

Computing Experiment-Constrained D-Optimal Designs

Aditya Pillai

H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30332,
apillai32@gatech.edu.

Gabriel Ponte

IOE Dept., University of Michigan and Federal University of Rio de Janeiro, gabponte@umich.edu

Marcia Fampa

COPPE, Federal University of Rio de Janeiro, fampa@cos.ufrj.br

Jon Lee

IOE Dept., University of Michigan, Ann Arbor, MI, jonxlee@umich.edu

Mohit Singh

H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30332,
mohitsinghr@gmail.com.

Weijun Xie

H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30332,
wxie@gatech.edu.

In optimal experimental design, the objective is to select a limited set of experiments that maximizes information about unknown model parameters based on factor levels. This work addresses the generalized D-optimal design problem, allowing for nonlinear relationships in factor levels. We develop scalable algorithms suitable for cases where the number of candidate experiments grows exponentially with the factor dimension, focusing on both first- and second-order models under design constraints. Particularly, our approach integrates convex relaxation with pricing-based local search techniques, which can provide upper bounds and performance guarantees. Unlike traditional local search methods, such as the “Fedorov exchange” and its variants, our method effectively accommodates arbitrary side constraints in the design space. Furthermore, it yields both a feasible solution and an upper bound on the optimal value derived from the convex relaxation. Numerical results highlight the efficiency and scalability of our algorithms, demonstrating superior performance compared to the state-of-the-art commercial software, JMP.

Key words: *D-optimal design; experimental design; convex relaxation; local search; approximation algorithm; determinant maximization*

1. Introduction. Optimal experimental design is a classical problem in statistics. The goal of optimal experimental design is to identify a limited number of experiments to obtain the maximum information about a vector of unknown model parameters in a given model based on levels of independent variables/factors. In this work, and as is common, we assume that the model is linear in the parameters. However, the model is often nonlinear in the factor levels. We

focus on one of the most popular optimal design criteria, D-optimality. In D-optimal experimental design, under i.i.d. Gaussian noise, we aim to identify a set of experiments to minimize the average covariance of the least-squares parameter estimator. In other words, we seek to minimize the volume of the standard confidence ellipsoid for the parameter centered on the least-squares parameter estimate. This is equivalent to maximizing the (logarithm of the) determinant of the Fisher information matrix (see [27] for more details; also see [7, 11, 12, 26, 28] for history and more information).

We let integer $d \geq 1$ denote the dimension of factors and let $\mathbf{x} := (x_1, \dots, x_d)^\top$ represent a vector of factors. In our most general setting, we consider a model based on a list of p distinct monomials $\mathcal{M} := \{m_1(\mathbf{x}), \dots, m_p(\mathbf{x})\}$ in the variables. Specifically,

$$y \approx \sum_{i=1}^p \theta_i m_i(\mathbf{x}),$$

where “ \approx ” indicates that the model incorporates a zero-mean Gaussian noise. The model is linear in the parameter vector $\boldsymbol{\theta} \in \mathbb{R}^p$, and so, by choosing a set of experiments $\mathbf{X} = \{\mathbf{x}^1, \dots, \mathbf{x}^k\}$ and their associated observations, y^1, \dots, y^k respectively, we can calculate the least-squares estimate $\hat{\boldsymbol{\theta}} = \arg \min_{\boldsymbol{\theta}} \sum_{j=1}^k ((\sum_{i=1}^p \theta_i m_i(\mathbf{x}^j)) - y^j)^2$. The challenge, of course, is to choose the best experiments, $\mathbf{X} = \{\mathbf{x}^1, \dots, \mathbf{x}^k\}$. Following the convention, we assume that experiments can be repeated.

For there to be a unique least-squares solution to the linear regression problem associated with a set of experiments (i.e., the model is “identifiable” given the experiments), we need a set of design points of rank p . Assuming that the model is identifiable, given a multi-set of experiments \mathbf{X} , the D-optimality criterion is the determinant of the information matrix of the design matrix, which we wish to maximize.

Before continuing, it is useful to settle on some terminology. If the maximum degree of a monomial in \mathcal{M} is ℓ , then we say that we have an ℓ -th order model. In an ℓ -th order model, if we have all possible monomials of maximum degree ℓ , then we say that it is a full ℓ -th order model. We refer to models that are not full as *partial*. In what follows, we will be mostly interested in full first-order models (having $d + 1$ monomials), full second-order models (having $\binom{d}{2} + d + 1$ monomials), and partial second-order models.

The *design point* associated with an experiment $\mathbf{x} \in \mathbb{R}^d$ is simply $\mathbf{p}(\mathbf{x}) := (m_1(\mathbf{x}), \dots, m_p(\mathbf{x}))^\top \in \mathbb{Z}^p$. That is, it is the vector of evaluations of each monomial on the experiment \mathbf{x} . We assume that we have a set of $L \geq 2$ levels $\{0, 1, \dots, L - 1\}$, that are the possible values of the factors. So, without additional structural requirements, we have a universe of L^d possible experiments, one for each $\mathbf{x} \in \{0, 1, \dots, L - 1\}^d$. It is common for one to have $L \geq \ell + 1$ for a full ℓ -th order model. For example, to capture a second-order effect at least in one variable, we would need at least three

levels (this is the motivation for many second-order designs in the literature that use three levels, such as the well-known Box-Behnken design; see [19, Sec. 5.3.3.6.2] and [4]).

We note that modeling and algorithms associated with constrained design regions are a central topic in the design of experiments literature; see [3, 6, 18] for references. Specifically, we consider the constraints of form $\mathbf{A}\mathbf{x} \leq \mathbf{b}$, for $\mathbf{x} \in \{0, 1, \dots, L-1\}^d$. So, we define our universe of allowable experiments as

$$\mathcal{Y} := \{\mathbf{x} \in \{0, 1, \dots, L-1\}^d : \mathbf{A}\mathbf{x} \leq \mathbf{b}\}.$$

Finally, we impose a “budget restriction” on the size of \mathbf{X} , the chosen multi-set of experiments, to be k .

With all of this, we are now ready to formally state our *experiment-constrained D-Optimal design problem*:

$$\phi := \max \left\{ \log \det \left(\sum_{\mathbf{x} \in \mathcal{Y}} \lambda(\mathbf{x}) \mathbf{p}(\mathbf{x}) \mathbf{p}(\mathbf{x})^\top \right) : \sum_{\mathbf{x} \in \mathcal{Y}} \lambda(\mathbf{x}) = k; \lambda(\mathbf{x}) \in \mathbb{Z}_+, \text{ for } \mathbf{x} \in \mathcal{Y} \right\}. \quad (\text{ECD-Opt})$$

The solution selects the multi-set of experiments \mathbf{X} , by simply having $\lambda(\mathbf{x})$ copies of \mathbf{x} in \mathbf{X} . We can further associate matrices \mathbf{P} and \mathbf{S} where \mathbf{P} has a row $\mathbf{p}(\mathbf{x})$ for each $\mathbf{x} \in \mathbf{X}$ and matrix $\mathbf{S} = \mathbf{P}^\top \mathbf{P}$. In this way, we have $\mathbf{S} = \mathbf{P}^\top \mathbf{P} = \sum_{\mathbf{x} \in \mathbf{X}} \mathbf{p}(\mathbf{x}) \mathbf{p}(\mathbf{x})^\top$, and we can rewrite the objective function as $\log \det(\mathbf{S}) = \log \det(\mathbf{P}^\top \mathbf{P})$.

In what follows, we let the collection of feasible design points be

$$\mathcal{P} := \{\mathbf{p}(\mathbf{x}) \in \mathbb{R}^p : \mathbf{x} \in \mathcal{Y}\}.$$

Therefore, ECD-Opt is equivalent to

$$\max \left\{ \log \det \left(\sum_{\mathbf{v} \in \mathcal{P}} \lambda(\mathbf{v}) \mathbf{v} \mathbf{v}^\top \right) : \sum_{\mathbf{v} \in \mathcal{P}} \lambda(\mathbf{v}) = k; \lambda(\mathbf{v}) \in \mathbb{Z}_+, \text{ for } \mathbf{v} \in \mathcal{P} \right\}. \quad (1)$$

The main computational challenge is that the set of vectors \mathcal{P} consists of exponentially many vectors in the dimension p , and in practice, may not be explicitly defined.

As studied in the literature and available software, when the set of allowable experiments \mathcal{Y} is given explicitly, as a list, it is interesting to have linear constraints on the integer variable vector λ , whereupon the associated D-optimal design problem has been referred to as CD-Opt (see [25]). In this paper, set \mathcal{Y} is given by a set of allowable experiment vectors \mathbf{x} , described via a range of levels for each factor with optionally additional linear constraints. To distinguish it from CD-Opt, we call our problem ECD-Opt.

There is a great deal of literature on heuristic algorithms for ECD-Opt, mainly focusing on the case where there are no side constraints $\mathbf{A}\mathbf{x} \leq \mathbf{b}$, based on the usual greedy and local search

ideas; see, for example, the references in [14, Sec. 1]. The popular software `JMP` embraces local search ideas described in [17], see [10, Page 125] and is capable of handling additional constraints as well as settings where design points are exponential in size. In Section 5, we give a detailed comparison of our approach to `JMP`. We remark that `JMP` does not produce any gaps to optimality and focuses only on heuristic based approaches, in contrast to our setting. Various branch-and-bound (B&B) based approaches aiming to find exact optimal solutions were also investigated in [13, 14, 23, 24, 29]. None of these handle the case in which the set of feasible design points is very large and not explicitly given. Our work fills this gap and aims to develop scalable computational techniques and fast approximation algorithms with performance guarantees, and to solve difficult instances of ECD-Opt.

1.1. Our results and contributions. In this work, we consider algorithms for cases of ECD-Opt for which the number of candidate experiments is exponential in the number of factors, d . We develop algorithms for ECD-Opt and study its continuous convex relaxation. For the sake of practicality, we focus on the cases of full first-order models and second-order models. We would like to emphasize that, in contrast to elementary/classical local search methods (known as “Fedorov exchange” and its variants), our approach effectively handles arbitrary side constraints $\mathbf{A}\mathbf{x} \leq \mathbf{b}$. Moreover, our approach provides not only a solution for the problem but also an upper bound for the optimal value based on the convex relaxation of the problem.

Pricing-based local search. Consider a given feasible solution $\lambda(\mathbf{x})$, for each $\mathbf{x} \in \mathcal{Y}$ of ECD-Opt. We can view it as a multi-set of experiments \mathbf{X} , where $\mathbf{x} \in \mathcal{Y}$ occurs with multiplicity $\lambda(\mathbf{x})$ in \mathbf{X} . A local search step, in our context, exchanges one experiment $\mathbf{x}' \in \mathbf{X}$ with an experiment $\mathbf{x} \in \mathcal{Y} \setminus \{\mathbf{x}'\}$, in a way that the objective of ECD-Opt increases. The classical “Fedorov exchange method” (i.e., local search step), typically applied when there are no or only very simple side constraints, treats set \mathcal{Y} as an input list. To relieve the computational burden, such a local search would only consider $\mathbf{x} \in \mathbf{X} \setminus \{\mathbf{x}'\}$ such that $|\mathbf{x} - \mathbf{x}'|$ is a standard unit vector. In contrast, we treat set \mathcal{Y} as given by constraints, and we formulate the “pricing problem for local search”, which seeks an improving \mathbf{x} using mathematical optimization techniques. This is akin to the well-known column-generation method for a wide variety of problems (which is why we call it a “pricing problem”, even though we do not use “prices” in the sense of dual variables). Our pricing problem effectively searches over a much larger local neighborhood than the simple local search described above, since it has no limitation on $|\mathbf{x} - \mathbf{x}'|$. For the case of an ℓ -th order model, the pricing problem has a polynomial objective function of degree 2ℓ , linear constraints, and integer variables. We will see that this is quite tractable for first-order models and for partial second-order models.

With respect to first-order models, the pricing problem for our local search becomes a maximization program with a convex quadratic objective function. This situation can be directly handled with, for example, `Gurobi`; or, using standard lifting methods, we can reformulate it linearly in the case where $L = 2$. If $L > 2$, such a linear reformulation is still possible (by the so-called “binarization”). See [15] for the use of similar techniques for a completely different problem. For second-order models, we proceed similarly but with a greater computational burden.

Pricing for the continuous relaxation. Finally, we also study the continuous relaxation of ECD-Opt obtained after removing the integrality constraints, which is a convex program. In connection with this, we let the optimal value of that be ϕ_R . Solving the continuous relaxation provides a useful upper bound to indicate the quality of any feasible integer solution to ECD-Opt. Our techniques are very similar to those for the local search pricing problem. The main difference is that the objective function for the continuous relaxation pricing problem has a different polynomial objective function of degree 2ℓ . We will refer then to the “pricing problem for the continuous relaxation” to distinguish it from the “pricing problem for local search”. The other differences between the two pricing approaches, which are rather significant, are in the computational details. In particular, we have found computational advantages by: (i) including randomly-generated experiments besides the best ones found by solving the pricing problem, (ii) working with the dual of the continuous relaxation if necessary, and (iii) reducing the number of experiments via a sparsification procedure over the iterations.

Theoretical results. We also theoretically analyze the performance of our pricing-based local search procedure. The polynomial optimization problem encountered at every local search step is, in general, NP-hard, but good approximation algorithms can be developed under certain settings. Building on the results [16], we provide an approximation algorithm at each local search step to obtain an approximate local search procedure with the following performance guarantee.

Theorem 1 *Suppose that there exists a polynomial-time $\frac{1}{\rho}$ -approximation algorithm for some $\rho > 1$ to the problem $\max_{\mathbf{v} \in \mathcal{P}} \mathbf{v}^\top \mathbf{G} \mathbf{v}$ for every positive semidefinite matrix $\mathbf{G} \succeq 0$. Then the pricing-based local search algorithm presented in Section 2 returns a feasible solution $\lambda(\mathbf{x})$ of ECD-Opt with the associated multi-set of experiments \mathbf{X} , where $\mathbf{x} \in \mathcal{Y}$ occurs with multiplicity $\lambda(\mathbf{x})$ in \mathbf{X} , and associated matrix \mathbf{S} satisfying*

$$\det(\mathbf{S}) \geq e^\phi \left(\frac{k-p+1}{k} \frac{p}{p+k(\rho-1)} \right)^p$$

where ϕ is the optimal objective of the ECD-Opt.

An interesting case where there exists an approximation algorithm for the sub-problem $\max_{\mathbf{v} \in \mathcal{P}} \mathbf{v}^\top \mathbf{G} \mathbf{v}$ for every $\mathbf{G} \succeq 0$ is when $\mathcal{P} = \{\mathbf{x} \in \{0, 1\}^d : x_1 = 1\}$. This corresponds to the full first-order model with two levels, where the first monomial is the constant, i.e., $m_1(\mathbf{x}) := 1$. Indeed, Nesterov [20], gives a $\rho = \frac{\pi}{2}$ approximation algorithm for the subproblem $\max_{\mathbf{v} \in \mathcal{P}} \mathbf{v}^\top \mathbf{G} \mathbf{v}$. Applying Nesterov's results along with Theorem 1, we obtain the following corollary.

Corollary 1 *For the full first-order model with two levels and no side constraints (i.e., $\mathcal{P} = \mathcal{Y}$), the pricing-based local search algorithm returns an approximate local optimum \mathbf{X} to ECD-Opt with associated matrix \mathbf{S} satisfying*

$$\det(\mathbf{S}) \geq e^\phi \left(\frac{k-p+1}{k} \frac{p}{p+k(\frac{\pi}{2}-1)} \right)^p$$

and additionally, satisfying

$$\det(\mathbf{S}) \geq \frac{k^p}{2^{2(p-1)}} \left(\frac{k-p+1}{k} \frac{p}{p+k(\frac{\pi}{2}-1)} \right)^p.$$

Here, the equality for ϕ_R follows using the symmetry of the convex relaxation.

1.2. Organization. In Section 2, we introduce the pricing based local search algorithm and discuss the ideas used in its implementation. Next, in Section 3, we analyze the local search theoretically by proving an approximation guarantee given an approximation algorithm for the pricing problem. We also give an efficient approximation algorithm for pricing problem in the first-order model and also show a hardness result for the pricing problem in the second-order model. In Section 4, we discuss the convex relaxation and how we implement a pricing based column generation to solve the relaxation. Next, in Section 5, we discuss numerical results of local search and column generation.

2. Local search algorithm for D-Optimal design. In this section, we discuss the local search algorithm for ECD-Opt. We are interested in instances of the problem where the set of vectors $\mathcal{P} := \{\mathbf{p}(\mathbf{x}) \in \mathbb{R}^p : \mathbf{x} \in \mathcal{Y}\}$ is not explicitly given and is too big to enumerate all its elements. We consider that we have a feasible solution $\hat{\lambda}(\mathbf{x})$, for $\mathbf{x} \in \mathcal{Y}$, to ECD-Opt, but we only store the vectors $\mathbf{p}(\mathbf{x})$ for which $\hat{\lambda}(\mathbf{x}) > 0$. We let $\mathcal{Y}^{>0} := \{\mathbf{x} \in \mathcal{Y} : \hat{\lambda}(\mathbf{x}) > 0\}$ and therefore, the associated matrix $\mathbf{S} = \sum_{\mathbf{x} \in \mathcal{Y}^{>0}} \hat{\lambda}(\mathbf{x}) \mathbf{p}(\mathbf{x}) \mathbf{p}(\mathbf{x})^\top$, and we assume that $\text{rank}(\mathbf{S}) = p$. Observe that $|\mathcal{Y}^{>0}| \leq p$ and thus can be explicitly stored. The simplest local search move decreases $\hat{\lambda}(\mathbf{x}')$ by one, for some $\mathbf{x}' \in \mathcal{Y}^{>0}$, and increases $\hat{\lambda}(\mathbf{x})$ by one, for some $\mathbf{x} \in \mathcal{Y}$ and $\mathbf{x} \neq \mathbf{x}'$. We can explicitly iterate over the experiments in the set $\mathcal{Y}^{>0}$ to find \mathbf{x}' . Since we do not explicitly generate and store all the vectors of \mathcal{P} , for a given $\mathbf{x}' \in \mathcal{Y}^{>0}$, we will seek to select an experiment $\mathbf{x} \in \mathcal{Y}$ and generate the corresponding vector $\mathbf{p}(\mathbf{x})$ of \mathcal{P} , so that $\det(\mathbf{S} - \mathbf{p}(\mathbf{x}') \mathbf{p}(\mathbf{x}')^\top + \mathbf{p}(\mathbf{x}) \mathbf{p}(\mathbf{x})^\top) > \det(\mathbf{S})$, if one exists.

We define $\mathbf{S}_{\mathbf{x}'} := \mathbf{S} - \mathbf{p}(\mathbf{x}')\mathbf{p}(\mathbf{x}')^\top$ for all $\mathbf{x}' \in \mathcal{Y}^{>0}$ and, for a $p \times p$ matrix \mathbf{S} and $k < p$, we define $\det_k(\mathbf{S})$ as the product of the k largest eigenvalues of \mathbf{S} . Then, for all $\mathbf{x}' \in \mathcal{Y}^{>0}$ and $\mathbf{x} \in \mathcal{Y}$, if $\text{rank}(\mathbf{S}_{\mathbf{x}'}) = p$, we have

$$\det(\mathbf{S}_{\mathbf{x}'} + \mathbf{p}(\mathbf{x})\mathbf{p}(\mathbf{x})^\top) = \det(\mathbf{S}_{\mathbf{x}'}) (1 + \mathbf{p}(\mathbf{x})^\top \mathbf{S}_{\mathbf{x}'}^{-1} \mathbf{p}(\mathbf{x})).$$

Otherwise, if $\mathbf{S}_{\mathbf{x}'}$ is rank deficient, we have

$$\det(\mathbf{S}_{\mathbf{x}'} + \mathbf{p}(\mathbf{x})\mathbf{p}(\mathbf{x})^\top) = \det(\mathbf{S}_{\mathbf{x}'}) \mathbf{p}(\mathbf{x})^\top (\mathbf{I}_p - \mathbf{S}_{\mathbf{x}'}^\dagger \mathbf{S}_{\mathbf{x}'}) \mathbf{p}(\mathbf{x}),$$

where $\mathbf{S}_{\mathbf{x}'}^\dagger$ is the Moore-Penrose pseudoinverse of $\mathbf{S}_{\mathbf{x}'}$ (see [22]).

Finally, for a given $\mathbf{x}' \in \mathcal{Y}^{>0}$, generating $\mathbf{p}(\mathbf{x}) \in \mathcal{P}$ that maximizes $\det(\mathbf{S}_{\mathbf{x}'} + \mathbf{p}(\mathbf{x})\mathbf{p}(\mathbf{x})^\top)$ is equivalent to the following pricing problem:

$$\max_{\mathbf{x} \in \mathcal{Y}} \mathbf{p}(\mathbf{x})^\top \mathbf{G} \mathbf{p}(\mathbf{x}), \quad (\text{Sub})$$

where $\mathbf{G} := \mathbf{S}_{\mathbf{x}'}^{-1}$, if $\text{rank}(\mathbf{S}_{\mathbf{x}'}) = p$, and $\mathbf{G} := \mathbf{I}_p - \mathbf{S}_{\mathbf{x}'}^\dagger \mathbf{S}_{\mathbf{x}'}$, otherwise.

It is easy to see (considering the \mathbf{x} associated with \mathbf{x}') that, if $\text{rank}(\mathbf{S}_{\mathbf{x}'}) = p$, the optimal objective value of (Sub) is at least $\det(\mathbf{S})/\det(\mathbf{S}_{\mathbf{x}'}) - 1$, and any objective value above that gives a local move with improving objective-function value. Otherwise, the optimal objective value of (Sub) is at least $\det(\mathbf{S})/\det(\mathbf{S}_{\mathbf{x}'}) - 1$.

In Section 2.1, we discuss how to solve the problem (Sub).

2.1. Solving the pricing problem for the local search. For the case of an ℓ -th order model, the pricing problem (Sub) maximizes a polynomial objective function of degree 2ℓ , subject to linear constraints and integrality constraints over the variables. For any order, this is a nonconvex optimization problem, and a common approach to solving it consists of reformulating it as an integer linear program by lifting the problem to a higher-dimensional space using McCormick inequalities. In the basic approach, we would instantiate a new variable y_i for each monomial $m_i(\mathbf{x})$, and we describe or approximate the convex hull of the graph of $y_i = m_i(\mathbf{x})$ on the domain $\{0, 1, \dots, L-1\}^d$, using linear inequalities. This can be done quite effectively. For example, when $L = 2$, we can assume that the degree of each variable in each monomial is 0 or 1. Therefore, we can define that $m_i(\mathbf{x}) = \prod_{j \in \mathcal{I}} x_j$. Then the inequalities $0 \leq y_i \leq x_j$ for $j \in \mathcal{I}$ and $y_i \geq \sum_{j \in \mathcal{I}} x_j - |\mathcal{I}| + 1$ define the convex hull of the graph of $y_i = m_i(\mathbf{x})$. We note that one way to handle more than two levels is through “binarization,” representing each integer variable $0 \leq x_i \leq L-1$ with a linear function of $\lceil \log_2 L \rceil$ binary variables (e.g., see [5, 21]).

Although our approach can be applied to solve (Sub) for any ℓ -th order model, we are particularly interested in solving two special cases— the full first- and second-order models.

For a full first-order model with d factors, a design point is given by $\mathbf{p}(\mathbf{x}) = (1, x_1, \dots, x_d)^\top \in \mathbb{Z}^{d+1}$. This particularizes (Sub) to the quadratic optimization problem with integer variables $\alpha_0, \dots, \alpha_d$:

$$\begin{aligned} \max_{\boldsymbol{\alpha}} \quad & \boldsymbol{\alpha}^\top \mathbf{G} \boldsymbol{\alpha}, \\ \text{s.t.} \quad & \alpha_0 = 1, \\ & \alpha_i \in \{0, 1, \dots, L-1\}, \quad \forall i = 1, \dots, d, \\ & \mathbf{A}(\alpha_1, \dots, \alpha_d)^\top \leq \mathbf{b}. \end{aligned} \quad (\text{Quad_Sub})$$

Note that this is a maximization program with convex quadratic objective function. For the “unconstrained” case (i.e., without side constraints $\mathbf{A}(\alpha_1, \dots, \alpha_d)^\top \leq \mathbf{b}$), we may relax the integrality constraints, and an optimal solution will be at a vertex of the box constraints. In this case, there are only 2^d possible solutions to check, rather than L^d possible integer solutions. This means that we will only generate experiments using the extreme levels, 0 and $L-1$. On the other hand, for the first-order model, we have $L=2$, and all levels are extreme. Therefore, in this unconstrained case, the problem (Quad_Sub) is equivalent to a “boolean quadratic program.” For modest values of d (e.g., $d=12$), it is feasible to solve (Quad_Sub) by enumeration. However, for larger d (e.g., $d=20$), enumeration becomes impractical as the solution space grows too large to explore in a reasonable time. To address this, we formulate (Quad_Sub) as a linear integer program by linearizing the quadratic terms, allowing us to efficiently solve it using *Gurobi*. In addition, we provide primal heuristics that speed up the search for the optimal solution for (Quad_Sub).

For the second-order model, we arrive at a degree four integer optimization problem over integer points in a d -dimensional polytope. We address this by linearizing all higher degree terms in the objective and solving the resulting higher-dimensional linear integer program using *Gurobi*.

Finally, we also consider solving (Sub) heuristically before solving it exactly using a solver. We employ a local search method with an initial solution $\bar{\boldsymbol{\alpha}} = (1; \mathbf{x}')$, where the criterion for the improvement of a given solution is the increase in the value of the objective function of (Sub) and different neighborhoods are explored, defined by flipping over one bit of the current solution, termed “Bit Flip,” or swapping two bits, termed “Bit Swap.” More specifically, for bit flip, the neighborhoods of a given solution $\bar{\boldsymbol{\alpha}}$ with $\bar{\alpha}_0 = 1$ are defined as

$$\begin{aligned} \mathcal{N}_1(\bar{\boldsymbol{\alpha}}) := \{ \boldsymbol{\alpha} \in \mathbb{Z}^{d+1} : \alpha_0 = 1, 0 \leq \alpha_i \leq L-1, \text{ for all } i \in [d], \mathbf{A}(\alpha_1, \dots, \alpha_d)^\top \leq \mathbf{b}, \\ (\alpha_1, \dots, \alpha_d)^\top - (\bar{\alpha}_1, \dots, \bar{\alpha}_d)^\top = \pm \mathbf{e}_i, \text{ for some } i \in [d] \}, \end{aligned} \quad (\text{Bit Flip})$$

and for bit swap, they are defined as

$$\begin{aligned} \mathcal{N}_2(\bar{\boldsymbol{\alpha}}) := \{ \boldsymbol{\alpha} \in \mathbb{Z}^{d+1} : \alpha_0 = 1, 0 \leq \alpha_i \leq L-1, \text{ for all } i \in [d], \mathbf{A}(\alpha_1, \dots, \alpha_d)^\top \leq \mathbf{b}, \\ (\alpha_1, \dots, \alpha_d)^\top - (\bar{\alpha}_1, \dots, \bar{\alpha}_d)^\top = \mathbf{e}_i - \mathbf{e}_j, \text{ for some } i, j \in [d], i \neq j \}, \end{aligned} \quad (\text{Bit Swap})$$

where \mathbf{e}_i is the i -th standard basis.

In the next section, we provide a theoretical analysis of a polynomial-time implementation for an approximate local search algorithm.

3. Theoretical results for D-optimal design. In this section, we prove Theorem 1 and Corollary 1. The first result gives a guarantee of the output of the local search algorithm if the polynomial optimization problem (Sub) described in Section 2.1 can be solved approximately. We also show that this result can be applied to obtain an approximation guarantee when set $\mathcal{Y} = \{-1, 1\}^d$. In addition, we prove the hardness for the sub-problem $\max_{\mathbf{v} \in \mathcal{P}} \mathbf{v}^\top \mathbf{G} \mathbf{v}$ arising from a second-order problem.

3.1. Local search with approximate local optima. In this section, we analyze the local search algorithm when it finds an approximate local solution that can be efficiently implemented in polynomial time, as stated in the following Theorem 1. More specifically, in Theorem 1, we prove that instead of identifying the best improving move by solving an integer program (IP), an algorithm that yields an approximately improving solution still achieves an approximately locally optimal solution. Besides, we show that for the special case when the set of allowed vectors is $\mathcal{Y} = \{\mathbf{x} \in \{0, 1\}^d : x_1 = 1\}$, Nesterov's result [20] can be used to give an efficient polynomial time algorithm for implementing the improvement step approximately. We also provide a performance guarantee on the approximate locally optimal solution.

We first prove Theorem 1. The proof follows the outline of Madan et al. [16] where they analyze the local search algorithm. In Theorem 1, we assume that the improvement step in the local search algorithm can be approximated within a factor of $\frac{1}{\rho}$.

A crucial ingredient in the analysis is the convex relaxation of (1) and its Lagrangian dual as defined below. First, the continuous relaxation of the D-opt design problem (1) is given by

$$\phi_R = \max_{\boldsymbol{\lambda}} \left\{ \ln \det \left(\sum_{i=1}^n \lambda_i \mathbf{v}^i \mathbf{v}^{i^\top} \right) : \sum_{i=1}^n \lambda_i = k, \lambda_i \geq 0, \forall i \in [n] \right\}. \quad (2)$$

Next, let us consider its equivalent form

$$\phi_R = \max_{\boldsymbol{\lambda}, \mathbf{Y}} \left\{ \ln \det (\mathbf{Y}) : \sum_{i=1}^n \lambda_i \mathbf{v}^i \mathbf{v}^{i^\top} \succeq \mathbf{Y}, \sum_{i=1}^n \lambda_i = k, \lambda_i \geq 0, \forall i \in [n] \right\},$$

which admits the following Lagrangian dual with dual multiplier $\boldsymbol{\Lambda}$ associated with constraint $\sum_{i=1}^n \lambda_i \mathbf{v}^i \mathbf{v}^{i^\top} \succeq \mathbf{Y}$ and dual multiplier ν associated with constraint $\sum_{i=1}^n \lambda_i = k$:

$$\phi_R^D = \min_{\boldsymbol{\Lambda} \succeq \mathbf{0}, \nu} \left\{ k\nu - \ln \det (\boldsymbol{\Lambda}) - p : \mathbf{v}^{i^\top} \boldsymbol{\Lambda} \mathbf{v}^i \leq \nu, \forall i \in [n] \right\}, \quad (3)$$

whose optimal value is equal to that of the continuous relaxation (2) due to the strong duality of the convex program with the Slater condition [2].

Theorem 1 Suppose that there exists a polynomial-time $\frac{1}{\rho}$ -approximation algorithm for some $\rho > 1$ to the problem $\max_{\mathbf{v} \in \mathcal{P}} \mathbf{v}^\top \mathbf{G} \mathbf{v}$ for every positive semidefinite matrix $\mathbf{G} \succeq 0$. Then the pricing-based local search algorithm presented in Section 2 returns a feasible solution $\lambda(\mathbf{x})$ of ECD-Opt with the associated multi-set of experiments \mathbf{X} , where $\mathbf{x} \in \mathcal{Y}$ occurs with multiplicity $\lambda(\mathbf{x})$ in \mathbf{X} , and associated matrix \mathbf{S} satisfying

$$\det(\mathbf{S}) \geq e^\phi \left(\frac{k-p+1}{k} \frac{p}{p+k(\rho-1)} \right)^p$$

where ϕ is the optimal objective of the ECD-Opt.

Proof: To prove the theorem, we show a stronger inequality by showing the given guarantee using ϕ_R instead of ϕ , which immediately implies the theorem since $\phi_R \geq \phi$. First, we show that the $(\frac{1}{\rho})$ -approximation algorithm for solving the problem $\max_{\mathbf{v} \in \mathcal{P}} \mathbf{v}^\top \mathbf{G} \mathbf{v}$ can be used to implement the local search algorithm where each local step can be solved approximately within a factor of $1/\rho$ as shown in the following lemma.

Lemma 1 Given some $\rho > 1$, if there exists a $\frac{1}{\rho}$ -approximation algorithm to the problem $\max_{\mathbf{v} \in \mathcal{P}} \mathbf{v}^\top \mathbf{G} \mathbf{v}$ for all $\mathbf{G} \succeq 0$, then there is a $\frac{1}{\rho}$ -approximation algorithm to the problem $\max_{\mathbf{v}^i \in \widehat{\mathcal{S}}, \mathbf{v}^j \in \mathcal{Y}} \det(\mathbf{S}_{-i} + \mathbf{v}^j \mathbf{v}^{j\top})$.

Proof: For each $i \in \widehat{\mathcal{S}}$, if \mathbf{S}_{-i} has an inverse, then $\det(\mathbf{S}_{-i} + \mathbf{v}^j \mathbf{v}^{j\top}) = \det(\mathbf{S}_{-i})(1 + \mathbf{v}^{j\top} \mathbf{S}_{-i}^{-1} \mathbf{v}^j)$. Thus, we can use the approximation algorithm to solve $\max_{\mathbf{v} \in \mathcal{P}} \mathbf{v}^\top \mathbf{S}_{-i}^{-1} \mathbf{v}$ to get a vector \mathbf{v}' satisfying $\det(\mathbf{S}_{-i} + \mathbf{v}' \mathbf{v}'^\top) = \det(\mathbf{S}_{-i})(1 + \mathbf{v}'^\top \mathbf{S}_{-i}^{-1} \mathbf{v}') \geq \det(\mathbf{S}_{-i})(1 + \frac{\mathbf{v}^{*\top} \mathbf{S}_{-i}^{-1} \mathbf{v}^*}{\rho}) \geq \frac{\det(\mathbf{S}_{-i})}{\rho}(1 + \mathbf{v}^{*\top} \mathbf{S}_{-i}^{-1} \mathbf{v}^*) = \frac{1}{\rho} \max_j \det(\mathbf{S}_{-i} + \mathbf{v}^j \mathbf{v}^{j\top})$. Similarly, if \mathbf{S}_{-i} does not have an inverse, then we have $\det(\mathbf{S}_{-i} + \mathbf{v}^j \mathbf{v}^{j\top}) = \det(\mathbf{S}_{-i}) \mathbf{v}^{j\top} (\mathbf{I}_p - \mathbf{S}_{-i}^\dagger \mathbf{S}_{-i}) \mathbf{v}^j$ and can apply the approximation algorithm to $\mathbf{I}_p - \mathbf{S}_{-i}^\dagger \mathbf{S}_{-i}$. Since we can get a $\frac{1}{\rho}$ -approximation for each fixed i we can also get the same approximation when we maximize over i . \diamond

We show the bound in the theorem by constructing a feasible dual solution $(\bar{\Lambda}, \nu)$ to the dual program (3), whose objective value is comparable to the objective value of the output of the approximate local search algorithm.

Let $(\bar{\Lambda}, \nu)$ where $\bar{\Lambda} = (\alpha \mathbf{S})^{-1}$ for $\alpha = \left(\frac{p}{k} \frac{k-p+1}{p+k(\rho-1)} \right)$ and $\nu = \frac{p}{k}$ be a candidate dual solution. This allows us to relate the objective value of \mathbf{S} to ϕ_R since

$$\phi_R \leq k\nu - \ln(\det((\alpha \mathbf{S})^{-1})) - p.$$

The inequality follows since the right-hand side is the objective value of the dual solution $(\bar{\Lambda}, \nu)$ and the dual is a minimization problem.

We are ready to show the approximation guarantee. Note that

$$\frac{\phi_R}{p} \leq \ln(\alpha) + \frac{\ln(\det(\mathbf{S}))}{p} + \frac{k\nu}{p} - 1 = \ln \left(\frac{p}{k} \frac{k-p+1}{p+k(\rho-1)} \right) + \frac{\ln(\det(\mathbf{S}))}{p}.$$

The first equality follows since the objective value of the dual solution $(\bar{\Lambda}, \nu)$ is given by $k\nu - \ln(\det(\bar{\Lambda})) - p$ and the second equality is from the definition of ν and α .

Thus, it remains to show that the candidate dual $(\bar{\Lambda}, \nu)$ is feasible to the dual problem (3). First, observe that $\bar{\Lambda} \succeq 0$ since $\mathbf{S} \succeq 0$ and $\alpha > 0$. We now verify the constraint $\mathbf{v}^{i^\top} \bar{\Lambda} \mathbf{v}^i \leq \nu$ for each i for which we need the following notation.

- (I) For $1 \leq i \leq n$, $\tau_i := \mathbf{v}^{i^\top} \mathbf{S}^{-1} \mathbf{v}^i$
- (II) For $1 \leq i, j \leq n$, $\tau_{ij} := \mathbf{v}^{j^\top} \mathbf{S}^{-1} \mathbf{v}^j$.

The next two lemmas are useful to prove the feasibility of the dual solution.

Lemma 2 *For any $i \in I$ and $1 \leq j \leq n$ we have $\tau_j - \tau_i \tau_j + \tau_{ij} \tau_{ji} \leq \tau_i + \rho - 1$.*

Proof: Since \mathbf{S} is an approximate locally optimal solution, we have that $\det(\mathbf{S} - \mathbf{v}^i \mathbf{v}^{i^\top} + \mathbf{v}^j \mathbf{v}^{j^\top}) \leq \rho \det(\mathbf{S})$. Following exactly the calculations of lemma 9 in [16], we get that $\rho \det(\mathbf{S}) \geq \det(\mathbf{S})(1 + \mathbf{v}^{j^\top} \mathbf{S}^{-1} \mathbf{v}^j)(1 - \mathbf{v}^{i^\top} (\mathbf{S} + \mathbf{v}^j \mathbf{v}^{j^\top})^{-1} \mathbf{v}^i)$ which is equivalent to $(1 + \tau_j)(1 - \tau_i + \frac{\tau_{ij} \tau_{ji}}{1 + \tau_j}) \leq \rho$. This concludes the proof since the last inequality is equivalent to the inequality of the claim. \diamond

Lemma 3 *For any $1 \leq j \leq n$, we have $\tau_j \leq \frac{p+k(\rho-1)}{k-p+1}$.*

Proof: Summing the inequality for all $i \in I$ from Lemma 2, we have

$$\sum_{i \in \widehat{S}} (\tau_j - \tau_i \tau_j + \tau_{ij} \tau_{ji}) \leq \sum_{i \in \widehat{S}} \tau_i + k(\rho - 1) = p + k(\rho - 1).$$

We also have the following expression for the left-hand side: $\sum_{i \in I} (\tau_j - \tau_i \tau_j + \tau_{ij} \tau_{ji}) = k\tau_j - p\tau_j + \tau_j$. Combining this with the upper bound concludes the proof. \diamond

To show $\mathbf{v}^{i^\top} \bar{\Lambda} \mathbf{v}^i \leq \nu$ for all $i \in [n]$, it is equivalent to show $\max_{i=1}^n \mathbf{v}^{i^\top} \bar{\Lambda} \mathbf{v}^i \leq \nu = \frac{p}{k}$. We have,

$$\begin{aligned} \max_{i=j}^n \mathbf{v}^{j^\top} \bar{\Lambda} \mathbf{v}^i &= \frac{1}{\alpha} \max_{j=1}^n \tau_j \\ &\leq \frac{p}{k}. \end{aligned}$$

The first equality follows by the definition of τ_j , and the inequality follows by applying the upper bound in Lemma 3. \square

We note that as in Madan et al. [16], the guarantee given in the above lemma is more meaningful when $k \gg p$. Now we apply Theorem 1 to prove Corollary 1.

Corollary 1 *For the full first-order model with two levels and no side constraints (i.e., $\mathcal{P} = \mathcal{Y}$), the pricing-based local search algorithm returns an approximate local optimum \mathbf{X} to ECD-Opt with associated matrix \mathbf{S} satisfying*

$$\det(\mathbf{S}) \geq e^\phi \left(\frac{k-p+1}{k} \frac{p}{p+k(\frac{\pi}{2}-1)} \right)^p$$

and additionally, satisfying

$$\det(\mathbf{S}) \geq \frac{k^p}{2^{2(p-1)}} \left(\frac{k-p+1}{k} \frac{p}{p+k(\frac{\pi}{2}-1)} \right)^p.$$

Proof: We prove the corollary by using Theorem 1 when $\mathcal{Y}' = \{-1, 1\}^d$ and then transforming the $\{-1, 1\}$ solution to a $\{0, 1\}$ solution. Nesterov's algorithm gives a $\frac{2}{\pi}$ -approximation to the problem $\max_{\mathbf{x} \in \mathcal{Y}'} \mathbf{x}^\top \mathbf{G} \mathbf{x}$ for any $\mathbf{G} \succeq 0$ with $\mathbf{G} \in \mathbb{R}^{d \times d}$. This satisfies the requirements of Theorem 1 and allows us to apply it with $\rho = \frac{2}{\pi}$. Let ϕ' be the optimal (ln det) value of the integral solution. Then, from Theorem 1, we get a solution \mathbf{S}' where $\mathbf{S} = \mathbf{V}'^\top \mathbf{V}'$ for $\mathbf{V}' \in \{-1, 1\}^{p \times k}$ satisfying

$$\det(\mathbf{S}') \geq e^{\phi'} \left(\frac{k-p+1}{k} \frac{p}{p+k(\frac{\pi}{2}-1)} \right)^p \quad (4)$$

We will show that the $\{-1, 1\}$ solution \mathbf{S}' can be transformed to a $\{0, 1\}$ solution \mathbf{S} such that $\det(\mathbf{S}) = \frac{\det(\mathbf{S}')}{2^{2(p-1)}}$. This transformation and the given guarantee immediately proves the first inequality of the corollary, since it also implies that $\phi = \frac{\phi'}{2^{2(p-1)}}$. Finally, we show that $\phi_R = \frac{\phi'_R}{2^{2(p-1)}}$ where ϕ'_R is the optimal value of the convex relaxation for the set \mathcal{Y}' which allows us to get a stronger version of the first inequality of the corollary where ϕ is replaced with ϕ_R . The transformation follows via a series of matrix operations that are standard for the Hadamard problem.

We will transform the matrix \mathbf{V}' to a matrix $\mathbf{V} \in \{0, 1\}^{p \times k}$ such that the first row is all-one. We can do this with the following steps:

Step 1. Multiply the columns by of \mathbf{V}' by $-1, 1$ so that the first row of \mathbf{V}' is all ones.

Step 2. Add the first row of \mathbf{V}' to all other rows.

Step 3. Divide all rows except the first row by 2.

We note that the first step can be achieved by multiplying by a $(k \times k)$ diagonal matrix \mathbf{D} on the right, and that steps 2 and 3 can be achieved by multiplying \mathbf{V}' on the left by $p-1$ matrices $\mathbf{L}_1, \dots, \mathbf{L}_{p-1}$. Thus, we have $\mathbf{V} = \mathbf{L}_{p-1} \dots \mathbf{L}_1 \mathbf{V}' \mathbf{D}$ so that $\det(\mathbf{L}_i) = 1/2$ for all $i = 1, \dots, p-1$. Thus, we can get a solution with vectors in $\mathcal{Y} = \{0, 1\}^d$ by setting $\mathbf{S} = \mathbf{V} \mathbf{V}^\top$ with objective value give by

$$\begin{aligned} \det(\mathbf{S}) &= \det(\mathbf{L}_{p-1} \mathbf{L}_2 \dots \mathbf{L}_1 \mathbf{V}' \mathbf{D}^2 \mathbf{V}'^\top \mathbf{L}_1^\top \dots \mathbf{L}_{p-1}^\top) \\ &= \det(\mathbf{L}_{p-1} \mathbf{L}_2 \dots \mathbf{L}_1 \mathbf{V}' \mathbf{V}'^\top \mathbf{L}_1^\top \dots \mathbf{L}_{p-1}^\top) \\ &= \frac{\det(\mathbf{S}')}{2^{2(p-1)}}. \end{aligned}$$

Then the proof of the corollary is concluded by the following lemma.

Lemma 4 When $\mathcal{Y} = \{\mathbf{x} \in \{0, 1\}^d : x_1 = 1\}$, an optimal solution to the continuous relaxation (2) is $\lambda_i^* = k/2^{p-1}$ for each $i \in [2^{p-1}]$ with optimal value $e^{\phi_R} = \frac{k^p}{2^{2(p-1)}}$.

Proof: The proof is to construct primal and dual solutions that have the same objective value. We begin by constructing the primal solution. We define the primal solution by letting $\lambda_i = \frac{k}{2^{p-1}}$ for $i = 1, \dots, 2^{p-1}$. Let \mathbf{A} be the matrix with all 2^{p-1} vectors from \mathcal{P} as its columns and \mathbf{D} be a diagonal matrix where $D_{ii} = \lambda_i$ for $i \in [2^{p-1}]$. Then the objective value of the primal solution is given by $\ln(\det(\mathbf{A}\mathbf{D}\mathbf{A}^\top))$. We now calculate $\det(\mathbf{A}\mathbf{D}\mathbf{A}^\top)$. We observe that $\mathbf{A}\mathbf{D}\mathbf{A}^\top = k \begin{bmatrix} 1 & \mathbf{1}^\top/2 \\ \mathbf{1}/2 & (\mathbf{J}_{p-1} + \mathbf{I}_{p-1})/4 \end{bmatrix}$ where $\mathbf{1}$ is the all ones vector in dimension $p-1$ and \mathbf{J}_{p-1} is the all ones matrix in dimension $p-1$. Then we have that $\det(\mathbf{A}\mathbf{D}\mathbf{A}^\top) = k^p \det((\mathbf{J}_{p-1} + \mathbf{I}_{p-1})/4 - \mathbf{1}(\mathbf{1}^\top)/4) = k^p \det(\mathbf{I}_{p-1}/4) = \frac{k^p}{2^{2(p-1)}}$. Now we construct a feasible dual solution with the same objective value. We set $\mathbf{\Lambda} = (\mathbf{A}\mathbf{D}\mathbf{A}^\top)^{-1}$ and $\nu = p/k$. Then the objective value of the dual solution, $k\nu - \ln(\det(\mathbf{\Lambda})) - p$, is equal to the primal objective by our calculation of $\det(\mathbf{A}\mathbf{D}\mathbf{A}^\top)$. We now verify the feasibility of the dual solution. The matrix $\mathbf{\Lambda}$ is PSD since its inverse is PSD, so it remains to show that $\mathbf{v}^\top \mathbf{\Lambda} \mathbf{v} \leq \nu = p/k$ for all $\mathbf{v} \in \mathcal{P}$. We can use the block structure of $\mathbf{A}\mathbf{D}\mathbf{A}^\top$ to get that $\mathbf{\Lambda} = \frac{1}{k} \begin{bmatrix} p & -2\mathbf{1}^\top \\ -2\mathbf{1} & 4\mathbf{I}_{p-1} \end{bmatrix}$. Then it is easy to verify that $\mathbf{v}^\top \mathbf{\Lambda} \mathbf{v} = p/k = \nu$ for all $\mathbf{v} \in \mathcal{P}$. \diamond

□

3.2. Hardness of polynomial optimization problem. So far, we have discussed two approaches to solve the problem $\max_{\mathbf{v} \in \mathcal{P}} \mathbf{v}^\top \mathbf{G} \mathbf{v}$ for $\mathbf{G} \succeq 0$: in Section 2.1, we describe how to solve the problem using an integer programming approach exactly whose running time can be exponential in the input size. In Corollary 1, we show that for the special case of a full first-order model when $\mathcal{P} = \mathcal{Y} = \{\mathbf{x} \in \{0, 1\}^d : x_1 = 1\}$, we have an efficient $\frac{2}{\pi}$ -approximation algorithm. This algorithm is based on the $\frac{2}{\pi}$ -approximation by Nesterov [20] to solve the problem $\max_{\mathbf{v} \in \mathcal{P}} \mathbf{v}^\top \mathbf{G} \mathbf{v}$ for $\mathbf{G} \succeq 0$ where $\mathcal{P} = \mathcal{Y} = \{-1, 1\}^d$. A natural question is whether such approximation algorithms also hold for general \mathcal{P} , in particular for second-order models.

In this section, we show a negative result and prove that it is NP-hard to get an approximation better than $1/2$ when the set \mathcal{P} is defined by a special case of the second-order model. In this case, hardness even follows when the matrix \mathbf{G} is diagonally dominant and \mathcal{M} has all degree-one monomials and a limited number of degree-two monomials. The diagonally dominant matrix is of special interest, since Goemans and Williamson [8] show that their algorithm for max-cut also achieves a 0.878-approximation for the problem $\max_{\mathbf{x} \in \{-1, 1\}^d} \mathbf{x}^\top \mathbf{G} \mathbf{x}$. The Goemans and Williamson result was the basis for the algorithm by Nesterov [20] for the general positive semi-definite matrix \mathbf{G} . For a diagonally dominant matrix \mathbf{G} our goal is to solve the problem

$$\max_{\mathbf{v} \in \mathcal{P}} \mathbf{v}^\top \mathbf{G} \mathbf{v} = \max_{\mathbf{x} \in \{-1, 1\}^d} \mathbf{p}(\mathbf{x})^\top \mathbf{G} \mathbf{p}(\mathbf{x}). \quad (5)$$

where $\mathbf{p}(\mathbf{x})$ has coordinates corresponding to all monomials of degree 1 and some monomials of degree 2.

We use the problem MAX – 3LIN for a reduction. In MAX – 3LIN we are given m equations over variables z_1, \dots, z_p of the form $z_i + z_j + z_k \equiv 0 \pmod{2}$ or $z_i + z_j + z_k \equiv 1 \pmod{2}$ and we need to select values of $z_i \in \{0, 1\}$ for each $i \in [p]$ to maximize the number of equations that are satisfied. The instance *cannot* contains two distinct equations involving the same set of variables. It has been shown that MAX – 3LIN is NP-hard to approximate better than $1/2$ [9].

Proposition 1 *There exists an approximation preserving polynomial-time reduction from MAX – 3LIN to Problem (5), and thus, it is NP-hard to approximate it better than a factor of $\frac{1}{2}$.*

Proof: We observe that $z_i + z_j + z_k \equiv 0 \pmod{2}$ is satisfied if and only if $\{z_i, z_j, z_k\}$ contains an even number of ones. Similarly, $z_i + z_j + z_k \equiv 1 \pmod{2}$ if and only if $\{z_i, z_j, z_k\}$ contains an odd number of ones. Given the set of m equations, we now construct an instance of problem (5) by defining the set of monomials \mathcal{M} and the matrix \mathbf{G} . Our reduction will have the following property:

- For all $\mathbf{x} \in \{-1, 1\}^d$ the value $\mathbf{p}(\mathbf{x})^\top \mathbf{G} \mathbf{p}(\mathbf{x})$ is equal to the number of equations satisfied by the MAX – 3LIN solution $z_i = \frac{x_i+1}{2}$ for all $i \in [p]$.

Firstly, we let \mathcal{M} contain all degree-one monomials meaning $m_i(\mathbf{x}) = x_i$ for $i = 1, \dots, d$. For every equation $z_i + z_j + z_k$ with $i < j < k$ we add a degree-two monomial corresponding to pair $\{j, k\}$ meaning $m_{jk}(\mathbf{x}) = x_j x_k$. Let \mathcal{C} be a set of subsets that defines the set of all monomials meaning for all $F \in \mathcal{C}$, $m_F(\mathbf{x}) = \prod_{i \in F} x_i \in \mathcal{M}$. We will use the elements \mathcal{C} to index entries of the matrix \mathbf{G} , which will have dimensions $|\mathcal{C}| \times |\mathcal{C}|$. For every equation $z_i + z_j + z_k \equiv 0 \pmod{2}$ with $i < j < k$ we set the entry $G_{i,\{j,k\}} = G_{\{j,k\},i} = -1/4$ and for $z_i + z_j + z_k \equiv 1 \pmod{2}$ we set the entry $G_{i,\{j,k\}} = G_{\{j,k\},i} = 1/4$. All other off-diagonal entries are set to 0. We define the diagonal entries by setting $G_{F_1, F_1} = \sum_{F_2 \in \mathcal{C}, F_2 \neq F_1} |G_{F_1, F_2}|$ for all $F_1 \in \mathcal{C}$. Thus, by construction, \mathbf{G} is symmetric, diagonally dominant, and therefore also positive semi-definite. We define sets $\mathcal{C}_1, \mathcal{C}_2$ where \mathcal{C}_1 contains all (i, j, k) such that $z_i + z_j + z_k \equiv 0 \pmod{2}$ is an equation in the MAX – 3LIN instance and \mathcal{C}_2 contains all (i, j, k) such that $z_i + z_j + z_k \equiv 1 \pmod{2}$ is an equation. For any $\mathbf{x} \in \{-1, 1\}^d$, we have

$$\mathbf{p}(\mathbf{x})^\top \mathbf{G} \mathbf{p}(\mathbf{x}) = \frac{1}{2} \sum_{(i,j,k) \in \mathcal{C}_2} (1 + x_i x_j x_k) + \frac{1}{2} \sum_{(i,j,k) \in \mathcal{C}_1} (1 - x_i x_j x_k).$$

Then we observe that

$$\frac{1 + x_i x_j x_k}{2} = \begin{cases} 1 & \text{if } x_i, x_j, x_k \text{ has an odd number of } -1 \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

and

$$\frac{1 - x_i x_j x_k}{2} = \begin{cases} 1 & \text{if } x_i, x_j, x_k \text{ has an even number of } -1 \\ 0 & \text{otherwise} \end{cases}. \quad (7)$$

Then the vector $\mathbf{x} = (x_1, \dots, x_d)$ can be mapped to a MAX – 3LIN solution $\mathbf{z} = (z_1, \dots, z_d)$ by setting $z_i = \frac{x_i+1}{2}$. The mapping from \mathbf{x} and \mathbf{z} has the property that $p(\mathbf{x})^\top \mathbf{G} p(\mathbf{x})$ is equal to the number of equations \mathbf{x} satisfies in the MAX – 3LIN instance. This follows by equations (6), (7) and the fact that $x_i = -1$ implies $z_i = 0$ and $x_i = 1$ implies $z_i = 1$. \square

4. Continuous relaxation. In this section, we consider the continuous relaxation (2) discussed in Section 3 and give an algorithm to solve it to optimality using column generation. We use both the primal (2) and dual (3) to solve the problem efficiently.

The algorithm works by maintaining a set of candidate vectors $\mathcal{P}' \subseteq \mathcal{P}$. We initialize the set \mathcal{P}' as a random subset of \mathcal{P} of small cardinality. We solve the restricted primal obtained by setting all variables for vectors not in \mathcal{P}' to zero and the corresponding restricted dual to optimality. We give the restricted primal and dual below, and for a set \mathcal{P}' , we refer to the restricted program as $R(\mathcal{P}')$.

$$\begin{aligned} \max_{\boldsymbol{\lambda}} \quad & \ln(\det(\sum_{\mathbf{v} \in \mathcal{P}'} \lambda_{\mathbf{v}} \mathbf{v} \mathbf{v}^\top)) & \min_{\boldsymbol{\Lambda}, \nu} \quad & k\nu - \ln(\det(\boldsymbol{\Lambda})) - p \\ \text{s.t.} \quad & \sum_{\mathbf{v} \in \mathcal{P}'} \lambda_{\mathbf{v}} = k & \text{s.t.} \quad & \mathbf{v}^\top \boldsymbol{\Lambda} \mathbf{v} \leq \nu, \quad \forall \mathbf{v} \in \mathcal{P}' \\ & \lambda_{\mathbf{v}} \geq 0, \quad \forall \mathbf{v} \in \mathcal{P}' & & \boldsymbol{\Lambda} \succeq 0. \end{aligned} \tag{8}$$

In the following lemma, we show that it is straightforward to obtain an optimal dual solution from an optimal primal solution.

Lemma 5 *Let $\lambda_{\mathbf{v}}$ for all $\mathbf{v} \in \mathcal{P}'$ be an optimal solution for the restricted primal with set of vectors, \mathcal{P}' . Then the dual solution $\boldsymbol{\Lambda} = (\sum_{\mathbf{v} \in \mathcal{P}'} \lambda_{\mathbf{v}} \mathbf{v} \mathbf{v}^\top)^{-1}$ and $\nu = \max_{\mathbf{v} \in \mathcal{P}'} \mathbf{v}^\top \boldsymbol{\Lambda} \mathbf{v}$ is an optimal solution for the restricted dual with set of vectors \mathcal{P}' .*

Proof: The proof follows by taking the Lagrangian multipliers for the primal program. We reformulate the primal program as follows,

$$\begin{aligned} \max_{\boldsymbol{\lambda}, \mathbf{X}} \quad & \ln(\det(\mathbf{X})) \\ \text{s.t.} \quad & \sum_{\mathbf{v} \in \mathcal{P}} \lambda_{\mathbf{v}} \mathbf{v} \mathbf{v}^\top \preceq \mathbf{X} \\ & \sum_{\mathbf{v} \in \mathcal{P}} \lambda_{\mathbf{v}} = k \\ & \lambda_{\mathbf{v}} \geq 0, \forall \mathbf{v} \in \mathcal{P}'. \end{aligned}$$

We use the multiplier Λ for the first constraint and ν for the equality constraint. Then, the Lagrangian is given by

$$\max_{\mathbf{X}, \boldsymbol{\lambda} \geq 0, \boldsymbol{\lambda}^\top \mathbf{1} = k} \min_{\nu, \boldsymbol{\Lambda} \succeq 0} \ln(\det(\mathbf{X})) + \nu(k - \sum_{\mathbf{v} \in \mathcal{P}} \lambda_{\mathbf{v}}) + \sum_{\mathbf{v} \in \mathcal{P}} \lambda_{\mathbf{v}} \mathbf{v}^\top \boldsymbol{\Lambda} \mathbf{v} - \langle \boldsymbol{\Lambda}, \mathbf{X} \rangle.$$

Taking the derivative with respect to \mathbf{X} gives us $(\mathbf{X}^{-1})^\top - \mathbf{\Lambda} = 0$ implying that $\mathbf{\Lambda} = \mathbf{X}^{-1}$. Hence, there exists an optimal dual solution that must also satisfy $\mathbf{X} = \sum_{v \in \mathcal{P}} \lambda_v v v^\top$ and $\nu = \max_{v \in \mathcal{P}'} v^\top \mathbf{\Lambda} v = p/k$. This concludes the proof of the lemma. \square

After solving the restricted primal and dual programs (8) to optimality, in the next step, the algorithm checks whether the optimal dual solution to the restricted dual program is feasible for the original dual problem. If it is, we obtain the result that the current primal and dual solutions are also optimal for the original problem. Otherwise, we find a violating constraint to dual program corresponding to some vector $v \in \mathcal{P} \setminus \mathcal{P}'$. We then add v to \mathcal{P}' and iterate in solving the updated formulated restricted primal and dual programs, which means that we move from the restricted program $R(\mathcal{P}')$ to $R(\mathcal{P}' \cup \{v\})$. The main challenge is to solve the *pricing problem* to find out whether there is a violating constraint. Given an optimal dual solution $(\bar{\mathbf{\Lambda}}, \bar{\nu})$ to the restricted dual, the problem of checking the feasibility of this solution to the original program is equivalent to solving

$$\max_{v \in \mathcal{P}} v^\top \bar{\mathbf{\Lambda}} v$$

and checking if the maximum is at most $\bar{\nu}$ or not. This problem turns out to be exactly the problem (Sub) that we have discussed in Section 2.1, and we apply the same methodology to solve it by reformulating it as a polynomial optimization problem. In the next proposition, we show that solving the pricing problem optimally also gives us a feasible dual solution. This is useful since it provides an upper bound for the value of the unrestricted program $R(\mathcal{P})$ and therefore, it is also an upper bound of the value of the integral optimum solution.

Proposition 2 *Let $(\bar{\mathbf{\Lambda}}, \bar{\nu})$ be a feasible solution to the restricted dual $R(\mathcal{P}')$ for a subset $\mathcal{P}' \subset \mathcal{P}$ and $\alpha = \max_{v \in \mathcal{P}} v^\top \bar{\mathbf{\Lambda}} v$. Then $(\bar{\mathbf{\Lambda}}, \alpha)$ is a feasible dual solution to the unrestricted dual program $R(\mathcal{P})$.*

Proof: This follows immediately, since $\bar{\mathbf{\Lambda}} \succeq 0$ and $v^\top \bar{\mathbf{\Lambda}} v \leq \alpha$ for all $v \in \mathcal{P}$ by definition of α . \square
We compute α in Proposition 2 every time we solve the pricing problem optimally when running the column generation algorithm. In Section 5, where we present our experimental results, we utilize Proposition 2 to derive an upper bound for large instances where obtaining the optimal value of $R(\mathcal{P})$ is computationally intractable.

To further speed up the implementation, we use additional steps that we describe in the following section.

4.1. Further implementation steps. In this section, we describe some extra steps that speed up the column generation algorithm. The full algorithm is described in Algorithm 1.

Sparsification. To ensure that the size of set \mathcal{Y}' does not increase too much over iterations, we describe a sparsification procedure that reduces the size of set \mathcal{Y}' without any loss to the objective value. This is useful since the formulation (8) depends on the size of \mathcal{P}' . In particular, whenever $|\mathcal{P}'| > \binom{p+1}{2}$, we can find a subset $\mathcal{P}'' \subset \mathcal{P}'$ of strictly smaller size such that the restricted problem $R(\mathcal{P}'')$ has the same objective as $R(\mathcal{P}')$. The search for \mathcal{P}'' can be reduced to solving a linear program described below. Given a feasible primal solution λ_v^* for all $v \in \mathcal{P}'$ to the restricted problem (8), we formulate a linear programming feasibility problem with variables λ_v for all $v \in \mathcal{P}'$. We ensure the matrix constraint that $\sum_{v \in \mathcal{P}'} \lambda_v v v^\top = \sum_{v \in \mathcal{P}'} \lambda_v^* v v^\top$ must hold, which is simply a set of linear constraints for variables λ_v . Additionally, we impose that λ is feasible for $R(\mathcal{P}')$. Then, we obtain the following LP feasibility problem:

$$\begin{aligned} \min_{\lambda} \quad & 0 \\ \text{s.t.} \quad & \sum_{v \in \mathcal{P}'} \lambda_v v v^\top = \sum_{v \in \mathcal{P}'} \lambda_v^* v v^\top \\ & \sum_{v \in \mathcal{P}'} \lambda_v = k \\ & \lambda_v \geq 0, \forall v \in \mathcal{P}'. \end{aligned} \tag{9}$$

We have the following theoretical bound on the size of the support of a continuous primal solution when we obtain a basic feasible solution. We then set $\mathcal{P}'' = \{v : \lambda_v > 0\}$ and continue.

Proposition 3 For any set of vectors \mathcal{P}' and weights $\lambda_v^* \geq 0$ for all $v \in \mathcal{P}'$ satisfying $\sum_{v \in \mathcal{P}'} \lambda_v^* = k$ there exists λ_v for all $v \in \mathcal{P}'$ satisfying $\sum_{v \in \mathcal{P}'} \lambda_v = k$ such that the following two hold

1. $\sum_{v \in \mathcal{P}'} \lambda_v v v^\top = \sum_{v \in \mathcal{P}'} \lambda_v^* v v^\top$
2. $|\{\lambda_v > 0 | v \in \mathcal{P}'\}| \leq \binom{p}{2} + p + 1$.

Proof: The proof of this lemma follows by setting λ to any basic feasible solution of LP (9) which exists since λ^* is a feasible solution to LP (9). The first property holds due to the first constraint in the LP, and the second property holds since the LP has $\binom{p}{2} + p + 1$ constraints apart from non-negativity constraints. Thus, any basic feasible solution would have no more than $\binom{p}{2} + p + 1$ nonzero variables. \square

Addition of random columns. If the optimal value of (8) is strictly less than the optimal value ϕ_R of the relaxation (3), solving the pricing problem (Sub) with an integer program guarantees that we obtain a violated column corresponding to a vector that is added to \mathcal{P}' . However, in practice, it can be slow to solve an integer program every iteration. A simple alternative approach is to add a small subset of random vectors from \mathcal{P} to \mathcal{P}' in addition to the vector that violates the

dual constraint. This procedure can accelerate the solution procedure by saving many iterations via preemptively adding vectors without solving an integer program.

Thus, the high level idea of our algorithm can be summarized with the following steps.

- (i) Solve the continuous relaxation of the restricted primal or dual problem (8) to obtain a dual solution (Λ, ν) . Check if the current dual solution is feasible to the original problem (3) by solving the pricing problem (Sub). If not, then add a subset of random vectors from \mathcal{P} to \mathcal{P}' along with the vector that violates the dual constraint.
- (ii) Sparsify the support of the current primal solution if the support is significantly larger than p^2 . Return to Step 1.

The details of the algorithm are described in Algorithm 1.

5. Experimental results. In this section, we show the efficacy and scalability of our approach for various classes of D -optimal design problems, in particular, first-order and second-order design, with knapsack constraints. We also compare our results with JMP [10], a popular commercial statistical software that can also compute the designs of the variants we study. In particular, we numerically compare the following variants of the D -optimal design problem.

- (i) Full first-order model with two levels and a cardinality constraint. In this model, we have all vectors with no more than r ones for some integer $r \leq d$, i.e., $\mathcal{P} = \mathcal{Y} = \{\mathbf{x} \in \{0, 1\}^d : \mathbf{1}^\top \mathbf{x} \leq r, x_1 = 1\}$.
- (ii) Full first-order model with two levels and two knapsack constraints. The knapsack constraints are given by two non-negative vectors $\mathbf{a}_1, \mathbf{a}_2 \in \mathbb{R}_+^d$ and two numbers $b_1, b_2 \geq 0$. Then the set of vectors is defined by $\mathcal{P} = \mathcal{Y} = \{\mathbf{x} \in \{0, 1\}^d : x_1 = 1, \mathbf{a}_1^\top \mathbf{x} \leq b_1, \mathbf{a}_2^\top \mathbf{x} \leq b_2\}$.
- (iii) We also consider the second-order model mixed with knapsack constraints. We recall the second-order model is defined by a set of monomials of degree at most two, given by set \mathcal{M} . That is, each design point is given by $\mathbf{p}(\mathbf{x}) = (m_1(\mathbf{x}), m_2(\mathbf{x}), \dots, m_p(\mathbf{x}))^\top$, where $m_j(\mathbf{x})$ is a monomial with degrees at most two. For vectors $\mathbf{a}_1, \mathbf{a}_2 \in \mathbb{R}_+^d$ and two numbers $b_1, b_2 \geq 0$ get the following set of feasible vectors:

$$\mathcal{P} = \{(m_1(\mathbf{x}), m_2(\mathbf{x}), \dots, m_p(\mathbf{x}))^\top \in \{0, 1\}^p : \mathbf{x} \in \{0, 1\}^d, x_1 = 1, \mathbf{a}_1^\top \mathbf{x} \leq b_1, \mathbf{a}_2^\top \mathbf{x} \leq b_2\}.$$

Experimental setup. We ran the experiments on a Windows machine with an AMD EPYC Processor that has 16 cores. For each of the variants described at the beginning of this section, we ran both the local search algorithm described in Section 2 and the column generation algorithm described in Section 4.1. We gathered various statistics on the runs of the two algorithms, such as total time, objective values, and time spent in Gurobi. Finally, we compared the results of our local search algorithm with JMP. We show a comparison between the objective values and also

Algorithm 1 Pricing-based algorithm for the continuous relaxation of ECD-Opt

Require: $\mathcal{P}' \subset \mathcal{P}$ which is a set of vectors and $\mathbf{S} = \sum_{v \in \mathcal{P}'} \mathbf{v}\mathbf{v}^\top$ satisfying $\det(\mathbf{S}) > 0$. Tolerances

$\delta, \epsilon, \gamma > 0$

- 1: Solve the primal program (8) on vectors \mathcal{P}' for to get primal solution λ .
- 2: Compute a feasible dual solution Λ, ν from the primal solution.
- 3: **while** True **do**
- 4: Let $\mathbf{v} = \arg \max_{v \in \mathcal{Y}} \mathbf{v}^\top \Lambda \mathbf{v}$ where the optimization problem is solved using a heuristic.
- 5: **if** $\mathbf{v}^\top \Lambda \mathbf{v} < (1 + \delta)\nu$ **then**
- 6: Let $\mathbf{v} = \arg \max_{v \in \mathcal{Y}} \mathbf{v}^\top \Lambda \mathbf{v}$ where the optimization problem is solved using an IP.
- 7: **end if**
- 8: **if** $\mathbf{v}^\top \Lambda \mathbf{v} \leq (1 + \epsilon)\nu$ and \mathbf{v} was obtained from the IP **then return** \mathcal{P}'
- 9: **end if**
- 10: $\mathcal{P}' \leftarrow \mathcal{P}' \cup \{\mathbf{v}\}$
- 11: **if** (Λ, ν) was obtained by solving the primal **then**
- 12: Add $p - 1$ random vectors from \mathcal{Y} to $\hat{\mathcal{S}}$
- 13: **else**
- 14: Add $2(p - 1)^2$ random vectors to $\hat{\mathcal{S}}$
- 15: **end if**
- 16: **if** $|\mathcal{P}'| > \lceil \frac{p^2}{3} \rceil$ and (Λ, ν) was obtained by solving the primal **then**
- 17: Solve an LP to get $\mathcal{P}'' \subseteq \mathcal{P}'$ such that $\sum_{v \in \mathcal{P}''} \lambda_v \mathbf{v} \mathbf{v}^\top = \sum_{v \in \mathcal{P}'} \lambda_v \mathbf{v} \mathbf{v}^\top$.
- 18: $\mathcal{P}' \leftarrow \mathcal{P}''$
- 19: **end if**
- 20: Compute (Λ, ν) using the new solution \mathcal{P}' .
- 21: **if** the improvement to the objective in the previous step is less than a factor of γ **then**
- 22: Solve only the dual in future iterations.
- 23: **end if**
- 24: **end while**

show that our local search algorithm can improve the JMP solution by using it as a starting point. When comparing the performance of our algorithm with JMP, we consistently ensure that the JMP is allowed to run for at least as long as the runtime of our algorithm. The implementation is done with Python and is available online on GitHub [1].

5.1. First-order model with the cardinality constraints We begin with the simplest of the three variants—the first-order model with the cardinality constraint. For the cardinality constraint,

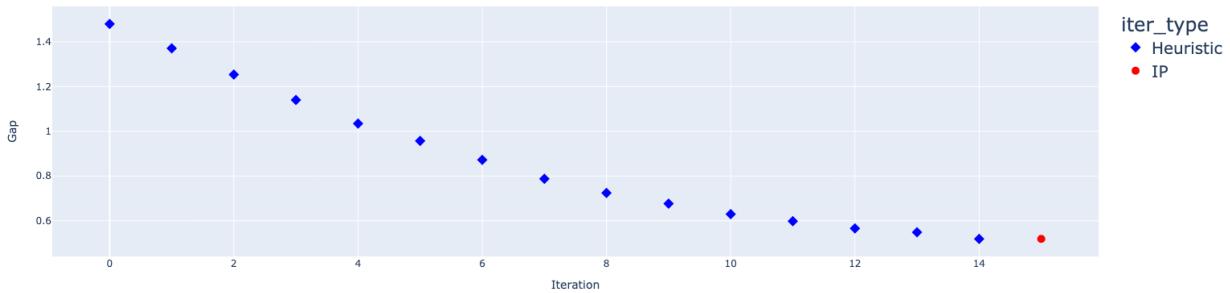


FIGURE 1. A single run of the local search algorithm for solving the first-order model with the cardinality constraints for $d = 18$. Heuristic means the local search found a locally improving move using the bit flip/swap heuristic, and IP means the algorithm failed to find an improving move with the heuristic and had to run *Gurobi*. In this example, the algorithm used the heuristic to find an improving move in every iteration, and the IP is only used in the last iteration to certify that the solution is a local optima.

we use $r = \lfloor \frac{d}{3} \rfloor$. First, we begin with a single run of the algorithm for $d = 18$ in Figure 1. The plot shows that for the simple cardinality constraint, the (Bit Flip) and (Bit Swap) heuristics work well to find a local improvement. The figure shows that the algorithm only needed to solve the IP in the last iteration to verify that the solution is a local optima.

We also compare the results of the local search algorithm to *JMP* in Table 1. These objective values are for $\ln \det(\cdot)$ and thus, a difference of 0.3 in the objective value amounts to $e^{0.3} \simeq 1.35$ factor improving in the determinant objective. Additionally, we are also able to improve the solution returned by *JMP* by using it as a starting point for our local search algorithm.

TABLE 1. A comparison of different methods for solving the first-order model with the cardinality constraint. In this table, Relaxation Value is the value of the convex relaxation (3), and values are $\ln \det$. LS Value is the value of the solution the local search algorithm returned with a random starting solution, *JMP* Value is the value of the solution *JMP* returned, and *JMP* + LS gap is the value of the solution returned by the local search algorithm with the *JMP* solution as a starting point.

d	LS Value	<i>JMP</i> Value	<i>JMP</i> + LS Value	Upper Bound (Convex Relaxation Value)
11	13.641	13.299	13.641	14.189
12	18.968	18.938	18.964	19.270
13	20.785	20.521	20.800	21.085
14	22.641	22.129	22.645	22.897
15	27.451	27.122	27.434	27.781
16	29.437	29.095	29.415	29.895
17	31.374	30.910	31.422	32.003
18	36.325	35.878	36.360	36.844
19	38.639	38.014	38.694	39.189
20	41.115	40.221	41.056	41.528

Table 2 shows the time it took to run the local search algorithm with a random starting solution. We can see from the table that the vast majority of the total time is spent on `Gurobi`. Specifically, the last iteration usually takes the most time, since it always requires k , the total number of points in the design, `Gurobi` calls to verify that a solution is locally optimal.

d	Total Time	Gurobi Time	Gurobi Iterations	Number of Iterations
11	11.15	10.43	1	1
12	11.21	10.40	1	2
13	18.02	16.82	2	8
14	21.54	20.33	1	7
15	26.30	23.97	1	12
16	27.54	25.79	1	13
17	29.73	28.12	1	9
18	30.57	28.15	1	16
19	30.65	28.71	1	17
20	82.53	78.89	2	32

TABLE 2. Local search timing statistics for solving first-order model with the cardinality constraint. Each row of the table corresponds to a different single run of the algorithm for dimension d , and the columns give information on each particular run. Total Time is the total time the algorithm took from beginning to end in seconds, Gurobi Time is the total time the algorithm spent in running `Gurobi`, Gurobi Iterations is the total number of iterations where the algorithm called `Gurobi` to search for an improvement, and Iterations is the total number of iterations for that run.

5.2. First-order model with knapsack constraints In the next experiment, we study the first-order model with two knapsack constraints. To generate the two knapsack constraints, we first generated two vectors $\mathbf{a}_1, \mathbf{a}_2 \in \mathbb{R}^d$ where $a_{i1} = 0$ and 80% of the entries are sampled uniformly at random from the set $\{0, \dots, 5\}$ and the remaining entries are uniformly at random from $\{20, \dots, 30\}$. The right-hand side the two constraints $b_i = \mathbb{E}[\mathbf{a}_i^\top \mathbf{x}] = 1/2 \sum_{j=2}^d a_{ij}$ for \mathbf{x} drawn uniformly at random from the set of vectors $\mathcal{P} = \mathcal{Y} = \{\mathbf{x} \in \{0, 1\}^d : x_1 = 1\}$ for $i = 1, 2$.

Figure 2 illustrates a single run of the local search. In this case, we solve the integer program in multiple iterations, in contrast to the first-order model with the cardinality constraint. In Figure 2, we can see that the local search algorithm would have been stuck in iteration 10 using only (Bit Flip, Bit Swap) and that using the IP helps to return a better solution.

Table 3 compares the results of `JMP` and our algorithm. `JMP` usually gets a better objective compared to our algorithm with a random starting solution. Nonetheless, for most instances, we can further improve the `JMP` solution using our local search algorithm when we use the `JMP` solution as an initial solution. Figure 3 illustrates a single run in which we used the `JMP` solution as a starting point. Here, the first iteration was an IP iteration, since the `JMP` solution was locally optimal in the bit flip/swaps neighborhood.

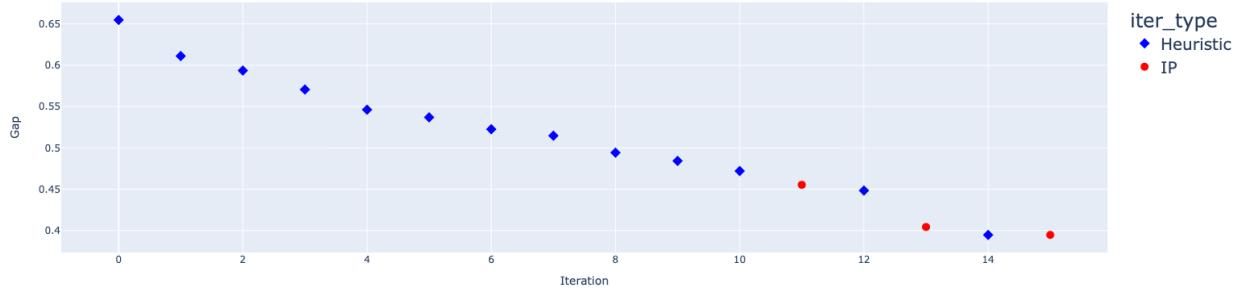


FIGURE 2. A single run of the local search algorithm for solving first-order model with knapsack constraints for $d = 13$. Heuristic means the local search found a locally improving move using the bit flip/swap heuristic, and IP means the algorithm failed to find an improving move with the heuristic and had to run `Gurobi` to solve the IP.

TABLE 3. A comparison of different methods for solving the first-order model with knapsack constraints.

d	LS Value	JMP Value	JMP + LS Value	Relaxation Value
11	17.677	17.717	17.717	18.007
12	20.675	20.767	20.767	21.068
13	23.407	23.501	23.501	23.802
14	26.192	26.264	26.274	26.678
15	29.912	29.987	29.987	30.378
16	32.668	32.636	32.767	33.262
17	35.971	36.027	36.044	36.663
18	39.090	39.152	39.152	39.729
19	42.467	42.518	42.598	43.193
20	45.928	45.924	45.924	46.661

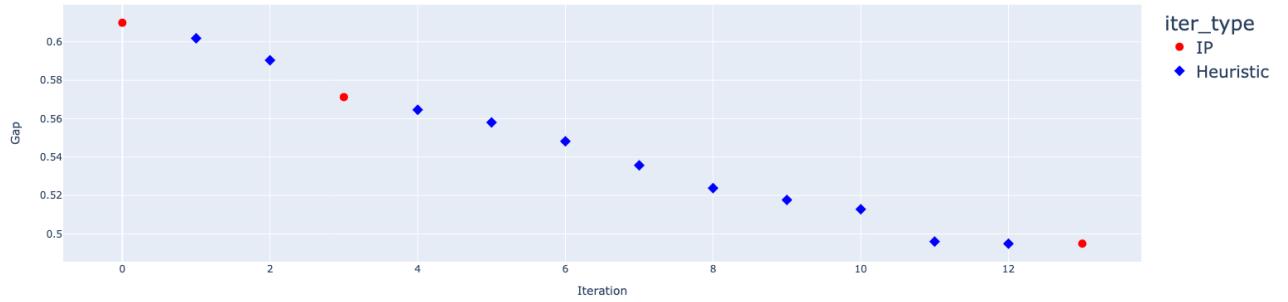


FIGURE 3. A single run of the local search algorithm for solving first-order model with knapsack constraints for $d = 16$. The local search is initialized with a solution given by JMP, and the IP is needed in the first iteration to find an improving move. A single run of the local search algorithm. The Heuristic means the local search found a locally improving move using the bit flip/swap heuristic, and IP means the algorithm failed to find an improving move with the heuristic and had to run `Gurobi` to solve IP. The Gap is the difference between the value of the convex relaxation and the value of the solution.

Table 4 displays the time it took to run the local search algorithm. Overall, the times are a bit higher in some runs than those with the cardinality constraint, since the Gurobi was executed in multiple iterations to solve IP.

TABLE 4. Local search timing statistics for solving first-order model with knapsack constraints.

d	Total Time	Gurobi Time	Gurobi Iterations	Iterations
11	9.10	8.09		1
12	13.61	12.22		1
13	32.78	30.42		3
14	28.06	26.84		2
15	20.89	19.79		1
16	44.77	42.65		2
17	24.64	22.49		1
18	27.41	25.50		1
19	55.36	52.26		2
20	73.51	69.19		2
				35

Table 5 shows timings for solving the convex relaxation (3). A large portion of the time taken to solve the convex relaxation is spent solving the restricted dual program (8) with Mosek. In Figure 4, we show the objective improvements returned by the column generation algorithm for a single instance. When we solve the separation problem (Sub) optimally, we can use the value to construct a feasible dual solution, as shown in Proposition 2, to get an upper bound on ϕ_R which is called IP bound in the figure.

TABLE 5. Timing statistics (in seconds) of the column generation algorithm for solving the first-order model with knapsack constraints. Each row corresponds to timings for a single run of the column generation for dimension d . Total Time is the total time the algorithm took from beginning to end, Gurobi Time is the total amount of time the algorithm spent in Gurobi, Mosek Time is the total time the algorithm spent in Mosek, and Iterations is the total number of iterations it took.

d	Total Time	Gurobi Time	Mosek Time	IPs solved	Iterations
11	11.45	3.44	0.13	17	17
12	20.32	8.08	1.19	33	41
13	17.64	4.29	0.07	16	18
14	53.99	17.36	0.11	44	67
15	113.31	32.34	43.32	82	103
16	151.70	63.94	27.31	130	201
17	3350.19	112.5	3173.13	181	260
18	3392.73	118.96	3194.11	181	218
19	22123.85	225.19	21799.30	255	340
20	45500.13	307.95	45083.08	236	365

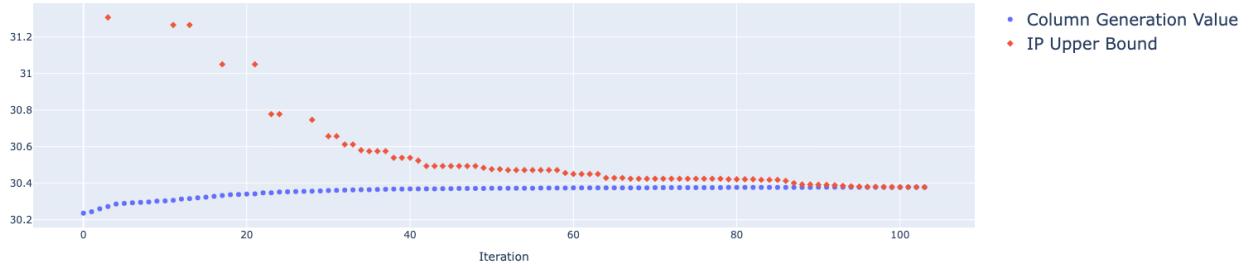


FIGURE 4. A single run of the column generation algorithm for the first-order model with knapsack constraints for $d = 15$. The IP upper bound is the value of a dual feasible solution, and the Column Generation Value is the value of the column generation algorithm.

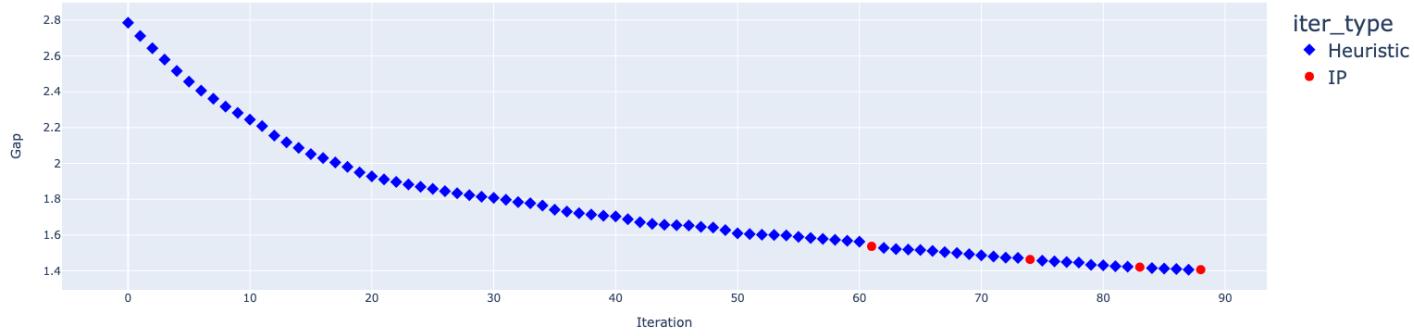


FIGURE 5. A single run of the local search algorithm for solving second-order model with knapsack constraints for $d = 16$. Heuristic means the local search found a locally improving move using the bit flip/swap heuristic and IP means the algorithm failed to find an improving move with the heuristic and had to run *Gurobi*. Using only the heuristic, the local search would have gotten stuck at iteration 60 with a gap of 1.562. Using the IP allows the algorithm to continue by improving the gap to 1.535 at iteration 61 and finally terminate at iteration 88 with a gap of 1.406.

5.3. Second-order model with knapsack constraints In this experiment, we let the first d monomials as degree one monomials, i.e., we let $m(\mathbf{x}) = x_i$ for $i \in [d]$ and the rest are the first $(\lfloor \frac{d}{2} \rfloor)$ pairs out of all $\binom{d}{2}$ pairs. This means that we also incorporate all the monomials in the form of $x_\ell x_j$ for all pairs $\{\ell, j\}$ from the set $\{2, 3, \dots, \lfloor \frac{d}{2} \rfloor + 1\}$ with $\ell \neq j$. The two knapsack constraints are the same as those used in Section 5.2.

Figure 5 illustrates a single run of the local search algorithm with $d = 20$. We can see that the IP is much more important for finding an improving move for this variant of the design problem. The heuristic failed to find an improving solution during earlier iterations, and the IP was used in more iterations.

Table 6 shows the run times of the local search algorithm with a random starting point. The times for this variant are significantly higher than those in the previous variants since the IP was used in each run more frequently, and the dimension of the design points is much larger due to degree-two monomials.

Table 7 shows a comparison between the value of the solution returned by the local search algorithm and JMP. For this variant, we are able to almost always get a better value than JMP and also improve the JMP solution using the local search.

TABLE 6. Local search timing statistics for second-order model with knapsack constraints.

d	Total Time	Gurobi Time	Gurobi Iterations	Iterations
11	71.26	67.88	1	9
12	147.75	144.44	1	30
13	147.80	141.45	2	30
14	364.51	355.45	3	38
15	684.80	672.79	3	63
16	1690.40	1664.10	4	89
17	1480.17	1422.48	2	97
18	1373.81	1316.34	1	98
19	2048.41	1976.67	2	141
20	6926.83	6722.65	6	247

TABLE 7. Local search values for second-order model with knapsack constraints. Then relaxation values listed for $d = 19, 20$ are the best upper bounds found for the instances using Proposition 2. For $d = 19$, the relaxation value of the optimal solution is in the interval [132.547, 132.634] and for $d = 20$ the relaxation value of the optimal solution is in the interval [164.070, 164.273]. The lower bound comes is the value of a feasible primal solution, and the upper bound is the value of a feasible dual solution.

d	LS Value	JMP Value	JMP + LS Value	Relaxation Value
11	32.804	32.824	32.824	33.367
12	46.819	46.638	46.707	47.484
13	45.761	45.505	45.627	46.426
14	68.768	68.744	69.067	69.866
15	73.762	73.513	73.712	74.798
16	90.827	90.272	90.681	92.233
17	98.448	98.137	98.531	100.451
18	120.817	119.719	120.901	122.743
19	130.339	130.183	130.384	132.634*
20	161.347	161.097	161.298	164.273*

6. Conclusion We developed local search and column generation algorithms for D -optimal design problems that are centered around solving a quadratic optimization problem (Sub). For the local search algorithm, we established a theoretical approximation guarantee, assuming that

(Sub) can be solved using an algorithm with a known approximation factor. Additionally, we showed that achieving an approximation ratio better than $1/2$ for (Sub) in the second-order model is NP-Hard.

We implemented the local search and column generation for both the first-order and second-order models with knapsack constraints. Our implementation solves (Sub) with an integer program, and the experimental results in Section 5 highlight the importance of the integer program in obtaining high-quality solutions. As the model complexity increases, basic heuristics such as (Bit Flip) and (Bit Swap) fail to find local improvements, making the integer program essential to find improving moves in larger neighborhoods. Notably, our local search starting from a random solution often outperforms JMP , and using the solution returned by JMP as a starting point for local search typically yields further improvements.

In Section 3, we established a negative result for the pricing problem (Sub) in the special case where the matrix G is diagonally dominant for the second-order model. A natural theoretical follow-up question is what kind of approximation guarantee can be achieved in the general case, where G is any positive semi-definite matrix. Answering this would allow us to apply Theorem 1 and obtain a result similar to Corollary 1 for the second-order model. From an empirical perspective, it would be interesting to use local search and column generation algorithms in a branch-and-bound framework to find an integral optimal solution.

Acknowledgments. A. Pillai and M. Singh were supported in part by NSF grant CCF-2106444 and NSF grant CCF-1910423. G. Ponte was supported in part by CNPq GM-GD scholarship 161501/2022-2 and AFOSR grant FA9550-22-1-0172. M. Fampa was supported in part by CNPq grant 307167/2022-4. J. Lee was supported in part by AFOSR grant FA9550-22-1-0172. W. Xie was supported in part by NSF grant 2246417 and ONR grant N00014-24-1-2066.

References

- [1] GitHub - adityapillai/DOPTDesign: column generation and local search algorithms for DOPT Design — [github.com](https://github.com/adityapillai/DOPTDesign). <https://github.com/adityapillai/DOPTDesign>.
- [2] Ahron Ben-Tal and Arkadi Nemirovski. *Lectures on modern convex optimization: analysis, algorithms, and engineering applications*, volume 2. SIAM, 2001.
- [3] George Box and Ian Hau. Constrained experimental designs Part I: Construction of projection designs, 1990. Technical report, <http://digital.library.wisc.edu/1793/69117>.
- [4] George E.P. Box and Donald W. Behnken. Some new three level designs for the study of quantitative variables. *Technometrics*, 2(4):455–475, 1960.
- [5] Christoph Buchheim and Giovanni Rinaldi. Efficient reduction of polynomial zero-one optimization to the quadratic case. *SIAM Journal on Optimization*, 18(4):1398–1413, 2008.
- [6] Dennis Cook and Valerii V. Fedorov. Constrained optimization of experimental design (invited with discussion). *Statistics*, 26:129–178, 1995.
- [7] Valerii V. Fedorov. *Theory of Optimal Experiments*. Academic Press, New York-London, 1972. Translated from the Russian and edited by W. J. Studden and E. M. Klimko.
- [8] Michel X Goemans and David P Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM (JACM)*, 42(6):1115–1145, 1995.
- [9] Johan Håstad. Some optimal inapproximability results. *Journal of the ACM*, 48(4):798–859, 2001.
- [10] JMP, A Business Unit of SAS. Design of experiments guide. Cary, NC: SAS Institute Inc, 2010.
- [11] Jack Kiefer. On the nonrandomized optimality and randomized nonoptimality of symmetrical designs. *The Annals of Mathematical Statistics*, 29(3):675–699, 1958.
- [12] Jack Kiefer and Jacob Wolfowitz. Optimum designs in regression problems. *The Annals of Mathematical Statistics*, 30(2):271–294, 1959.
- [13] Chun-Wa Ko, Jon Lee, and Kevin Wayne. A spectral bound for D-optimality, 1994. Manuscript <https://marciafampa.com/pdf/KLW1994.pdf>.
- [14] Chun-Wa Ko, Jon Lee, and Kevin Wayne. Comparison of spectral and Hadamard bounds for D-optimality. In *MODA 5*, pages 21–29. Physica, Heidelberg, 1998.
- [15] Jon Lee. In situ column generation for a cutting-stock problem. *Computers & Operations Research*, 34(8):2345–2358, 2007.
- [16] Vivek Madan, Mohit Singh, Uthaipon Tantipongpipat, and Weijun Xie. Combinatorial algorithms for optimal design. In *Conference on Learning Theory*, pages 2210–2258. PMLR, 2019.
- [17] Ruth K. Meyer and Christopher J. Nachtsheim. The coordinate-exchange algorithm for constructing exact optimal experimental designs. *Technometrics*, 37(1):60–69, 1995.

- [18] Douglas C. Montgomery, Elvira N. Loredo, Duangporn Jearkpaporn, and Murat Caner Testik. Experimental designs for constrained regions. *Quality Engineering*, 14(4):587–601, 2002.
- [19] National Institute of Standards and Technology. NIST/SEMATECH e-Handbook of Statistical Methods, 2012. <https://doi.org/10.18434/M32189>.
- [20] Yuri Nesterov and Vladimir Spokoiny. Random gradient-free minimization of convex functions. *Foundations of Computational Mathematics*, 17:527–566, 2017.
- [21] Jonathan Owen and Sanjay Mehrotra. On the value of binary expansions for general mixed-integer linear programs. *Operations Research*, 50:810–819, 10 2002.
- [22] Roger Penrose. A generalized inverse for matrices. *Mathematical Proceedings of the Cambridge Philosophical Society*, 51:406–413, 1955.
- [23] Gabriel Ponte, Marcia Fampa, and Jon Lee. Exact and heuristic solution approaches for the D-optimality problem, 2022. In: Proceedings of the LIV Brazilian Symposium on Operations Research (SOBRAPO 2022), https://proceedings.science/proceedings/100311/_papers/157447/download/fulltext_file2.
- [24] Gabriel Ponte, Marcia Fampa, and Jon Lee. Branch-and-bound for D-optimality with fast local search and variable-bound tightening, 2023. preprint: <https://arxiv.org/abs/2302.07386>.
- [25] Gabriel Ponte, Marcia Fampa, and Jon Lee. Convex relaxation for the generalized maximum-entropy sampling problem. In L. Liberti, editor, *Proceedings of SEA 2024 (22nd International Symposium on Experimental Algorithms)*, volume 301 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 25:1–25:14, 2024. Long version at: <http://arxiv.org/abs/2404.01390>.
- [26] Friedrich Pukelsheim. *Optimal Design of Experiments*, volume 50 of *Classics in Applied Mathematics*. SIAM, 2006. Reprint of the 1993 original.
- [27] Ralph C. St. John and Norman R. Draper. D-optimality for regression designs: A review. *Technometrics*, 17(1):15–23, 1975.
- [28] Abraham Wald. On the efficient design of statistical investigations. *Annals of Mathematical Statistics*, 14:134–140, 1943.
- [29] William J. Welch. Branch-and-bound search for experimental designs based on D-optimality and other criteria. *Technometrics*, 24(1):41–48, 1982.