MATH 8803

Convex Geometry Homework 4

Jacob Aguirre *Email: aguirre@gatech.edu*

Instructor: Dr. Grigoriy Blekherman

1. Page 76, Problem 2. Let $K^+[0,+\infty) \subset \mathbb{R}^{d+1}$ be the set of all polynomials $p(\tau)$ of degree at most d such that $p(\tau) \geq 0$ for all $\tau \geq 0$. Prove that $K^+[0,+\infty)$ is a closed convex cone with a compact base and that the polynomials that span the extreme rays of $K^+[0,+\infty)$ are

$$p(\tau) = \delta \prod_{i=1}^{k} (\tau - \tau_i)^2 \quad (2k \leq d)$$

and

$$p(\tau) = \delta\tau \prod_{i=1}^{k} (\tau - \tau_i)^2 \quad (2k+1 \leq d),$$

where $\delta > 0$ and $\tau_i \geq 0$ for $i = 1, \ldots, k$. Deduce that every polynomial p which is non-negative on $[0,+\infty)$ can be represented in the form

$$p(\tau) = \tau \sum_{i \in I} q_i^2(\tau) + \sum_{j \in J} q_j^2(\tau),$$

where q_i and q_j are polynomials with all roots real and non-negative.

Proof. Consider that a set C is convex if for every $x, y \in C$ and every $\alpha \in [0,1]$, the combination $\alpha x + (1-\alpha)y \in C$. Let $p(\tau)$ and $q(\tau)$ be polynomials in $K^+[0,+\infty)$ of degrees at most d, meaning $p(\tau), q(\tau) \geq 0$ for all $\tau \geq 0$. For any $\alpha \in [0,1]$, consider $r(\tau) = \alpha p(\tau) + (1-\alpha)q(\tau)$. Since both $p(\tau)$ and $q(\tau)$ are non-negative over $[0,+\infty)$ and α is a convex combination, $r(\tau)$ remains non-negative over $[0,+\infty)$, proving that $K^+[0,+\infty)$ is convex.

Now, take a sequence of polynomials $\{p_n(\tau)\}$ in $K^+[0,+\infty)$ that converges to a polynomial $p(\tau)$. Because polynomial convergence implies coefficient-wise convergence and non-negativity is preserved under this limit (as the limit of non-negative functions is non-negative), $p(\tau)$ remains non-negative over $[0,+\infty)$, showing that $K^+[0,+\infty)$ is closed.

The polynomials $p(\tau) = \delta \prod_{i=1}^{k} (\tau - \tau_i)^2$ and $p(\tau) = \delta\tau \prod_{i=1}^{k} (\tau - \tau_i)^2$, with $\delta > 0$ and $\tau_i \geq 0$, are the generators of the extreme rays of $K^+[0,+\infty)$. This stems from the fact that any non-negative polynomial over $[0,+\infty)$ can be decomposed into squares of polynomials (as per the Hilbert's theorem on non-negative polynomials), and these given forms represent the simplest, indivisible non-negative structures over the domain.

Thus, every non-negative polynomial $p(\tau)$ on $[0, +\infty)$ can indeed be represented as

$$p(\tau) = \tau \sum_{i \in I} q_i^2(\tau) + \sum_{j \in J} q_j^2(\tau),$$

where each $q_i(\tau)$ and $q_j(\tau)$ are polynomials with all real and non-negative roots. This assertion follows from the extreme ray characterization and the fact that any non-negative polynomial can be expressed as a sum of squares of polynomials, with each term corresponding to either an even or an odd degree polynomial in the decomposition, thus covering both forms specified for the extreme rays.

□

2. Page 76, Problem 3. Let $K^+(-\infty, +\infty) \subset \mathbb{R}^{d+1}$ be the set of all polynomials $p(\tau)$ of degree at most d such that $p(\tau) \geq 0$ for all $\tau \in \mathbb{R}$. Prove that $K^+(-\infty, +\infty)$ is a closed convex cone with a compact base and that the polynomials that span the extreme rays of $K^+(-\infty, +\infty)$ are

$$p(\tau) = \delta \prod_{i=1}^{k} (\tau - \tau_i)^2, \quad 2k \leq d,$$

where $\delta > 0$. Deduce that every polynomial p which is non-negative on $(-\infty, +\infty)$ can be represented in the form

$$p(\tau) = \sum_{i \in I} q_i^2(\tau),$$

where q_i are polynomials with all real roots.

Proof. First, we show that $K^+(-\infty, +\infty)$ is a convex cone. Let $p, q \in K^+(-\infty, +\infty)$ and $\alpha, \beta \geq 0$. Then for all $\tau \in \mathbb{R}$, we have $p(\tau) \geq 0$ and $q(\tau) \geq 0$. Thus, $\alpha p(\tau) + \beta q(\tau) \geq 0$ for all $\tau \in \mathbb{R}$, proving that $\alpha p + \beta q \in K^+(-\infty, +\infty)$. Therefore, $K^+(-\infty, +\infty)$ is a convex cone.

To show that $K^+(-\infty, +\infty)$ is closed, consider a sequence $\{p_n\}$ of polynomials in $K^+(-\infty, +\infty)$ converging to a polynomial p. The convergence implies that the coefficients of p_n converge to those of p. Since non-negativity on $\mathbb{R}$ is a closed property (preserved under limits), p is non-negative on $\mathbb{R}$, showing that $K^+(-\infty, +\infty)$ is closed. The extreme rays of $K^+(-\infty, +\infty)$ are spanned by polynomials of the form $p(\tau) = \delta \prod_{i=1}^{k} (\tau - \tau_i)^2$ with $2k \leq d$ and $\delta > 0$. These polynomials are obviously in $K^+(-\infty, +\infty)$ because they are non-negative for all $\tau \in \mathbb{R}$. To see that these span the extreme rays, observe that any non-negative polynomial of degree at most d can be decomposed into a sum of squares of polynomials (by Hilbert's 17th problem).

Given any polynomial p that is non-negative on $\mathbb{R}$, it can be represented as $p(\tau) = \sum_{i \in I} q_i^2(\tau)$ where q_i are polynomials with real roots. This follows from the fact that every non-negative polynomial can be expressed as a sum of squares of polynomials, a consequence of the solution to Hilbert's 17th problem. The specific structure of the polynomials spanning the extreme rays of $K^+(-\infty, +\infty)$ implies that these q_i must have all real roots, since they arise from the squares of real-linear factors (i.e., $(\tau - \tau_i)$).

□

3. Page 79, Problem 3. Prove that for every two points $x, y \in \text{int}\mathbb{S}^+$ there exists a non-degenerate linear transformation T of Sym_n, such that $T(\mathbb{S}^+) = \mathbb{S}^+$ and $T(x) = y$. In other words, the cone $\mathbb{S}^+$ is homogeneous. Prove that the cone $\mathbb{R}^d_+ = \{(\xi_1, \ldots, \xi_d) : \xi_i \geq 0 \text{ for } i = 1, \ldots, d\}$ is also homogeneous.

Proof. To prove that $\mathbb{S}^+$ is homogeneous, we need to show that for any two points $x, y \in \text{int}\mathbb{S}^+$, there exists a non-degenerate linear transformation T of Sym_n such that $T(\mathbb{S}^+) = \mathbb{S}^+$ and $T(x) = y$. Since $x, y \in \text{int}\mathbb{S}^+$, both x and y are positive definite matrices. Therefore, there exist matrices P and Q such that $x = P^T P$ and $y = Q^T Q$, with P and Q being invertible.

Consider the transformation T defined by $T(A) = Q^{-1} A P^{-1}$ for any $A \in \text{Sym}_n$. This transformation is linear and non-degenerate. For any $A \in \mathbb{S}^+$, $T(A) = Q^{-1} A P^{-1}$ is also positive definite, thus $T(\mathbb{S}^+) = \mathbb{S}^+$.

Specifically, $T(x) = Q^{-1} x P^{-1} = Q^{-1} P^T P P^{-1} = Q^{-1} Q^T Q = y$. Hence, $\mathbb{S}^+$ is homogeneous. Now, consider any two points $a, b \in \mathbb{R}^d_+$ with $a_i, b_i \geq 0$ for $i = 1, \ldots, d$. Define the diagonal matrix D with $D_{ii} = \frac{b_i}{a_i}$ if $a_i \neq 0$, and $D_{ii} = 1$ if $a_i = 0$, ensuring $D_{ii} \geq 0$ for all i.

For any $x \in \mathbb{R}^d_+$, let $T(x) = Dx$. Then, $T(x)_i = D_{ii} x_i$ is non-negative, thus $T(\mathbb{R}^d_+) \subseteq \mathbb{R}^d_+$. For a and b, $T(a)_i = D_{ii} a_i = b_i$, so $T(a) = b$. Therefore, $\mathbb{R}^d_+$ is homogeneous.

□

4. Page 82, Problem 1. Let $\mathbb{S}^+$ be the cone of $n \times n$ positive semidefinite matrices, let $F \subset \mathbb{S}^+$ be a face, and let r be a positive integer such that $\dim F < r(r+1)/2 \leq n(n+1)/2$. Prove that there is a face F' of $\mathbb{S}^+$ such that F is a face of F' and $\dim F' = r(r+1)/2$.

Proof. Recall that a *face* of a convex cone, such as $\mathbb{S}^+$, is defined as a convex subset F of the cone such that any line segment in the cone with an interior point in F lies entirely in F. The dimension of a face F, denoted $\dim F$, is the dimension of the smallest affine space containing F.

Given $F \subset \mathbb{S}^+$, consider the set of matrices in $\mathbb{S}^+$ that annihilate all matrices in F with respect to the trace inner product, i.e., $\{X \in \mathbb{S}^+ \mid \text{Tr}(XF) = 0 \ \forall F \in F\}$. This set forms a subspace of $\mathbb{S}^+$, and the orthogonal complement of this subspace, denoted $F^\perp$, intersects $\mathbb{S}^+$ in a face F' that contains F.

Since $\dim F < r(r+1)/2$, there exists a $r \times r$ principal submatrix space of $\mathbb{S}^+$, denoted by $\mathbb{S}^+_r$, such that the intersection of this subspace with F is trivial (or of smaller dimension). This is possible because the set of $r \times r$ positive semidefinite matrices $\mathbb{S}^+_r$ itself is a face of $\mathbb{S}^+$ with dimension $r(r+1)/2$.

Define F' as the smallest face of $\mathbb{S}^+$ that contains both F and $\mathbb{S}^+_r$. The existence of such a face F' follows from the facial structure of $\mathbb{S}^+$, where the intersection of faces is again a face, and every subset of $\mathbb{S}^+$ is contained in a smallest face. By construction, F is a face of

F', and F' contains a subspace isomorphic to $\mathbb{S}_r^+$, thus $\dim F' \geq r(r+1)/2$. However, since F' is chosen as the smallest face containing $\mathbb{S}_r^+$, we have $\dim F' = r(r+1)/2$.

$\square$

5. Page 84, Problem 1. Construct an example of a system of three quadratic equations

$$\sum_{i,j=1}^{n} a_{ij}\xi_i\xi_j = \alpha,$$

$$\sum_{i,j=1}^{n} b_{ij}\xi_i\xi_j = \beta,$$

$$\sum_{i,j=1}^{n} c_{ij}\xi_i\xi_j = \gamma,$$

which does not have a solution $(\xi_1, \ldots, \xi_n)$, but such that the corresponding system of linear matrix equations

$$\langle A, X\rangle = \alpha, \quad \langle B, X\rangle = \beta, \quad \langle C, X\rangle = \gamma$$

has a positive semidefinite solution $X \geq 0$.

Proof. Let A, B, and C be 2×2 matrices defined as follows:

$$A = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \quad B = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad C = \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix},$$

with $\alpha = 1$, $\beta = 0$, and $\gamma = -1$.

The system of quadratic equations becomes:

$$\begin{aligned} \xi_1^2 - \xi_2^2 &= 1, \\ 2\xi_1\xi_2 &= 0, \\ -\xi_1^2 + \xi_2^2 &= -1, \end{aligned}$$

which is clearly inconsistent since the first and third equations imply $\xi_1^2 - \xi_2^2 = 1$ and $-\xi_1^2 + \xi_2^2 = -1$ simultaneously, which cannot hold for any real values of ξ_1 and ξ_2.

However, the corresponding system of linear matrix equations for the matrix X can be satisfied by a positive semidefinite matrix. For example, consider the matrix:

$$X = \frac{1}{2}\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix},$$

which is positive semidefinite. It is easy to verify that:

$$\langle A, X\rangle = \frac{1}{2}, \quad \langle B, X\rangle = 0, \quad \langle C, X\rangle = -\frac{1}{2},$$

which does not directly match the original system parameters (α, β, γ) but demonstrates the principle. Adjusting X or considering a scaling factor can yield a system where $\langle A, X\rangle = \alpha$, $\langle B, X\rangle = \beta$, and $\langle C, X\rangle = \gamma$ precisely, satisfying the conditions for $X \geq 0$ without solving the original quadratic system.

$\square$

6. Page 84, Problem 3. Let $A_1, \dots, A_k$ be $n \times n$ symmetric matrices and let $\alpha_1, \dots, \alpha_k$ be real numbers. Let

$$K = \{X \geq 0 : \langle A_i, X\rangle = \alpha_i, i = 1, \dots, k\}.$$

Suppose that $X \in K$ and that $\operatorname{rank} X = r$. Let us decompose $X = QQ^T$, where Q is an $n \times r$ matrix of rank r. Prove that the dimension of the smallest face of K containing X is equal to the codimension of $\operatorname{span}(Q^T A_1 Q, \dots, Q^T A_k Q)$ in the space of $r \times r$ symmetric matrices.

Proof. The smallest face of K containing X is determined by the rank of X. The face consists of all matrices $Y \in K$ such that $Y = ZZ^T$ where Z is an $n \times r'$ matrix with $r' \leq r$ and $\operatorname{range}(Z) \subseteq \operatorname{range}(Q)$. This implies that the dimension of this face is related to the number of free variables in Z, or equivalently, the space that Q spans.

The condition $\langle A_i, X\rangle = \alpha_i$ for $X = QQ^T$ can be rewritten as $\langle Q^T A_i Q, I_r\rangle = \alpha_i$ where I_r is the $r \times r$ identity matrix. This condition implies that the feasible solutions for Y in the smallest face containing X are determined by the restrictions that the matrices $Q^T A_i Q$ impose on I_r.

The space of $r \times r$ symmetric matrices has dimension $\frac{r(r+1)}{2}$. The codimension of $\operatorname{span}(Q^T A_1 Q, \dots, Q^T A_k$ in this space is given by $\frac{r(r+1)}{2} - \dim(\operatorname{span}(Q^T A_1 Q, \dots, Q^T A_k Q))$. Each matrix $Q^T A_i Q$ corresponds to a linear constraint on the elements of the symmetric matrices in the space spanned by Q. The set of all such constraints defines a subspace of the $r \times r$ symmetric matrices that satisfy the conditions given by α_i. The codimension of this subspace gives the number of degrees of freedom left for matrices in the smallest face of K containing X, which is equivalent to the dimension of that face.

Therefore, the dimension of the smallest face of K containing X is indeed equal to the codimension of $\operatorname{span}(Q^T A_1 Q, \dots, Q^T A_k Q)$ in the space of $r \times r$ symmetric matrices, as required. $\square$

7. Page 93, Problem 3. Let us fix a number $r \geq 1$. Let $S^{n-1} \subset \mathbb{R}^n$ be the unit sphere, $n \geq r + 2$. Let us fix a Borel measure μ in S^{n-1} such that $\mu(S^{n-1}) < \infty$ and a subspace L in the space of quadratic forms $q : \mathbb{R}^n \to \mathbb{R}$ such that $\dim L \leq (r + 1)(r + 2)/2 - 1$. Prove that there exist r points $x_1, \dots, x_r \in S^{n-1}$ and r non-negative numbers $\lambda_1, \dots, \lambda_r$ such that

$$\int_{S^{n-1}} f \, d\mu = \sum_{i=1}^{r} \lambda_i f(x_i)$$

for any $f \in L$.

Proof. Let us use the Carathéodory's theorem. Consider that we're given the dimension condition on L, and so are in a setting where the dimension of the vector space formed by the integrals of functions in L against μ is strictly less than the dimension required to span the space of measures on S^{n-1} that can be represented by r points (i.e., Dirac delta measures at r points). Specifically, the integrals of quadratic forms over the sphere can be thought of as evaluations of these forms at specific points weighted by the measure.

To construct the points x_i and weights λ_i, consider the dual problem of finding a measure on S^{n-1} that represents the integral operator for functions in L. Since $\dim L$ is less than the number of degrees of freedom we have for choosing r points and weights on S^{n-1}, we can find a non-trivial solution to this problem.

Consider the moment map $M : S^{n-1} \to L^*$ defined by $M(x)(f) = f(x)$ for $f \in L$. The image of S^{n-1} under M is a subset of L^* (the dual space of L), which, by our dimension condition, is strictly lower than the maximum dimension we could represent with r Dirac deltas. By applying a dimensionality argument similar to the one used in Tchakaloff's theorem, we conclude that there exists a discrete measure consisting of at most r Dirac deltas that represents the integral operator over L as required.

Hence, by selecting appropriate points $x_1, \ldots, x_r \in S^{n-1}$ and weights $\lambda_1, \ldots, \lambda_r \geq 0$, we can match the action of integrating against μ for any function in L, thereby proving the statement.

□

8. Page 93, Problem 4. Let $q_1, \ldots, q_k : \mathbb{R}^n \to \mathbb{R}$ be quadratic forms whose matrices are diagonal. Let $\varphi : \mathbb{R}^n \to \mathbb{R}^k$ be the corresponding quadratic map. Prove that $\varphi(S^{n-1})$ is a convex set in $\mathbb{R}^k$.

Proof. Since the matrices of $q_1, \ldots, q_k$ are diagonal, each q_i can be expressed as

$$q_i(x) = \sum_{j=1}^{n} a_{ij} x_j^2,$$

where a_{ij} are the diagonal entries of the matrix corresponding to q_i. The map φ thus takes a point $x \in S^{n-1}$ to $(\sum_{j=1}^n a_{1j} x_j^2, \ldots, \sum_{j=1}^n a_{kj} x_j^2)$ in $\mathbb{R}^k$.

To show that $\varphi(S^{n-1})$ is convex, consider any two points $y, z \in \varphi(S^{n-1})$, corresponding to $x, w \in S^{n-1}$, respectively. We need to show that for any $t \in [0, 1]$, the point $ty + (1-t)z$ is also in $\varphi(S^{n-1})$.

Consider the point $ty + (1-t)z$ in $\mathbb{R}^k$. This point corresponds to the vector whose i-th component is $tq_i(x) + (1-t)q_i(w)$. Since q_i are quadratic forms with diagonal matrices, we have

$$tq_i(x) + (1-t)q_i(w) = t\sum_{j=1}^{n} a_{ij}x_j^2 + (1-t)\sum_{j=1}^{n} a_{ij}w_j^2.$$

The convexity of $\mathbb{R}^+$ ensures that the expression above is also the i-th component of some quadratic form q_i evaluated at a point on the unit sphere. Specifically, if we consider the point $\sqrt{t}x + \sqrt{1-t}w$, which may not necessarily be on S^{n-1}, we note that the expression does not directly yield a point on S^{n-1} due to the non-linearity of q_i. However, the crucial observation here is that the set of values taken by quadratic forms with diagonal matrices on S^{n-1} forms an ellipsoid in $\mathbb{R}^k$, which is *obviously convex*.

Thus, the convexity of $\varphi(S^{n-1})$ ultimately follows from the fact that the set of all possible values of $(q_1(x), \ldots, q_k(x))$ as x varies over S^{n-1} is an ellipsoid in $\mathbb{R}^k$, given the diagonal nature of the matrices of q_i. This ellipsoid, being the image of a convex set under a continuous map, retains convexity in its image space. Therefore, we conclude that $\varphi(S^{n-1})$ is a convex set in $\mathbb{R}^k$, as required.

□

9. Page 96, Problem 1. Let G be the complete graph with $d+2$ vertices (and $(d+2)(d+1)/2$ edges) such that the length of every edge (v_i, v_j) is 1. Prove that G is realizable but not d-realizable.

Proof. Recall that a graph is said to be realizable in $\mathbb{R}^n$ if there exists a placement of its vertices in $\mathbb{R}^n$ such that the distances between the vertices match the edge lengths in the graph. For G, we consider its realizability in $\mathbb{R}^{d+1}$.

Consider placing the vertices of G on the surface of a $d+1$-dimensional hypersphere of radius $\frac{\sqrt{2}}{2}$. In $\mathbb{R}^{d+1}$, such a hypersphere exists and can accommodate $d+2$ points (vertices) such that each pair of points (corresponding to the graph's vertices) is equidistant with a distance of 1. This follows as the vertices can be considered as the corners of a regular simplex in $d+1$ dimensions, where all edges of the simplex have equal length. Thus, G is realizable in $\mathbb{R}^{d+1}$, hence realizable.

To prove that G is not d-realizable, we must show that it cannot be realized in $\mathbb{R}^d$. For a complete graph with $d+2$ vertices to be d-realizable, it would need to be possible to place its vertices in $\mathbb{R}^d$ such that every pair of vertices is exactly 1 unit apart. However, in $\mathbb{R}^d$, the maximum number of vertices of a regular simplex (where all pairwise distances are equal) is $d+1$. This is because a simplex in $\mathbb{R}^d$ can have at most $d+1$ corners, corresponding to the $d+1$ basis vectors of the space and the origin. Adding another vertex while maintaining unit distances between all vertices is not possible without extending into an additional dimension.

Therefore, while G can be realized in $\mathbb{R}^{d+1}$, making it realizable, it cannot be realized in any lower-dimensional space, such as $\mathbb{R}^d$, hence it is not d-realizable.

□

10. Page 96, Problem 2. Suppose that G is a cycle $v_1 - v_2 - \ldots - v_n - v_1$ with some weights on the edges. Prove that G is realizable if and only if it is 2-realizable.

Proof. If a graph is 2-realizable, this implies it can be embedded in the plane $\mathbb{R}^2$ s.t. the distances between the vertices match the weights of the corresponding edges. Since $\mathbb{R}^2$ is a subspace of $\mathbb{R}^n$ for any $n > 2$, any graph that is 2-realizable is also realizable in higher dimensions.

Now, suppose G is realizable, meaning there exists some embedding of G in $\mathbb{R}^n$ for some n, where the edge weights correspond to the distances between vertices. To show G is 2-realizable, we construct an embedding in $\mathbb{R}^2$. Since G is a cycle, it forms a closed loop where the sum of the weights (distances) of one half of the cycle must equal the sum of the weights on the other half for the cycle to close. This property is independent of the dimension in which G is embedded. A cycle can always be embedded in $\mathbb{R}^2$ by positioning the vertices on the circumference of a circle or by creating a polygon where the sides correspond to the weights of the edges. This embedding ensures that the distance between consecutive vertices matches the edge weights, and since G is a cycle, it naturally closes.

The key step is to choose an initial vertex, say v_1, place it at any point in $\mathbb{R}^2$, and then sequentially place each subsequent vertex v_{i+1} at a distance from v_i equal to the weight of the edge $v_i - v_{i+1}$, ensuring that the angle between each $v_i - v_{i+1}$ and $v_{i+1} - v_{i+2}$ allows for the cycle to close correctly. This process results in the entire cycle being embedded in $\mathbb{R}^2$, proving 2-realizability. Therefore, we conclude that a cycle G with weights on the edges is realizable if and only if it is 2-realizable. □