
Elementary Symmetric Polynomials for Optimal Experimental Design

Zelda Mariet

Massachusetts Institute of Technology
Cambridge, MA 02139
zelda@csail.mit.edu

Suvrit Sra

Massachusetts Institute of Technology
Cambridge, MA 02139
suvrit@mit.edu

Abstract

We revisit the classical problem of optimal experimental design (OED) under a new mathematical model grounded in a geometric motivation. Specifically, we introduce models based on elementary symmetric polynomials; these polynomials capture “partial volumes” and offer a graded interpolation between the widely used A-optimal design and D-optimal design models, obtaining each of them as special cases. We analyze properties of our models, and derive both greedy and convex-relaxation algorithms for computing the associated designs. Our analysis establishes approximation guarantees on these algorithms, while our empirical results substantiate our claims and demonstrate a curious phenomenon concerning our greedy method. Finally, as a byproduct, we obtain new results on the theory of elementary symmetric polynomials that may be of independent interest.

1 Introduction

Optimal Experimental Design (OED) develops the theory of selecting experiments to perform in order to estimate a hidden parameter as well as possible. It operates under the assumption that experiments are costly and cannot be run as many times as necessary or run even once without tremendous difficulty [33]. OED has been applied in a large number of experimental settings [35, 9, 28, 46, 36], and has close ties to related machine-learning problems such as outlier detection [15, 22], active learning [19, 18], Gaussian process driven sensor placement [27], among others.

We revisit the classical setting where each experiment depends linearly on a hidden parameter $\theta \in \mathbb{R}^m$. We assume there are n possible experiments whose outcomes $y_i \in \mathbb{R}$ can be written as

$$y_i = x_i^\top \theta + \epsilon_i \quad 1 \leq i \leq n,$$

where the $x_i \in \mathbb{R}^m$ and ϵ_i are independent, zero mean, and homoscedastic noises. OED seeks to answer the question: how to choose a set S of k experiments that allow us to estimate θ without bias and with minimal variance?

Given a feasible set S of experiments (i.e., $\sum_{i \in S} x_i x_i^\top$ is invertible), the Gauss-Markov theorem shows that the lowest variance for an unbiased estimate $\hat{\theta}$ satisfies $\text{Var}[\hat{\theta}] = (\sum_{i \in S} x_i x_i^\top)^{-1}$. However, $\text{Var}[\hat{\theta}]$ is a matrix, and matrices do not admit a total order, making it difficult to compare different designs. Hence, OED is cast as an optimization problem that seeks an optimal design S^*

$$S^* \in \underset{S \in [n], |S| \leq k}{\text{argmin}} \Phi \left(\left(\sum_{i \in S} x_i x_i^\top \right)^{-1} \right), \quad (1.1)$$

where Φ maps positive definite matrices to \mathbb{R} to compare the variances for each design, and may help elicit different properties that a solution should satisfy, either statistical or structural.

Elfving [16] derived some of the earliest theoretical results for the linear dependency setting, focusing on the case where one is interested in reconstructing a predefined linear combination of the

underlying parameters $c^\top \theta$ (C-optimal design). Kiefer [26] introduced a more general approach to OED, by considering matrix means on positive definite matrices as a general way of evaluating optimality [33, Ch. 6], and Yu [48] derived general conditions for a map Φ under which a class of multiplicative algorithms for optimal design has guaranteed monotonic convergence.

Nonetheless, the theory of OED branches into multiple variants of (1.1) depending on the choice of Φ , among which A-optimal design ($\Phi = \text{trace}$) and D-optimal design ($\Phi = \text{determinant}$) are probably the two most popular choices. Each of these choices has a wide range of applications as well as statistical, algorithmic, and other theoretical results. We refer the reader to the classic book [33], which provides an excellent overview and introduction to the topic; see also the summaries in [1, 35].

For A-optimal design, recently Wang et al. [44] derived greedy and convex-relaxation approaches; [11] considers the problem of constrained adaptive sensing, where θ is supposed sparse. D-optimal design has historically been more popular, with several approaches to solving the related optimization problem [17, 38, 31, 20]. The dual problem of D-optimality, Minimum Volume Covering Ellipsoid (MVCE) is also a well-known and deeply studied optimization problem [3, 34, 43, 41, 14, 42]. Experimental design has also been studied in more complex settings: [8] considers Bayesian optimal design; under certain conditions, non-linear settings can be approached with linear OED [13, 25].

Due to the popularity of A- and D-optimal design, the theory surrounding these two sub-problems has diverged significantly. However, both the trace and the determinant are special cases of fundamental spectral polynomials of matrices: *elementary symmetric polynomials (ESP)*, which have been extensively studied in matrix theory, combinatorics, information theory, and other areas due to their importance in the theory of polynomials [24, 30, 21, 6, 23, 4].

These considerations motivate us to derive a broader view of optimal design which we call *ESP-Design*, where Φ is obtained from an elementary symmetric polynomial. This allows us to consider A-optimal design and D-optimal design as special cases of ESP-design, and thus treat the entire ESP-class in a unified manner. Let us state the key contributions of this paper more precisely below.

Contributions

- We introduce ESP-design, a new, general framework for OED that leverages geometric properties of positive definite matrices to interpolate between A- and D-optimality. ESP-design offers an intuitive setting in which to gradually scale between A-optimal and D-optimal design.
- We develop a convex relaxation as well as greedy algorithms to compute the associated designs. As a byproduct of our convex relaxation, we prove that ESPs are geodesically log-convex on the Riemannian manifold of positive definite matrices; this result may be of independent interest.
- We extend a result of Avron and Boutsidis [2] on determinantal column-subset selection to ESPs; as a consequence we obtain a greedy algorithm with provable optimality bounds for ESP-design.

Experiments on synthetic and real data illustrate the performance of our algorithms and confirm that ESP-design can be used to obtain designs with properties that scale between those of both A- and D-optimal designs, allowing users to tune trade-offs between their different benefits (e.g. predictive error, sparsity, etc.). We show that our greedy algorithm generates designs of equal quality to the famous Fedorov exchange algorithm [17], while running in a fraction of the time.

2 Preliminaries

We begin with some background material that also serves to set our notation. We omit proofs for brevity, as they can be found in standard sources such as [6].

We define $[n] \triangleq \{1, 2, \dots, n\}$. For $S \subseteq [n]$ and $M \in \mathbb{R}^{n \times m}$, we write M_S the $|S| \times m$ matrix created by keeping only the rows of M indexed by S , and $M[S|S']$ the submatrix with rows indexed by S and columns indexed by S' ; by $x_{(i)}$ we denote the vector x with its i -th component removed. For a vector $v \in \mathbb{R}^m$, the *elementary symmetric polynomial (ESP)* of order $\ell \in \mathbb{N}$ is defined by

$$e_\ell(v) \triangleq \sum_{1 \leq i_1 < \dots < i_\ell \leq m} \prod_{j=1}^{\ell} v_{i_j} = \sum_{I \subseteq [m], |I|=\ell} \prod_{j \in I} v_j, \quad (2.1)$$

where $e_\ell \equiv 0$ for $\ell = 0$ and $\ell > m$. Let \mathbb{S}_m^+ (\mathbb{S}_m^{++}) be the cone of positive semidefinite (positive definite) matrices of order m . We denote by $\lambda(M)$ the eigenvalues (in decreasing order) of a symmetric matrix M . Def. (2.1) extends to matrices naturally; ESPs are *spectral functions*, as we set

$E_\ell(M) \triangleq e_\ell \circ \lambda(M)$; additionally, they enjoy another representation that allows us to interpret them as “partial volumes”, namely,

$$E_\ell(M) = \sum_{S \subseteq [n], |S|=\ell} \det(M[S|S]). \quad (2.2)$$

The following proposition captures basic properties of ESPs that we will require in our analysis.

Proposition 2.1. *Let $M \in \mathbb{R}^{m \times m}$ be symmetric and $1 \leq \ell \leq m$; also let $A, B \in \mathbb{S}_m^+$. We have the following properties: (i) If $A \succeq B$ in Löwner order, then $E_\ell(A) \geq E_\ell(B)$; (ii) If M is invertible, then $E_\ell(M^{-1}) = \det(M^{-1})E_{m-\ell}(M)$; (iii) $\nabla e_\ell(\lambda) = [e_{\ell-1}(\lambda_{(i)})]_{1 \leq i \leq m}$.*

3 ESP-design

A-optimal design uses $\Phi \equiv \text{tr}$ in (1.1), and thus selects designs with low average variance. Geometrically, this translates into selecting confidence ellipsoids whose bounding boxes have a small diameter. Conversely, D-optimal design uses $\Phi \equiv \det$ in (1.1), and selects vectors that correspond to the ellipsoid with the smallest volume; as a result it is more sensitive to outliers in the data¹. We introduce a natural model that scales between A- and D-optimal design. Indeed, by recalling that both the trace and the determinant are special cases of ESPs, we obtain a new model as fundamental as A- and D-optimal design, while being able to interpolate between the two in a graded manner.

Unless otherwise indicated, we consider that we are selecting experiments *without* repetition.

3.1 Problem formulation

Let $X \in \mathbb{R}^{n \times m}$ ($m \ll n$) be a design matrix with full column rank, and $k \in \mathbb{N}$ be the budget ($m \leq k \leq n$). Define $\Gamma_k = \{S \subseteq [n] \text{ s.t. } |S| \leq k, X_S^\top X_S \succ 0\}$ to be the set of *feasible* designs that allow unbiased θ estimates. For $\ell \in \{1, \dots, m\}$, we introduce the *ESP-design* model:

$$\min_{S \in \Gamma_k} f_\ell(S) \triangleq \frac{1}{\ell} \log E_\ell((X_S^\top X_S)^{-1}). \quad (3.1)$$

We keep the $1/\ell$ -factor in (3.1) to highlight the homogeneity (E_ℓ is a polynomial of degree ℓ) of our design criterion, as is advocated in [33, Ch. 6].

For $\ell = 1$, (3.1) yields A-optimal design, while for $\ell = m$, it yields D-optimal design. For $1 < \ell < m$, ESP-design interpolates between these two extremes. Geometrically, we may view it as seeking an ellipsoid with the smallest average volume for ℓ -dimensional slices (taken across sets of size ℓ). Alternatively, ESP-design can be also be interpreted as a regularized version of D-optimal design via Prop. 2.1-(ii). In particular, for $\ell = m - 1$, we recover a form of regularized D-optimal design:

$$f_{m-1}(S) = \frac{1}{m-1} [\log \det((X_S^\top X_S)^{-1}) + \log \|X_S\|_2^2].$$

(3.1) is a known hard combinatorial optimization problem (in particular for $\ell = m$), which precludes an exact optimal solution. However, its objective enjoys remarkable properties that help us derive efficient algorithms for its approximate solution. The first one of these is based on a natural convex relaxation obtained below.

3.2 Continuous relaxation

We describe below a traditional approach of relaxing (3.1) by relaxing the constraint on S , allowing elements in the set to have fractional multiplicities. The new optimization problem takes the form

$$\min_{z \in \Gamma_k^c} \frac{1}{\ell} \log E_\ell\left((X^\top \text{Diag}(z)X)^{-1}\right), \quad (3.2)$$

where we Γ_k^c denotes the set of vectors $\{z \in \mathbb{R}^n \mid 0 \leq z_i \leq 1\}$ such that $X^\top \text{Diag}(z)X$ remains invertible and $\mathbf{1}^\top z \leq k$. The following is a direct consequence of Prop 2.1-(i):

Proposition 3.1. *Let z^* be the optimal solution to (3.2). Then $\|z^*\|_1 = k$.*

Convexity of f_ℓ on Γ_k^c (where by abuse of notation, f_ℓ also denotes the continuous relaxation in (3.2)) can be obtained as a consequence of [32]; however, we obtain it as a corollary Lemma 3.3, which shows that $\log E_\ell$ is geodesically convex; this result seems to be new, and is *stronger* than convexity of f_ℓ ; hence it may be of independent interest.

¹For a more in depth discussion of the geometric interpretation of various optimal designs, refer to e.g. [7, Section 7.5].

Definition 3.2 (geodesic-convexity). A function $f : \mathbb{S}_m^{++} \rightarrow \mathbb{R}$ defined on the Riemannian manifold \mathbb{S}_m^{++} is called *geodesically convex* if it satisfies

$$f(P\#_tQ) \leq (1-t)f(P) + tf(Q), \quad t \in [0, 1], \text{ and } P, Q \succ 0.$$

where we use the traditional notation $P\#_tQ := P^{1/2}(P^{-1/2}QP^{-1/2})^tP^{1/2}$ to denote the *geodesic* between P and $Q \in \mathbb{S}_m^{++}$ under the Riemannian metric $g_P(X, Y) = \text{tr}(P^{-1}XP^{-1}Y)$.

Lemma 3.3. *The function E_ℓ is geodesically log-convex on the set of positive definite matrices.*

Corollary 3.4. *The map $M \mapsto E_\ell^{1/\ell}((X^\top MX)^{-1})$ is log-convex on the set of PD matrices.*

For further details on the theory of geodesically convex functions on \mathbb{S}_m^+ and their optimization, we refer the reader to [40]. We prove Lemma 3.3 and Corollary 3.4 in Appendix A.

From Corollary 3.4, we immediately obtain that (3.2) is a convex optimization problem, and can therefore be solved using a variety of efficient algorithms. Projected gradient descent turns out to be particularly easy to apply because we only require projection onto the intersection of the cube $0 \leq z \leq 1$ and the plane $\{z \mid z^\top \mathbf{1} = k\}$ (as a consequence of Prop 3.1). Projection onto this intersection is a special case of the so-called continuous quadratic knapsack problem, which is a very well-studied problem and can be solved essentially in linear time [10, 12].

Remark 3.5. The convex relaxation remains log-convex when points can be chosen with multiplicity, in which case the projection step is also significantly simpler, requiring only $z \geq 0$.

We conclude the analysis of the continuous relaxation by showing a bound on the support of its solution under some mild assumptions:

Theorem 3.6. *Let ϕ be the mapping from \mathbb{R}^m to $\mathbb{R}^{m(m+1)/2}$ such that $\phi(x) = (\xi_{ij}x_ix_j)_{1 \leq i, j \leq m}$ with $\xi_{ij} = 1$ if $i = j$ and 2 otherwise. Let $\tilde{\phi}(x) = (\phi(x), 1)$ be the affine version of ϕ . If for any set of $m(m+1)/2$ distinct rows of X , the mapping under $\tilde{\phi}$ is independent, then the support of the optimum z^* of (3.2) satisfies $\|z^*\|_0 \leq k + \frac{m(m+1)}{2}$.*

The proof is identical to that of [44, Lemma 3.5], which shows such a result for A-optimal design; we relegate it to Appendix B.

4 Algorithms and analysis

Solving the convex relaxation (3.2) does not directly provide a solution to (3.1); first, we must round the relaxed solution $z^* \in \Gamma_k^c$ to a discrete solution $S \in \Gamma_k$. We present two possibilities: (i) rounding the solution of the continuous relaxation (§4.1); and (ii) a greedy approach (§4.2).

4.1 Sampling from the continuous relaxation

For conciseness, we concentrate on sampling without replacement, but note that these results extend with minor changes to with replacement sampling (see [44]). Wang et al. [44] discuss the sampling scheme described in Alg. 1) for A-optimal design; the same idea easily extends to ESP-design. In particular, Alg. 1, applied to a solution of (3.2), provides the same asymptotic guarantees as those proven in [44, Lemma 3.2] for A-optimal design.

Algorithm 1: Sample from z^*

Data: budget k , $z^* \in \mathbb{R}^n$

Result: S of size k

$S \leftarrow \emptyset$

while $|S| < k$ **do**

 Sample $i \in [n] \setminus S$ uniformly at random

 Sample $x \sim \text{Bernoulli}(z_i^*)$

if $x = 1$ **then** $S \leftarrow S \cup \{i\}$

return S

Theorem 4.1. *Let $\Sigma_* = X^\top \text{Diag}(z^*)X$. Suppose $\|\Sigma_*^{-1}\|_2 \kappa(\Sigma_*) \|X\|_\infty^2 \log m = \mathcal{O}(1)$. The subset S constructed by sampling as above verifies with probability $p = 0.8$*

$$E_\ell \left((X_S^\top X_S)^{-1} \right)^{1/\ell} \leq \mathcal{O}(1) \cdot E_\ell \left((X_{S^*}^\top X_{S^*})^{-1} \right)^{1/\ell}.$$

Theorem 4.1 shows that under reasonable conditions, we can probabilistically construct a good approximation to the optimal solution in linear time, given the solution z^* to the convex relaxation.

4.2 Greedy approach

In addition to the solution based on convex relaxation, ESP-design admits an intuitive greedy approach, despite not being a submodular optimization problem in general. Here, elements are removed one-by-one from a base set of experiments; this greedy removal, as opposed to greedy addition, turns out to be much more practical. Indeed, since f_ℓ is not defined for sets of size smaller than k , it is hard to greedily add experiments to the empty set and then bound the objective function after k items have been added. This difficulty precludes analyses such as [45, 39] for optimizing non-submodular set functions by bounding their “curvature”.

Algorithm 2: Greedy algorithm

Data: matrix X , budget k , initial set S_0

Result: S of size k

$S \leftarrow S_0$

while $|S| > k$ **do**

Find $i \in S$ such that $S \setminus \{i\}$ is feasible and i minimizes $f_\ell(S \setminus \{i\})$
 $S \leftarrow S \setminus \{i\}$

return S

Bounding the performance of Algorithm 2 relies on the following lemma.

Lemma 4.2. *Let $X \in \mathbb{R}^{n \times m}$ ($n \geq m$) be a matrix with full column rank, and let k be a budget $m \leq k \leq n$. Let S of size k be subset of $[n]$ drawn with probability $\mathcal{P} \propto \det(X_S^\top X_S)$. Then*

$$\mathbb{E}_{S \sim \mathcal{P}} \left[E_\ell \left((X_S^\top X_S)^{-1} \right) \right] \leq \prod_{i=1}^\ell \frac{n-m+i}{k-m+i} \cdot E_\ell \left((X^\top X)^{-1} \right), \quad (4.1)$$

with equality if $X_S^\top X_S \succ 0$ for all subsets S of size k .

Lemma 4.2 extends a result from [2, Lemma 3.9] on column-subset selection via volume sampling to all ESPs. In particular, it follows that removing one element (by volume sampling a set of size $n-1$) will in expectation decrease f by a multiplicative factor which is clearly also attained by a greedy minimization. This argument then entails the following bound on Algorithm 2’s performance. Proofs of both results are in Appendix C.

Theorem 4.3. *Algorithm 2 initialized with a set S_0 of size n_0 produces a set S^+ of size k such that*

$$E_\ell \left((X_{S^+}^\top X_{S^+})^{-1} \right) \leq \prod_{j=1}^\ell \frac{n_0 - m + j}{k - m + j} \cdot E_\ell \left((X_{S_0}^\top X_{S_0})^{-1} \right) \quad (4.2)$$

As Wang et al. [44] note regarding A-optimal design, (4.2) provides a trivial optimality bound on the greedy algorithm when initialized with $S_0 = \{1, \dots, n\}$ (S^* denotes the optimal set):

$$E_\ell \left((X_{S^+}^\top X_{S^+})^{-1} \right)^{1/\ell} \leq \frac{n-m+\ell}{k-m+1} f(\{1, \dots, n\}) \leq \frac{n-m+\ell}{k-m+1} E_\ell \left((X_{S^*}^\top X_{S^*})^{-1} \right)^{1/\ell} \quad (4.3)$$

However, this naive initialization can be replaced by the support $\|z^*\|_0$ of the convex relaxation solution; in the common scenario described by Theorem 3.6, we then obtain the following result:

Theorem 4.4. *Let $\tilde{\phi}$ be the mapping defined in 3.6, and assume that all choices of $m(m+1)/2$ distinct rows of X always have their mapping independent mappings for $\tilde{\phi}$. Then the outcome of the greedy algorithm initialized with the support of the solution to the continuous relaxation verifies*

$$f_\ell(S^+) \leq \log \left(\frac{k + m(m-1)/2 + \ell}{k - m + 1} \right) + f_\ell(S^*).$$

4.3 Computational considerations

Computing the ℓ -th elementary symmetric polynomial on a vector of size m can be done in $\mathcal{O}(m \log^2 \ell)$ using Fast Fourier Transform for polynomial multiplication, due to the construction introduced by Ben-Or (see [37]); hence, computing $f_\ell(S)$ requires $\mathcal{O}(nm^2)$ time, where the cost is dominated by computing $X_S^\top X_S$. Alg. 1 runs in expectation in $\mathcal{O}(n)$; Alg. 2 costs $\mathcal{O}(m^2 n^3)$.

5 Further Implications

We close our theoretical presentation by discussing a potentially important geometric problem related to ESP-design. In particular, our motivation here is the dual problem of D-optimal design (i.e., dual to the convex relaxation of D-optimal design): this is nothing but the well-known *Minimum Volume Covering Ellipsoid (MVCE)* problem, which is a problem of great interest to the optimization community in its own right—see the recent book [42] for an excellent account.

With this motivation, we develop the dual formulation for ESP-design now. We start by deriving $\nabla E_\ell(A)$, for which we recall that $E_\ell(\cdot)$ is a spectral function, whereby the spectral calculus of Lewis [29] becomes applicable, saving us from intractable multilinear algebra [23]. More precisely, say $U^\top \Lambda U$ is the eigendecomposition of A , with U unitary. Then, as $E_\ell(A) = e_\ell \circ \lambda(A)$,

$$\nabla E_\ell(A) = U^\top \text{Diag}(\nabla e_\ell(\Lambda))U = U^\top \text{Diag}(e_{\ell-1}(\Lambda^{-\langle i \rangle}))U. \quad (5.1)$$

We can now derive the dual of ESP-design (we consider only $z \geq 0$); in this case problem (3.2) is

$$\sup_{A \succ 0, z \geq 0} \inf_{\mu \in \mathbb{R}, H} -\frac{1}{\ell} \log E_\ell(A) - \text{tr}(H(A^{-1} - X^\top \text{Diag}(z)X)) - \mu(\mathbf{1}^\top z - k),$$

which admits as dual

$$\inf_{\mu \in \mathbb{R}, H} \sup_{A \succ 0, z \geq 0} \underbrace{-\frac{1}{\ell} \log E_\ell(A) - \text{tr}(HA^{-1})}_{g(A)} + \text{tr}(HX^\top \text{Diag}(z)X) - \mu(\mathbf{1}^\top z - k). \quad (5.2)$$

We easily show that $H \succeq 0$ and that g reaches its maximum on \mathbb{S}_m^{++} for A such that $\nabla g = 0$. Rewriting $A = U^\top \Lambda U$, we have

$$\nabla g(A) = 0 \iff \Lambda \text{Diag}(e_{\ell-1}(\Lambda_{\langle i \rangle}))\Lambda = e_\ell(\Lambda)UHU^\top.$$

In particular, H and A are co-diagonalizable, with $\Lambda \text{Diag}(e_{\ell-1}(\Lambda_{\langle i \rangle}))\Lambda = \text{Diag}(h_1, \dots, h_m)$. The eigenvalues of A must thus satisfy the system of equations

$$\lambda_i^2 e_{\ell-1}(\lambda_1, \dots, \lambda_{i-1}, \lambda_{i+1}, \dots, \lambda_m) = h_i e_\ell(\lambda_1, \dots, \lambda_m), \quad 1 \leq i \leq m.$$

Let $a(H)$ be such a matrix (notice, $a(H) = \nabla g^*(0)$). Since f_ℓ is convex, $g(a(H)) = f_\ell^*(-H)$ where f_ℓ^* is the Fenchel conjugate of f_ℓ . Finally, the dual optimization problem is given by

$$\sup_{x_i^\top H x_i \leq 1, H \succeq 0} f_\ell^*(-H) = \sup_{x_i^\top H x_i \leq 1, H \succeq 0} \frac{1}{\ell} \log E_\ell(a(H))$$

Details of the calculation are provided in Appendix D. In the general case, deriving $a(H)$ or even $E_\ell(a(H))$ does not admit a closed form that we know of. Nevertheless, we recover the well-known duals of A-optimal design and D-optimal design as special cases.

Corollary 5.1. *For $\ell = 1$, $a(H) = \text{tr}(H^{1/2})H^{1/2}$ and for $\ell = m$, $a(H) = H$. Consequently, we recover the dual formulations of A- and D-optimal design.*

6 Experimental results

We compared the following methods to solving (3.1):

- UNIF / UNIFFDV: k experiments are sampled uniformly / with Fedorov exchange
- GREEDY / GREEDYFDV: greedy algorithm (relaxed init.) / with Fedorov exchange
- SAMPLE: sampling (relaxed init.) as in Algorithm 1.

We also report the results for solution of the continuous relaxation (RELAX); the convex optimization was solved using projected gradient descent, the projection being done with the code from [12].

6.1 Synthetic experiments: optimization comparison

We generated the experimental matrix X by sampling n vectors of size m from the multivariate Gaussian distribution of mean 0 and sparse precision Σ^{-1} (density d ranging from 0.3 to 0.9). Due to the runtime of Fedorov methods, results are reported for only one run; results averaged over multiple iterations (as well as for other distributions over X) are provided in Appendix E.

As shown in Fig. 1, the greedy algorithm applied to the convex relaxation’s support outperforms sampling from the convex relaxation solution, and does as well as the usual Fedorov algorithm UNIFFDV; GREEDYFDV marginally improves upon the greedy algorithm and UNIFFDV. Strikingly, GREEDY provides designs of comparable quality to UNIFFDV; furthermore, as very few local exchanges improve upon its design, running the Fedorov algorithm with GREEDY initialization is much faster (Table 1); this is confirmed by Table 2, which shows the number of experiments in common for different algorithms: GREEDY and GREEDYFDV only differ on very few elements. As the budget k increases, the difference in performances between SAMPLE, GREEDY and the continuous relaxation decreases, and the simpler SAMPLE algorithm becomes competitive. Table 3 reports the support of the continuous relaxation solution for ESP-design with $\ell = 10$.

Table 1: Runtimes (s) ($\ell = 10, d = 0.6$)

k	40	80	120	160	200
GREEDY	$2.8 \cdot 10^1$	$2.7 \cdot 10^1$	$3.1 \cdot 10^1$	$4.0 \cdot 10^1$	$5.2 \cdot 10^1$
GREEDYFDV	$6.6 \cdot 10^1$	$2.2 \cdot 10^2$	$3.2 \cdot 10^2$	$1.2 \cdot 10^2$	$1.3 \cdot 10^2$
UNIFFDV	$1.6 \cdot 10^3$	$4.1 \cdot 10^3$	$6.0 \cdot 10^3$	$6.2 \cdot 10^3$	$4.7 \cdot 10^3$

Table 2: Common items between solutions ($\ell = 10, d = 0.6$)

k	40	80	120	160	200
$ \text{GREEDY} \cap \text{UNIFFDV} $	26	76	114	155	200
$ \text{GREEDY} \cap \text{GREEDYFDV} $	40	78	117	160	200
$ \text{UNIFFDV} \cap \text{GREEDYFDV} $	26	75	113	155	200

Table 3: $\|z^*\|_0$ ($\ell = 10, d = 0.6$)

k	40	80	120	160	200
$d = 0.3$	93 ± 3	117 ± 3	148 ± 2	181 ± 3	213 ± 2
$d = 0.6$	92 ± 7	117 ± 4	145 ± 4	180 ± 3	214 ± 4
$d = 0.9$	88 ± 3	116 ± 3	147 ± 4	179 ± 3	214 ± 1

6.2 Real data

We used the Concrete Compressive Strength dataset [47] (with column normalization) from the UCI repository to evaluate ESP-design on real data; this dataset consists in 1030 possible experiments to model concrete compressive strength as a linear combination of 8 physical parameters. In Figure 2 (a), OED chose k experiments to run to estimate θ , and we report the normalized prediction error on the remaining $n - k$ experiments. The best choice of OED for this problem is of course A-optimal design, which shows the smallest predictive error. In Figure 2 (b), we report the fraction of non-zero entries in the design matrix X_S ; higher values of ℓ correspond to increasing sparsity. This confirms that OED allows us to scale between the extremes of A-optimal design and D-optimal design to tune desirable side-effects of the design; for example, sparsity in a design matrix can indicate not needing to tune a potentially expensive experimental parameter, which is instead left at its default value.

7 Conclusion and future work

We introduced the family of ESP-design problems, which evaluate the quality of an experimental design using elementary symmetric polynomials, and showed that typical approaches to optimal design such as continuous relaxation and greedy algorithms can be extended to this broad family of problems, which covers A-optimal design and D-optimal design as special cases.

We derived new properties of elementary symmetric polynomials: we showed that they are geodesically log-convex on the space of positive definite matrices, enabling fast solutions to solving the relaxed ESP optimization problem. We furthermore showed in Lemma 4.2 that volume sampling, applied to the columns of the design matrix X has a constant multiplicative impact on the objective function $E_\ell((X_S^\top X_S)^{-1})$, extending Avron and Boutsidis [2]’s result from the trace to all el-

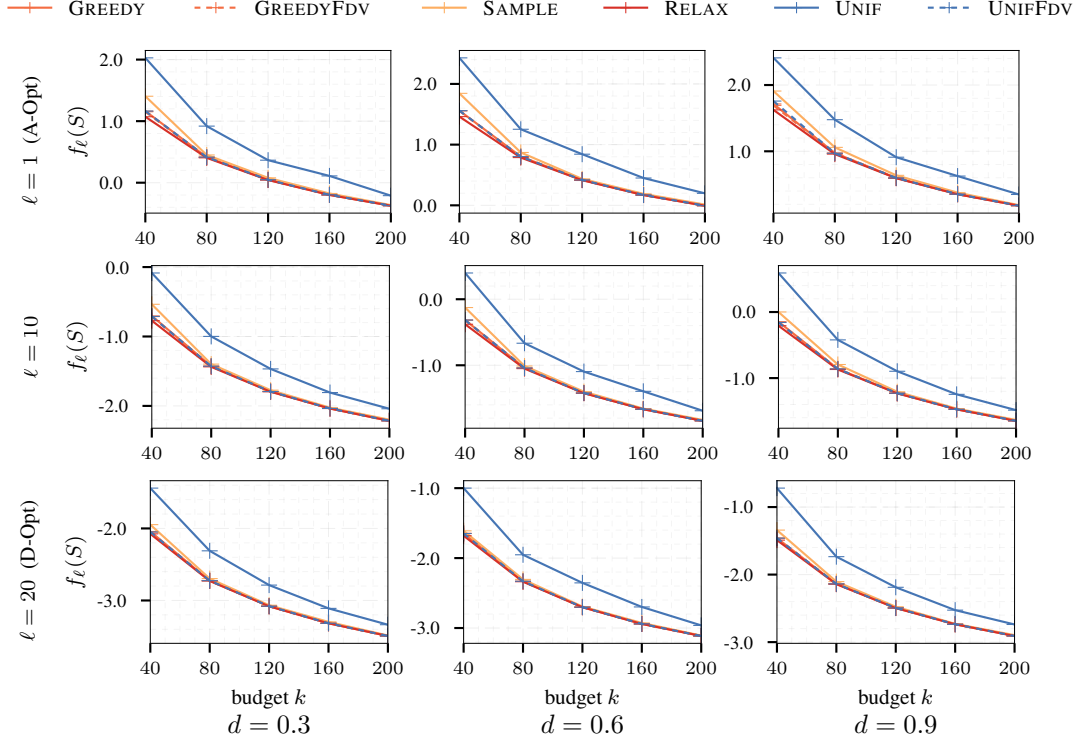


Figure 1: Synthetic experiments, $n = 500$, $m = 30$. The greedy algorithm performs as well as the classical Fedorov approach; as k increases, all designs except UNIF converge towards the continuous relaxation, making SAMPLE the best approach for large designs.




Figure 2: Predicting concrete compressive strength via the greedy method; higher ℓ increases the sparsity of the design matrix X_S , at the cost of marginally decreasing predictive performance.

elementary symmetric polynomials. This allows us to derive a greedy algorithm with performance guarantees, which empirically performs as well as Fedorov exchange, in a fraction of the runtime.

However, our work still includes some open questions: in deriving the Lagrangian dual of the optimization problem, we had to introduce the function $a(H)$ which maps \mathbb{S}_m^{++} ; however, although $a(H)$ is known for $\ell = 1, m$, its form for other values of ℓ is unknown, making the dual form a purely theoretical object in the general case. Whether the closed form of a can be derived, or whether $E_\ell(a(H))$ can be obtained with only knowledge of H , remains an open problem. Due to the importance of the dual form of D-optimal design as the Minimum Volume Covering Ellipsoid, we believe that further investigation of the general dual form of ESP-design will provide valuable insight, both into optimal design and for the general theory of optimization.

ACKNOWLEDGEMENTS

Suvrit Sra acknowledges support from NSF grant IIS-1409802 and DARPA Fundamental Limits of Learning grant W911NF-16-1-0551.

References

- [1] A. Atkinson, A. Donev, and R. Tobias. *Optimum Experimental Designs, With SAS*. Oxford Statistical Science Series. OUP Oxford, 2007.
- [2] H. Avron and C. Boutsidis. Faster subset selection for matrices and applications. *SIAM J. Matrix Analysis Applications*, 34(4):1464–1499, 2013.
- [3] E. R. Barnes. An algorithm for separating patterns by ellipsoids. *IBM Journal of Research and Development*, 26:759–764, 1982.
- [4] H. H. Bauschke, O. Güler, A. S. Lewis, and H. S. Sendov. Hyperbolic polynomials and convex analysis. *Canad. J. Math.*, 53(3):470–488, 2001.
- [5] R. Bhatia. *Matrix Analysis*. Springer, 1997.
- [6] R. Bhatia. *Positive Definite Matrices*. Princeton University Press, 2007.
- [7] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [8] K. Chaloner and I. Verdinelli. Bayesian experimental design: A review. *Statistical Science*, 10:273–304, 1995.
- [9] D. A. Cohn. Neural network exploration using optimal experiment design. In *Neural Networks*, pages 679–686. Morgan Kaufmann, 1994.
- [10] R. Cominetti, W. F. Mascarenhas, and P. J. S. Silva. A newton’s method for the continuous quadratic knapsack problem. *Mathematical Programming Computation*, 6(2):151–169, 2014.
- [11] M. A. Davenport, A. K. Massimino, D. Needell, and T. Woolf. Constrained adaptive sensing. *IEEE Transactions on Signal Processing*, 64(20):5437–5449, 2016.
- [12] T. A. Davis, W. W. Hager, and J. T. Hungerford. An efficient hybrid algorithm for the separable convex quadratic knapsack problem. *ACM Trans. Math. Softw.*, 42(3):22:1–22:25, 2016.
- [13] H. Dette, V. B. Melas, and W. K. Wong. Locally d -optimal designs for exponential regression models. *Statistica Sinica*, 16(3):789–803, 2006.
- [14] A. N. Dolia, T. De Bie, C. J. Harris, J. Shawe-Taylor, and D. M. Titterton. The minimum volume covering ellipsoid estimation in kernel-defined feature spaces. In *European Conference on Machine Learning*, pages 630–637, 2006.
- [15] E. N. Dolia, N. M. White, and C. J. Harris. D -optimality for minimum volume ellipsoid with outliers. In *In Proceedings of the Seventh International Conference on Signal/Image Processing and Pattern Recognition*, pages 73–76, 2004.
- [16] G. Elfving. Optimum allocation in linear regression theory. *Ann. Math. Statist.*, 23(2):255–262, 1952.
- [17] V. Fedorov. *Theory of optimal experiments*. Probability and mathematical statistics. Academic Press, 1972.
- [18] Y. Gu and Z. Jin. Neighborhood preserving d -optimal design for active learning and its application to terrain classification. *Neural Computing and Applications*, 23(7):2085–2092, 2013.
- [19] X. He. Laplacian regularized d -optimal design for active learning and its application to image retrieval. *IEEE Trans. Image Processing*, 19(1):254–263, 2010.
- [20] T. Horel, S. Ioannidis, and S. Muthukrishnan. *Budget Feasible Mechanisms for Experimental Design*, pages 719–730. Springer Berlin Heidelberg, 2014.
- [21] R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, 1985.
- [22] D. A. Jackson and Y. Chen. Robust principal component analysis and outlier detection with ecological data. *Environmetrics*, 15(2):129–139, 2004.
- [23] T. Jain. Derivatives for antisymmetric tensor powers and perturbation bounds. *Linear Algebra and its Applications*, 435(5):1111 – 1121, 2011.
- [24] R. Jozsa and G. Mitchison. Symmetric polynomials in information theory: Entropy and subentropy. *Journal of Mathematical Physics*, 56(6), 2015.
- [25] A. I. Khuri, B. Mukherjee, B. K. Sinha, and M. Ghosh. Design issues for generalized linear models: A review. *Statist. Sci.*, 21(3):376–399, 2006.
- [26] J. Kiefer. Optimal design: Variation in structure and performance under change of criterion. *Biometrika*, 62:277–288, 1975.

- [27] A. Krause, A. Singh, and C. Guestrin. Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies. *Journal of Machine Learning Research*, 9(Feb):235–284, 2008.
- [28] F. S. Lasheras, J. V. Vilán, P. G. Nieto, and J. del Coz Díaz. The use of design of experiments to improve a neural network model in order to predict the thickness of the chromium layer in a hard chromium plating process. *Mathematical and Computer Modelling*, 52(78):1169 – 1176, 2010.
- [29] A. S. Lewis. Derivatives of spectral functions. *Math. Oper. Res.*, 21(3):576–588, 1996.
- [30] I. G. Macdonald. *Symmetric functions and Hall polynomials*. Oxford university press, 1998.
- [31] A. J. Miller and N.-K. Nguyen. A fedorov exchange algorithm for d -optimal design. *Applied Statistics*, 43:669–677, 1994.
- [32] W. W. Muir. Inequalities concerning the inverses of positive definite matrices. *Proceedings of the Edinburgh Mathematical Society*, 19(2):109113, 1974.
- [33] F. Pukelsheim. *Optimal Design of Experiments*. Society for Industrial and Applied Mathematics, 2006.
- [34] P. J. Rousseeuw and A. M. Leroy. *Robust Regression and Outlier Detection*. John Wiley & Sons, Inc., 1987.
- [35] G. Sagnol. *Optimal design of experiments with application to the inference of traffic matrices in large networks: second order cone programming and submodularity*. Theses, École Nationale Supérieure des Mines de Paris, 2010.
- [36] A. Schein and L. Ungar. A-Optimality for Active Learning of Logistic Regression Classifiers, 2004.
- [37] A. Shpilka and A. Wigderson. Depth-3 arithmetic circuits over fields of characteristic zero. *computational complexity*, 10(1):1–27, 2001.
- [38] S. Silvey, D. Titterington, and B. Torsney. An algorithm for optimal designs on a design space. *Communications in Statistics - Theory and Methods*, 7(14):1379–1389, 1978.
- [39] J. D. Smith and M. T. Thai. Breaking the bonds of submodularity: Empirical estimation of approximation ratios for monotone non-submodular greedy maximization. *CoRR*, abs/1702.07002, 2017.
- [40] S. Sra and R. Hosseini. Conic Geometric Optimization on the Manifold of Positive Definite Matrices. *SIAM J. Optimization (SIOPT)*, 25(1):713–739, 2015.
- [41] P. Sun and R. M. Freund. Computation of minimum-volume covering ellipsoids. *Operations Research*, 52(5):690–706, 2004.
- [42] M. Todd. *Minimum-Volume Ellipsoids*. Society for Industrial and Applied Mathematics, 2016.
- [43] L. Vandenberghe, S. Boyd, and S.-P. Wu. Determinant maximization with linear matrix inequality constraints. *SIAM J. Matrix Anal. Appl.*, 19(2):499–533, 1998.
- [44] Y. Wang, A. W. Yu, and A. Singh. On computationally tractable selection of experiments in regression models, 2016.
- [45] Z. Wang, B. Moran, X. Wang, and Q. Pan. Approximation for maximizing monotone non-decreasing set functions with a greedy method. *J. Comb. Optim.*, 31(1):29–43, 2016.
- [46] T. C. Xygkis, G. N. Korres, and N. M. Manousakis. Fisher information based meter placement in distribution grids via the d -optimal experimental design. *IEEE Transactions on Smart Grid*, PP(99), 2016.
- [47] I.-C. Yeh. Modeling of strength of high-performance concrete using artificial neural networks. *Cement and Concrete Research*, 28(12):1797 – 1808, 1998.
- [48] Y. Yu. Monotonic convergence of a general algorithm for computing optimal designs. *Ann. Statist.*, 38(3):1593–1606, 2010.