DETERMINANT MAXIMIZATION WITH LINEAR MATRIX INEQUALITY CONSTRAINTS[†]

LIEVEN VANDENBERGHE[‡], STEPHEN BOYD[§], AND SHAO-PO WU[§]

Abstract. The problem of maximizing the determinant of a matrix subject to linear matrix inequalities arises in many fields, including computational geometry, statistics, system identification, experiment design, and information and communication theory. It can also be considered as a generalization of the semidefinite programming problem.

We give an overview of the applications of the determinant maximization problem, pointing out simple cases where specialized algorithms or analytical solutions are known. We then describe an interior-point method, with a simplified analysis of the worst-case complexity and numerical results that indicate that the method is very efficient, both in theory and in practice. Compared to existing specialized algorithms (where they are available), the interior-point method will generally be slower; the advantage is that it handles a much wider variety of problems.

1. Introduction. We consider the optimization problem

$$\begin{array}{ll} \text{minimize} & c^T x + \log\det G(x)^{-1} \\ \text{subject to} & G(x) \succ 0 \\ & F(x) \succeq 0, \end{array} \tag{1.1}$$

where the optimization variable is the vector $x \in \mathbf{R}^m$. The functions $G : \mathbf{R}^m \to \mathbf{R}^{l\times l}$ and $F : \mathbf{R}^m \to \mathbf{R}^{n\times n}$ are affine:

$$G(x) = G_0 + x_1 G_1 + \cdots + x_m G_m,$$
$$F(x) = F_0 + x_1 F_1 + \cdots + x_m F_m,$$

where $G_i = G_i^T$ and $F_i = F_i^T$. The inequality signs in (1.1) denote matrix inequalities, *i.e.*, $G(x) \succ 0$ means $z^T G(x) z > 0$ for all nonzero z and $F(x) \succeq 0$ means $z^T F(x) z \geq 0$ for all z. We call $G(x) \succ 0$ and $F(x) \succeq 0$ (strict and nonstrict, respectively) *linear matrix inequalities* (LMIs) in the variable x. We will refer to problem (1.1) as a max-det problem, since in many cases the term $c^T x$ is absent, so the problem reduces to maximizing the determinant of $G(x)$ subject to LMI constraints.

The max-det problem is a convex optimization problem, *i.e.*, the objective function $c^T x + \log\det G(x)^{-1}$ is convex (on $\{x \mid G(x) \succ 0\}$), and the constraint set is convex. Indeed, LMI constraints can represent many common convex constraints, including linear inequalities, convex quadratic inequalities, and matrix norm and eigenvalue constraints (see Alizadeh[1], Boyd, El Ghaoui, Feron and Balakrishnan[13], Lewis and Overton[47], Nesterov and Nemirovsky[51, §6.4], and Vandenberghe and Boyd[69]).

In this paper we describe an interior-point method that solves the max-det problem very efficiently, both in worst-case complexity theory and in practice. The method we describe shares many features of interior-point methods for linear and semidefinite programming. In particular, our computational experience (which is limited to problems of moderate size — several hundred variables, with matrices up to 100×100)

[†]Research supported in part by AFOSR (under F49620-95-1-0318), NSF (under ECS-9222391 and EEC-9420565), MURI (under F49620-95-1-0525), and a NATO Collaborative Research Grant (CRG-941269). Associated software is available at URL `http://www-isl.stanford.edu/people/boyd` and from anonymous FTP to `isl.stanford.edu` in `pub/boyd/maxdet`.

[‡]Electrical Engineering Department, University of California, Los Angeles CA 90095-1594. E-mail: `vandenbe@ee.ucla.edu`.

[§]Information Systems Laboratory, Electrical Engineering Department, Stanford University, Stanford CA 94305. E-mail: `boyd@isl.stanford.edu`, `clive@isl.stanford.edu`.

indicates that the method we describe solves the max-det problem (1.1) in a number of iterations that hardly varies with problem size, and typically ranges between 5 and 50; each iteration involves solving a system of linear equations.

Max-det problems arise in many fields, including computational geometry, statistics, and information and communication theory, so the duality theory and algorithms we develop have wide application. In some of these applications, and for very simple forms of the problem, the max-det problems can be solved by specialized algorithms or, in some cases, analytically. Our interior-point algorithm will generally be *slower* than the specialized algorithms (when the specialized algorithms can be used). The *advantage* of our approach is that it is much more general; it handles a much wider variety of problems. The analytical solutions or specialized algorithms, for example, cannot handle the addition of (convex) constraints; our algorithm for general max-det problems does.

In the remainder of §1, we describe some interesting special cases of the max-det problem, such as semidefinite programming and analytic centering. In §2 we describe examples and applications of max-det problems, pointing out analytical solutions where they are known, and interesting extensions that can be handled as general max-det problems. In §3 we describe a duality theory for max-det problems, pointing out connections to semidefinite programming duality. Our interior-point method for solving the max-det problem (1.1) is developed in §4–§9. We describe two variations: a simple 'short-step' method, for which we can prove polynomial worst-case complexity, and a 'long-step' or adaptive step predictor-corrector method which has the same worst-case complexity, but is much more efficient in practice. We finish with some numerical experiments. For the sake of brevity, we omit most proofs and some important numerical details, and refer the interested reader to the technical report [70]. A C implementation of the method described in this paper is also available [76].

Let us now describe some special cases of the max-det problem.

Semidefinite programming. When $G(x) = 1$, the max-det problem reduces to

$$\begin{array}{ll} \text{minimize} & c^T x \\ \text{subject to} & F(x) \succeq 0, \end{array} \tag{1.2}$$

which is known as a *semidefinite program* (SDP). Semidefinite programming unifies a wide variety of convex optimization problems, *e.g.*, linear programming,

$$\begin{array}{ll} \text{minimize} & c^T x \\ \text{subject to} & Ax \leq b \end{array}$$

which can be expressed as an SDP with $F(x) = \mathbf{diag}(b - Ax)$. For surveys of the theory and applications of semidefinite programming, see [1], [13], [46], [47], [51, §6.4], and [69].

Analytic centering. When $c = 0$ and $F(x) = 1$, the max-det problem (1.1) reduces to

$$\begin{array}{ll} \text{minimize} & \log\det G(x)^{-1} \\ \text{subject to} & G(x) \succ 0, \end{array} \tag{1.3}$$

which we call the *analytic centering* problem. We will assume that the feasible set $\{x \mid G(x) \succ 0\}$ is nonempty and bounded, which implies that the matrices G_i, $i = 1, \ldots, m$, are linearly independent, and that the objective $\phi(x) = \log\det G(x)^{-1}$

is strictly convex (see, *e.g.*, [69] or [12]). Since the objective function grows without bound as x approaches the boundary of the feasible set, there is a unique solution $x^\star$ of (1.3). We call $x^\star$ the *analytic center* of the LMI $G(x) \succ 0$. The analytic center of an LMI generalizes the analytic center of a set of linear inequalities, introduced by Sonnevend[64, 65].

Since the constraint cannot be active at the analytic center, $x^\star$ is characterized by the optimality condition $\nabla\phi(x^\star) = 0$:

$$(\nabla\phi(x^\star))_i = -\mathbf{Tr} G_i G(x^\star)^{-1} = 0, \quad i = 1, \ldots, m \tag{1.4}$$

(see for example Boyd and El Ghaoui[12]).

The analytic center of an LMI is important for several reasons. We will see in §5 that the analytic center can be computed very efficiently, so it can be used as an easily computed robust solution of the LMI. Analytic centering also plays an important role in interior-point methods for solving the more general max-det problem (1.1). Roughly speaking, the interior-point methods solve the general problem by solving a sequence of analytic centering problems.

Parametrization of LMI feasible set. Let us restore the term $c^T x$:

$$\begin{array}{ll} \text{minimize} & c^T x + \log\det G(x)^{-1} \\ \text{subject to} & G(x) \succ 0, \end{array} \tag{1.5}$$

retaining our assumption that the feasible set $\mathbf{X} = \{x \mid G(x) \succ 0\}$ is nonempty and bounded, so the matrices G_i are linearly independent and the objective function is strictly convex. Thus, problem (1.5) has a unique solution $x^\star(c)$, which satisfies the optimality conditions $c + \nabla\phi(x^\star(c)) = 0$, *i.e.*,

$$\mathbf{Tr} G_i G(x^\star(c))^{-1} = c_i, \quad i = 1, \ldots, m.$$

Thus for each $c \in \mathbf{R}^m$, we have a (readily computed) point $x^\star(c)$ in the set $\mathbf{X}$.

Conversely, given a point $x \in \mathbf{X}$, define $c \in \mathbf{R}^m$ by $c_i = \mathbf{Tr} G(x)^{-1} G_i$, $i = 1, \ldots, m$. Evidently we have $x = x^\star(c)$. In other words, there is a one-to-one correspondence between vectors $c \in \mathbf{R}^m$ and feasible vectors $x \in \mathbf{X}$: the mapping $c \mapsto x^\star(c)$ is a parametrization of the feasible set $\mathbf{X}$ of the strict LMI $G(x) \succ 0$, with parameter $c \in \mathbf{R}^m$. This parametrization of the set $\mathbf{X}$ is related to the *Legendre transform* of the convex function $\log\det G(x)^{-1}$, defined by

$$\mathcal{L}(y) = -\inf\{-y^T x + \log\det G(x)^{-1} \mid G(x) \succ 0\}.$$

Maximal lower bounds in the positive definite cone. Here we consider a simple example of the max-det problem. Let $A_i = A_i^T$, $i = 1, \ldots, L$, be positive definite matrices in $\mathbf{R}^{p \times p}$. A matrix X is a lower bound of the matrices A_i if $X \preceq A_i$, $i = 1, \ldots, L$; it is a maximal lower bound if there is no lower bound Y with $Y \neq X$, $Y \succeq X$.

Since the function $\log\det X^{-1}$ is monotone decreasing with respect to the positive semidefinite cone, *i.e.*,

$$0 \prec X \preceq Y \Longrightarrow \log\det Y^{-1} \leq \log\det X^{-1},$$

we can compute a maximal lower bound $A_{\rm mlb}$ by solving

$$\begin{array}{ll} \text{minimize} & \log\det X^{-1} \\ \text{subject to} & X \succ 0 \\ & X \preceq A_i, \quad i = 1, \ldots, L. \end{array} \tag{1.6}$$

This is a max-det problem with $p(p+1)/2$ variables (the elements of the matrix X), and L LMI constraints $A_i - X \succeq 0$, which we can also consider as diagonal blocks of one single block diagonal LMI

$$\mathbf{diag}(A_1 - X, A_2 - X, \ldots, A_L - X) \succeq 0.$$

Of course there are other maximal lower bounds; replacing $\log\det X^{-1}$ by any other monotone decreasing matrix function, *e.g.*, $-\mathbf{Tr}X$ or $\mathbf{Tr}X^{-1}$, will also yield (other) maximal lower bounds. The maximal lower bound A_{mlb} obtained by solving (1.6), however, has the property that it is invariant under congruence transformations, *i.e.*, if the matrices A_i are transformed to TA_iT^T, where $T \in \mathbf{R}^{p\times p}$ is nonsingular, then the maximal lower bound obtained from (1.6) is $TA_{\mathrm{mlb}}T^T$.

2. Examples and applications. In this section we catalog examples and applications. The reader interested only in duality theory and solution methods for the max-det problem can skip directly to §3.

2.1. Minimum volume ellipsoid containing given points. Perhaps the earliest and best known application of the max-det problem arises in the problem of determining the minimum volume ellipsoid that contains given points $x^1, \ldots, x^K$ in $\mathbf{R}^n$ (or, equivalently, their convex hull $\mathbf{Co}\{x^1, \ldots, x^K\}$). This problem has applications in cluster analysis (Rosen[58], Barnes[9]), and robust statistics (in ellipsoidal peeling methods for outlier detection; see Rousseeuw and Leroy[59, §7]).

We describe the ellipsoid as $\mathcal{E} = \{x \mid \|Ax + b\| \le 1\}$, where $A = A^T \succ 0$, so the volume of $\mathcal{E}$ is proportional to $\det A^{-1}$. Hence the minimum volume ellipsoid that contains the points x^i can be computed by solving the convex problem

$$\begin{array}{ll} \text{minimize} & \log\det A^{-1} \\ \text{subject to} & \|Ax^i + b\| \le 1, \; i = 1, \ldots, K \\ & A = A^T \succ 0, \end{array} \tag{2.1}$$

where the variables are $A = A^T \in \mathbf{R}^{n\times n}$ and $b \in \mathbf{R}^n$. The norm constraints $\|Ax^i + b\| \le 1$, which are just convex quadratic inequalities in the variables A and b, can be expressed as LMIs

$$\left[\begin{array}{cc} I & Ax^i + b \\ (Ax^i + b)^T & 1 \end{array}\right] \succeq 0.$$

These LMIs can in turn be expressed as one large block diagonal LMI, so (2.1) is a max-det problem in the variables A and b.

Nesterov and Nemirovsky[51, §6.5], and Khachiyan and Todd[43] describe interior-point algorithms for computing the maximum-volume ellipsoid in a polyhedron described by linear inequalities (as well as the minimum-volume ellipsoid covering a polytope described by its vertices).

Many other geometrical problems involving ellipsoidal approximations can be formulated as max-det problems. References [13, §3.7], [16] and [68] give several examples, including the maximum volume ellipsoid contained in the intersection or in the sum of given ellipsoids, and the minimum volume ellipsoid containing the sum of given ellipsoids. For other ellipsoidal approximation problems, suboptimal solutions can be computed via max-det problems.

Ellipsoidal approximations of convex sets are used in control theory and signal processing in *bounded-noise* or *set-membership* techniques. These techniques were first introduced for state estimation (see, *e.g.*, Schweppe[62, 63], Witsenhausen[75], Bertsekas and Rhodes[11], Chernousko[16, 17]), and later applied to system identification (Fogel and Huang[35, 36], Norton[52, 53, §8.6], Walter and Piet-Lahanier[71], Cheung, Yurkovich and Passino[18]), and signal processing Deller[23]. (For a survey emphasizing signal processing applications, see Deller *et al.* [24]).

Other applications include the *method of inscribed ellipsoids* developed by Tarasov, Khachiyan, and Erlikh[66], and design centering (Sapatnekar[60]).

2.2. Matrix completion problems.

Positive definite matrix completion. In a positive definite matrix completion problem we are given a symmetric matrix $A_f \in \mathbf{R}^{n\times n}$, some entries of which are fixed; the remaining entries are to be chosen so that the resulting matrix is positive definite.

Let the positions of the free (unspecified) entries be given by the index pairs (i_k, j_k), (j_k, i_k), $k = 1, \ldots, m$. We can assume that the diagonal elements are fixed, *i.e.*, $i_k \neq j_k$ for all k. (If a diagonal element, say the (l,l)th, is free, we take it to be very large, which makes the lth row and column of A_f irrelevant.) The positive definite completion problem can be cast as an SDP feasibility problem:

$$\begin{array}{ll} \text{find} & x \in \mathbf{R}^m \\ \text{such that} & A(x) \stackrel{\Delta}{=} A_f + \sum_{k=1}^{m} x_k \left(E_{i_k j_k} + E_{j_k i_k}\right) \succ 0, \end{array}$$

where E_{ij} denotes the matrix with all elements zero except the (i,j) element, which is equal to one. Note that the set $\{x \mid A(x) \succ 0\}$ is bounded since the diagonal elements of $A(x)$ are fixed.

Maximum entropy completion. The analytic center of the LMI $A(x) \succ 0$ is sometimes called the *maximum entropy completion* of A_f. From the optimality conditions (1.4), we see that the maximum entropy completion $x^\star$ satisfies

$$2\mathbf{Tr} E_{i_k j_k} A(x^\star)^{-1} = 2\left(A(x^\star)^{-1}\right)_{i_k j_k} = 0, \; k = 1, \ldots, m,$$

i.e., the matrix $A(x^*)^{-1}$ has a zero entry in every location corresponding to an unspecified entry in the original matrix. This is a very useful property in many applications; see, for example, Dempster[27], or Dewilde and Ning[30].

Parametrization of all positive definite completions. As an extension of the maximum entropy completion problem, consider

$$\begin{array}{ll} \text{minimize} & \mathbf{Tr} CA(x) + \log\det A(x)^{-1} \\ \text{subject to} & A(x) \succ 0, \end{array} \tag{2.2}$$

where $C = C^T$ is given. This problem is of the form (1.5); the optimality conditions are

$$A(x^*) \succ 0, \quad \left(A(x^*)^{-1}\right)_{i_k j_k} = C_{i_k j_k}, \;\; k = 1, \ldots, m, \tag{2.3}$$

i.e., the inverse of the optimal completion matches the given matrix C in every free entry. Indeed, this gives a parametrization of all positive definite completions: a positive definite completion $A(x)$ is uniquely characterized by specifying the elements of its inverse in the free locations, *i.e.*, $\left(A(x)^{-1}\right)_{i_k j_k}$. Problem (2.2) has been studied by Bakonyi and Woerdeman[8].

Contractive completion. A related problem is the contractive completion problem: given a (possibly nonsymmetric) matrix A_f and m index pairs (i_k, j_k), $k = 1, \ldots, m$, find a matrix

$$A(x) = A_f + \sum_{k=1}^{m} x_k E_{i_k, j_k}.$$

with spectral norm (maximum singular value) less than one.

This can be cast as a semidefinite programming feasibility problem [69]: find x such that

$$\begin{bmatrix} I & A(x) \\ A(x)^T & I \end{bmatrix} \succ 0. \tag{2.4}$$

One can define a maximum entropy solution as the solution that maximizes the determinant of (2.4), *i.e.*, solves the max-det problem

$$\begin{array}{ll} \text{maximize} & \log\det(I - A(x)^T A(x)) \\ \text{subject to} & \begin{bmatrix} I & A(x) \\ A(x)^T & I \end{bmatrix} \succ 0. \end{array} \tag{2.5}$$

See Nævdal and Woerdeman[50], Helton and Woerdeman[38]. For a statistical interpretation of (2.5), see §2.3.

Specialized algorithms and references. Very efficient algorithms have been developed for certain specialized types of completion problems. A well known example is the maximum entropy completion of a positive definite banded Toeplitz matrix (Dym and Gohberg[31], Dewilde and Deprettere[29]). Davis, Kahan, and Weinberger[22] discuss an analytic solution for a contractive completion problem with a special (block matrix) form. The methods discussed in this paper solve the *general* problem efficiently, although they are slower than the specialized algorithms where they are applicable. Moreover they have the advantage that other convex constraints, *e.g.*, upper and lower bounds on certain entries, are readily incorporated.

Completion problems, and specialized algorithms for computing completions, have been discussed by many authors, see, *e.g.*, Dym and Gohberg[31], Grone, Johnson, Sá and Wolkowicz[37], Barrett, Johnson and Lundquist[10], Lundquist and Johnson[49], Dewilde and Deprettere[29], Dembo, Mallows, and Shepp[26]. Johnson gives a survey in [41]. An interior-point method for an approximate completion problem is discussed in Johnson, Kroschel, and Wolkowicz[42].

We refer to Boyd *et al.* [13, §3.5], and El Ghaoui[32], for further discussion and additional references.

2.3. Risk-averse linear estimation. Let $y = Ax + w$ with $w \sim \mathcal{N}(0, I)$ and $A \in \mathbf{R}^{q \times p}$. Here x is an unknown quantity that we wish to estimate, y is the measurement, and w is the measurement noise. We assume that $p \leq q$ and that A has full column rank.

A linear estimator $\widehat{x} = My$, with $M \in \mathbf{R}^{p \times q}$, is unbiased if $\mathbf{E}\widehat{x} = x$ where $\mathbf{E}$ means expected value, *i.e.*, the estimator is unbiased if $MA = I$. The minimum-variance unbiased estimator is the unbiased estimator that minimizes the error variance

$$\mathbf{E}\|My - x\|^2 = \mathbf{Tr} MM^T = \sum_{i=1}^{p} \sigma_i^2(M),$$

where $\sigma_i(M)$ is the ith largest singular value of M. It is given by $M = A^+$, where $A^+ = (A^TA)^{-1}A^T$ is the pseudo-inverse of A. In fact the minimum-variance estimator is optimal in a stronger sense: it not only minimizes $\sum_i \sigma_i^2(M)$, but each singular value $\sigma_i(M)$ separately:

$$MA = I \Longrightarrow \sigma_i(A^+) \le \sigma_i(M), \quad i = 1, \dots, p. \tag{2.6}$$

In some applications estimation errors larger than the mean value are more costly, or less desirable, than errors less than the mean value. To capture this idea of *risk aversion* we can consider the objective or cost function

$$2\gamma^2 \log \mathbf{E} \exp\left(\frac{1}{2\gamma^2}\|My - x\|^2\right) \tag{2.7}$$

where the parameter γ is called the *risk-sensitivity parameter*. This cost function was introduced by Whittle in the more sophisticated setting of stochastic optimal control; see [72, §19]. Note that as $\gamma \to \infty$, the risk-sensitive cost (2.7) converges to the cost $\mathbf{E}\|My - x\|^2$, and is always larger (by convexity of exp). We can gain further insight from the first terms of the series expansion in $1/\gamma^2$:

$$\begin{aligned} 2\gamma^2 \log \mathbf{E} \exp\left(\frac{1}{2\gamma^2}\|\widehat{x} - x\|^2\right) &\simeq \mathbf{E}\|\widehat{x} - x\|^2 + \frac{1}{4\gamma^2}\left(\mathbf{E}\|\widehat{x} - x\|^4 - \left(\mathbf{E}\|\widehat{x} - x\|^2\right)^2\right) \\ &= \mathbf{E}z + \frac{1}{4\gamma^2}\mathrm{var}\ z, \end{aligned}$$

where $z = \|\widehat{x} - x\|^2$ is the squared error. Thus for large γ, the risk-averse cost (2.7) augments the mean-square error with a term proportional to the variance of the squared error.

The unbiased, risk-averse optimal estimator can be found by solving

$$\begin{array}{ll} \text{minimize} & 2\gamma^2 \log \mathbf{E} \exp\left(\frac{1}{2\gamma^2}\|My - x\|^2\right) \\ \text{subject to} & MA = I, \end{array}$$

which can be expressed as a max-det problem. The objective function can be written as

$$\begin{aligned} &2\gamma^2 \log \mathbf{E} \exp\left(\frac{1}{2\gamma^2}\|My - x\|^2\right) \\ &= 2\gamma^2 \log \mathbf{E} \exp\left(\frac{1}{2\gamma^2} w^T M^T M w\right) \\ &= \begin{cases} 2\gamma^2 \log\det(I - (1/\gamma^2)M^TM)^{-1/2} & \text{if } M^TM \prec \gamma^2 I \\ \infty & \text{otherwise} \end{cases} \\ &= \begin{cases} \gamma^2 \log\det \begin{bmatrix} I & \gamma^{-1}M^T \\ \gamma^{-1}M & I \end{bmatrix}^{-1} & \text{if } \begin{bmatrix} I & \gamma^{-1}M^T \\ \gamma^{-1}M & I \end{bmatrix} \succ 0 \\ \infty & \text{otherwise} \end{cases} \end{aligned}$$

so the unbiased risk-averse optimal estimator solves the max-det problem

$$
\begin{array}{ll}
\text{minimize} & \gamma^2 \log\det \begin{bmatrix} I & \gamma^{-1}M^T \\ \gamma^{-1}M & I \end{bmatrix}^{-1} \\
\text{subject to} & \begin{bmatrix} I & \gamma^{-1}M^T \\ \gamma^{-1}M & I \end{bmatrix} \succ 0 \\
& MA = I.
\end{array} \tag{2.8}
$$

This is in fact an analytic centering problem, and has a simple analytic solution: the least squares estimator $M = A^+$. To see this we express the objective in terms of the singular values of M:

$$
\gamma^2 \log\det \begin{bmatrix} I & \gamma^{-1}M^T \\ \gamma^{-1}M & I \end{bmatrix}^{-1} = \begin{cases} -\gamma^2 \sum_{i=1}^{p} \log(1-\sigma_i^2(M)/\gamma^2)^{-1} & \text{if } \sigma_1(M) < \gamma \\ \infty & \text{otherwise.} \end{cases}
$$

It follows from property (2.6) that the solution is $M = A^+$ if $\|A^+\| < \gamma$, and that the problem is infeasible otherwise. (Whittle refers to the infeasible case, in which the risk-averse cost is always infinite, as 'neurotic breakdown'.)

In the simple case discussed above, the optimal risk-averse and the minimum-variance estimators coincide (so there is certainly no advantage in a max-det problem formulation). When additional convex constraints on the matrix M are added, *e.g.*, a given sparsity pattern, or triangular or Toeplitz structure, the optimal risk-averse estimator can be found by including these constraints in the max-det problem (2.8) (and will not, in general, coincide with the constrained minimum-variance estimator).

2.4. Experiment design.

Optimal experiment design. As in the previous section, we consider the problem of estimating a vector x from a measurement $y = Ax + w$, where $w \sim \mathcal{N}(0, I)$ is measurement noise. The error covariance of the minimum-variance estimator is equal to $A^+(A^+)^T = (A^TA)^{-1}$. We suppose that the rows of the matrix $A = [a_1 \ \dots \ a_q]^T$ can be chosen among M possible test vectors $v^{(i)} \in \mathbf{R}^p$, $i = 1, \dots, M$:

$$
a_i \in \{v^{(1)}, \dots, v^{(M)}\}, \quad i = 1, \dots, q.
$$

The goal of experiment design is to choose the vectors a_i so that the error covariance $(A^TA)^{-1}$ is 'small'. We can interpret each component of y as the result of an experiment or measurement that can be chosen from a fixed menu of possible experiments; our job is to find a set of measurements that (together) are maximally informative.

We can write $A^TA = q\sum_{i=1}^{M} \lambda_i v^{(i)} {v^{(i)}}^T$, where λ_i is the fraction of rows a_k equal to the vector $v^{(i)}$. We ignore the fact that the numbers λ_i are integer multiples of $1/q$, and instead treat them as continuous variables, which is justified in practice when q is large. (Alternatively, we can imagine that we are designing a random experiment: each experiment a_i has the form $v^{(k)}$ with probability λ_k.)

Many different criteria for measuring the size of the matrix $(A^TA)^{-1}$ have been proposed. For example, in E-optimal design, we minimize the norm of the error covariance, $\lambda_{\max}((A^TA)^{-1})$, which is equivalent to maximizing the smallest eigenvalue of A^TA. This is readily cast as the SDP

$$\begin{array}{ll}\text{maximize} & t \\ \text{subject to} & \sum_{i=1}^{M} \lambda_i v^{(i)} v^{(i)T} \succeq tI \\ & \sum_{i=1}^{M} \lambda_i = 1 \\ & \lambda_i \geq 0, \quad i = 1, \ldots, M\end{array}$$

in the variables $\lambda_1, \ldots, \lambda_M$, and t. Another criterion is A-optimality, in which we minimize $\mathbf{Tr}(A^TA)^{-1}$. This can be cast as an SDP:

$$\begin{array}{ll}\text{minimize} & \sum_{i=1}^{p} t_i \\ \text{subject to} & \begin{bmatrix} \sum_{i=1}^{M} \lambda_i v^{(i)} v^{(i)T} & e_i \\ e_i^T & t_i \end{bmatrix} \succeq 0, \quad i = 1, \ldots, p, \\ & \lambda_i \geq 0, \quad i = 1, \ldots, M, \\ & \sum_{i=1}^{M} \lambda_i = 1,\end{array}$$

where e_i is the ith unit vector in $\mathbf{R}^p$, and the variables are λ_i, $i = 1, \ldots, M$, and t_i, $i = 1, \ldots, p$.

In D-optimal design, we minimize the determinant of the error covariance $(A^TA)^{-1}$, which leads to the max-det problem

$$\begin{array}{ll}\text{minimize} & \log\det\left(\sum_{i=1}^{M} \lambda_i v^{(i)} v^{(i)T}\right)^{-1} \\ \text{subject to} & \lambda_i \geq 0, \; i = 1, \ldots, M \\ & \sum_{i=1}^{M} \lambda_i = 1.\end{array} \tag{2.9}$$

In §3 we will derive an interesting geometrical interpretation of the D-optimal matrix A, and show that A^TA determines the minimum volume ellipsoid, centered at the origin, that contains $v^{(1)}, \ldots, v^{(M)}$.

Fedorov[33], Atkinson and Donev[7], Pukelsheim[55], and Cook and Fedorov[19] give surveys and additional references on optimal experiment design. Wilhelm[74, 73] discusses nondifferentiable optimization methods for experiment design. Jávorzky *et al.* [40] describe an application in frequency domain system identification, and compare the interior-point method discussed later in this paper with conventional algorithms. Lee *et al.* [45, 5] discuss a non-convex experiment design problem and a relaxation solved by an interior-point method.

Extensions of D-optimal experiment design. The formulation of D-optimal design as an max-det problem has the advantage that one can easily incorporate additional useful convex constraints. For example, one can add linear inequalities $c_i^T\lambda \leq \alpha_i$, which can reflect bounds on the total cost of, or time required to carry out, the experiments.

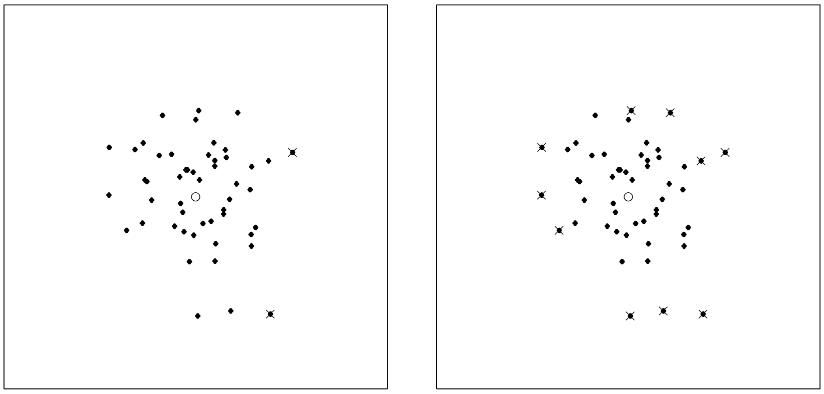

FIG. 2.1. *A D-optimal experiment design involving 50 test vectors in $\mathbf{R}^2$, with and without the 90-10 constraint. The circle is the origin; the dots are the test vectors that are not used in the experiment (*i.e.*, have a weight $\lambda_i = 0$); the crosses are the test vectors that are used (*i.e.*, have a weight $\lambda_i > 0$). Without the 90-10 constraint, the optimal design allocates all measurements to only two test vectors. With the constraint, the measurements are spread over ten vectors, with no more than 90% of the measurements allocated to any group of five vectors. See also Figure 2.2.*

We can also consider the case where each experiment yields several measurements, *i.e.*, the vectors a_i and $v^{(k)}$ become matrices. The max-det problem formulation (2.9) remains the same, except that the terms $v^{(k)}v^{(k)T}$ can now have rank larger than one. This extension is useful in conjunction with additional linear inequalities representing limits on cost or time: we can model discounts or time savings associated with performing groups of measurements simultaneously. Suppose, for example, that the cost of simultaneously making measurements $v^{(1)}$ and $v^{(2)}$ is less than the sum of the costs of making them separately. We can take $v^{(3)}$ to be the matrix

$$v^{(3)} = \left[\begin{array}{cc} v^{(1)} & v^{(2)} \end{array} \right]$$

and assign costs c_1, c_2, and c_3 associated with making the first measurement alone, the second measurement alone, and the two simultaneously, respectively.

Let us describe in more detail another useful additional constraint that can be imposed: that no more than a certain fraction of the total number of experiments, say 90%, is concentrated in less than a given fraction, say 10%, of the possible measurements. Thus we require

$$\sum_{i=1}^{\lfloor M/10 \rfloor} \lambda_{[i]} \leq 0.9, \tag{2.10}$$

where $\lambda_{[i]}$ denotes the ith largest component of λ. The effect on the experiment design will be to spread out the measurements over more points (at the cost of increasing the determinant of the error covariance). (See Figures 2.1 and 2.2.)

The constraint (2.10) is convex; it is satisfied if and only if there exists $x \in \mathbf{R}^M$ and t such that

$$\begin{array}{l} \lfloor M/10 \rfloor \, t + \sum_{i=1}^{M} x_i \leq 0.9 \\ t + x_i \geq \lambda_i, \quad i = 1, \ldots, M \\ x \geq 0 \end{array} \tag{2.11}$$

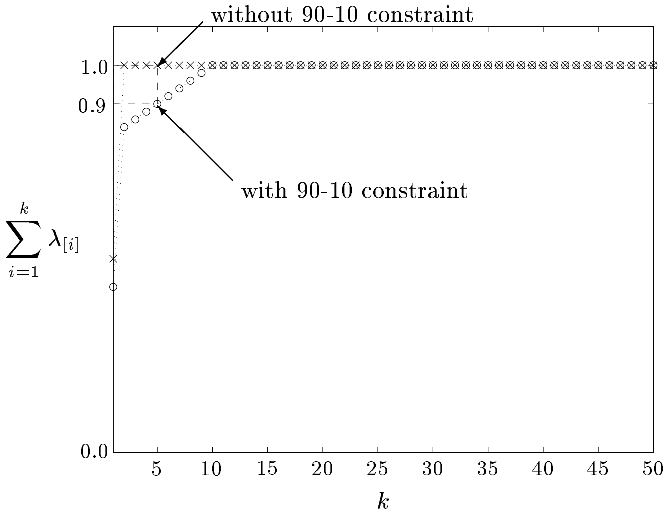

FIG. 2.2. *Experiment design of Figure 2.1. The curves show the sum of the largest k components of λ as a function of k, without the 90-10 constraint ('$\times$'), and with the constraint ('$\circ$'). The constraint specifies that the sum of the largest five components should be less than 0.9,* i.e., *the curve should avoid the area inside the dashed rectangle.*

(see [14, p.318]). One can therefore compute the D-optimal design subject to the 90-10 constraint (2.10) by adding the linear inequalities (2.11) to the constraints in (2.9) and solving the resulting max-det problem in the variables λ, x, t.

2.5. Maximum likelihood estimation of structured covariance matrices. The next example is the maximum likelihood (ML) estimation of structured covariance matrices of a normal distribution. This problem has a long history; see *e.g.*, Anderson[2, 3].

Let $y^{(1)}, \ldots, y^{(M)}$ be M samples from a normal distribution $\mathcal{N}(0, \Sigma)$. The ML estimate for Σ is the positive definite matrix that maximizes the log-likelihood function $\log \prod_{i=1}^{M} p(y^{(i)})$, where

$$p(x) = \left((2\pi)^p \det \Sigma\right)^{-1/2} \exp\left(-\frac{1}{2} x^T \Sigma^{-1} x\right).$$

In other words, Σ can be found by solving

$$\begin{array}{ll} \text{maximize} & \log\det \Sigma^{-1} - \frac{1}{M} \sum_{i=1}^{N} {y^{(i)}}^T \Sigma^{-1} y^{(i)} \\ \text{subject to} & \Sigma \succ 0. \end{array} \qquad (2.12)$$

This can be expressed as a max-det problem in the *inverse* $R = \Sigma^{-1}$:

$$\begin{array}{ll} \text{minimize} & \mathbf{Tr} SR + \log\det R^{-1} \\ \text{subject to} & R \succ 0, \end{array} \qquad (2.13)$$

where $S = \frac{1}{M} \sum_{i=1}^{n} y^{(i)} {y^{(i)}}^T$. Problem (2.13) has the straightforward analytical solution $R = S^{-1}$ (provided S is nonsingular).

It is often useful to impose additional structure on the covariance matrix Σ or its inverse R (Anderson[2, 3], Burg, Luenberger, Wenger[15], Scharf[61, §6.13], Dembo[25]). In some special cases (*e.g.*, Σ is circulant) analytical solutions are known; in other cases where the constraints can be expressed as LMIs in R, the ML estimate can be obtained from a max-det problem. To give a simple illustration, bounds on the variances Σ_{ii} can be expressed as LMIs in R:

$$\Sigma_{ii} = e_i^T R^{-1} e_i \leq \alpha \Longleftrightarrow \begin{bmatrix} R & e_i \\ e_i^T & \alpha \end{bmatrix} \succeq 0.$$

The formulation as a max-det problem is also useful when the matrix S is singular (for example, because the number of samples is too small) and, as a consequence, the max-det problem (2.13) is unbounded below. In this case we can impose constraints (*i.e.*, prior information) on Σ, for example lower and upper bounds on the diagonal elements of R.

2.6. Gaussian channel capacity.

The Gaussian channel and the water-filling algorithm. The entropy of a normal distribution $\mathcal{N}(\mu, \Sigma)$ is, up to a constant, equal to $\frac{1}{2}\log\det\Sigma$ (see Cover and Thomas[21, Chapter 9]). It is therefore not surprising that max-det problems arise naturally in information theory and communications. One example is the computation of channel capacity.

Consider a simple Gaussian communication channel: $y = x + v$, where y, x, and v are random vectors in $\mathbf{R}^n$; $x \sim \mathcal{N}(0, X)$ is the input; y is the output, and $v \sim \mathcal{N}(0, R)$ is additive noise, independent of x. This model can represent n parallel channels, or one single channel at n different time instants or n different frequencies.

We assume the noise covariance R is known and given; the input covariance X is the variable to be determined, subject to constraints (such as power limits) that we will describe below. Our goal is to maximize the *mutual information* between input and output, given by

$$\frac{1}{2}(\log\det(X + R) - \log\det R) = \frac{1}{2}\log\det(I + R^{-1/2} X R^{-1/2})$$

(see [21]). The *channel capacity* is defined as the maximum mutual information over all input covariances X that satisfy the constraints. (Thus, the channel capacity depends on R and the constraints.)

The simplest and most common constraint is a limit on the average total power in the input, *i.e.*,

$$\mathbf{E}x^T x/n = \mathbf{Tr}X/n \leq P. \tag{2.14}$$

The information capacity subject to this average power constraint is the optimal value of

$$\begin{array}{ll} \text{maximize} & \frac{1}{2}\log\det(I + R^{-1/2} X R^{-1/2}) \\ \text{subject to} & \mathbf{Tr}X \leq nP \\ & X \succeq 0 \end{array} \tag{2.15}$$

(see [21, §10]). This is a max-det problem in the variable $X = X^T$.

There is a straightforward solution to (2.15), known in information theory as the *water-filling* algorithm (see [21, §10], [20]). Let $R = V \Lambda V^T$ be the eigenvalue decomposition of R. By introducing a new variable $\widetilde{X} = V^T X V$, we can rewrite the problem as

$$\begin{array}{ll} \text{maximize} & \frac{1}{2}\log\det(I + \Lambda^{-1/2}\widetilde{X}\Lambda^{-1/2}) \\ \text{subject to} & \mathbf{Tr}\widetilde{X} \leq nP \\ & \widetilde{X} \succeq 0. \end{array}$$

Since the off-diagonal elements of $\widetilde{X}$ do not appear in the constraints, but decrease the objective, the optimal $\widetilde{X}$ is diagonal. Using Lagrange multipliers one can show that the solution is $\widetilde{X}_{ii} = \max(\nu - \lambda_i, 0)$, $i = 1, \ldots, n$, where the Lagrange multiplier ν is to be determined from $\sum \widetilde{X}_{ii} = nP$. The term 'water-filling' refers to a visual description of this procedure (see [21, §10], [20]).

Average power constraints on each channel. Problem (2.15) can be extended and modified in many ways. For example, we can replace the average *total* power constraint by an average power constraint on the individual channels, *i.e.*, we can replace (2.14) by $\mathbf{E}x_k^2 = X_{kk} \leq P$, $k = 1, \ldots, n$. The capacity subject to this constraint can be determined by solving the max-det problem

$$\begin{array}{ll} \text{maximize} & \frac{1}{2}\log\det\left(I + R^{-1/2}XR^{-1/2}\right) \\ \text{subject to} & X \succeq 0 \\ & X_{kk} \leq P, \quad k = 1, \ldots, n. \end{array}$$

The water-filling algorithm does not apply here, but the capacity is readily computed by solving this max-det problem in X. Moreover, we can easily add other constraints, such as power limits on subsets of individual channels, or an upper bound on the correlation coefficient between two components of x:

$$\frac{|X_{ij}|}{\sqrt{X_{ii}X_{jj}}} \leq \rho_{\max} \Longleftrightarrow \left[\begin{array}{cc} \sqrt{\rho_{\max}}X_{ii} & X_{ij} \\ X_{ij} & \sqrt{\rho_{\max}}X_{jj} \end{array}\right] \succeq 0.$$

Gaussian channel capacity with feedback. Suppose that the n components of x, y, and v are consecutive values in a time series. The question whether knowledge of the past values v_k helps in increasing the capacity of the channel is of great interest in information theory [21, §10.6]). In the Gaussian channel with feedback one uses, instead of x, the vector $\tilde{x} = Bv + x$ as input to the channel, where B is a strictly lower triangular matrix. The output of the channel is $y = \tilde{x} + v = x + (B + I)v$. We assume there is an average total power constraint: $\mathbf{E}\tilde{x}^T\tilde{x}/n \leq P$.

The mutual information between $\tilde{x}$ and y is

$$\frac{1}{2}\left(\log\det((B+I)R(B+I)^T + X) - \log\det R\right),$$

so we maximize the mutual information by solving

$$\begin{array}{ll} \text{maximize} & \frac{1}{2}\left(\log\det((B+I)R(B+I)^T + X) - \log\det R\right) \\ \text{subject to} & \mathbf{Tr}(BRB^T + X) \leq nP \\ & X \succeq 0 \\ & B \text{ strictly lower triangular} \end{array}$$

over the matrix variables B and X. To cast this problem as a max-det problem, we introduce a new variable $Y = (B + I)R(B + I)^T + X$ (*i.e.*, the covariance of y), and obtain

$$
\begin{array}{ll}
\text{maximize} & \log\det Y \\
\text{subject to} & \mathbf{Tr}(Y - RB^T - BR - R) \le nP \\
& Y - (B+I)R(B+I)^T \succeq 0 \\
& B \text{ strictly lower triangular.}
\end{array} \tag{2.16}
$$

The second constraint can be expressed as an LMI in B and Y,

$$
\begin{bmatrix} Y & B+I \\ (B+I)^T & R^{-1} \end{bmatrix} \succeq 0,
$$

so (2.16) is a max-det problem in B and Y.

Capacity of channel with cross-talk. Suppose the n channels are independent, *i.e.*, all covariances are diagonal, and that the noise covariance depends on X: $R_{ii} = r_i + a_i X_{ii}$, with $a_i > 0$. This has been used as a model of near-end cross-talk (see [6]). The capacity (with the total average power constraint) is the optimal value of

$$
\begin{array}{ll}
\text{maximize} & \dfrac{1}{2}\displaystyle\sum_{i=1}^{n} \log\left(1 + \frac{X_{ii}}{r_i + a_i X_{ii}}\right) \\
\text{subject to} & X_{ii} \ge 0, \quad i = 1, \ldots, n \\
& \displaystyle\sum_{i=1}^{n} X_{ii} \le nP,
\end{array}
$$

which can be cast as a max-det problem

$$
\begin{array}{ll}
\text{maximize} & \frac{1}{2}\displaystyle\sum_{i=1}^{n} \log(1 + t_i) \\
\text{subject to} & X_{ii} \ge 0, \; t_i \ge 0, \quad i = 1, \ldots, n, \\
& \begin{bmatrix} 1 - a_i t_i & \sqrt{r_i} \\ \sqrt{r_i} & a_i X_{ii} + r_i \end{bmatrix} \succeq 0, \quad i = 1, \ldots, n, \\
& \displaystyle\sum_{i=1}^{n} X_{ii} \le nP.
\end{array}
$$

The LMI is equivalent to $t_i \le X_{ii}/(r_i + a_i X_{ii})$. This problem can be solved using standard methods; the advantage of a max-det problem formulation is that we can add other (LMI) constraints on X, *e.g.*, individual power limits. As another interesting possibility, we could impose constraints that distribute the power across the channels more uniformly, *e.g.*, a 90-10 type constraint (see §2.4).

3. The dual problem. We associate with (1.1) the *dual* problem

$$
\begin{array}{ll}
\text{maximize} & \log\det W - \mathbf{Tr} G_0 W - \mathbf{Tr} F_0 Z + l \\
\text{subject to} & \mathbf{Tr} G_i W + \mathbf{Tr} F_i Z = c_i, \quad i = 1, ..., m, \\
& W = W^T \succ 0, \quad Z = Z^T \succeq 0.
\end{array} \tag{3.1}
$$

The variables are $W \in \mathbf{R}^{l \times l}$ and $Z \in \mathbf{R}^{n \times n}$. Problem (3.1) is also a max-det problem, and can be converted into a problem of the form (1.1) by elimination of the equality constraints.