



Fast algorithms for maximizing the minimum eigenvalue in fixed dimension

Adam Brown^{a,*}, Aditi Laddha^b, Mohit Singh^a

^a Georgia Institute of Technology, Atlanta, GA, 30318, USA

^b Yale University, New Haven, CT, 06520, USA

ARTICLE INFO

Keywords:

Convex optimization
Minimum eigenvalue
Partition matroids

ABSTRACT

In the minimum eigenvalue problem, we are given a collection of vectors in $\mathbb{R}^d$, and the goal is to pick a subset B to maximize the minimum eigenvalue of the matrix $\sum_{i \in B} v_i v_i^\top$. We give a $O\left(n^{O(d \log(d)/\epsilon^2)}\right)$ -time randomized algorithm that finds an assignment subject to a partition constraint whose minimum eigenvalue is at least $(1 - \epsilon)$ times the optimum, with high probability. As a byproduct, we also get a simple algorithm for an algorithmic version of Kadison-Singer problem.

1. Introduction

Subset selection problems with spectral objectives offer a natural model for studying problems in a variety of fields, including numerical linear algebra [1], graph theory [2], convex geometry [3,4], resource allocation [5,6], and optimal design of experiments [7–9], all under a single umbrella.

In this work, we consider the minimum eigenvalue problem. In an instance of a minimum eigenvalue problem, we are given a collection of n vectors $v_1, \dots, v_n \in \mathbb{R}^d$, and the goal is to pick a subset $B \subseteq [n]$ of given vectors to maximize the minimum eigenvalue of the matrix $\sum_{i \in B} v_i v_i^\top$. The selected set B must satisfy additional constraints such as cardinality, partition, or more generally matroid constraints. While much of the focus in previous works [7–9] has been on cardinality constraints, in this work, we consider partition constraints, but note that with some additional work our approach can be extended to general matroid constraints. Partition constraints are enough to model the algorithmic version of the Kadison-Singer problem [10,11] as well as the Max-min allocation problem [5,12,13,6,14] as a special case of the minimum eigenvalue problem.

The “discrepancy” formulation of the Kadison-Singer problem (shown to be equivalent to the original formulation in [15] and proved in [10]) states that given a set of vectors $v_1, \dots, v_m \in \mathbb{R}^d$ with $\|v_i\| \leq \alpha$ and $\sum_i v_i v_i^\top = I_d$, there exists a partition of $[m]$ into two subsets S_1, S_2 , such that for every j , $\sum_{i \in S_j} v_i v_i^\top$ spectrally approximates $I_d/2$ to an additive factor of $O(\alpha)$. Algorithmically, finding such a partition is

equivalent to solving an instance of the minimum eigenvalue maximization problem under partition matroid constraints (see Section 2.2).

Another classical application of the minimum eigenvalue problem arises in the area of optimal design of experiments in statistics [16,7]. The goal in the design of experiments is to select a subset of vectors S from a given list of vectors $\{v_1, \dots, v_n\}$ such that certain measures of the covariance matrix $(\sum_{i \in S} v_i v_i^\top)^{-1}$ are small. In particular, minimizing the maximum eigenvalue of the covariance matrix, classically known as the E -design problem in statistics, is exactly the minimum eigenvalue problem. While much of the previous work has focused on the case when the selected set of measurements S must satisfy cardinality constraints, our work generalizes this problem to be studied under general matroid constraints.

1.1. Our results and contributions

In this note, we present an approximation algorithm for the minimum eigenvalue problem for partition constraints.

Theorem 1. For any $\epsilon > 0$ there is an $O\left(n^{O(d \log(d)/\epsilon^2)}\right)$ -time algorithm which, given a collection of vectors $v_{ij} \in \mathbb{R}^d$ for $i \in [k]$ and $j \in P_i$ returns an assignment $\sigma : [n] \rightarrow [n]$ such that with probability at least $1 - d^{-4}$

$$\lambda_{\min} \left(\sum_{i=1}^k v_{i\sigma(i)} v_{i\sigma(i)}^\top \right)$$

* Corresponding author.

E-mail addresses: ajmbrown@gatech.edu (A. Brown), aditi.laddha@yale.edu (A. Laddha), msingh94@gatech.edu (M. Singh).

$$\geq (1 - \epsilon) \cdot \max_{\sigma^*} \lambda_{\min} \left(\sum_{i=1}^k v_{i\sigma^*(i)} v_{i\sigma^*(i)}^\top \right).$$

Our result generalizes to give a PTAS (for constant dimension) when the objective is a general matrix function satisfying certain technical properties. In particular, this implies that a similar result as in Theorem 1 is achievable when the objective is to maximize the determinant of $\sum_{i \in B} v_i v_i^\top$ or to minimize any norm of the eigenvalues of $(\sum_{i \in B} v_i v_i^\top)^{-1}$. We use $\mathbb{S}_+^d$ to denote the set of $d \times d$ positive semidefinite matrices.

Theorem 2. Suppose we are given a partition $\mathcal{P} = (P_1, \dots, P_k)$, and a collection $\mathcal{V}$ of vectors $v_{ij} \in \mathbb{R}^d$ for $i \in [k]$ and $j \in P_i$. Let $f : \mathbb{S}_+^d \rightarrow \mathbb{R}$ be a concave, monotone, and homogeneous function given with a value and first order oracle.

For any $\epsilon > 0$, there is an $O(n^{O(d \log(d)/\epsilon^2)})$ -time randomized algorithm, which takes $(\mathcal{V}, \mathcal{P}, f)$ as input and returns an assignment $\sigma : [k] \rightarrow [n]$ such that with probability at least $1 - d^{-4}$,

$$f \left(\sum_{i=1}^k v_{i\sigma(i)} v_{i\sigma(i)}^\top \right) \geq (1 - \epsilon) \cdot \max_{\sigma^*} f \left(\sum_{i=1}^k v_{i\sigma^*(i)} v_{i\sigma^*(i)}^\top \right).$$

Although Theorem 2 is stated in terms of maximizing concave functions, our algorithm can also be applied to minimize monotone and homogeneous convex functions (e.g., trace $((\sum_{i \in B} v_i v_i^\top)^{-1})$) by considering the natural convex relaxation of the function over the matroid base polytope and using the same rounding strategy. This allows us to model a variety of experimental design problems by selecting an appropriate objective function, and using a uniform matroid to model the cardinality constraint. For example, with the determinant or trace objective mentioned earlier, we can get a PTAS for the D- and A-design problem in constant dimension. The extension from the minimum eigenvalue to a more general objective is not a special property of our particular algorithm. For example, a similar connection was made in [7] in the context of experimental design.

Using the more general technique of pipage rounding, the result can be extended to all matroids. We give an overview of the approach in Section 2.1 and the full details are available in the arxiv version [17].

Theorem 3. Suppose we are given a collection $\mathcal{V}$ of vectors $v_i \in \mathbb{R}^d$ for $i \in [n]$ and a matroid $\mathcal{M} = ([n], I)$ with an independence oracle. Let $f : \mathbb{S}_+^d \rightarrow \mathbb{R}$ be a concave, monotone, and homogeneous function given with a value and first order oracle.

For any $\epsilon > 0$ there is an $O(n^{O(d \log(d)/\epsilon^2)})$ -time algorithm which, takes $(\mathcal{V}, \mathcal{M}, f)$ as input and returns a set $B \in I$ such that with probability at least $1 - d^{-4}$

$$f \left(\sum_{i \in B} v_i v_i^\top \right) \geq (1 - \epsilon) \cdot \max_{B^* \in I} f \left(\sum_{i \in B^*} v_i v_i^\top \right).$$

We also show the application of our result to obtain an algorithmic version of the Kadison-Singer problem [10] for constant dimension. We slightly improve the run time compared to the recent work [11].

Corollary 1. Suppose we are given collection of vectors $\mathcal{U} = (u_1, \dots, u_n) \in \mathbb{R}^d$ with $\|u_i\|^2 \leq \alpha$ for any $i \in [n]$ and $\sum_{i=1}^n u_i u_i^\top = I_d$ and a constant $c > 0$ such that there exists a set T^* satisfying

$$\left(\frac{1}{2} - c\sqrt{\alpha} \right) I_d \leq \sum_{i \in T^*} u_i u_i^\top \leq \left(\frac{1}{2} + c\sqrt{\alpha} \right) I_d.$$

For any $\epsilon > 0$, there exists a randomized algorithm such which given $\mathcal{U}$ and c as input, returns a set T such that

$$(1 - \epsilon) \cdot \left(\frac{1}{2} - c\sqrt{\alpha} \right) I_d \leq \sum_{i \in T} u_i u_i^\top \leq (1 + \epsilon) \cdot \left(\frac{1}{2} + c\sqrt{\alpha} \right) I_d,$$

with probability at least $1 - O(d^{-4})$. The run time of the algorithm is $O(n^{O(d \log d / \epsilon^2)})$.

1.2. Related work

The minimum eigenvalue problem with partition constraints can be interpreted as a generalization of the max-min allocation problem. In the case of cardinality constraints, it can also model problems from experimental design and spectral sparsification. We give an overview of prior work for these special cases.

Max-min allocation: In the max-min allocation problem, we are given a set $[d]$ of agents and a set $[n]$ of items where agent $j \in [d]$ has valuation $h_{ij} \geq 0$ for item i . The goal is to select an assignment $\sigma : [n] \rightarrow [d]$ which maximizes

$$\min_{j \in [d]} \sum_{i : \sigma(i)=j} h_{ij}.$$

This can be seen as a special case of the minimum eigenvalue problem with partition constraints.

Bansal and Sviridenko [18] introduced the configuration LP as a relaxation for the max-min allocation problem but showed that it has an integrality gap of $\Omega(\sqrt{n})$ [18]. Asadpour and Saberi [5] gave a rounding scheme for the same LP, which achieves an $O(\sqrt{n} \log^3 n)$ -approximation. This was later improved by Chakrabarty et al. [6] to an $\tilde{O}(n^\epsilon)$ -approximation for any $\epsilon \in \Omega(\log \log n / \log n)$ by iteratively constructing new instances with smaller integrality gap.

Experimental design (E-optimal design): Even with cardinality constraints (uniform matroid of rank k), the minimum eigenvalue problem is NP-hard [19]. Allen-Zhu et al. [7] showed that it is possible to deterministically find a $(1 - \epsilon)$ -approximation so long as $k \geq \Omega(d/\epsilon^2)$ by rounding the natural convex relaxation. They also conjectured that this requirement was necessary. This conjecture was confirmed in [8], where they showed an integrality gap instance for the convex relaxation. Recently Lau and Zhou [9] have built on the regret minimization framework from [7] to show that a modified local search algorithm with a “smoothed” objective works as long as there is a near-optimal solution with a good condition number.

Since our algorithm also gives approximations for other objectives we can also make comparisons to other experimental design problems. In D-design the objective is select a collection of k vectors to maximize $f : X \mapsto \det(X)$ and in A-design the goal is to minimize $f : X \mapsto \text{trace}(X^{-1})$. In both cases there are combinatorial $(1 + \epsilon)$ -approximations when $k \geq \Omega(d/\epsilon)$ [20,9]. This requirement on k is weaker than if prior results for E-design are extended directly. The approximations are much worse for k close to d . For example a d -approximation for A-design when $k = d$ [8].

Spectral sparsification and Kadison-Singer The problem of rounding the natural convex programming relaxation for the minimum eigenvalue problem is closely related to spectral sparsification [2] and the Kadison-Singer problem [10]. In spectral sparsification [2], the goal is to pick a small subset of vectors $S \subseteq [n]$ such that $\sum_{i \in S} w_i v_i v_i^\top$ spectrally approximates $\sum_{i \in [n]} v_i v_i^\top$ for some weights w_i . In the cardinality constrained minimum eigenvalue problem, rounding the convex programming solution involves finding a small set S , such that $\sum_{i \in S} v_i v_i^\top$ spectrally approximates $\sum_{i \in [n]} x_i v_i v_i^\top$, where the weights x_i form the solution to the convex relaxation. Indeed [7] essentially build on this connection to obtain their results for the E-design problem discussed earlier. The Kadison-Singer problem [10] is closely related to the minimum eigenvalue problem under a partition matroid constraint. We utilize this connection in Corollary 1 to give an algorithmic version of the Kadison-Singer problem for constant dimensions. More generally, the Kadison-Singer problem can be reformulated as showing that the integrality gap of the natural relaxation of the minimum eigenvalue problem

under partition matroid constraints is at most $1/(1 - \sqrt{\epsilon^2})$ if the length of each vector is at most $O(\epsilon)$. We discuss this connection in Section 3.

2. The algorithm for partition matroids

Suppose that our ground set is split into k parts $E = P_1 \cup \dots \cup P_k$ with each part containing n elements, and we have a collection of vectors v_{ij} for $i \in [k]$ and $j \in P_i$. The goal is to select an element $\sigma(i) \in P_i$ for each i to maximize $\lambda_{\min} \left(\sum_{i=1}^k v_{i\sigma(i)} v_{i\sigma(i)}^\top \right)$.

We can construct the natural convex relaxation of this problem as follows. For each $i \in [k]$ and $j \in P_i$, we add a decision variable x_{ij} which represents whether we select the vector v_j from part P_i , i.e., if $\sigma(i) = j$. Then we get the convex program

$$\begin{aligned} \max \quad & \lambda_{\min}(X) \\ X = \quad & \sum_{i=1}^k \sum_{j \in P_i} x_{ij} \cdot v_{ij} v_{ij}^\top \\ \sum_{j \in P_i} x_{ij} = 1, \quad & \forall i \in [k] \\ x \geq 0 \end{aligned} \quad (\text{CP})$$

The constraint $\sum_{j \in P_i} x_{ij} = 1$ ensures that we have a probability distribution over the possible assignments within each part in the optimal solution.

Given an optimal solution x^* with value OPT , a natural rounding strategy is to round independently within each part. Following this rounding strategy, we get a rank 1 random matrix M_i for each part P_i with

$$\Pr(M_i = v_{ij} v_{ij}^\top) = x_{ij}^*, \quad \forall j \in P_i.$$

The following matrix concentration inequality bounds the probability of failure of this rounding strategy.

Theorem 4. [21, Theorem 5.1.1] Consider independent random matrices $M_1, \dots, M_k \in \mathbb{S}_+^d$. Set

$$\mu_{\min} = \lambda_{\min} \left(\mathbb{E} \left[\sum_{i=1}^k M_i \right] \right).$$

If $\lambda_{\max}(M_i) \leq R$ for all $i \in [k]$ a.s. then

$$\Pr \left(\lambda_{\min} \left(\sum_{i=1}^k M_i \right) < (1 - \epsilon) \mu_{\min} \right) \leq d \cdot \exp \left(\frac{-\epsilon^2 \mu_{\min}}{2R} \right).$$

If we round according to the optimal solution x^* then

$$\mathbb{E} \left[\sum_{i=1}^k M_i \right] = \sum_{i=1}^k \sum_{j \in P_i} x_{ij}^* v_{ij} v_{ij}^\top.$$

So $\mu_{\min} = OPT$, and since for our particular case M_i are rank 1, $R = \max_i \lambda_{\max}(M_i) = \max_{i,j} \|v_{ij}\|^2$. To bound the failure probability, we want $R \approx \epsilon^2 / \log(d)$, which in turn requires that $\max_{i,j} \|v_{ij}\|^2 = O(\epsilon^2 / \log(d))$. This is a very strong assumption on an instance and thus the above approach does not work directly. We now discuss how to remedy it by the following procedure.

The plan is to “guess” a suitable change of basis such that all the vectors in the support of our optimal solution have a small norm. This will be useful because of the following standard, but slightly more flexible, version of the preceding matrix concentration inequality.

Corollary 2. Let A be a positive definite matrix and consider independent random matrices $M_1, \dots, M_k \in \mathbb{S}_+^d$. Define

$$\mu_{\min} := \lambda_{\min} \left(A^{-1/2} \mathbb{E} \left[\sum_{i=1}^k M_i \right] A^{-1/2} \right).$$

If $\lambda_{\max}(A^{-1/2} M_i A^{-1/2}) \leq R$ for all $i \in [k]$ a.s. then

$$\Pr \left(\sum_{i=1}^k M_i \not\geq (1 - \epsilon) \mu_{\min} \cdot A \right) \leq d \cdot \exp \left(\frac{-\epsilon^2 \mu_{\min}}{2R} \right).$$

Again, since M_i is rank 1 for our case, we have $R = \max_{i \in [k]} \lambda_{\max}(A^{-1/2} M_i A^{-1/2}) = \max_{i,j} v_{ij}^\top A^{-1} v_{ij}$. So, to use this corollary, we first need to find a matrix A such that $v_{ij}^\top A^{-1} v_{ij} = O(\epsilon^2 / \log(d))$ for all $[i] \in [k], j \in P_i$. We will only need to consider matrices A of a specific form that uses the input vectors.

Given a subset $S \subseteq E$, we define $A_S := \sum_{(i,j) \in S} v_{ij} v_{ij}^\top$, and consider the set of long vectors in the norm induced by A_S : $L(S) := \{(i,j) \in E \setminus S : v_{ij}^\top A_S^{-1} v_{ij} > \frac{\epsilon^2}{10 \log(d)}\}$. For a fixed set S , the following convex program ensures that S is included in the solution and no “long” vectors from $L(S)$ are included in the solution.

$$\begin{aligned} \max \quad & \lambda_{\min}(X) \\ X = \quad & \sum_{i=1}^k \sum_{j \in P_i} x_{ij} \cdot v_{ij} v_{ij}^\top \\ \sum_{j \in P_i} x_{ij} = 1, \quad & \forall i \in [k] \\ x_{ij} = 0, \quad & \forall (i,j) \in L(S) \\ x_{ij} = 1, \quad & \forall (i,j) \in S \\ x \geq 0 \end{aligned} \quad (\text{CP}(S))$$

The additional constraints exclude all long vectors outside of S , and force us to include all elements of S (which may also be long). Thus, all fractional variables correspond to short vectors and we can now use the flexible matrix concentration inequalities to randomly round the optimal solution.

But, it is not clear that there is a good choice of S for which the convex program (CP(S)) is still a relaxation of the original problem. Our main lemma shows that there exists a suitable set S that is not too large.

Lemma 1. For any set T and a set of vectors $\{v_i : i \in T\}$ in $\mathbb{R}^d$ such that $\sum_{i \in T} v_i v_i^\top$ is invertible, there exists a subset $S \subseteq T$ such that $|S| = O(d \log(d) / \epsilon^2)$, $A_S = \sum_{i \in S} v_i v_i^\top$ is invertible, and for all $i \in T \setminus S$,

$$v_i^\top A_S^{-1} v_i \leq \frac{\epsilon^2}{10 \log(d)}.$$

The proof of this lemma is inspired by the local search algorithm of [20] which they used to find an approximately optimal solution to the determinant maximization problem. Their algorithm is combinatorial, and while they never do any rounding, they do rely on a convex program for the analysis.

At first glance, it may not be apparent why a subset satisfying the conditions of Lemma 1 should exist. However, in the proof, we show that any subset of T that is locally optimal with respect to a local search criteria indeed satisfies the guarantees of Lemma 1.

Proof of Lemma 1. We consider the local search process of [20]. Starting with a set S of size $\ell := 10d \log(d) / \epsilon^2 + d - 1$ such that $A_S = \sum_{i \in S} v_i v_i^\top$ is invertible, we apply the following update rule. For any $j \in T \setminus S$ and $i \in S$, if $\det(A_S) < \det(A_S - v_i v_i^\top + v_j v_j^\top)$, update $S = (S \setminus \{i\}) \cup \{j\}$ and iterate.

Let $S \subseteq T$ be a locally optimal (under single element swaps) solution for this process (such an S corresponds to the locally optimal solution determinant maximization problem subject to the cardinality constraint $|S| \leq \ell$), and let $A = A_S = \sum_{i \in S} v_i v_i^\top$. More concretely, this means that for all $i \in S$ and $j \in T \setminus S$,

$$\det(A) \geq \det(A - v_i v_i^\top + v_j v_j^\top).$$

We calculate the ratio of right-hand side with the left-hand side using the matrix determinant lemma,

$$\begin{aligned} \det(A)^{-1} \cdot \det(A - v_i v_i^\top + v_j v_j^\top) \\ &= \det(A)^{-1} \cdot \det\left(A + \begin{bmatrix} v_i & v_j \end{bmatrix} \begin{bmatrix} -v_i & v_j \end{bmatrix}^\top\right) \\ &= \det\left(I_2 + \begin{bmatrix} -v_i & v_j \end{bmatrix}^\top A^{-1} \begin{bmatrix} v_i & v_j \end{bmatrix}\right) \\ &= (1 - v_i^\top A^{-1} v_i)(1 + v_j^\top A^{-1} v_j) + (v_i^\top A^{-1} v_j)^2. \end{aligned}$$

So local optimality implies that for every $i \in S$ and $j \notin S$, $(1 - v_i^\top A^{-1} v_i)(1 + v_j^\top A^{-1} v_j) + (v_i^\top A^{-1} v_j)^2 \leq 1$. Rearranging this inequality we get

$$\begin{aligned} v_j^\top A^{-1} v_j - (v_i^\top A^{-1} v_i) \cdot (v_j^\top A^{-1} v_j) + (v_i^\top A^{-1} v_j)^2 \\ \leq v_i^\top A^{-1} v_i. \end{aligned} \quad (1)$$

Note that

$$\begin{aligned} \sum_{i \in S} v_i^\top A^{-1} v_i &= \sum_{i \in S} \text{trace}(v_i^\top A^{-1} v_i) \\ &= \sum_{i \in S} \text{trace}(v_i v_i^\top A^{-1}) = \text{trace}(A A^{-1}) = d \end{aligned}$$

and, for all $j \in T$,

$$\begin{aligned} \sum_{i \in S} (v_i^\top A^{-1} v_j)^2 &= \sum_{i \in S} v_j^\top A^{-1} v_i^\top v_i A^{-1} v_j \\ &= v_j^\top A^{-1} A A^{-1} v_j = v_j^\top A^{-1} v_j. \end{aligned}$$

So for a fixed $j \in T \setminus S$, summing equation (1) over all $i \in S$ implies $\ell \cdot v_j^\top A^{-1} v_j - d \cdot v_j^\top A^{-1} v_j + v_j^\top A^{-1} v_j \leq d$. Rearranging, we see that for any $j \in T \setminus S$,

$$v_j^\top A^{-1} v_j \leq \frac{d}{\ell - d + 1} = \frac{\epsilon^2}{10 \log(d)},$$

where the last equality follows from the choice of

$$\ell = 10d \log(d)/\epsilon^2 + d - 1. \quad \square$$

We will apply this lemma to the set $T = \{v_{i\sigma^*(i)} : i \in [k]\}$ where σ^* is the choice function that maximizes the minimum eigenvalue, i.e., when T contains the vectors from an optimal integral assignment. In particular, we get the following corollary.

Lemma 2. *There is a subset $S \subseteq E$ such that $|S| = O(d \log(d)/\epsilon^2)$, and the convex program (CP(S)) is a relaxation for the minimum eigenvalue problem.*

As d is a constant, the size of the set S we search for is also constant. Thus, there are at most $O(n^{O(d \log(d)/\epsilon^2)})$ possible choices for S . We will consider each choice in turn to guess the correct set. Note that trying every set of the appropriate size will be the dominant factor in determining the algorithm's runtime.

The following lemma proves that for any fixed subset S , rounding the optimal solution to (CP(S)) gives a good approximation to the optimal value of (CP(S)).

Lemma 3. *Let $S \subseteq E$ be an independent set, and let x be a feasible solution to (CP(S)). Then rounding randomly in each part outputs an assignment $\sigma : [k] \rightarrow E$ with $\sigma(i) \in P_i$ such that*

$$\sum_{i=1}^k v_{i\sigma(i)} v_{i\sigma(i)}^\top \geq (1 - \epsilon) \cdot \sum_{i=1}^k \sum_{j \in P_i} x_{ij} v_{ij} v_{ij}^\top$$

with probability at least $1 - O(d^{-4})$.

Proof. Let $X = \sum_{i=1}^k \sum_{j \in P_i} x_{ij} v_{ij} v_{ij}^\top$. The matrix X contains $v_{ij} v_{ij}^\top$ with coefficient 1 for every $(i, j) \in S$. Thus $X \geq \sum_{i \in S} v_i v_i^\top$, so $v_j^\top X^{-1} v_j \leq$

$\frac{\epsilon^2}{10 \log(d)}$ for all $j \notin L \cup S$. For the purposes of the analysis, for every $j \in S$ we can replace the vector v_j with $(v_j^\top X^{-1} v_j) \cdot \sqrt{10 \log(d)/\epsilon^2}$ copies of the same vector scaled down to have squared-length at most $\epsilon^2/(10 \log(d))$ with respect to X . Since all elements of S get value 1 in x , we can similarly extend the vector x so that it has a 1 in all the copied entries. Since these values of x are deterministic, nothing changes about the resulting distribution over matrices, but we can now assume that $v_{ij}^\top X^{-1} v_{ij} < \epsilon^2/(10 \log(d))$ for all (i, j) in the support of x .

Next, define random matrices $M_1, \dots, M_k$ such that for any $i \in [k]$, $\Pr(M_i = v_{ij} v_{ij}^\top) = x_{ij}$ for all $j \in P_i$. We then apply Corollary 2 with $A = X$ on random matrices $M_1, \dots, M_k$. Since x is not supported on L ,

$$R = \max_i \lambda_{\max}(X^{-1/2} M_i X^{-1/2}) \leq \frac{\epsilon^2}{10 \log(d)}.$$

In addition, as $\mathbb{E}(\sum_i M_i) = X$ by definition, we have $\mu_{\min} = 1$.

Thus, if $\sigma : [k] \rightarrow E$ is the choice function obtained by independent rounding,

$$\begin{aligned} \Pr_\sigma \left[\sum_{i=1}^k v_{i\sigma(i)} v_{i\sigma(i)}^\top \not\geq (1 - \epsilon) X \right] \\ \leq d \cdot \exp(-5 \log(d)) = d^{-4}. \quad \square \end{aligned}$$

Combining this lemma with the earlier guarantee that there exists a set S of reasonable size such that (CP(S)) is a relaxation, we get the following algorithm: try all possible choices for the set S and return the solution with the best objective.

Algorithm 1 Algorithm to find an approximation to OPT .

-
- 1: **Input:** Partition matroid $\mathcal{M}$ with k parts $P_1, \dots, P_k$
 - 2: **for** each $S \subseteq [n]$ such that $|S| = 10d \log(d)/\epsilon^2 + d - 1$ **do**
 - 3: $x^* \leftarrow$ optimal solution of (CP(S)) for matroid $\mathcal{M}$
 - 4: For each $i \in [k]$, set $\sigma_S(i) = j$ with probability x_{ij}^*
 - 5: **end for**
 - 6: Return the assignment function σ_S which maximizes $\lambda_{\min} \left(\sum_i v_{i\sigma_S(i)} v_{i\sigma_S(i)}^\top \right)$ over all choices of S
-

We will prove Theorem 2, from which Theorem 1 follows as a special case. The only difference in the algorithm is that we solve the convex program (CP(S)) with a concave, monotone, homogeneous function f as the objective instead of $\lambda_{\min}$. We will be able to apply Lemma 2 since its proof makes no reference to the objective function, only the optimal solution, and we will use Lemma 3 without modification.

Proof of Theorem 2. By Lemma 2 there is a set $S \subseteq E$ with $|S| = O(d \log d/\epsilon^2)$ such that (CP(S)) is a relaxation.

Let x^* be the optimal solution of (CP(S)) and let σ^* be the choice function of the optimal basis for the integral problem (both with objective f). Since (CP(S)) is a relaxation, we have $f \left(\sum_{i=1}^k v_{i\sigma^*(i)} v_{i\sigma^*(i)}^\top \right) \leq f \left(\sum_{ij} x_{ij}^* v_{ij} v_{ij}^\top \right)$.

Lemma 3 implies that with high probability, the choice function obtained by rounding x^* , σ_S , is a good approximation to (CP(S)). So combining Lemma 3 with the previous inequality gives

$$\begin{aligned} f \left(\sum_{i=1}^k v_{i\sigma_S(i)} v_{i\sigma_S(i)}^\top \right) \\ \geq (1 - \epsilon) \cdot f \left(\sum_{ij} x_{ij}^* v_{ij} v_{ij}^\top \right) \\ \geq (1 - \epsilon) \cdot f \left(\sum_{i=1}^k v_{i\sigma^*(i)} v_{i\sigma^*(i)}^\top \right), \end{aligned}$$

with probability at least $1 - d^{-4}$. Since we iterate over all choice functions in step 6 of Algorithm 1, we will output a choice function σ which is at least as good as σ_S with the same probability. $\square$

2.1. General matroid constraints

In the general form of the problem, we are given a collection of vectors $v_1, \dots, v_n \in \mathbb{R}^d$ and a matroid $\mathcal{M} = ([n], \mathcal{I})$, and the goal is to find a basis $B \in \mathcal{I}$ which maximizes $\lambda_{\min}(\sum_{i \in B} v_i v_i^\top)$.

For a general matroid, the idea of finding a linear transformation under which all elements in the optimal solution have a small norm generalizes easily. So we can use the same approach of first guessing a set $S \subseteq E$ on a reasonable size and then solving the convex relaxation of the problem conditioned on S being included in the solution.

Given a subset $S \subseteq [n]$, we can again set $A_S = \sum_{i \in S} v_i v_i^\top$, and consider the set of long vectors:

$$L(S) = \left\{ i \in [n] \setminus S : v_i^\top A_S^{-1} v_i > \frac{\epsilon^2}{10 \log(d)} \right\}.$$

For a matroid $\mathcal{M}$, let $\mathcal{P}(\mathcal{M}) \subseteq [0, 1]^n$ denote the matroid base polytope. Then the following is the natural convex programming relaxation which excludes the “long” vectors.

$$\begin{aligned} \max \quad & \lambda_{\min}(X) \\ \text{s.t.} \quad & X = \sum_{i=1}^n x_i \cdot v_i v_i^\top \\ & x \in \mathcal{P}(\mathcal{M}) \\ & x_i = 0, \forall i \in L(S) \\ & x_i = 1, \forall i \in S \\ & x \geq 0 \end{aligned} \quad (\text{CP}(S))$$

This convex program can be solved in polynomial time using an independence oracle. Just like in the partition case, Lemma 1 guarantees that there is a set S for which (CP(S)) is a relaxation for the minimum eigenvalue problem. As before, after solving (CP(S)), we can guarantee that all the vectors in the fractional support of the optimal solution will have a small norm with respect to A_S .

The challenge in extending the earlier approach to general matroid constraints comes from the rounding step. For a partition matroid, we could simply round the fractional optimum of (CP(S)) independently in each part to obtain a basis. However, for more general constraints, it is not so clear how to round a fractional solution to a basis.

Instead of rounding independently, we will use the technique of pipage rounding to find a basis.

Algorithm 2 Randomized Pipage Rounding.

- 1: **Input:** Point $x \in \mathcal{P}(\mathcal{M})$, where $\mathcal{P}(\mathcal{M})$ is a matroid base polytope
 - 2: **while** x is not integral **do**
 - 3: $a, b \leftarrow$ distinct elements of $[n]$ s.t. $\exists \epsilon > 0$ with $x \pm \epsilon(e_a - e_b) \in \mathcal{P}(\mathcal{M})$
 - 4: $\ell \leftarrow \min\{y \geq 0 : x - y(e_a - e_b) \in \mathcal{P}(\mathcal{M})\}$
 - 5: $h \leftarrow \max\{y \geq 0 : x + y(e_a - e_b) \in \mathcal{P}(\mathcal{M})\}$
 - 6: $x \leftarrow \begin{cases} x - \ell(e_a - e_b) & \text{w.p. } \ell/(\ell + h) \\ x + h(e_a - e_b) & \text{w.p. } h/(\ell + h) \end{cases}$
 - 7: **end while**
 - 8: **Return** basis B of $\mathcal{M}$ with indicator vector $x \in \mathcal{P}(\mathcal{M})$
-

The following lemma is the lower-tail version of the same concentration inequality proved in [22]. The complete details of the proof can be found in the arxiv version [17].

Lemma 4. Let $\mathcal{P}(\mathcal{M})$ be a matroid base polytope and $x \in \mathcal{P}(\mathcal{M})$. Let $M_1, \dots, M_m$ be self-adjoint matrices that satisfy $\lambda_{\max}(M_i) \leq R$. Let $\mu = \lambda_{\min}(\sum_{i \in [n]} x_i M_i)$. If randomized pipage rounding (Algorithm 2) starts at x and outputs the characteristic vector of a basis B of $\mathcal{M}$, then we have

$$\Pr \left[\sum_{i \in B} M_i \leq (1 - \epsilon) \cdot \mu \right] \leq d \cdot \exp \left(-\frac{\epsilon^2 \mu}{2R} \right).$$

We use this lemma to generalize our earlier approach to all matroids.

Lemma 5. Let $S \subseteq E$ be an independent set in $\mathcal{M}$ and let x be a feasible solution to (CP(S)). Then pipage rounding starting at x^* outputs a basis B such that

$$\Pr \left[\sum_{i \in B} v_i v_i^\top \not\geq (1 - \epsilon) \sum_{i \in [n]} x_i v_i v_i^\top \right] < d^{-4}.$$

The proof is identical to that of Lemma 3, except we use the matrix concentration inequality from Lemma 4.

2.2. Application: algorithmic Kadison-Singer problem

The Kadison-Singer conjecture was resolved in [10] using the following theorem which can be interpreted as a generalization of Weaver's conjecture [15].

Theorem 5. [10, Corollary 1.5 with $r = 2$] Let $u_1, \dots, u_m \in \mathbb{R}^d$ be vectors such that $\sum_{i=1}^m u_i u_i^\top = I$ and $\|u_i\|^2 \leq \alpha$ for all i . There exists a set $T \subseteq [m]$ such that

$$\left(\frac{1}{2} - 3\sqrt{\alpha} \right) I_d \leq \sum_{i \in T} u_i u_i^\top \leq \left(\frac{1}{2} + 3\sqrt{\alpha} \right) I_d.$$

Their proof is based on analyzing interlacing families of polynomials and does not lead to an efficient algorithm to find such a subset T . In [11], they introduce an algorithmic form of the Kadison-Singer problem, which asks to find such a subset assuming it exists. For a constant $c > 0$ and a set of vectors $u_1, \dots, u_m \in \mathbb{R}^d$ such that $\|u_i\|^2 \leq \alpha$, $\sum_{i=1}^m u_i u_i^\top = I$ where there exists a subset $T \subseteq [m]$ satisfying

$$\left(\frac{1}{2} - c\sqrt{\alpha} \right) I \leq \sum_{i \in T} u_i u_i^\top \leq \left(\frac{1}{2} + c\sqrt{\alpha} \right) I, \quad (2)$$

the goal is actually to find a set $T \subseteq [m]$ which satisfies the above condition. This problem is FNP-hard when $c = 1/(4\sqrt{2})$ for general values of d [11, Theorem 2].

Their main result [11, Theorem 1] is an algorithm with running time $O\left(\binom{m}{k} \cdot \text{poly}(m, d)\right)$ where

$$k = O\left(\frac{d}{\epsilon^2} \log(d) \log\left(\frac{1}{c\sqrt{\alpha}}\right)\right)$$

which returns a set $T' \subseteq [m]$ such that

$$(1 - \epsilon) \left(\frac{1}{2} - c\sqrt{\alpha} \right) I \leq \sum_{i \in T'} u_i u_i^\top \leq (1 + \epsilon) \left(\frac{1}{2} + c\sqrt{\alpha} \right) I, \quad (3)$$

In this section, we will show how to use the rounding technique for partition matroids to give a simpler algorithm that achieves the same guarantee with the same run time, except we save the small dependence on $\log(1/c\sqrt{\alpha})$ in the exponent.

The main idea of the proof is not new; the central construction is essentially the same as the one used in [10] to prove Corollary 1.5 and deduce Weaver's conjecture from their main result. We make a few aesthetic modifications to more closely match our prior setup.

Proof of Corollary 1. Given vectors $u_1, \dots, u_m \in \mathbb{R}^d$, we construct an instance of the minimum eigenvalue with partition constraints as follows. Let $E = \{1, 2\} \times [m]$, with m parts $P_1, \dots, P_m$ so that $P_i = \{(i, 1), (i, 2)\}$ for $i \in [m]$. For each $i \in [m]$ define the vectors

$$v_{i1} = \begin{bmatrix} u_i \\ 0 \end{bmatrix} \in \mathbb{R}^{2d}, \text{ and } v_{i2} = \begin{bmatrix} 0 \\ u_i \end{bmatrix} \in \mathbb{R}^{2d}.$$

To see how v and u are related, note that for any $\delta \in [0, 1/2)$ there is a choice function $\sigma : [m] \rightarrow \{1, 2\}$ such that

$$\left(\frac{1}{2} - \delta\right) I_{2d} \leq \sum_{i=1}^m v_{i\sigma(i)} v_{i\sigma(i)}^\top \quad (4)$$

if and only if there is a set $T \subseteq [m]$ such that

$$\left(\frac{1}{2} - \delta\right) I \leq \sum_{i \in T} u_i u_i^\top \leq \left(\frac{1}{2} + \delta\right) I. \quad (5)$$

Given σ satisfying (4), let $X_1 := \sum_{i:\sigma(i)=1} u_i u_i^\top$ and $X_2 := \sum_{i:\sigma(i)=2} u_i u_i^\top$. Then X_1 and X_2 are respectively the first and second diagonal $d \times d$ block of $\sum_{i=1}^m v_{i\sigma(i)} v_{i\sigma(i)}^\top$. Therefore $\left(\frac{1}{2} - \delta\right) I_{2d} \leq \sum_{i=1}^m v_{i\sigma(i)} v_{i\sigma(i)}^\top$ if and only if $X_1 \geq \left(\frac{1}{2} - \delta\right) I_d$ and $X_2 \geq \left(\frac{1}{2} - \delta\right) I_d$. In addition, since $X_1 + X_2 = I_d$, this is equivalent to

$$\left(\frac{1}{2} - \delta\right) I_d \leq \sum_{i:\sigma(i)=1} u_i u_i^\top = I_d - X_2 \leq \left(\frac{1}{2} + \delta\right) I_d.$$

We then use Algorithm 1 to find a $(1 - \epsilon)$ approximate solution $\sigma : [m] \rightarrow \{1, 2\}$ to input $\mathcal{M}$ and vectors v_{ij} . Since we assume there is a set T satisfying (2), Theorem 1 implies that with probability at least $1 - O(d^{-4})$, Algorithm 1 will return a choice function σ^* such that $(1 - \epsilon) \left(\frac{1}{2} - c\sqrt{\alpha}\right) I_{2d} \leq \sum_{i=1}^m v_{i\sigma^*(i)} v_{i\sigma^*(i)}^\top$, and we will return the set $T' = \{i \in [m] : \sigma^*(i) = 1\}$.

From the equivalence between (4) and (5), the set $T' = \{i \in [m] : \sigma(i) = 1\}$ satisfies (3)

$$(1 - \epsilon) \left(\frac{1}{2} - c\sqrt{\alpha}\right) I_d \leq \sum_{i \in T'} u_i u_i^\top \leq (1 + \epsilon) \left(\frac{1}{2} + c\sqrt{\alpha}\right) I_d. \quad \square$$

3. Conclusion and remarks

The resolution of the Kadison-Singer problem in [10] using the interlacing families of polynomials implies the following existential result about maximizing the minimum eigenvalue under partition matroid constraints.

Theorem 6. [10, Theorem 1.4] For $\epsilon > 0$ and vectors $\{v_{ij}\}_{i \in [k], j \in [n]} \in \mathbb{R}^d$ with $\|v_{ij}\|^2 \leq \epsilon$ for all $i \in [k], j \in [n]$, if there exist $x_{ij} \geq 0$ such that

$$\sum_{i=1}^k \sum_{j=1}^n x_{ij} \cdot v_{ij} v_{ij}^\top = I_d \quad \text{and} \quad \sum_{j=1}^n x_{ij} = 1 \text{ for all } i \in [k],$$

then there exists a choice function $\sigma : [k] \rightarrow [n]$ such that

$$(1 - \sqrt{\epsilon})^2 \cdot I_d \leq \sum_{i=1}^k v_{i\sigma(i)} v_{i\sigma(i)}^\top \leq (1 + \sqrt{\epsilon})^2 \cdot I_d.$$

We can state this result equivalently as an “existential” rounding result. When $\|v_{ij}\|^2 \leq \epsilon$, Theorem 6 implies that the integrality gap of the natural convex relaxation (CP) for the minimum eigenvalue problem with partition constraints is only $1/(1 - \sqrt{\epsilon})^2$. It is an open problem to efficiently round the solution to the convex relaxation with comparable guarantees for any dimension d .

More generally, the problem of designing an approximation algorithm for the minimum eigenvalue problem under partition or matroid constraints in arbitrary dimensions remains wide open. However, checking whether there is a solution with a non-zero objective can be solved in polynomial time through matroid intersection. Recently, there has been significant progress in the case of maximizing the determinant [4,3,23–27], but it remains open whether those techniques can be applied to the minimum eigenvalue problem.

CRedit authorship contribution statement

Adam Brown: Writing – review & editing, Writing – original draft, Investigation, Conceptualization. **Aditi Laddha:** Writing – review &

editing, Writing – original draft, Investigation, Conceptualization. **Mohit Singh:** Writing – review & editing, Supervision, Conceptualization.

Data availability

No data was used for the research described in the article.

Acknowledgements

Adam Brown and Mohit Singh were supported in part by NSF CCF-2106444 and NSF CCF-1910423. Aditi Laddha was supported in part by the Institute for Foundations of Data Science at Yale and NSF CCF-2007443.

References

- [1] H. Avron, C. Boutsidis, Faster subset selection for matrices and applications, *SIAM J. Matrix Anal. Appl.* 34 (4) (2013) 1464–1499, <https://doi.org/10.1137/120867287>.
- [2] J.D. Batson, D.A. Spielman, N. Srivastava, Twice-ramanujan sparsifiers, in: *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing*, STOC '09, Association for Computing Machinery, New York, NY, USA, 2009, pp. 255–262.
- [3] M.D. Summa, F. Eisenbrand, Y. Faenza, C. Moldenhauer, On largest volume simplices and sub-determinants, in: *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '15, Society for Industrial and Applied Mathematics, USA, 2015, pp. 315–323.
- [4] A. Nikolov, Randomized rounding for the largest simplex problem, in: *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing*, STOC '15, Association for Computing Machinery, New York, NY, USA, 2015, pp. 861–870.
- [5] A. Asadpour, A. Saberi, An approximation algorithm for max-min fair allocation of indivisible goods, in: *Proceedings of the Thirty-Ninth Annual ACM Symposium on Theory of Computing*, STOC '07, Association for Computing Machinery, New York, NY, USA, 2007, pp. 114–121.
- [6] D. Chakraborty, J. Chuzhoy, S. Khanna, On allocating goods to maximize fairness, in: *Proceedings of the 2009 50th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '09, IEEE Computer Society, USA, 2009, pp. 107–116.
- [7] Z. Allen-Zhu, Y. Li, A. Singh, Y. Wang, Near-optimal discrete optimization for experimental design: a regret minimization approach, *Math. Program.* 186 (11) (2017), <https://doi.org/10.1007/s10107-019-01464-2>.
- [8] A. Nikolov, M. Singh, U.T. Tantipongipat, Proportional volume sampling and approximation algorithms for a-optimal design, in: *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, 2019, pp. 1369–1386.
- [9] L.C. Lau, H. Zhou, A local search framework for experimental design, in: *Proceedings of the Thirty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '21, Society for Industrial and Applied Mathematics, USA, 2021, pp. 1039–1058.
- [10] A.W. Marcus, D.A. Spielman, N. Srivastava, Interlacing families ii: mixed characteristic polynomials and the Kadison–Singer problem, *Ann. Math.* 182 (2015) 327–350.
- [11] B. Jourdan, P. Macgregor, H. Sun, Is the algorithmic Kadison-Singer problem hard?, *arXiv:2205.02161*, 2022, <https://arxiv.org/abs/2205.02161>.
- [12] A. Asadpour, U. Feige, A. Saberi, Santa claus meets hypergraph matchings, *ACM Trans. Algorithms* 8 (3) (jul 2012), <https://doi.org/10.1145/2229163.2229168>.
- [13] U. Feige, On allocations that maximize fairness, in: *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '08, Society for Industrial and Applied Mathematics, USA, 2008, pp. 287–293.
- [14] S. Davies, T. Rothvoss, Y. Zhang, A tale of santa claus, hypergraphs and matroids, in: *Proceedings of the Thirty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '20, Society for Industrial and Applied Mathematics, USA, 2020, pp. 2748–2757.
- [15] N. Weaver, The Kadison–Singer problem in discrepancy theory, *Discrete Math.* 278 (1–3) (2004) 227–239.
- [16] F. Pukelsheim, *Optimal Design of Experiments*, SIAM, 2006.
- [17] A. Brown, A. Laddha, M. Singh, Maximizing the minimum eigenvalue in constant dimension, *arXiv:2401.14317*, 2024.
- [18] N. Bansal, M. Sviridenko, The Santa claus problem, in: *Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '06, Association for Computing Machinery, New York, NY, USA, 2006, pp. 31–40.
- [19] A. Çivril, M. Magdon-Ismail, On selecting a maximum volume sub-matrix of a matrix and related problems, *Theor. Comput. Sci.* 410 (2009) 4801–4811.
- [20] V. Madan, M. Singh, U. Tantipongipat, W. Xie, Combinatorial algorithms for optimal design, in: *Proceedings of the Thirty-Second Conference on Learning Theory*, 2019, pp. 2210–2258.
- [21] J.A. Tropp, An introduction to matrix concentration inequalities, *Found. Trends Mach. Learn.* 8 (1–2) (2015) 1–230, <https://doi.org/10.1561/22000000048>.
- [22] N.J. Harvey, N. Olver, Pipe rounding, pessimistic estimators and matrix concentration, in: *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, SIAM, 2014, pp. 926–945.

- [23] A. Nikolov, M. Singh, Maximizing determinants under partition constraints, in: Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing, STOC '16, Association for Computing Machinery, New York, NY, USA, 2016, pp. 192–201.
- [24] N. Anari, S.O. Gharan, C. Vinzant, Log-concave polynomials, entropy, and a deterministic approximation algorithm for counting bases of matroids, in: 2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS), IEEE, 2018, pp. 35–46.
- [25] M. Singh, W. Xie, Approximate positive correlated distributions and approximation algorithms for D-optimal design, in: Proceedings of SODA, 2018.
- [26] V. Madan, A. Nikolov, M. Singh, U. Tantipongpipat, Maximizing determinants under matroid constraints, in: 2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS), 2020, pp. 565–576, arXiv:2004.07886.
- [27] A. Brown, A. Laddha, M. Pittu, M. Singh, P. Tetali, Determinant maximization via matroid intersection algorithms, in: 2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS), IEEE Computer Society, Los Alamitos, CA, USA, 2022, pp. 255–266, <https://doi.ieeecomputersociety.org/10.1109/FOCS54457.2022.00031>.