$ when $\mathcal{X}_N$ is the prefix of a LDS.

The packing radius (also designated by separating radius, or maximin-distance)

$$\text{PR}(\mathbf{Z}_n) = \min_{\mathbf{z}_i \neq \mathbf{z}_j \in \mathbf{Z}_n} \frac{1}{2} \|\mathbf{z}_i - \mathbf{z}_j\| \quad (n \geq 2),$$

is also often used as a space-filling characteristic of a design $\mathbf{Z}_n$ [5, 14, 28, 30], in particular due to the simplicity of its calculation when compared to $\text{CR}(\mathbf{Z}_n)$. Good designs should have a large packing radius. Notice, however, that $\text{PR}(\mathbf{Z}_n)$ is an even more local space-filling characteristic of $\mathbf{Z}_n$ than $\text{CR}(\mathbf{Z}_n)$, in the sense that moving a single point $\mathbf{z}_i$ of $\mathbf{Z}_n$ to make it coincide with another $\mathbf{z}_j$ sets $\text{PR}(\mathbf{Z}_n)$ to zero whereas $\text{CR}(\mathbf{Z}_n)$ only increases to $\text{CR}(\mathbf{Z}_n \setminus \{\mathbf{z}_i\})$.

We reserve the notation $\mathbf{X}_n = [\mathbf{x}_1, \dots, \mathbf{x}_n]$ to designs that are constructed incrementally, our objective being that all prefix designs $\mathbf{X}_m$, $m = 1, \dots, n$, should have good space-filling properties. A special notation is used for the Sobol' and Halton sequences, with $\mathbf{S}_{n,d}$ and $\mathbf{H}_{n,d}$ respectively denoting their first n points in dimension d .

When using Gaussian process modeling and kriging, the number of design points required for a given prediction accuracy can be related to the correlation length of the model [12, 29]. However, this requires prior knowledge on the behavior of the function to be approximated, or that preliminary experimentation be performed with an initial design. When this initial design is suitably ordered, function evaluations can be interrupted when the confidence in the identified model is deemed sufficient to proceed to a second design phase, where complementary design points are added to improve the precision of the model predictions. In this situation, the size n of the initial design is small, much smaller than 2^d when d is large, and in particular we can have in mind the “10 d ” rule of [19].

Interest in incremental constructions brings an important constraint to the optimization problem, since designs with minimum covering radius are not nested. The consequences are well known in the large n situation: for $d = 1$, remember for instance the low dispersion sequence of

Ruzsa whose discrepancy does not tend to zero as $n \rightarrow \infty$ [25, p. 154]. But the phenomenon is already present for small n : for $\mathcal{X} = \mathcal{C}_d$ the one-point CR-optimal design is $\mathbf{Z}_1^* = \{\mathbf{1}_d/2\}$ (the center of $\mathcal{C}_d$), whereas 2-point CR-optimal designs have the form $\mathbf{Z}_2^* = \{\mathbf{z}_1^*, \mathbf{z}_2^* = \mathbf{1}_d - \mathbf{z}_1^*\}$ with $\mathbf{z}_1^*$ having all its coordinates equal to $1/2$ except one which equals $1/4$; see the appendix. Any incremental construction is therefore already suboptimal for $\text{CR}(\mathbf{X}_n)$ at $n = 2$ (and $\text{CR}(\mathbf{X}_n)$ can remain equal to $\sqrt{d}/2$ for some iterations when the first design points are chosen near the vertices of $\mathcal{C}_d$; see [7]).

2 Performance guarantee for submodular function maximization

This section recalls the notion of submodular set functions and a well known theorem establishing, when they are non-decreasing, bounds on the optimality gap of their greedy maximization.

2.1 Submodularity and the greedy algorithm

Let $\mathcal{X}_C$ be a finite candidate set with $C = |\mathcal{X}_C|$ elements and denote by $2^{\mathcal{X}_C}$ its power set. A (scalar real valued) set-function is an application of $2^{\mathcal{X}_C}$ on $\mathbb{R}$. Space-filling criteria are set-functions. We say that a set-function f is non-decreasing when $f(\mathbf{A} \cup \{\mathbf{x}\}) \geq f(\mathbf{A})$ for any $\mathbf{A} \subset \mathcal{X}_C$ and any $\mathbf{x} \in \mathcal{X}_C$.

Definition 1 (Submodular function) *A set-function $f : 2^{\mathcal{X}_C} \rightarrow \mathbb{R}$ is submodular if and only if it satisfies the following three equivalent conditions, see e.g., [1].*

- (a) $f(\mathbf{A}) + f(\mathbf{B}) \geq f(\mathbf{A} \cup \mathbf{B}) + f(\mathbf{A} \cap \mathbf{B}), \quad \forall \mathbf{A}, \mathbf{B} \in 2^{\mathcal{X}_C};$
- (b) $f(\mathbf{A} \cup \{\mathbf{x}\}) - f(\mathbf{A}) \geq f(\mathbf{B} \cup \{\mathbf{x}\}) - f(\mathbf{B}), \quad \forall \mathbf{A} \subset \mathbf{B} \in 2^{\mathcal{X}_C}, \mathbf{x} \in \mathcal{X}_C \setminus \mathbf{B};$
- (c) $f(\mathbf{A} \cup \{\mathbf{x}\}) - f(\mathbf{A}) \geq f(\mathbf{A} \cup \{\mathbf{x}, \mathbf{y}\}) - f(\mathbf{A} \cup \{\mathbf{y}\}), \quad \forall \mathbf{A}, \mathbf{B} \in 2^{\mathcal{X}_C}, \mathbf{x}, \mathbf{y} \in \mathcal{X}_C \setminus \mathbf{A}.$

Inequality (3) is known as the *diminishing returns* property, stating that the increment resulting of the addition of an element $\mathbf{x}$ to a set is a decreasing function of the set to which it is added.

The problem of maximizing a set-function is in general NP-hard, its exact resolution requiring the evaluation of f over all the 2^C elements of the power set, which is infeasible except in trivial cases of little practical interest. In this paper we concentrate on the greedy (one-step-ahead) algorithm below to efficiently find approximate maximizers of size k of a given set-function f .

Algorithm 1 (Greedy Algorithm)

```

1: set  $\mathbf{X} = \emptyset$ 
2: while  $|\mathbf{X}| < k$  do
3:   find  $\mathbf{x}$  in  $\mathcal{X}_C$  such that  $f(\mathbf{X} \cup \{\mathbf{x}\})$  is maximal
4:    $\mathbf{X} \leftarrow \mathbf{X} \cup \{\mathbf{x}\}$ 
5: end while
6: return  $\mathbf{X}$ 

```

When several solutions exist at Step 3, a single one is selected (e.g. randomly). In spite of its simplicity, under mild conditions on f Algorithm 1 can be fairly efficient when applied to a non-decreasing submodular function as stated by the following theorem [24].

Theorem 1 *Let f be a non-decreasing submodular set function, then, for any given k , $1 \leq k \leq C$, Algorithm 1 returns a set $\mathbf{X}_k$ with bounded optimality gap*

$$\frac{f_k^* - f(\mathbf{X}_k)}{f_k^* - f(\emptyset)} \leq (1 - 1/k)^k \leq 1/e < 0.3679, \quad (4)$$

where $f_k^* = \max_{\mathbf{X} \subset \mathcal{X}_C: |\mathbf{X}| \leq k} f(\mathbf{X})$ and $e = \exp(1)$.

Note that (4) implies the efficiency bound

$$\frac{f(\mathbf{X}_k) - f(\emptyset)}{f_k^* - f(\emptyset)} \geq 1 - \frac{1}{e} > 63.2\%, \quad \forall k \geq 1. \quad (5)$$

2.2 The lazy-greedy algorithm

At each iteration of the greedy algorithm the maximum of $f(\mathbf{X} \cup \{\mathbf{x}\})$ over all $\mathbf{x} \in \mathcal{X}_C \setminus \mathbf{X}$ must be found. As presented in [21], the submodularity of f can be further exploited to restrict actual evaluation of f to a proper subset of $\mathcal{X}_C \setminus \mathbf{X}$ (i.e., strictly included in $\mathcal{X}_C \setminus \mathbf{X}$). Although the worst-case complexity of the modified algorithm is still $\mathcal{O}(kC)$, important gains in computational complexity are observed in practice; see Figure 1 for an illustration.

Denote by $\delta_{\mathbf{X}}(\mathbf{x}) = f(\mathbf{X} \cup \{\mathbf{x}\}) - f(\mathbf{X}) \geq 0$ the improvement of f when $\mathbf{x}$ is added to $\mathbf{X}$, and let $\mathbf{X}_n$ be the greedy solution at iteration n . Also, remark that maximization of $f(\mathbf{X}_n \cup \{\mathbf{x}\})$ in Step 3 of Algorithm 1 is equivalent to maximization of $\delta_{\mathbf{X}_n}(\mathbf{x})$.

By construction, $\mathbf{X}_i \subset \mathbf{X}_n$ for all $i < n$. Then, since f is submodular, for all $i < n$, $\delta_{\mathbf{X}_n}(\mathbf{x}) \leq \delta_{\mathbf{X}_i}(\mathbf{x})$; i.e., for each $\mathbf{x}$ the increments $\delta_{\mathbf{X}_i}(\mathbf{x})$ decrease from iteration to iteration, as the size of $\mathbf{X}_i$ grows. At the first iteration, with $\mathbf{X}_0 = \emptyset$, we compute all $\delta_{\mathbf{X}_0}(\mathbf{x})$ for all $\mathbf{x} \in \mathcal{X}_C$, establishing for each $\mathbf{x}$ an upper bound $\bar{\delta}(\mathbf{x})$ on $\delta_{\mathbf{X}_k}(\mathbf{x})$ at subsequent iterations. These upper bounds are updated as follows.

Consider, at iteration k , scanning of the set $\mathcal{X}_C \setminus \mathbf{X}_{k-1}$ to compute the solution of Step 3 of Algorithm 1. Let $\mathcal{L}_{k-1} \subset \mathcal{X}_C$ denote the set of points that are possible solutions of Step 3. Initialize $\mathcal{L}_{k-1} = \mathcal{X}_C \setminus \mathbf{X}_{k-1}$, and let $\mathbf{x}_k^{**}$ be its member with largest $\bar{\delta}(\mathbf{x})$. While $\mathcal{L}_{k-1} \neq \emptyset$, update $\bar{\delta}(\mathbf{x}_k^{**}) = \delta_{\mathbf{X}_{k-1}}(\mathbf{x}_k^{**})$ and remove from $\mathcal{L}_{k-1}$ all $\mathbf{x}$ whose upper bound is smaller than or equal to $\bar{\delta}(\mathbf{x}_k^{**})$. When $\mathcal{L}_{k-1} = \emptyset$, update $\mathbf{X}_{k-1}$ into $\mathbf{X}_k = \mathbf{X}_{k-1} \cup \{\mathbf{x}_k^{**}\}$.

This modified algorithm trades memory for computational power, requiring the storage of the most recently updated values of $\delta_{\mathbf{X}}(\mathbf{x})$ for all $\mathbf{x} \in \mathcal{X}_C \setminus \mathbf{X}$.

The left panel of Figure 1 shows the computational times² of the greedy (solid lines) and lazy-greedy (dashed lines) maximizations³ of the non-decreasing submodular criterion $\hat{I}_{B,q}$ that will be introduced in Section 3.3; see (10). Notice the tremendous acceleration provided by the lazy-greedy implementation, contrasting with the fast linear increase of computational cost with n for the greedy version. We can understand the lazy algorithm as inducing a decrease of the effective size of the candidate set $\mathcal{X}_C$ at each iteration, from C to γC , $\gamma < 1$. Denote by m_k the number of updates of $\mathbf{x}_k^{**}$ at iteration k (with therefore $m_1 = C$), and let $\gamma_k = m_k/C$: the effective size of $\mathcal{X}_C$ for the lazy-greedy algorithm is $m_k = C \gamma_k$ at iteration k and $\bar{\gamma}_n C = (C/n) \sum_{k=1}^n \gamma_k$ in n iterations (i.e., for a n -point design). The right panel of Figure 1 shows γ_k as a function of k for the maximization of $\hat{I}_{B,q}(\mathbf{Z}_n)$, with an average (dashed horizontal line) of $\bar{\gamma}_n \simeq 0.05$ for $n = 200$. The behaviour observed is typical.

²All calculations are made with Matlab, on a PC with a clock speed of 3.3 GHz and 8 GB RAM, and operations are vectorized whenever possible.

³Note that, whenever the solution at Step 3 is unique, the designs obtained in the two cases are identical.

Figure 1: Left: computational times of the greedy and lazy-greedy implementations; Right: effective size γ_k of $\mathcal{X}_C$ as a function of k (maximization of $\hat{I}_{B,q}(\mathbf{Z}_n)$ of Section 3.3 with $\mathcal{X}_Q = \mathcal{X}_C = \mathbf{S}_{2048,d}$, $q = 5$, $B = \sqrt{d}/2$; $d = 5$ on the right panel).

In the sequel, we shall always use the lazy-greedy algorithm to maximize a submodular set function.

3 Covering measures

3.1 Definitions

For any positive scalar r , we define the covering measure $\Phi_r(\mathbf{Z}_n)$ of $\mathbf{Z}_n$ by

$$\Phi_r(\mathbf{Z}_n) = \frac{\text{vol}\{\mathcal{X} \cap [\cup_{i=1}^n \mathcal{B}(\mathbf{z}_i, r)]\}}{\text{vol}(\mathcal{X})},$$

with $\mathcal{B}(\mathbf{z}, r)$ the closed ball with center $\mathbf{z}$ and radius r . For a given $\mathbf{Z}_n$, consider also the function $r \in \mathbb{R}^+ \rightarrow F_{\mathbf{Z}_n}(r) = \Phi_r(\mathbf{Z}_n)$. $F_{\mathbf{Z}_n}$ is non-decreasing and satisfies $F_{\mathbf{Z}_n}(0) = 0$ and $F_{\mathbf{Z}_n}(r) = 1$ for any $r \geq \text{CR}(\mathbf{Z}_n)$. If X is distributed with the uniform probability measure μ on $\mathcal{X}$, we have

$$\text{Prob}\{X \in \cup_{i=1}^n \mathcal{B}(\mathbf{z}_i, r)\} = \text{Prob}\{d(X, \mathbf{Z}_n) \leq r\} = \int_{\{\mathbf{x} \in \mathcal{X} : d(\mathbf{x}, \mathbf{Z}_n) \leq r\}} \mu(d\mathbf{x}) = F_{\mathbf{Z}_n}(r), \quad (6)$$

and $F_{\mathbf{Z}_n}$ is the c.d.f. of the random variable $d(X, \mathbf{Z}_n)$, supported on $[0, \text{CR}(\mathbf{Z}_n)]$. We will call $F_{\mathbf{Z}_n}$ the distance c.d.f. and denote by $Q_\alpha(\mathbf{Z}_n)$ its α -quantile of:

$$Q_\alpha(\mathbf{Z}_n) = \inf\{t : F_{\mathbf{Z}_n}(t) \geq \alpha\},$$

with $Q_1(\mathbf{Z}_n) = \text{CR}(\mathbf{Z}_n)$. The set function $\mathbf{Z}_n \subset \mathcal{X} \rightarrow \Phi_r(\mathbf{Z}_n)$ is non-decreasing and satisfies $\Phi_r(\emptyset) = 0$. Moreover, for any $\mathbf{x} \in \mathcal{X}$, the difference $\Phi_r(\mathbf{Z}_n \cup \{\mathbf{x}\}) - \Phi_r(\mathbf{Z}_n)$ is non-increasing with respect to $\mathbf{Z}_n$, so that Φ_r is submodular.

In [26], the authors also define design criteria based on the distribution of distances to design points, in particular the $(1 - \gamma)$ -covering radius which corresponds to $Q_{1-\gamma}(\mathbf{Z}_n)$. Using

approximations valid for large d (see Remark 2), they investigate the properties of $Q_{1-\gamma}(\mathbf{Z}_n)$ for different random designs, in particular designs formed by n points i.i.d. in $\mathcal{C}_d = [0, 1]^d$ with a beta distribution. In [27], they consider full-factorial and 2^{d-1} fractional factorial designs, also for large d . We follow a different route, considering the values of $F_{\mathbf{Z}_n}(r)$ taken at different r and exploiting the submodularity of Φ_r to construct deterministic incremental designs $\mathbf{X}_n$ which, thanks to Theorem 1, have guaranteed efficiency for all n for the criterion considered. In contrast with [26], our constructions are effective also for small d , as several examples will illustrate. In our numerical study, additionally to the covering radius, we will also assess performance through the covering α -quantile Q_α for α close to 1.

3.2 A c.d.f.-based submodular covering criterion

Denote by $f_{\mathbf{Z}_n}$ the probability density function (p.d.f.) corresponding to $F_{\mathbf{Z}_n}$. For any $B > 0$, $q > -1$ and $\mathbf{Z}_n \neq \emptyset$, define the integrated covering measure

$$I_{B,q}(\mathbf{Z}_n) = \int_0^B r^q F_{\mathbf{Z}_n}(r) dr \quad (7)$$

$$= \frac{1}{q+1} \left\{ B^{q+1} F_{\mathbf{Z}_n}(B) - \int_0^B r^{q+1} f_{\mathbf{Z}_n}(r) dr \right\}, \quad (8)$$

and set $I_{B,q}(\emptyset) = 0$.

The set function $I_{B,q} : \mathbf{Z}_n \rightarrow I_{B,q}(\mathbf{Z}_n)$ is non-decreasing and submodular, and satisfies $I_{B,q}(\emptyset) = 0$. As $I_{B,q}$ satisfies the conditions of Theorem 1, its greedy maximization with Algorithm 1 provides sequences of incremental designs $\mathbf{X}_n$ with a $I_{B,q}$ -efficiency of at least 63.2% for any n , see (5). Since $F_{\mathbf{Z}_n}(B) = 1$ for any $B \geq \text{CR}(\mathbf{Z}_n)$, maximizing $I_{B,q}(\mathbf{Z}_n)$ with respect to $\mathbf{Z}_n$ for $B \geq \text{diam}(\mathcal{X})$ is equivalent to minimizing $\int_0^B r^{q+1} f_{\mathbf{Z}_n}(r) dr = \mathbb{E}_n\{R^{q+1}\}$ in (8), where the random variable R has the p.d.f. $f_{\mathbf{Z}_n}$ and where

$$(\mathbb{E}_n\{R^{q+1}\})^{1/(q+1)} = E_{q+1}(\mathbf{Z}_n) \quad (9)$$

is the L^{q+1} -mean quantization error induced by $\mathbf{Z}_n$, see [11]. It satisfies $E_{q+1}(\mathbf{Z}_n) \nearrow \text{CR}(\mathbf{Z}_n)$ as $q \rightarrow \infty$. For B and q large enough, maximizing $I_{B,q}(\mathbf{Z}_n)$ should therefore provide designs with small values of $\text{CR}(\mathbf{Z}_n)$; moreover, the greedy maximization of $I_{B,q}(\mathbf{Z}_n)$ is equivalent to the greedy minimization of $E_{q+1}(\mathbf{Z}_n)$, which is proved in [20] to ensure that $E_{q+1}(\mathbf{Z}_n)$ tends to zero at rate $n^{-1/d}$. These observations motivate the investigations on the properties of $I_{B,q}$ and its numerical implementation presented in the paper.

3.3 Numerical implementation

To evaluate $I_{B,q}(\mathbf{Z}_n)$ we substitute the empirical c.d.f. $\hat{F}_{\mathbf{Z}_n}$ for $F_{\mathbf{Z}_n}(r)$ in (7), where $\hat{F}_{\mathbf{Z}_n}$ is obtained by replacing μ in (6) by the uniform measure μ_Q supported on a finite subset $\mathcal{X}_Q$ of $\mathcal{X}$,

$$\mathcal{X}_Q = \{\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(Q)}\}.$$

The set $\mathcal{X}_Q$ must be well spread over $\mathcal{X}$; it may be a regular grid, or correspond to the first Q points of a LDS in $\mathcal{X}$. The empirical distance c.d.f. $\hat{F}_{\mathbf{Z}_n}$ is based on the Q distances $d_j(\mathbf{Z}_n) =$

$d(\mathbf{x}^{(j)}, \mathbf{Z}_n)$, $j = 1, \dots, Q$. Denoting by $\bar{d}_j$ the truncated version of d_j , $\bar{d}_j(\mathbf{Z}_n) = \min\{d_j(\mathbf{Z}_n), B\}$, $j = 1, \dots, Q$, we obtain

$$I_{B,q}(\mathbf{Z}_n) \approx \hat{I}_{B,q}(\mathbf{Z}_n) = \frac{B^{q+1}}{q+1} - \frac{1}{Q(q+1)} \sum_{j=1}^Q \bar{d}_j^{q+1}(\mathbf{Z}_n). \quad (10)$$

Remark 1 (A connection with clustering) When $B \geq \text{diam}(\mathcal{X})$ ($B \geq \text{diam}(\mathcal{X})/2$ when $\mathcal{X}$ is convex and $\mathbf{x}_1$ is its Chebyshev center), the greedy maximization of $I_{B,q}(\mathbf{Z}_n)$, $q > -1$, selects

$$\mathbf{x}_{n+1} \in \text{Arg max}_{\mathbf{x} \in \mathcal{X}} \int_{\mathcal{C}(\mathbf{x})} [d^{q+1}(\mathbf{z}, \mathbf{X}_n) - \|\mathbf{x} - \mathbf{z}\|^{q+1}] \mu(d\mathbf{z})$$

at iteration n , with $\mathcal{C}(\mathbf{x})$ the cell with generator $\mathbf{x}$ in the Voronoi partition of $\mathcal{X}$ associated with $\mathbf{X}_n \cup \{\mathbf{x}\}$.

One may notice that (10) implies that, when B is large enough, the maximization of $\hat{I}_{B,q}(\mathbf{Z}_n)$ is equivalent to the center location problem [21] defined by the minimization of $\sum_i d_j^{q+1}(\mathbf{Z}_n)$. It also corresponds to the minimization of the L^{q+1} -mean quantization error induced by $\mathbf{Z}_n$, which, in this discrete setting where μ_Q is substituted for μ , can be solved by clustering algorithms, see [18], including **kmeans** and variants based of the Chebyshev centers of the Voronoi cells defined by the elements of $\mathbf{Z}_n$ [28]. $\square$

Iterative update of the $d(\mathbf{x}^{(j)}, \mathbf{X}_n)$. Consider the greedy maximization of $\hat{I}_{B,q}$. The criterion $\hat{I}_{B,q}(\mathbf{X}_n \cup \{\mathbf{x}\})$ optimized at step n of the incremental construction involves the distances $d_j(\mathbf{X}_n \cup \{\mathbf{x}^{(k)}\})$ for the current design $\mathbf{X}_n$ and all $\mathbf{x}^{(k)} \in \mathcal{X}_C$, the set of eligible design points. The $Q \times C$ matrix $\mathbf{D}_n$ with elements $\{\mathbf{D}_n\}_{j,k} = d_j(\mathbf{X}_n \cup \{\mathbf{x}^{(k)}\})$, $\mathbf{x}^{(k)} \in \mathcal{X}_C$, can be computed recursively as follows. At initialization of Algorithm 1, $\{\mathbf{D}_0\}_{j,k}$ is the $Q \times C$ matrix of inter-distances between the points of $\mathcal{X}_Q$ and those of $\mathcal{X}_C$. Then, when $\mathbf{X}_{n+1} = \mathbf{X}_n \cup \{\mathbf{x}^{(k')}\}$ for some $\mathbf{x}^{(k')} \in \mathcal{X}_C$, each column $\{\mathbf{D}_{n+1}\}_{\cdot,\ell}$ of $\mathbf{D}_{n+1}$ is given by $\min\{\{\mathbf{D}_n\}_{\cdot,\ell}, \{\mathbf{D}_0\}_{\cdot,k'}\}$, where the minimum is taken element-wise, for $\ell = 1, \dots, C$. In practise, it is computationally more efficient to calculate $\mathbf{D}_0^{(q+1)}$ and directly update $\mathbf{D}_n^{(q+1)}$, with $\{\mathbf{D}_n^{(q+1)}\}_{j,k} = [d_j(\mathbf{X}_n \cup \{\mathbf{x}^{(k)}\})]^{q+1}$ for any $n \geq 0$.

The complexity of the evaluation of $\hat{I}_{B,q}(\mathbf{X}_n \cup \{\mathbf{x}^{(k)}\})$ for an $\mathbf{x}^{(k)} \in \mathcal{X}_C$ is of order $\mathcal{O}(Q)$. The complexity of the greedy maximization of $\hat{I}_{B,q}(\mathbf{Z}_n)$ is thus of order $\mathcal{O}((\kappa + n)CQ)$, where κCQ counts for the initial calculation of $\mathbf{D}_0^{(q+1)}$; it becomes $\mathcal{O}((\kappa + n\bar{\gamma}_n)CQ)$ for the lazy-greedy algorithm, with $\bar{\gamma}_n \ll 1$, see Section 2.2.

Remark 2 (Large d approximation) Using an Edgeworth expansion of the Central Limit Theorem, in [40] the authors derive approximations of $F_{\mathbf{x}}(r)$ that are valid for large d , with $F_{\mathbf{x}}$ the distance c.d.f. for the one-point design $\{\mathbf{x}\}$. For X distributed with the uniform probability measure on $\mathcal{X}$ and $\mathbf{Z}_n$ a random design with n i.i.d. elements $\mathbf{z}_i$, we have $F_{\mathbf{Z}_n}(r) = \text{Prob}\{d(X, \mathbf{Z}_n) \leq r\} = 1 - \prod_{i=1}^n [1 - F_{\mathbf{z}_i}(r)]$; this is also exploited in [26]. However, the approximation of $F_{\mathbf{x}}$ is not directly exploitable for the estimation of $I_{B,q}(\mathbf{Z}_n)$ for a given $\mathbf{Z}_n$ since in that case $F_{\mathbf{Z}_n}(r) \neq 1 - \prod_{i=1}^n [1 - F_{\mathbf{z}_i}(r)]$. $\square$

Remark 3 (Initialization) When one is not interested in the performance of the smallest designs, with n less than some $n_{\min}$, it seems preferable to directly choose the first $n_{\min}$ points by

minimizing $\text{CR}(\mathbf{Z}_{n_{\min}})$, and initialize the greedy algorithm at this $n_{\min}$ -point design. The batch optimization of $\text{CR}(\mathbf{Z}_{n_{\min}})$ is a difficult task, but several methods available are able to produce designs with reasonably good performance; see, e.g., [28]. The investigation of properties and performance of incremental constructions initialized in this way is not considered in the present paper, where we consider the entire nested sequence, starting at $n = 1$. This will be the subject of future studies, together with backwards constructions that start with a nearly optimal $n_{\max}$ -point design and iteratively eliminate design points until an $n_{\min}$ -points design is reached, for some $n_{\min} < n_{\max}$. $\square$

3.4 Parameter setting

Choice of B . Large enough values of B have no influence on the maximization of $I_{B,q}(\mathbf{Z}_n)$. Indeed, take $B_2 \geq B_1 \geq \text{CR}(\mathbf{Z}_n)$; then $I_{B_2,q}(\mathbf{Z}_n) = I_{B_1,q}(\mathbf{Z}_n) + (B_2^{q+1} - B_1^{q+1})/(q+1)$ since $F_{\mathbf{Z}_n}(r) = 1$ for $B_1 \leq r \leq B_2$. The simplest rule sets $B = \text{diam}(\mathcal{X})$, guaranteeing that $F_{\mathbf{Z}_n}(B) = 1$ for any $\mathbf{Z}_n$. A less simplistic rule is to choose B of the order of magnitude of the largest possible value of $\text{CR}(\mathbf{Z}_n)$ over the range $[n_{\min}, n_{\max}]$ of design sizes of interest, i.e., $\text{CR}(\mathbf{Z}_{n_{\min}})$.

When $\mathcal{X} = \mathcal{C}_d$ and $B \geq \sqrt{d}/2$, the first design point is necessarily taken at $\mathbf{c} = (1/2, \dots, 1/2)$, the center of $\mathcal{C}_d$. The unit cube $\mathcal{C}_d$ can be covered by m^d hypercubes with side $1/m$, and therefore by m^d balls with radius $\sqrt{d}/(2m)$. Let CR_n^* denote the minimum value of $\text{CR}(\mathbf{Z}_n)$ for designs of size n . Taking $m = \lfloor n^{1/d} \rfloor$, the largest integer smaller than or equal to $n^{1/d}$, we have $n \geq m^d$ and therefore the following upper bound holds:

$$\text{CR}_n^* \leq \text{CR}_{m^d}^* \leq R^*(n, d) = \frac{\sqrt{d}}{2 \lfloor n^{1/d} \rfloor}. \quad (11)$$

This suggests taking $B = R^*(n_{\min}, d)$ instead of $B = \sqrt{d}/2$. The upper bound (11) is very pessimistic though, with $R^*(1, d) = \sqrt{d}/2$ for all d and the k th jump downwards of $R^*(n, d)$ occurring at $n = (k+1)^d$. For example, the effect of choosing $B = R^*(n_{\min}, d)$ will only be effective for $n_{\min} \geq 32$ when $d = 5$ and $n_{\min} \geq 1024$ when $d = 10$. Our numerical studies indicate that it sometimes slightly improves the efficiency of the resulting designs when $n \geq 2^d$ (that is, for very large designs when d is large). More interestingly, the first points selected are then no longer uniquely defined, which can be exploited to implement multi-start strategies.

Figures 2 and 3 illustrate the impact of the choice of B when $d = 2$, $\mathcal{X}_C = \mathcal{X}_Q = \mathbf{S}_{2048,2}$ and $n_{\max} = 20$. Figure 2 presents the exact value of $\text{CR}(\mathbf{X}_n)$ for n between 1 and $n_{\max}$ (red solid curve with $\star$) and its under approximation $\text{CR}_{[\mathcal{X}_Q]}(\mathbf{X}_n)$ given by (1) (black dashed curve with $\circ$). The blue dashed curve (+) presents the estimated⁴ values of CR_n^* ; the coloured area shows the region between CR_n^* and $\text{CR}(\mathbf{X}_n)$. The magenta curve (∇) shows the empirical value of $E_{q+1}(\mathbf{X}_n)$ (a lower bound on $\text{CR}_{[\mathcal{X}_Q]}(\mathbf{X}_n)$, see (9) in Section 3.2), for the uniform measure μ_Q on $\mathcal{X}_Q$. Although the value $q = 5$ is too small for $E_{q+1}(\mathbf{X}_n)$ to provide a good approximation of $\text{CR}(\mathbf{X}_n)$, we can see that the greedy algorithm manages to ensure a reasonable decrease of $\text{CR}(\mathbf{X}_n)$ along iterations. This decrease is more regular when $B = R^*(n_{\min}, d)$ than when $B = \sqrt{d}/2$. Figure 3 shows the nested designs that are constructed and helps understanding

⁴Exact for $n = 1, 2$, best (smallest) values obtained for each n between 3 and 20 by running a `kmeans`-type clustering algorithm on a 50×50 regular grid in $\mathcal{X}$; see [28]. The values plotted are therefore not necessarily equal to the true values of CR_n^* , but we believe that the overestimation is negligible.

the different behaviors of $\text{CR}(\mathbf{X}_n)$: when $B = \sqrt{d}/2$ (left), $\mathbf{x}_1$ corresponds to the optimal one-point design and is at the center of $\mathcal{C}_2$, the next two points $\mathbf{x}_2$ and $\mathbf{x}_3$ are close to those for the optimal 3-point design $\{\mathbf{x}_1, \hat{\mathbf{x}}_2, \hat{\mathbf{x}}_3\}$ with $\mathbf{x}_1$ fixed at $(1/2, 1/2)$; see the appendix. The situation is different when $B = R^*(n_{\min}, d)$ (right panel), and we see that sacrificing optimality at a given n may be at the benefit of a more regular decrease of $\text{CR}(\mathbf{X}_n)$: here the choice of the initial points $\mathbf{x}_1$, $\mathbf{x}_2$ and $\mathbf{x}_3$ is much worse than with $B = \sqrt{d}/2$ (with $\text{CR}(\mathbf{X}_1) = 1.0054$ and $\text{CR}(\mathbf{X}_2) = \text{CR}(\mathbf{X}_3) = 0.7442$), but the situation then improves, leading in particular to significantly smaller $\text{CR}(\mathbf{X}_n)$ for $n \in \{5, \dots, 8\}$.

Figure 2: $\text{CR}(\mathbf{X}_n)$ (red $\star$), $\text{CR}_{[\mathcal{X}_Q]}(\mathbf{X}_n)$ (black $\circ$), CR_n^* (blue $+$), empirical value of $E_{q+1}(\mathbf{X}_n)$ (9) (magenta ∇), for $\mathbf{X}_n$ obtained by greedy maximization of $\hat{I}_{B,q}(\mathbf{X}_n)$.

Figure 3: $\mathbf{X}_{20}$ and circles centered at design points with radius $\text{CR}(\mathbf{X}_{20})$; the order of selection of the points is indicated.

Remark 4 (Truncation of B to $\text{CR}(\mathbf{X}_n)$) For $B \geq \text{CR}(\mathbf{X}_n)$, we have $I_{B,q}(\mathbf{X}_n \cup \{\mathbf{x}\}) = I_{\text{CR}(\mathbf{X}_n),q}(\mathbf{X}_n \cup \{\mathbf{x}\}) + \{B^{q+1} - [\text{CR}(\mathbf{X}_n)]^{q+1}\}/(q+1)$, where the second term does not depend on

$\mathbf{x}$. We may thus select $\mathbf{x}_{n+1}$ as the maximizer of $I_{\min\{B, \text{CR}(\mathbf{X}_n)\}, q}(\mathbf{X}_n \cup \{\mathbf{x}\})$. When q is large, the constant $C_n = \{B^{q+1} - [\text{CR}(\mathbf{X}_n)]^{q+1}\}/(q+1)$ may become much larger than most terms in the sum (10), so that ignoring it significantly improves numerical stability. Note that the contribution of C_n to $I_{B,q}(\mathbf{X}_n)$ must be suitably accounted for in the calculation of the bounds $\bar{\delta}(\mathbf{x}^{(k)})$ on the increments in the lazy-greedy algorithm of Section 2.2. $\square$

Remark 5 (A positive lower integration limit in $I_{B,q}$) We discussed above the possibility of restricting the largest distances taken into account in $I_{B,q}$. We can also think of restricting the contribution of the smallest distances by using a strictly positive lower integration limit b . The criterion is then more focussed on the influence of points of $\mathcal{X}_Q$ at large distances from the design. Parameter b must be taken larger than the smallest covering radius for the design sizes of interest, and we suggest to relate b to a lower bound on $\text{CR}_{n_{\max}}^*$. Consider designs of size n in $\mathcal{X} = \mathcal{C}_d$. Since the n balls $\mathcal{B}(\mathbf{x}_i, \text{CR}_n^*)$ centered at the optimal design points cover $\mathcal{X}$, $nV_d(\text{CR}_n^*)^d \geq \text{vol}(\mathcal{X}) = 1$, with V_d the volume of the d -dimensional unit ball $\mathcal{B}(\mathbf{0}, 1)$, $V_d = \pi^{d/2}/\Gamma(d/2 + 1)$. This implies the following classical lower bound [25, p. 150]:

$$\text{CR}_n^* \geq R_*(n, d) = (nV_d)^{-1/d}. \quad (12)$$

We may note that V_d is maximum for $d = 5$ (with $V_5 \simeq 5.2638$) and that $n^{1/d} R_*(n, d) = \sqrt{d}/\sqrt{2\pi e} + \mathcal{O}(1/\sqrt{d})$, $d \rightarrow \infty$. A reasonable choice is then $b = R_*(n_{\max}, d)$. In practice, our numerical investigations indicate that this modification has marginal influence on performance, although it may slightly simplify the computations since the summation in (10) involves a smaller number of terms. It will not be used in the numerical experiments of Section 5. $\square$

Choice of q . The discussion in Section 3.2 suggests that performance, as measured by the covering radius, should improve as q increases. This is indeed the case for the criterion $\text{CR}_{[\mathcal{X}_Q]}(\mathbf{Z}_n)$ defined by (1). Whenever $B > \text{CR}(\mathbf{Z}_n)$, when q tends to infinity a design $\mathbf{Z}_n$ that maximizes $\hat{I}_{B,q}(\mathbf{Z}_n)$ tends to minimize $\text{CR}_{[\mathcal{X}_Q]}(\mathbf{Z}_n)$ as $\hat{I}_{B,q}(\mathbf{Z}_n)$ is focussed on the furthest points from $\mathbf{Z}_n$ in $\mathcal{X}_Q$, irrespectively of the presence of points in $\mathcal{X} \setminus \mathcal{X}_Q$ further away from $\mathbf{Z}_n$. The connection with a clustering problem for the L^{q+1} -mean quantization error noticed in Remark 1 suggests that avoiding values of q too large may lead to better generalization properties of the statistics on which $I_{B,q}$ is based when the number of data points on which it is evaluated is finite. Rather intensive numerical studies show that performance is robust with respect to the choice of q and indicate that a value $q \in [5, 25]$ is a convenient choice (see Figure 9 for an illustration). Figure 4 illustrates the impact of q , comparing the performances obtained for $q = 5$ (red) and $q = 50$ (blue) when $\mathcal{X} = \mathcal{C}_5$ ($d = 5$), with $\mathcal{X}_C = \mathcal{X}_Q = \mathbf{S}_{2048,d}$ and $B = \sqrt{d}/2$. We normalize $\text{CR}(\mathbf{X}_n)$ by dividing it by the lower bound (12) on the optimal (minimum) value CR_n^* . On the left panel, $\text{CR}(\mathbf{X}_n)$ is approximated by $\text{CR}_{[\mathcal{X}_Q]}(\mathbf{X}_n)$, on the right panel it is approximated by $\text{CR}_{[\mathcal{X}_N]}(\mathbf{X}_n)$, with $\mathcal{X}_N$ given by 2^{18} points of a scrambled Sobol' sequence complemented by a 2^d full factorial design (which gives $N = 262176$). The choice $q = 50$ is preferable to $q = 5$ in terms of $\text{CR}_{[\mathcal{X}_Q]}(\mathbf{X}_n)$, but is significantly worse for $\text{CR}_{[\mathcal{X}_N]}(\mathbf{X}_n)$, which is of course a better approximation of $\text{CR}(\mathbf{X}_n)$.

Choices of $\mathcal{X}_C$ and $\mathcal{X}_Q$. The problem considered may impose constraints on the set $\mathcal{X}_C$ of eligible points. For instance, in Section 5 we force the algorithm to choose points within a Latin hypercube design, and use $I_{B,q}$ to define an order in this design. When no such constraints are

Figure 4: Left: $\text{CR}_{[\mathcal{X}_Q]}(\mathbf{X}_n)/R_*(n, d)$; Right: $\text{CR}_{[\mathcal{X}_N]}(\mathbf{X}_n)/R_*(n, d)$, $n = 1, \dots, 250$; $q = 5$ (red $\star$), $q = 50$ (blue $+$).

enforced, any finite set evenly distributed on $\mathcal{X}$ can be used and in the paper we most often use points from a LDS.

A set of points evenly distributed on $\mathcal{X}$ is required for $\mathcal{X}_Q$, with Q large being more crucial than having C large. In particular, $\mathcal{X}_Q$ must sample the regions expected to be the furthest away from designs with good CR. For that reason, when $\mathcal{X} = \mathcal{C}_d$, $\mathcal{X}_Q$ should contain the vertices of $\mathcal{C}_d$. Given the factor CQ in the algorithmic complexity of the construction and the size $Q \times C$ of the matrices involved, see Section 3.3, computational constraints impose an upper bound on the product CQ . Reducing the size of $\mathcal{X}_C$ to a minimal but acceptable size then allows to take a larger set $\mathcal{X}_Q$, which is beneficial for the quality of the approximation of $I_{B,q}(\mathbf{Z}_n)$ by $\hat{I}_{B,q}(\mathbf{Z}_n)$. Our numerical experiments suggest using $\mathcal{X}_C \subset \mathcal{X}_Q$ with the rule of thumb $Q = 2C$.

4 Spacings and coffee-house design

Coffee-house design. The algorithmic construction presented in the previous section is based on the distance c.d.f. $F_{\mathbf{Z}_n}$. In this section we present a second family of incremental algorithms, directly related to the most common geometrical criteria, the covering and packing radii. Even if these criteria are not submodular, and thus Theorem 1 cannot be invoked to guarantee a lower bound on the efficiency of their greedy optimization, an efficiency of 50% can still be guaranteed.

Coffee-house designs [22], [23, Chap. 4] greedily maximize $\text{PR}(\mathbf{Z}_n)$ for $n > 1$; see [15] for an early suggestion. When $\mathcal{X}$ is convex, the first design point $\mathbf{x}_1$ is usually chosen at the Chebyshev center of $\mathcal{X}$ (the center of the minimal-radius ball enclosing $\mathcal{X}$); then, at any $n \geq 1$, $\mathbf{x}_{n+1} \in \text{Arg max}_{\mathbf{x} \in \mathcal{X}} d(\mathbf{x}, \mathbf{Z}_n)$. Denote by PR_n^* the optimal (maximum) value of $\text{PR}(\mathbf{Z}_n)$ for a n -point design in $\mathcal{X}$. The two ratios $\text{CR}_n^*/\text{CR}(\mathbf{Z}_n)$ and $\text{PR}(\mathbf{Z}_n)/\text{PR}_n^*$ are less than one by construction and quantify the efficiency of a design $\mathbf{Z}_n$ in terms of its covering and packing radii. It is remarkable that the simple greedy coffee-house construction ensures the following property (irrespectively of the choice made for $\mathbf{x}_1$):

$$\frac{1}{2} \leq \frac{\text{CR}_n^*}{\text{CR}(\mathbf{Z}_n)} \leq 1 \quad (n \geq 1) \quad \text{and} \quad \frac{1}{2} \leq \frac{\text{PR}(\mathbf{Z}_n)}{\text{PR}_n^*} \leq 1 \quad (n \geq 2). \quad (13)$$

Indeed, by construction $\text{PR}(\mathbf{X}_{n+1}) = d(\mathbf{x}_{n+1}, \mathbf{X}_n)/2 = \text{CR}(\mathbf{X}_n)/2$ for all $n \geq 1$. Take any n -point design $\mathbf{Z}_n = \{\mathbf{z}_1, \dots, \mathbf{z}_n\}$ in $\mathcal{X}$. From the pigeonhole principle, one of the balls $\mathcal{B}(\mathbf{z}_i, \text{CR}(\mathbf{Z}_n))$ must contain two points $\mathbf{x}_i, \mathbf{x}_j$ of $\mathbf{X}_{n+1}$, implying that $\text{PR}(\mathbf{X}_{n+1}) \leq \|\mathbf{x}_i - \mathbf{x}_j\|/2 \leq \text{CR}(\mathbf{Z}_n)$. Therefore, $\text{CR}_n^* \geq \text{PR}(\mathbf{X}_{n+1}) = \text{CR}(\mathbf{X}_n)/2$. Similarly, $\text{PR}_{n+1}^* \leq \text{CR}_n^* \leq \text{CR}(\mathbf{X}_n) = 2 \text{PR}(\mathbf{X}_{n+1})$. The original proof is given in [10].

As for the maximization of $I_{B,q}(\mathbf{Z}_n)$ in Section 3, the implementation of Algorithm 1 is much facilitated when $\mathbf{x}_{n+1} \in \text{Argmax}_{\mathbf{x} \in \mathcal{X}_C} d(\mathbf{x}, \mathbf{X}_n)$ with $\mathcal{X}_C$ a finite set of candidates. The efficiencies given in (13) remain valid provided that $\text{CR}(\mathbf{X}_n)$ is approximated by $\text{CR}_{[\mathcal{X}_C]}(\mathbf{X}_n)$ and that CR_n^* and PR_n^* are relative to optimal n -point designs in $\mathcal{X}_C$. Contrary to the maximization of $\hat{I}_{B,q}(\mathbf{Z}_n)$ in Section 3.3, the greedy maximization of $\text{PR}(\mathbf{Z}_n)$ does not require computing a $Q \times C$ matrix of inter-distances, but only the update of the C distances $d(\mathbf{x}^{(i)}, \mathbf{Z}_n)$ for the $\mathbf{x}^{(i)}$ in $\mathcal{X}_C$. For that reason, the size C of the candidate set can be taken much larger than in Section 3.4.

Denote by $\rho(\mathbf{Z}_n)$ the mesh-ratio of $\mathbf{Z}_n$, defined by

$$\rho(\mathbf{Z}_n) = \frac{\text{CR}(\mathbf{Z}_n)}{\text{PR}(\mathbf{Z}_n)} \quad (n \geq 2).$$

If a sequence of n -points designs $\mathbf{Z}_n$ is such that $\rho(\mathbf{Z}_n)$ is uniformly bounded, then the sequence is said to be quasi-uniform; see, e.g., [4]. The inverse of $\rho(\mathbf{Z}_n)$ is sometimes called uniformity measure [5]. When $\mathcal{X}$ is convex it cannot be covered by two or more non-overlapping balls having their centers in $\mathcal{X}$, and $\text{PR}(\mathbf{Z}_n) < \text{CR}(\mathbf{Z}_n)$, implying that $\rho(\mathbf{Z}_n) > 1$. When $\mathbf{X}_n$ is a coffee-house design, $\text{PR}(\mathbf{X}_{n+1}) = \text{CR}(\mathbf{X}_n)/2$ for $n \geq 1$, and since $\text{CR}(\mathbf{Z}_n) \geq \text{CR}(\mathbf{Z}_{n+1})$ for all n and any nested designs $\mathbf{Z}_n \subset \mathbf{Z}_{n+1}$, $\rho(\mathbf{X}_n)$ satisfies

$$1 \leq \rho(\mathbf{X}_n) \leq 2 \quad (n \geq 2).$$

That $\rho(\mathbf{X}_n)$ is upper-bounded by 2 is a very strong guarantee of coffee-house design, which does not hold for our c.d.f.-based construction. This powerful property is certainly related to the bounded optimality gaps for both the covering and packing radii. The method presented in the previous section, tailored to control CR only, is unable to guarantee a large spacing between design points, which can translate into large values⁵ of $\rho(\mathbf{X}_n)$.

Spacings and boundary-phobic coffee-house design. By construction, a coffee-house design places design points on the boundary $\partial\mathcal{X}$ of $\mathcal{X}$, in contradiction with the objective of constructing designs with low values of $\text{CR}(\mathbf{Z}_n)$. For that reason, we consider below a modified method that forces design points to stay away from $\partial\mathcal{X}$.

Following [13]⁶, we define the maximal spacing $S(\mathbf{Z}_n)$ as the radius of the largest ball contained in $\mathcal{X}$ and not intersecting $\mathbf{Z}_n$,

$$S(\mathbf{Z}_n) = \sup \{r : \exists \mathbf{x} \text{ such that } \mathbf{x} \oplus r \mathcal{B}(\mathbf{0}, 1) \subset \mathcal{X} \setminus \mathbf{Z}_n\},$$

with $\oplus$ denoting the Minkowski sum (for two subsets $\mathcal{A}_1$ and $\mathcal{A}_2$ of $\mathbb{R}^d$, $\mathcal{A}_1 \oplus \mathcal{A}_2 = \{\mathbf{x}_1 + \mathbf{x}_2 : \mathbf{x}_1 \in \mathcal{A}_1, \mathbf{x}_2 \in \mathcal{A}_2\}$.) We slightly extend this notion by introducing a parameter β that controls

⁵Values larger than 3 have been observed in Example 1 of Section 5.

⁶In [13], $\mathcal{X}$ is only assumed to be bounded and bounded convex sets other than $\mathcal{B}(\mathbf{0}, 1)$ are also considered.

the ratio between the distance to the design and the distance⁷ to $\partial\mathcal{X}$, and define the β -spacing of $\mathbf{Z}_n$, for $\beta > 0$, as

$$S_\beta(\mathbf{Z}_n) = \sup \left\{ r : \exists \mathbf{x} \in \mathcal{X} \text{ such that } d(\mathbf{x}, \mathbf{Z}_n) \geq r \text{ and } d(\mathbf{x}, \partial\mathcal{X}) \geq \frac{r}{\beta} \right\} = \sup_{\mathbf{x} \in \mathcal{X}} D_\beta(\mathbf{x}, \mathbf{Z}_n),$$

where

$$D_\beta(\mathbf{x}, \mathbf{Z}_n) = \min \{ d(\mathbf{x}, \mathbf{Z}_n), \beta d(\mathbf{x}, \partial\mathcal{X}) \}, \quad \mathbf{x} \in \mathcal{X}, \quad (14)$$

and $d(\mathbf{x}, \partial\mathcal{X}) = \inf_{\mathbf{z} \in \partial\mathcal{X}} \|\mathbf{x} - \mathbf{z}\|$. Figure 5 shows the variations of $D_\beta(\mathbf{x}, \mathbf{Z}_n)$ for a 4-point design $\mathbf{Z}_n$ in the square $[0, 1]^2$ with two different values of β . We have $S_1(\mathbf{Z}_n) = S(\mathbf{Z}_n)$ and we define $S_\infty(\mathbf{Z}_n) = \text{CR}(\mathbf{Z}_n)$. We also define

$$P_\beta(\mathbf{Z}_n) = \min_{\mathbf{z}_i \neq \mathbf{z}_j \in \mathbf{Z}_n} \frac{1}{2} \min \{ \|\mathbf{z}_i - \mathbf{z}_j\|, \beta d(\mathbf{z}_i, \partial\mathcal{X}) \} \text{ and } \rho_\beta(\mathbf{Z}_n) = \frac{S_\beta(\mathbf{Z}_n)}{P_\beta(\mathbf{Z}_n)} \quad (n \geq 2), \quad (15)$$

and set $P_\infty(\mathbf{Z}_n) = \text{PR}(\mathbf{Z}_n)$, so that $\rho_\infty(\mathbf{Z}_n) = \rho(\mathbf{Z}_n)$.

Figure 5: Boundary-phobic compromise distance $D_\beta(\mathbf{x}, \mathbf{Z}_n)$ for a 4-point design (yellow stars) in $\mathcal{X} = [0, 1]^2$.

The coffee-house algorithm can be straightforwardly extended to the greedy maximization of $P_\beta(\mathbf{Z}_n)$, $\beta > 0$, using $\mathbf{x}_{n+1} \in \text{Arg max}_{\mathbf{x} \in \mathcal{X}} D_\beta(\mathbf{x}, \mathbf{X}_n)$ for any $n \geq 1$ (and if we use the convention $\mathbf{X}_0 = \emptyset$ and $d(\mathbf{x}, \emptyset) = +\infty$ for any $\mathbf{x}$, the same rule sets $\mathbf{x}_1$ at the Chebyshev center of $\mathcal{X}$ when $\mathcal{X}$ is convex and to a point of the medial axis [3] at maximum distance from $\partial\mathcal{X}$ in general). Following the same steps as for coffee-house design, we obtain the following property, whose proof is given in the appendix.

⁷When $\mathcal{X} = \mathcal{C}_d$, an anonymous referee suggested to penalize differently distances to the vertices and distances to the boundary of $\mathcal{C}_d$ (or more generally to the d' -dimensional faces of $\mathcal{C}_d$, $1 < d' < d$), using different metrics; we have not explored this possibility here.

Theorem 2 Let $\mathcal{X}$ be a compact and convex subset of $\mathbb{R}^d$ and let $\mathbf{X}_n$ be a design obtained with the following greedy construction: $\mathbf{X}_0 = \emptyset$ and $\mathbf{x}_{k+1} \in \text{Arg max}_{\mathbf{x} \in \mathcal{X}} D_\beta(\mathbf{x}, \mathbf{X}_k)$, $k = 0, 1, \dots, n-1$, where D_β is defined in (14), $\beta > 0$, and $d(\mathbf{x}, \emptyset) = +\infty$. Then $P_\beta(\mathbf{X}_{n+1}) = (1/2) S_\beta(\mathbf{X}_n)$ for $n \geq 1$ and

$$\frac{1}{2} \leq \frac{S_{\beta,n}^*}{S_\beta(\mathbf{X}_n)} \leq 1 \quad (n \geq 1) \quad \text{and} \quad \frac{1}{2} \leq \frac{P_\beta(\mathbf{X}_n)}{P_{\beta,n}^*} \leq 1, \quad 1 \leq \rho_\beta(\mathbf{X}_n) \leq 2 \quad (n \geq 2), \quad (16)$$

where $S_{\beta,n}^* = \min_{\mathbf{Z}_n \subset \mathcal{X}} S_\beta(\mathbf{Z}_n)$ and $P_{\beta,n}^* = \max_{\mathbf{Z}_n \subset \mathcal{X}} P_\beta(\mathbf{Z}_n)$.

Algorithm 2 of [37], also involving the domain boundary, corresponds to the greedy maximization of $P_\beta(\mathbf{Z}_n)$ in $\mathcal{X} = \mathcal{C}_d$, for the value $\beta = 2\sqrt{2d}$ (chosen in [37] by trial and error). Considering that the covering radius is of special importance, we recommend a different choice for β when $\mathcal{X} = \mathcal{C}_d$. Since $\mathbf{x}_1 = \mathbf{1}_d/2$, the center of $\mathcal{X}$, when $\beta = \infty$ the second point $\mathbf{x}_2$ coincides with a vertex of $\mathcal{X}$. We suggest to choose β depending on the maximum target design size $n_{\max}$, such that $\mathbf{x}_2$ be at distance $R_\star(n_{\max}, d)$ from a vertex of $\mathcal{X}$, with $R_\star(n_{\max}, d)$ the lower bound on $\text{CR}_{n_{\max}}^*$ given by (12). This implies $d(\mathbf{x}_2, \partial\mathcal{X}) = R_\star(n_{\max}, d)/\sqrt{d} = \|\mathbf{x}_2 - \mathbf{x}_1\|/\beta$ and gives

$$\beta = \beta_\star(n_{\max}, d) = \frac{d}{2R_\star(n_{\max}, d)} - \sqrt{d}. \quad (17)$$

The left panel of Figure 13 in the appendix shows $\beta_\star(n, d)$ as a function of d for $n = 50$ (red $\star$), $n = 100$ (blue ∇) and $n = 200$ (black $\circ$); the curve with magenta $\times$ corresponds to the values suggested in [37]. We see that as n grows $\beta_\star(n, d)$ increases (since we expect designs that better cover $\mathcal{X}$), and as d becomes large the value of $\beta_\star(n, d)$ exhibits a very slow growth after an initial fast decay (being almost constant for large values of n).

5 Numerical study

This section assesses the performance of the design algorithms presented in the previous sections. Designs are compared in terms of the covering radius $\text{CR}(\mathbf{X}_n)$ and of the covering α -quantile $Q_\alpha(\mathbf{X}_n)$, with $\alpha = 0.99$; we will show plots of the evolution of these criteria as the greedy constructions unfold. We stress that for the incremental design methods studied in this paper, the entire evolution of the performance indicators over a target range of design sizes is of interest, and not just their value for a given final design size. Since $\text{CR}(\mathbf{X}_n)$ and $Q_\alpha(\mathbf{X}_n)$ ideally decrease as $n^{-1/d}$ all our plots will show normalized versions, $\text{CR}(\mathbf{X}_n)/R_\star(n, d)$, see (12), and $n^{1/d} Q_\alpha(\mathbf{X}_n)$. An ideal design method should lead to trajectories close to an horizontal line at a small value.

We compare the performance of designs $\mathbf{X}_n^*$ that incrementally maximize $\widehat{I}_{B,q}(\mathbf{X}_n)$ to three families of constructions: (i) prefixes of low discrepancy Halton ($\mathbf{H}_{n,d}$) and Sobol' ($\mathbf{S}_{n,d}$) sequences; (ii) $\mathbf{X}_n^{CH,\beta}$ obtained by greedy maximization of $P_\beta(\mathbf{Z}_n)$ given by (15), with $\beta \in \{\beta_\star(n_{\max}, d), 2\sqrt{2}, \infty\}$; and (iii) two incremental constructions (designs $\mathbf{X}_n^{VD}$ and $\mathbf{X}_n^{RD}$) based on a direct relaxation of the covering radius recently proposed in [32], see precise definitions in the appendix.

Two case-studies are considered. In Example 1 the design space has a simple geometry embedded in a relatively large dimensional space, $\mathcal{X} = \mathcal{C}_{10} = [0, 1]^{10}$, while Example 2 considers

a more challenging topology involving a non-convex (annular) domain $\mathcal{X} \subset \mathbb{R}^2$. In both examples, the parameters of $\hat{I}_{B,q}(\mathbf{X}_n)$ are $B = \text{diam}(\mathcal{X})$ and $q = 10$ (we also use $q = 10$ for $\mathbf{X}_n^{VD}$ and $\mathbf{X}_n^{RD}$); design sizes exceed the range prescribed by the $n = 10d$ rule of [19]⁸.

Example 1 ($\mathcal{X} = \mathcal{C}_{10}$). In this example, the sets $\mathcal{X}_C$ and $\mathcal{X}_Q$ are always of the form $\mathcal{X}_C = \mathbf{S}_{C,d}$ and $\mathcal{X}_Q = \mathbf{S}_{Q',d} \cup \{0,1\}^d$ (for $\mathbf{X}_n^{VD}$ and $\mathbf{X}_n^{RD}$ a scrambled Sobol' sequence is used for $\mathcal{X}_C$, as their constructions require that $\mathcal{X}_C \cap \mathcal{X}_Q = \emptyset$). We use $C = 8192$ and $Q' = 16384$ for all methods. Note that the size of the candidate set for the construction of $\mathbf{X}_n^{CH,\beta}$ exceeds the one suggested in [37] ($C = 1000d + 2n_{\max}$). The value of $\text{CR}(\mathbf{Z}_n)$ is approximated by $\text{CR}_{[\mathcal{X}_N]}(\mathbf{Z}_n)$ given by (1), with $\mathcal{X}_N$ consisting of 2^{18} points of a scrambled Sobol' sequence complemented by the 2^d full factorial design $\{0,1\}^d$ ($N = 263168$), and $Q_\alpha(\mathbf{X}_n)$ is approximated by its empirical estimate based on $\mathcal{X}_N$. All plots are normalized and consider $n \in [10, 200]$.

It is noticeable that for all designs constructed in this example $\text{CR}_{[\mathcal{X}_N]}(\mathbf{X}_n) = \text{CR}_{\{0,1\}^d}(\mathbf{X}_n)$, indicating that only the vertices of $\mathcal{C}_d$ matter when evaluating CR for these designs. This dominant effect of the vertices on the values of $\text{CR}(\mathbf{X}_n)$ for large d and $n \ll 2^d$ flags the limitation of the covering radius as a measure of the global space-fillingness of designs in the hypercube.

(i) Figure 6 shows the comparison of $\mathbf{X}_n^*$ (red) with $\mathbf{H}_{n,d}$ (blue ∇) and $\mathbf{S}_{n,d}$ (black $\times$). To illustrate the impact of the presence of the vertices of $\mathcal{C}_d$ in $\mathcal{X}_Q$, the performance of the c.d.f.-based method is presented both for $\mathcal{X}_Q = \mathbf{S}_{Q',d} \cup \{0,1\}^d$ (solid red line with $\star$) and $\mathcal{X}_Q = \mathbf{S}_{Q',d}$ (dotted red lines with $+$). The left panel plots the evolution of $\text{CR}(\mathbf{X}_n)$, showing that the greedy maximization of $\hat{I}_{B^*,q}(\mathbf{Z}_n)$ yields designs with significantly smaller covering radii than the two classical LDS. The right panel shows the evolution of $Q_\alpha(\mathbf{X}_n)$, revealing an even stronger dominance of $\mathbf{X}_n^*$ for this criterion. Notice the significant increase of CR when the vertices of $\mathcal{C}_d$ are not included in $\mathcal{X}_Q$ (dotted red line). Contrariwise, this results in a smaller covering α -quantile, since the empirical c.d.f. criterion ignores then the immediate neighborhoods of the vertices.

(ii) Figure 7 shows the comparison of $\mathbf{X}_n^*$ (red $\star$, as in (i)) with the designs $\mathbf{X}_n^{CH,\beta}$ produced by three variants of coffee-house. The following code is used: $\beta = \infty$ (standard coffee-house design) in black $\times$, $\beta = 2\sqrt{2}d$ in magenta $\circ$, and $\beta = \beta_*(100, d)$ given by (17) in blue ∇ . To show the sensitivity of standard coffee-house with respect to the candidate set, we also plot in this figure (black $+$) the performance of coffee-house design $\mathbf{X}_n^{CH,\infty}$ for a candidate set extended by the vertices of $\mathcal{C}_{10}$, $\mathcal{X}_C = \mathbf{S}_{4096,10} \cup \{0,1\}^d$. The design $\mathbf{X}_n^*$ has smaller CR (see left panel) than all versions of coffee-house for all design sizes considered⁹. Except for very small design sizes, the boundary-phobic versions of coffee-house have better Q_α values (right panel) than $\mathbf{X}_n^*$. Note that $\mathbf{X}_n^{CH,\infty}$ has much larger Q_α than its competitors and is very sensitive to the choice of the candidate set: when $\mathcal{X}_C$ includes the vertices of $\mathcal{C}_{10}$, those are systematically selected at all iterations considered ($n \leq 200$), resulting in both large CR and Q_α .

(iii) Figure 8 shows the comparison of $\mathbf{X}_n^*$ (red $\star$, as in (i) and (ii)) with the designs $\mathbf{X}_n^{VD}$ (blue ∇) and $\mathbf{X}_n^{RD}$ (black $\times$). As before we can see (left panel) that $\mathbf{X}_n^*$ dominates $\mathbf{X}_n^{VD}$ and $\mathbf{X}_n^{RD}$ for CR for all design sizes. Note the very good performance of $\mathbf{X}_n^{RD}$ for the covering

⁸In Example 1, very small designs with $n \leq 10$ are not meaningful and are excluded from the analysis.

⁹Our c.d.f.-based method thus inherits the ranking established in [37] where it is shown that $\mathbf{X}_n^{CH,2\sqrt{2}d}$ performs better in terms of CR than sliced and nested Latin hypercube designs proposed in the literature, which are only batch-incremental (see the references in [37]).

Figure 6: Left: $CR(\mathbf{X}_n)/R_*(n_{\max}, d)$; Right: $n^{1/d} Q_\alpha(\mathbf{X}_n)$, $n = 10, \dots, 200$, $d = 10$: $\mathbf{X}_n^*$ (red $\star$, $q = 10$), $\mathbf{H}_{n,10}$ (blue ∇), $\mathbf{S}_{n,10}$ (black $\times$). The dotted red lines with + correspond to $\mathbf{X}_n^*$ when $\mathcal{X}_Q$ is restricted to $\mathbf{S}_{Q',d}$.

Figure 7: Left: $CR(\mathbf{X}_n)/R_*(n_{\max}, d)$; Right: $n^{1/d} Q_\alpha(\mathbf{X}_n)$, $n = 10, \dots, 200$, $d = 10$: $\mathbf{X}_n^*$ (red $\star$, $q = 10$), $\mathbf{X}_n^{CH, 2\sqrt{2}d}$ (magenta $\circ$), $\mathbf{X}_n^{CH, \beta_*(n,d)}$ with $\beta_*(n, d)$ given by (17) (blue ∇) and $\mathbf{X}_n^{CH, \infty}$ (black $\times$, and dotted line with black + for the candidate set $\mathbf{S}_{4096,10} \cup \{0, 1\}^d$).

quantile (right panel).

Figure 9 confirms the robustness of the c.d.f.-based method with respect to q , displaying the evolution of both CR and Q_α for $q = 5, 10$ and 25 . It shows that increasing q reduces the covering radius of $\mathbf{X}_n^*$ by giving more weight to points in $\mathcal{X}_Q$ far from the design (i.e., the vertices of $\mathcal{C}_d$), at the detriment of Q_α . Note the larger impact of increasing q from 5 to 10 than from 10 to 25, suggesting that further increase of q will not result in significantly different performances.

Table 1 gives the computational time of methods considered, with $\mathbf{X}_n^{LRD}$ the design obtained by lazy-greedy minimization of $\Psi_q(\mathbf{Z}_n, \mu_Q)$ given by (19). The coffee-house variants are by far

Figure 8: Left: $CR(\mathbf{X}_n)/R_*(n_{\max}, d)$; Right: $n^{1/d} Q_\alpha(\mathbf{X}_n)$, $n = 10, \dots, 200$, $d = 10$; $\mathbf{X}_n^*$ (red $\star$), $\mathbf{X}_n^{VD}$ (blue ∇) and $\mathbf{X}_n^{RD}$ (black $\times$); $q = 10$.

Figure 9: Left: $CR(\mathbf{X}_n^*)/R_*(n_{\max}, d)$; Right: $n^{1/d} Q_\alpha(\mathbf{X}_n^*)$, $n = 10, \dots, 200$, $d = 10$; $q = 5$ (blue ∇), $q = 10$ (red $\star$), and $q = 25$ (black $\times$).

the fastest, followed by the c.d.f.-based method. Whenever numerical complexity must be kept to a minimum, one of the variants of coffee-house design based on spacings generates designs with slightly increased covering radii than $\mathbf{X}_n^*$ but at a much smaller cost.

Table 1: Computational time of $\mathbf{X}_{200}$ (in s, 10 repetitions).

$\mathbf{X}_n^*$	$\mathbf{X}_n^{VD}$	$\mathbf{X}_n^{RD}$	$\mathbf{X}_n^{LRD}$	$\mathbf{X}_n^{CH, \infty}$	$\mathbf{X}_n^{CH, 2\sqrt{2}d}$	$\mathbf{X}_n^{CH, \beta_*(100, d)}$
18.0	32.9	303.8	126.2	0.5	0.9	0.9

We assess now the impact of properties of the set of candidate points by repeating the comparison between $\mathbf{X}_n^*$ and the coffee-house designs $\mathbf{X}_n^{CH, \beta}$ when $\mathcal{X}_C = \mathbf{Z}_{Lh, 100}$, a Latin

hypercube design with maximum packing radius; see <https://spacefillingdesigns.nl/>. We define in this way an order in $\mathbf{Z}_{Lh,100}$, hoping that all prefix designs $\mathbf{Z}_{Lh,1:n} = \{\mathbf{z}_i\}_{i=1}^n$, $\mathbf{z}_i \in \mathbf{Z}_{Lh,100}$, $n \leq 100$, have small covering radius.

Figure 10 shows CR (left panel) and Q_α (right panel) of designs $\mathbf{X}_n^*$ along with those of $\mathbf{X}_n^{CH,\beta}$. The figure shows that for this extremely constrained candidate set $\mathbf{X}_n^*$ (red) still outperforms the other methods¹⁰. This example confirms that the incremental optimization of the c.d.f.-based criterion produces, when applied to a final design with good space-filling properties, a sequence of nested designs with good overall space-filling properties (notice however that better CR values are obtained without the Latin hypercube constraint).

Figure 10: Left: $CR(\mathbf{X}_n)/R_*(n_{\max}, d)$; Right: $n^{1/d} Q_\alpha(\mathbf{X}_n)$, $n = 1, \dots, n_{\max} = 100$, $d = 10$, candidate set $\mathcal{X}_{100} = \mathbf{Z}_{Lh,100}$: $\mathbf{X}_n^*$ (red $\star$, $q = 10$) and $\mathbf{X}_n^{CH,\beta}$ with $\beta = 2\sqrt{2d}$ (magenta $\circ$), $\beta = \beta_*(200, d)$ given by (17) (blue ∇) and $\beta = \infty$ (black $\times$).

Example 2. We investigate in this example the impact of non-convex domains $\mathcal{X}$, by considering an annular geometry, letting $\mathcal{X} = \{\mathbf{x} \in \mathbb{R}^2 : 1/2 \leq \|\mathbf{x}\| \leq 1\}$ (for simplicity, we still base the covering and packing radii on the Euclidean distance, although geodesics could have been used as well). We compare the maximization of the c.d.f.-based criterion to (i) use of prefixes of LDS and (ii) to coffee-house design with $\beta = \infty$. We take $\mathcal{X}_C = \mathcal{X}_Q$ given by the first $C = 2048$ points of a Sobol' sequence of points in $\mathcal{C}_2$ — renormalized to $[-1, 1]^2$ — falling inside $\mathcal{X}$. An analogous construction is used for the set $\mathcal{X}_N$ used to approximate $CR(\mathbf{Z}_n)$, by retaining the first $N = 2^{18}$ points of a renormalized scrambled Sobol' sequence that fall inside $\mathcal{X}$. The LDS design $\mathbf{S}_{n,d}$ gathers the first n elements of $\mathcal{X}_C$.

Spacings are difficult to handle for non-convex domains, explaining why only $\beta = \infty$ is studied in this example. Instead, we consider two distinct candidate sets in the construction of coffee-house designs: $\mathcal{X}_C$, the points of the Sobol' renormalized and clipped sequence, and, in an effort to enforce boundary avoidance, we also use a different candidate set $\mathcal{X}'_C$ which is the restriction of the renormalized Sobol' sequence to the eroded annulus $\mathcal{X}' = \{\mathbf{x} \in \mathbb{R}^2 : 1/2 + r \leq \|\mathbf{x}\| \leq 1 - r\} \subset \mathcal{X}$, with $r = R_*(100, d)/2$.

¹⁰The magenta and blue curves, corresponding to $\beta = 2\sqrt{2d}$ and $\beta = \beta_*(100, d)$, coincide.

Figure 11 shows the 100-point designs obtained. The left panel shows $\mathbf{X}_n^*$ (red ■) and $\mathbf{X}_n^S$ (black ●), and the right panel displays the coffee-house designs $\mathbf{X}_n^{CH,\infty}$ (magenta ■) and $\mathbf{X}_n^{CH',\infty}$ (blue ●) obtained for the candidate sets $\mathcal{X}$ and $\mathcal{X}'$, respectively. Visual inspection shows that they all fill $\mathcal{X}$ reasonably well, although $\mathbf{S}_{n,d}$ has a few pairs of nearly coincident points.

Figure 11: Left: $\mathbf{X}_n^*$ (red ■) and $\mathbf{S}_{n,2}$ (black ●); Right: $\mathbf{X}_n^{CH,\infty}$ (magenta ■) and $\mathbf{X}_n^{CH',\infty}$ (blue ●).

Figure 12 shows the performance of the four designs, in terms of normalized CR (left panel) and Q_α (right panel). The three greedy constructions exhibit comparable performance and clearly outperform the Sobol' sequence for both criteria. $\mathbf{X}_n^*$ and $\mathbf{X}_n^{CH,\infty}$ have slightly smaller covering radii than $\mathbf{X}_n^{CH',\infty}$ for $n \gtrsim 50$: the absence of sharp corners and edges seems to soften the importance of staying away from the boundaries of the domain; notice that $\mathbf{X}_n^{CH,\infty}$ has many points near the boundary while $\mathbf{X}_n^*$ has few, suggesting that boundaries are not necessarily the main concern here.

6 Conclusions

This paper proposes a novel space-filling c.d.f.-based criterion $I_{B,q}$ tailored to the minimization of the covering radius. This set function satisfies the conditions of the celebrated Nemhauser's Theorem on the greedy maximization of submodular functions, thereby opening the possibility to challenge the simple but 50%-efficient coffee-house design methodology. Although the performance guarantees provided by the submodularity of $I_{B,q}$ do not translate into performance guarantees for the covering radius, we verified experimentally, over a larger set of case studies than those presented in the paper, that the design sequences obtained by its greedy maximization have smaller CR values than competing state-of-the-art incremental design methods. The fact that it almost always outperforms coffee-house which has a guaranteed 50% efficiency suggests that an even stronger guarantee may hold for our method.

By fully taking into account how the points of $\mathbf{X}_n$ are collectively embedded in the domain $\mathcal{X}$, $I_{B,q}$ overcomes the limitations due to the locality of $\text{CR}(\mathbf{X}_n)$ mentioned in the introduction. It involves a regularization parameter q and a range parameter B , whose precise values do not significantly impact its ability to express the space-fillingness of a design. The method also

Figure 12: $n^{1/d} \text{CR}(\mathbf{X}_n)$ (left) and $n^{1/d} Q_{\alpha, [\mathcal{X}_N]}(\mathbf{X}_n)$ (right) for the four designs in Figure 11: $\mathbf{X}_n^*$ (red $\star$), $\mathbf{S}_{n,2}$ (black $\times$), $\mathbf{X}_n^{CH, \infty}$ (magenta $\circ$) and $\mathbf{X}_n^{CH', \infty}$ (blue ∇).

requires the specification of a set of points $\mathcal{X}_Q$ to compute the distance c.d.f. and of a candidate set $\mathcal{X}_C$ of eligible design points. For finite q , the method greedily minimizes the L^{q+1} -mean quantization error evaluated on $\mathcal{X}_Q$, which, when $B = \text{diam}(\mathcal{X})$ and $q \rightarrow \infty$, tends to an underestimate of $\text{CR}(\mathbf{X}_n)$ (since based on $\mathcal{X}_Q$ only). We verified that setting $q \in [5, 25]$ and $B = \text{diam}(\mathcal{X})$ yields good designs in a wide variety of situations. The choice of $\mathcal{X}_C$ is not critical, even in the challenging situation where $\mathcal{X}$ is the hypercube $\mathcal{C}_d$, for which the presence of vertices is easily problematic for large d : the method never tries to select points near the vertices and, more generally, is not sensitive to the proximity of points of $\mathcal{X}_C$ to the boundary of $\mathcal{X}$.

The coffee-house algorithm only requires the specification of a candidate set $\mathcal{X}_C$ and runs much faster. However, when $\mathcal{X} = \mathcal{C}_d$ it naturally tends to select points as near as possible to the vertices of $\mathcal{X}$, thereby exhibiting a lack of robustness to the choice of $\mathcal{X}_C$. Boundary avoidance can be enforced through the notion of spacings, resulting in simple and efficient algorithms when $\mathcal{X}$ has a simple enough geometry.

Overall, the geometry of $\mathcal{X}$ appears as being a key factor, including in terms of the relevance of CR to measure the space-filling quality of a design. We thus concur with the authors of [26] on the importance of using a more global characteristic like the covering α -quantile Q_α . Incremental design constructed with $I_{B,q}$ are also competitive in terms of Q_α compared to other incremental design methods.

Note finally that the development of greedy algorithms for the minimization of $Q_\alpha(\mathbf{X}_n)$ itself forms another challenging but promising objective towards the incremental construction of efficient space-filling designs, in particular in the hypercube $\mathcal{C}_d$.

Appendix

Two-point CR-optimal design in $\mathcal{C}_d$. The CR-optimal one-point design is $\mathbf{Z}_1^* = \mathbf{1}_d/2$, i.e., the center of $\mathcal{C}_d$.

Consider the design $\mathbf{Z}_2(\alpha) = \{\mathbf{z}_1(\alpha), \mathbf{z}_2(\alpha) = \mathbf{1}_d - \mathbf{z}_1(\alpha)\}$ with $\mathbf{z}_1(\alpha) = (1/2, \dots, 1/2, \alpha)$, $\alpha < 1/2$. It defines a partition of $\mathcal{C}_d$ into two polyhedral Voronoi cells $\mathcal{C}_1$ and $\mathcal{C}_2$ which are separated by the bisecting hyperplane $\mathcal{H}$ of the segment joining $\mathbf{z}_1(\alpha)$ and $\mathbf{z}_2(\alpha)$, $\mathcal{H} = \{\mathbf{z} \in \mathcal{C}_d : z_d = 1/2\}$. We have $\text{CR}(\mathbf{Z}_2(\alpha)) = \|\mathbf{x}^* - \mathbf{z}_1(\alpha)\|$ for $\mathbf{x}^*$ a vertex of $\mathcal{C}_d$ when $\alpha \geq 1/4$ and for $\mathbf{x}^*$ having all its coordinates in $\{0, 1\}$ except the last one equal to $1/2$ otherwise, implying $\text{CR}(\mathbf{Z}_2(\alpha)) = \max\{[(d-1)/4 + \alpha^2]^{1/2}, [(d-1)/4 + (1/2 - \alpha)^2]^{1/2}\}$, whose minimum value is reached for $\alpha = 1/4$, with $\text{CR}(\mathbf{Z}_2(1/4)) = (1/2) \sqrt{d-3/4} < \sqrt{d}/2$.

Let $\mathbf{Z}_2^* = \{\mathbf{z}_1^*, \mathbf{z}_2^*\}$ denote a 2-point CR-optimal design, and denote by $\mathcal{C}_1^*$ and $\mathcal{C}_2^*$ the two cells of the Voronoi tessellation of $\mathcal{C}_d$ that $\mathbf{Z}_2^*$ generates. Each of them contains exactly 2^{d-1} vertices of $\mathcal{C}_d$, since otherwise one of them, $\mathcal{C}_1^*$ say, would contain two opposite vertices, implying $\text{CR}(\mathbf{Z}_2^*) \geq \text{diam}(\mathcal{C}_1^*)/2 = \sqrt{d}/2 > \text{CR}(\mathbf{Z}_2(1/4))$. We have $\text{CR}(\mathbf{Z}_2^*) \geq \max\{\text{diam}(\mathcal{C}_1^*)/2, \text{diam}(\mathcal{C}_2^*)/2\}$. The minimum is obtained when the two cells have the same diameter and are such that this diameter is minimal, which gives $\text{CR}(\mathbf{Z}_2^*) = \text{CR}(\mathbf{Z}_2(1/4))$.

Consider now an incremental greedy construction, one-step-ahead optimal. Since any 2-point design of the form $\mathbf{Z}_2 = \{\mathbf{1}_d/2, \mathbf{z}\}$ is such that $\text{CR}(\mathbf{Z}_2) = \sqrt{d}/2$, we necessarily have $\text{CR}(\mathbf{X}_2) = \sqrt{d}/2 > \text{CR}(\mathbf{Z}_2^*)$. Note that if we could add two points at a time, we could reach $\text{CR}(\{\mathbf{1}_d/2, \mathbf{x}_2, \mathbf{x}_3\}) = \text{CR}(\mathbf{Z}_2^*)$ by taking $\mathbf{x}_2 = \mathbf{z}_1^*$, $\mathbf{x}_3 = \mathbf{z}_2^*$, and even $\text{CR}(\{\mathbf{1}_d/2, \hat{\mathbf{x}}_2, \hat{\mathbf{x}}_3\}) = (1/2) \sqrt{d-8/9}$ by choosing $\hat{\mathbf{x}}_2 = (1/2, \dots, 1/2, 1/6)$ and $\hat{\mathbf{x}}_3 = \mathbf{1}_d - \hat{\mathbf{x}}_2$.

Evaluation of $\text{CR}(\mathbf{x}_n)$ when $\mathcal{X} = \mathcal{C}_d$. For $d \leq 4$, the exact values of $\text{CR}(\mathbf{X}_n)$ are calculated using Voronoi tessellations; see [28]; for larger d , we underestimate $\text{CR}(\mathbf{X}_n)$ by $\text{CR}_{[\mathcal{X}_N]}(\mathbf{X}_n)$ given by (1), with $\mathcal{X}_N$ a finite subset of $\mathcal{X}$. When $\mathcal{X}_N$ corresponds to the first N points of a (t, s) -sequence in base β in $\mathcal{X}$, we have

$$\text{CR}(\mathcal{X}_N) < \overline{\text{CR}}(\mathcal{X}_N) = \frac{\sqrt{d} \beta^{1+t/d}}{N^{1/d}}, \quad (18)$$

see [25, Th. 6.11]. The values of t for the sequences of Sobol' (for which $\beta = 2$) are given in Table 2; see [38]. The sequences of Faure [8] have $t = 0$ but β is the first prime number larger than or equal to d , making the upper bound in (18) larger than for Sobol'. Figure 13 shows $\overline{\text{CR}}(\mathcal{X}_N)$ as a function of d for both types of sequences when $N = 10^6$, showing that the inequality $\text{CR}(\mathbf{X}_n) \leq \text{CR}_{[\mathcal{X}_Q]}(\mathbf{X}_n) + \sqrt{d} \beta^{1+t/d} N^{-1/d}$ implied by (2) is of limited practical use. For that reason, for $d > 4$ we have chosen to only indicate the value of $\text{CR}_{[\mathcal{X}_N]}(\mathbf{X}_n)$ for $\mathcal{X}_N$ given by the union of $\mathbf{S}_{N_S}$, the first N_S points of a (possibly scrambled) Sobol' sequence, and a 2^d full factorial design (i.e., the vertices of $\mathcal{X}$), with thus $N = N_S + 2^d$.

Table 2: Values of $t(d)$ for Sobol' sequence, for $2 \leq d \leq 13$.

d	2	3	4	5	6	7	8	9	10	11	12	13
t	0	1	3	5	8	11	15	19	23	27	31	35

Proof of Theorem 2. We first prove by induction that $P_\beta(\mathbf{X}_{n+1}) = (1/2) S_\beta(\mathbf{X}_n)$ for $n \geq 1$. Since $d(\mathbf{x}_2, \partial\mathcal{X}) \leq d(\mathbf{x}_1, \partial\mathcal{X})$, we have

$$P_\beta(\mathbf{X}_2) = \frac{1}{2} \min\{\beta d(\mathbf{x}_1, \partial\mathcal{X}), \beta d(\mathbf{x}_2, \partial\mathcal{X}), \|\mathbf{x}_1 - \mathbf{x}_2\|\} = \frac{1}{2} \min\{\beta d(\mathbf{x}_2, \partial\mathcal{X}), \|\mathbf{x}_1 - \mathbf{x}_2\|\} = \frac{1}{2} S_\beta(\mathbf{X}_1),$$

Figure 13: Left: $\beta_*(n, d)$ given by (17) as a function of d for $n = 50, 100, 200$; the magenta curve with $\times$ is for the value $\beta = 2\sqrt{2d}$ suggested in [37]. Right: upper bound $CR(\mathcal{X}_N)$ given by (18) as a function of d for the sequences of Faure and Sobol' for $N = 10^6$.

which proves the property for $n = 1$. Assume that $P_\beta(\mathbf{X}_n) = (1/2) S_\beta(\mathbf{X}_{n-1})$. By construction, we have $P_\beta(\mathbf{X}_{n+1}) \leq (1/2) S_\beta(\mathbf{X}_n)$; we show that $P_\beta(\mathbf{X}_{n+1}) \geq (1/2) S_\beta(\mathbf{X}_n)$. The greedy construction of $\mathbf{X}_{n+1}$ implies that $d(\mathbf{x}_{n+1}, \mathbf{X}_n) \geq S_\beta(\mathbf{X}_n)$ and $\beta d(\mathbf{x}_{n+1}, \partial\mathcal{X}) \geq S_\beta(\mathbf{X}_n)$. Also, any $\mathbf{x}_i \neq \mathbf{x}_j \in \mathbf{X}_n$ satisfy

$$\|\mathbf{x}_i - \mathbf{x}_j\| \geq 2 P_\beta(\mathbf{X}_n) = S_\beta(\mathbf{X}_{n-1}) \geq S_\beta(\mathbf{X}_n),$$

where the first inequality comes from the definition of P_β , the equality comes from the induction hypothesis and the second inequality from the fact that S_β is nonincreasing. Similarly, $\beta d(\mathbf{x}_i, \partial\mathcal{X}) \geq 2 P_\beta(\mathbf{X}_n) \geq S_\beta(\mathbf{X}_n)$ for all $\mathbf{x}_i \in \mathbf{X}_n$. Altogether, this gives $P_\beta(\mathbf{X}_{n+1}) \geq (1/2) S_\beta(\mathbf{X}_n)$.

Let $\mathbf{Z}_n^*$ and $\mathbf{Z}_n^{**}$ be two n -point designs in $\mathcal{X}$ such that $S_\beta(\mathbf{Z}_n^*) = S_{\beta,n}^* = \min_{\mathbf{Z}_n \subset \mathcal{X}} S_\beta(\mathbf{Z}_n)$ ($n \geq 1$) and $P_\beta(\mathbf{Z}_n^{**}) = P_{\beta,n}^* = \max_{\mathbf{Z}_n \subset \mathcal{X}} P_\beta(\mathbf{Z}_n)$ ($n \geq 2$). Denote by $\mathcal{X}^-(r)$ the erosion of $\mathcal{X}$ by $\mathcal{B}(\mathbf{0}, r)$, $\mathcal{X}^-(r) = \mathcal{X} \ominus \mathcal{B}(\mathbf{0}, r) = \{\mathbf{x} \in \mathcal{X} : \mathbf{x} + \mathbf{z} \in \mathcal{X}, \forall \mathbf{z} \in \mathcal{B}(\mathbf{0}, r)\}$, with $\mathcal{X}^-(r) \neq \emptyset$ for any $r \leq \text{diam}(\mathcal{X})/2$.

Suppose that the n balls $\mathcal{B}(\mathbf{x}_i, S_\beta(\mathbf{X}_n))$ do not cover $\mathcal{X}^-(S_\beta(\mathbf{X}_n)/\beta)$. This would imply the existence of $\mathbf{x}^* \in \mathcal{X}^-(S_\beta(\mathbf{X}_n)/\beta)$ such that $\beta d(\mathbf{x}^*, \partial\mathcal{X}) > S_\beta(\mathbf{X}_n)$ and $d(\mathbf{x}^*, \mathbf{X}_n) > S_\beta(\mathbf{X}_n)$, and therefore $D_\beta(\mathbf{x}^*, \mathbf{X}_n) > S_\beta(\mathbf{X}_n)$. This is impossible since $S_\beta(\mathbf{X}_n) = \max_{\mathbf{x} \in \mathcal{X}} D_\beta(\mathbf{x}^*, \mathbf{X}_n)$.

We prove that $S_\beta(\mathbf{X}_n) \leq 2 S_{\beta,n}^*$. Take any $\mathbf{x}_i \in \mathbf{X}_{n+1}$. It satisfies $\mathbf{x}_i \in \mathcal{X}^-(2 P_\beta(\mathbf{X}_{n+1})/\beta)$ from the definition of P_β , and thus $\mathbf{x}_i \in \mathcal{X}^-(S_\beta(\mathbf{X}_n)/\beta)$ since $P_\beta(\mathbf{X}_{n+1}) = (1/2) S_\beta(\mathbf{X}_n)$, and $\mathbf{x}_i \in \mathcal{X}^-(S_{\beta,n}^*/\beta)$ since $S_{\beta,n}^* \leq S_\beta(\mathbf{X}_n)$. Since the n balls $\mathcal{B}(\mathbf{z}_i^*, S_\beta(\mathbf{X}_n))$, $\mathbf{z}_i^* \in \mathbf{Z}_n^*$, cover $\mathcal{X}^-(S_{\beta,n}^*/\beta)$, this implies the existence of a $\mathbf{z}_\ell^* \in \mathbf{Z}_n^*$ and of $\mathbf{x}_j \neq \mathbf{x}_k \in \mathbf{X}_{n+1}$ such that $\mathbf{x}_j, \mathbf{x}_k \in \mathcal{B}(\mathbf{z}_\ell^*, S_{\beta,n}^*)$. Therefore, $S_\beta(\mathbf{X}_n) = 2 P_\beta(\mathbf{X}_{n+1}) \leq \|\mathbf{x}_i - \mathbf{x}_j\| \leq 2 S_{\beta,n}^*$.

We now prove that $P_{\beta,n+1}^* \leq 2 P_\beta(\mathbf{X}_{n+1})$. Take any $\mathbf{z}_i^{**} \in \mathbf{Z}_{n+1}^{**}$; $\mathbf{z}_i^{**} \in \mathcal{X}^-(2 P_{\beta,n+1}^*/\beta)$ from the definition of P_β , $\mathbf{z}_i^{**} \in \mathcal{X}^-(2 P_\beta(\mathbf{X}_{n+1})/\beta) = \mathcal{X}^-(S_\beta(\mathbf{X}_n)/\beta)$ since $P_{\beta,n+1}^* \geq P_\beta(\mathbf{X}_{n+1}) = (1/2) S_\beta(\mathbf{X}_n)$, and $\mathbf{z}_i^{**} \in \mathcal{X}^-(S_{\beta,n}^*/\beta)$ since $S_{\beta,n}^* \leq S_\beta(\mathbf{X}_n)$. Therefore, there exist $\mathbf{z}_\ell^* \in \mathbf{Z}_n^*$ and $\mathbf{z}_j^{**} \neq \mathbf{z}_k^{**} \in \mathbf{Z}_{n+1}^{**}$ such that $\mathbf{z}_j^{**}, \mathbf{z}_k^{**} \in \mathcal{B}(\mathbf{z}_\ell^*, S_{\beta,n}^*)$. It implies that $2 P_{\beta,n+1}^* \leq 2 S_{\beta,n}^* \leq 2 S_\beta(\mathbf{X}_n) = 4 P_\beta(\mathbf{X}_{n+1})$.

Finally, we show that $1 \leq \rho_\beta(\mathbf{X}_n) \leq 2$. Since $\mathcal{X}$ is convex, for any $\mathbf{X}_n$, $n \geq 2$, the n balls $\mathcal{B}(\mathbf{x}_i, P_\beta(\mathbf{X}_n))$ do not cover $\mathcal{X}^-(P_\beta(\mathbf{X}_n)/\beta)$, which implies the existence of a $\mathbf{x} \in \mathcal{X}^-(P_\beta(\mathbf{X}_n)/\beta) \setminus \cup_{i=1}^n \mathcal{B}(\mathbf{x}_i, P_\beta(\mathbf{X}_n))$. It satisfies $d(\mathbf{x}, \partial\mathcal{X}) \geq P_\beta(\mathbf{X}_n)/\beta$ and $d(\mathbf{x}, \mathbf{X}_n) \geq P_\beta(\mathbf{X}_n)/\beta$; therefore, $D_\beta(\mathbf{x}, \mathbf{X}_n) \geq P_\beta(\mathbf{X}_n)$ and $S_\beta(\mathbf{X}_n) \geq P_\beta(\mathbf{X}_n)$, so that $\rho_\beta(\mathbf{X}_n) \geq 1$. Any design obtained with the greedy algorithm satisfies $S_\beta(\mathbf{X}_{n+1}) \leq S_\beta(\mathbf{X}_n) = 2P_\beta(\mathbf{X}_{n+1})$, $n \geq 1$, so that $\rho_\beta(\mathbf{X}_{n+1}) \leq 2$. $\blacksquare$

Minimization of a relaxed version of $\text{CR}(\mathbf{Z}_n)$. In [32], the following relaxed version of the coverage criterion of [34] is considered:

$$\Psi_q(\mathbf{Z}_n) = \Psi_q(\mathbf{Z}_n; \mu) = \left[\int_{\mathcal{X}} \left(\frac{1}{n} \sum_{i=1}^n \|\mathbf{z}_i - \mathbf{x}\|^{-q} \right)^{-1} \mu(d\mathbf{x}) \right]^{1/q}, \quad q \neq 0, \quad (19)$$

with μ the uniform probability measure on $\mathcal{X}$. For any n -point design $\mathbf{Z}_n$, it satisfies $\Psi_q(\mathbf{Z}_n) \rightarrow \text{CR}(\mathbf{Z}_n)$ as $q \rightarrow \infty$. If $\xi_n = (1/n) \sum_{i=1}^n \delta_{\mathbf{z}_i}$ denotes the empirical measure associated with $\mathbf{Z}_n$, we can define $\psi_q(\xi_n) = \Psi_q(\mathbf{Z}_n)$ and more generally

$$\psi_q(\xi) = \psi_q(\xi; \mu) = \left[\int_{\mathcal{X}} \left(\int_{\mathcal{X}} \|\mathbf{z} - \mathbf{x}\|^{-q} \xi(d\mathbf{z}) \right)^{-1} \mu(d\mathbf{x}) \right]^{1/q}, \quad q \neq 0,$$

for any probability measure ξ on $\mathcal{X}$. It is shown in [32] that $\psi_q^q(\cdot)$ is a convex functional for $q > 0$ (strictly convex for $q \in (0, d)$); for reasons explained in that paper, q should be taken larger than $\max\{0, d-2\}$ and we use $q = d$ in the examples below. The criterion $\psi_q(\xi)$ can be minimized by the conditional (or constrained) gradient method of [9]; when the step size at iteration k equals $1/(k+1)$, the method corresponds to Wynn's Vertex-Direction (VD) method [39] of approximate design theory (in practice, a discrete measure μ_Q supported on $\mathcal{X}_Q \subset \mathcal{X}$ is substituted for μ , and the minimization of $\psi_q(\xi)$ corresponds to an A -optimal design problem). Taking the initial measure $\xi^{(1)}$ as the one-point measure supported at $\mathbf{z}_1 \in \mathcal{X}$ and using iteratively $\xi^{(k+1)} = [k/(k+1)] \xi^{(k)} + \delta_{\mathbf{z}_{k+1}}/(k+1)$, then, for each $n \geq 1$, $\xi^{(n)}$ has n support points which define an incremental design $\mathbf{X}_n$. Here, $\mathbf{z}_1 \in \text{Arg min}_{\mathbf{z} \in \mathcal{X}} \int_{\mathcal{X}} \|\mathbf{z} - \mathbf{x}\|^q \mu(d\mathbf{x})$ and $\mathbf{z}_{k+1}$ minimizes the directional derivative of $\psi_q^q(\cdot)$ at $\xi^{(k)}$ in the direction of the delta measure $\delta_{\mathbf{z}}$ with respect to $\mathbf{z}$. When a finite candidate set $\mathcal{X}_C$ is used, this corresponds to

$$\mathbf{z}_{k+1} \in \text{Arg max}_{\mathbf{z} \in \mathcal{X}_C} \int_{\mathcal{X}} \left[\|\mathbf{z} - \mathbf{x}\|^{-q} \left(\int_{\mathcal{X}} \|\mathbf{y} - \mathbf{x}\|^{-q} \xi^{(k)}(d\mathbf{y}) \right)^{-2} \right] \mu(d\mathbf{x});$$

see [32]. We shall denote by $\mathbf{X}_n^{VD}$ the designs constructed in this manner, with μ_Q substituted for μ .

Straightforward calculations indicate that, for $q > 0$, $(1/n) \Psi_q^q(\mathbf{Z}_n, \mu_Q)$, a discrete unnormalized form of $\Psi_q^q(\mathbf{Z}_n; \mu)$ defined in (19), is non-increasing and supermodular (i.e., $-(1/n) \Psi_q^q(\mathbf{Z}_n, \mu_Q)$ is submodular). Theorem 1 thus applies to the greedy maximization of $-(1/n) \Psi_q^q(\mathbf{Z}_n, \mu_Q)$; we shall denote by $\mathbf{X}_n^{RD}$ a design obtained by greedy minimization of $\Psi_q(\mathbf{Z}_n, \mu_Q)$, a Relaxed and Discretized version of $\text{CR}(\mathbf{Z}_n)$. The lazy-greedy algorithm of Section 2.2 can also be applied, see Table 1.

Declarations

Funding

This work was partly supported by project INDEX (INcremental Design of EXperiments) ANR-18-CE91-0007 of the French National Research Agency (ANR).

Conflict of interest

On behalf of all authors, the corresponding author states that there are no conflicts of interest.

Code availability

The matlab scripts implementing the greedy and lazy greedy algorithms for the maximization of the covering criterion $I_{B,q}$ are available at <https://sdb3.i3s.unice.fr/anrindex>.

References

- [1] Bach, F.: Learning with submodular functions: a convex optimization perspective. *Foundations and Trends in Machine Learning* **6**(2-3), 145–373 (2013)
- [2] Beck, J., Guillas, S.: Sequential design with mutual information for computer experiments (MICE): emulation of a tsunami model. *SIAM/ASA Journal on Uncertainty Quantification* **4**, 739–766 (2016)
- [3] Blum, H.: A transformation for extracting new descriptors of shape. In: W. Whalen-Dunn (ed.) *Models for the Perception of Speech and Visual Forms*, pp. 362–380. MIT Press, Cambridge (1967)
- [4] Bondarenko, A., Hardin, D., Saff, E.: Mesh ratios for best-packing and limits of minimal energy configurations. *Acta Mathematica Hungarica* **142**(1), 118–131 (2014)
- [5] De Marchi, S.: On optimal center locations for radial basis function interpolation: computational aspects. *Rend. Sem. Math. Torino* **61**(3), 343–358 (2003)
- [6] Fang, K.T., Li, R., Sudjianto, A.: *Design and Modeling for Computer Experiments*. Chapman & Hall/CRC, Boca Raton (2006)
- [7] Cabral Farias, R., Pronzato, L., Rendas, M.J.: Incremental construction of nested designs based on two-level fractional factorial designs. HAL preprint hal-02483004 (2020)
- [8] Faure, H.: Discrepances de suites associées à un système de numération (en dimension s). *Acta Arithmetica* **41**, 337–351 (1982)
- [9] Frank, M., Wolfe, P.: An algorithm for quadratic programming. *Naval Res. Logist. Quart.* **3**, 95–110 (1956)
- [10] Gonzalez, T.: Clustering to minimize the maximum intercluster distance. *Theoretical Computer Science* **38**, 293–306 (1985)

- [11] Graf, S., Luschgy, H.: Foundations of Quantization for Probability Distributions. Springer, Berlin (2000)
- [12] Harari, O., Bingham, D., Dean, A., Higdon, D.: Computer experiments: prediction accuracy, sample size and model complexity revisited. *Statistica Sinica* **28**, 899–919 (2018)
- [13] Janson, S.: Maximal spacings in several dimensions. *The Annals of Probability* **15**(1), 274–280 (1987)
- [14] Johnson, M., Moore, L., Ylvisaker, D.: Minimax and maximin distance designs. *Journal of Statistical Planning and Inference* **26**, 131–148 (1990)
- [15] Kennard, R., Stone, L.: Computer aided design of experiments. *Technometrics* **11**(1), 137–148 (1969)
- [16] Korte, B., Vygen, J.: Combinatorial Optimization: Theory and Algorithms (5 ed.). Springer, New York (2012)
- [17] Krause, A., Singh, A., Guestrin, C.: Near-optimal sensor placements in gaussian processes: theory, efficient algorithms and empirical studies. *Journal of Machine Learning Research* **9**, 235–284 (2008)
- [18] Lekivetz, R., Jones, B.: Fast flexible space-filling designs for nonrectangular regions. *Quality and Reliability Engineering International* **31**(5), 829–837 (2015)
- [19] Loeppky, J., Sacks, J., Welch, W.: Choosing the sample size of a computer experiment: a practical guide. *Journal of the American Statistical Association* **51**(4), 366–376 (2009)
- [20] Luschgy, H., Pagès, G.: Greedy vector quantization. *Journal of Approximation Theory* **198**, 111–131 (2015)
- [21] Minoux, M.: Accelerated greedy algorithms for maximizing submodular set functions. In: *Proc. 8th IFIP Conference, Wurzburg (Part 2)*, pp. 234–243. Springer, New-York (1977)
- [22] Müller, W.: Coffee-house designs. In: A. Atkinson, B. Bogacka, A. Zhigljavsky (eds.) *Optimum Design 2000*, chap. 21, pp. 241–248. Kluwer, Dordrecht (2001)
- [23] Müller, W.: *Collecting Spatial Data*. Springer, Berlin (2007). [3rd ed.]
- [24] Nemhauser, G., Wolsey, L., Fisher, M.: An analysis of approximations for maximizing submodular set functions–I. *Mathematical Programming* **14**(1), 265–294 (1978)
- [25] Niederreiter, H.: *Random Number Generation and Quasi-Monte Carlo Methods*. SIAM, Philadelphia (1992)
- [26] Noonan, J., Zhigljavsky, A.: Non-lattice covering and quantization of high dimensional sets. *arXiv preprint arXiv:2006.02705v1* (2020)
- [27] Noonan, J., Zhigljavsky, A.: Efficient quantization and weak covering of high dimensional cubes. In: *Black Box Optimization, Machine Learning, and No-Free Lunch Theorems*. Springer (2021). *arXiv preprint arXiv:2005.07938*

- [28] Pronzato, L.: Minimax and maximin space-filling designs: some properties and methods for construction. *Journal de la Société Française de Statistique* **158**(1), 7–36 (2017)
- [29] Pronzato, L.: Sensitivity analysis via Karhunen-Loève expansion of a random field model: estimation of Sobol’ indices and experimental design. *Reliability Engineering and System Safety* **187**, 93–109 (2019). Hal-01545604v2
- [30] Pronzato, L., Müller, W.: Design of computer experiments: space filling and beyond. *Statistics and Computing* **22**, 681–701 (2012)
- [31] Pronzato, L., Rendas, J.: A submodular criterion for space-filling design. *SIAM Conference on Uncertainty Quantification*, EPFL, Lausanne, Switzerland, April 5-8 (2016)
- [32] Pronzato, L., Zhigljavsky, A.: Measures minimizing regularized dispersion. *Journal of Scientific Computing* **78**(3), 1550–1570 (2019)
- [33] Pronzato, L., Zhigljavsky, A.: Bayesian quadrature, energy minimization and space-filling design. *SIAM/ASA J. Uncertainty Quantification* **8**(3), 959–1011 (2020)
- [34] Royle, J., Nychka, D.: An algorithm for the construction of spatial coverage designs with implementation in SPLUS. *Computers and Geosciences* **24**(5), 479–488 (1998)
- [35] Santner, T., Williams, B., Notz, W.: *The Design and Analysis of Computer Experiments*. Springer, Heidelberg (2003)
- [36] Schaback, R.: Error estimates and condition numbers for radial basis function interpolation. *Advances in Computational Mathematics* **3**(3), 251–264 (1995)
- [37] Shang, B., Apley, D.: Full-sequential space-filling design algorithms for computer experiments. *Journal of Quality Technology* (2020). <https://doi.org/10.1080/00224065.2019.1705207>
- [38] Sobol’, I.: On the distribution of points in a cube and the approximate evaluation of integrals. *USSR Computational Mathematics and Mathematical Physics* **7**(4), 86–112 (1967)
- [39] Wynn, H.: The sequential generation of D -optimum experimental designs. *Annals of Math. Stat.* **41**, 1655–1664 (1970)
- [40] Zhigljavsky, A., Noonan, J.: Covering of high-dimensional cubes and quantization. *SN Operations Research Forum* **1**(3), 1–32 (2020). arXiv preprint arXiv:2002.06118