Lecture notes – CS492 – Spring 2008 – Andreas Holmsen

A generalization of Carathéodory's theorem. One of the most fundamental theorems of convexity is Carathéodory's theorem. It states that the origin is contained in the convex hull of a set S in d-dimensional Euclidean space if and only if the origin is contained in a simplex spanned by points of S. Here we will prove a partitioned version of Carathéodory's theorem.

Theorem 1. *Let $A_1, \ldots, A_{d+1}$ be disjoint, non-empty, finite point sets in $\mathbb{R}^d$ such that $A_1 \cup \cdots \cup A_{d+1}$ is in general position with respect to the origin. If the origin is contained in* $\text{conv}(A_i \cup A_j)$ *for all $1 \leq i < j \leq d+1$, then the origin is contained in some simplex S with $|S \cap A_i| = 1$ for every $1 \leq i \leq d+1$.*

Here the term *"in general position with respect to the origin"* means that any k-tuple of points ($1 \leq k \leq d$) together with the origin span a k-dimensional simplex. Equivalently, it means that the vectors spanned by any k-tuple are linearly independent.

Figure 1: The set of points on the left are not in general position with respect to the origin since there is a pair of points whose connecting segment pass through the origin. The set of points on the right, however, are in general position with respect to the origin.

Let us first prove some special cases. For $d = 1$, Theorem 1 is trivial, so let us consider the case $d = 2$. In this case we have a set of red points, a set of blue points, and a set of green points, in the plane, which together are in general position with respect to the origin. Theorem 1 claims that if the origin is contained in the convex hull of the red and blue points, the red and green points, and the blue and green points, then there is a triangle with one red vertex, one blue vertex, and one green vertex which contains the origin. (Check that Theorem 1 holds for the point sets above. In fact, in the plane it doesn't really matter whether or not the points are in general position with respect to the origin).

Now to prove Theorem 1 for $d = 2$. Let us first note that if the origin is contained in the convex hull of, say, the red and blue points, then there is a triangle spanned by the red and blue points, that contains the origin, and which has at least one red vertex and one blue vertex (Excercise 1a). Let us assume that there is a red point r and blue points b_1 and b_2 such that the triangle rb_1b_2 contains the origin. Let L_i be the line that passes through b_i and the origin ($i = 1, 2$). The point r is contained in a unique open quadrant, Q, defined by the lines L_1 and L_2. (See the figure below). Let us suppose, for a contradiction, that there is no red/blue/green triangle that contains the origin. It is not hard to see (really?) that if there is a green point g that is not contained in Q, then one of the triangles rb_1g or rb_2g will contain the origin. So every green point must be contained in Q. Now, for any green point g, the triangle gb_1b_2 contains the origin. So, by the same argument as above, all the red points must be contained in Q. This means that the convex hull of the red and the green points is contained in Q, and therefore it cannot contain the origin. This contradiction establishes Theorem 1 for $d = 2$.

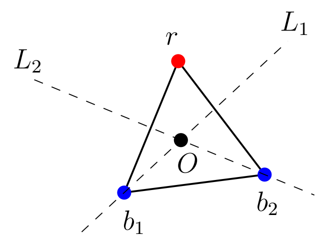

Figure 2: The lines L_1 and L_2 cut the plane into four quadrants. If there is no red/blue/green triangle containing the origin, then the green points must be contained in the open quadrant that contains the red point.

As we see, the case $d = 2$ is quite simple. Why can't we just apply the same proof for $d > 2$? Well, it is in fact not clear what to do. Let us illustrate this by considering the case $d = 3$. It is true that if the origin is contained in the convex hull of a set of red and blue points, then there is a *tetrahedron* spanned by the red and blue points, that contains the origin, and which has a at least one red vertex and one blue vertex (Excercise 1b). It is here things become more complicated: The tetrahedron could have one red vertex and three blue vertices, *or*, it could have two red vertices and two blue vertices. If we accept the situation that we need to consider both cases (and even more as d grows!), we are still stuck with the trouble of finding an analogue of the lines L_1 and L_2. At this point it should be clear that it will be difficult to prove Theorem 1, for arbitrary $d > 2$, by naively mimicking the proof of the planar case. It turns out that one way to prove Theorem 1 is to give a different higher dimensional interpretation of the fact that in the plane we can find a red-blue triangle that contains the origin (this is essentially Lemma 2, below). But before we get to that we will reformulate the problem in *spherical* terms.

Let S^{d-1} denote the $(d-1)$-dimensional unit sphere in $\mathbb{R}^d$ centered at the origin. We say that a finite point set $A \subset S^{d-1}$ is in *general position* if any $1 \leq k \leq d$ points of A span a k-dimensional linear subspace of $\mathbb{R}^d$. This is the spherical equivalent to saying that a point set is in general position with respect to the origin.

The fact that a point set is in general position on S^{d-1} implies that any d points of A are contained in some *open* hemisphere $H \subset S^{d-1}$. For $0 \leq k \leq d$, it makes sense to speak of a k-simplex of A, i.e. the *spherical convex hull* of some $k+1$ points of A, which is denoted by conv_S. More generally, if X is contained in some open hemisphere, then $\text{conv}_S X$ is the intersection of all open hemispheres that contain X.

Notice that there is a fundamental difference between the *spherical* convex hull and the regular *affine* convex hull: We only defined the spherical convex hull for point sets that are contained in some open hemisphere. The connection between the affine convex hull and the spherical convex hull is the following: *A set of points in general position on S^{d-1} contain the origin in their affine convex hull if and only if the set of points is not contained in any open hemisphere*. So a set of points on S^{d-1} that contains the origin in its affine convex hull is *precisely* the kind which do not have a spherical convex hull!

It is also important to notice the following: Given a point set in $\mathbb{R}^d$ which is in general position with respect to the origin, we can project each point from the origin to a unique point on the unit sphere to obtain a point set on S^{d-1}. The resulting point set will be in general position on the S^{d-1}. Moreover, a subset of the original points in $\mathbb{R}^d$ contains the origin in its convex hull *if and only if* the corresponding point set on S^{d-1} is not contained in any open hemisphere.

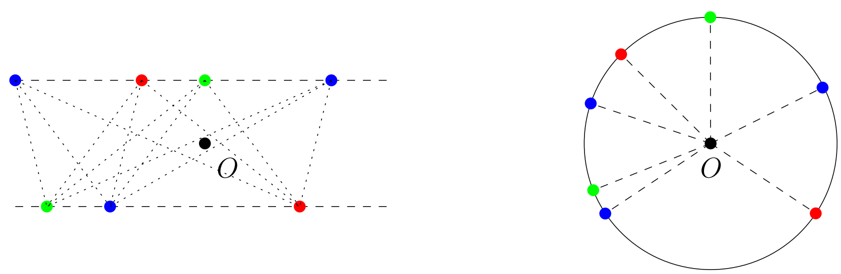

Figure 3: A point set in $\mathbb{R}^2$ in general position with respect to the origin (on the left) can be projected to a point set in general position on S^1 (on the right). The *triangles* in $\mathbb{R}^2$ that contain the origin correspond to *triples* of points on S^1 that are not contained in an open halfcircle.

It should now be clear that to prove Theorem 1, it is enough to prove it for point sets that are in general position on S^{d-1}. Here is the main lemma that we need.

Lemma 2. *For $d \geq 2$, let $A_1, \dots, A_d$ be disjoint, non-empty, finite point sets in S^{d-1} such that $A_1 \cup \dots \cup A_d$ is in general position, and suppose $A_i \cup A_j$ is not contained in any open hemisphere, for any $1 \leq i < j \leq d$. Let U denote the collection of all $(d-1)$-simplices spanned by the d-tuples consisting of a single point from each A_i. If U does not cover S^{d-1}, then for some $1 \leq i \leq d$ there exists an open hemisphere H such that $A_i \subset H \subset U$.*

Before proving Lemma 2, let us consider some cases for small d. For $d = 2$ this is a statement concerning red and blue points on the circle. This case is quite simple and in fact it was implicitly proved above. So let us consider the case $d = 3$.

Here we have red, blue, and green points in general position on S^2. The set U now consists of *spherical triangles* with a vertex of each color. Clearly the union of these triangles is a closed set, and if U does not cover S^2, then U must have a boundary. The boundary of U consists of pieces of arcs (spherical 1-simplices) connecting points of distinct colors.

Let us consider a point p on the boundary of U which lies in the relative interior of a unique arc with endpoints of distinct colors. Clearly such a point must exist. (Excercise 2). Let us suppose that r and b are the endpoints of the arc containg p, where r is red and b is blue. The points r and b are contained in a unique *great circle* H which bounds two opposite open hemispheres H^+ and H^-. We claim that it is impossible that there are green points in both H^+ and H^-. If there are green points $g^+ \in H^+$ and $g^- \in H^-$, then the spherical triangles rbg^+ and rbg^- lie in opposite closed hemispheres, share the common arc rb, and they both belong to U. But clearly this cannot happen since p is a boundary point. So we may assume that the green points are contained in H^+.

Next we show that there must at least one red point and one blue point in H^-. Suppose there were no red points in H^-. Since there already exists a red and a blue point on H this would mean that the remaining red points (if there are any!) must lie in H^+. But then the red and the green points are contained in $H^+ \cup H$, and by the general position assumption we can find an open hemisphere that contains all the red and green points. The same argument applies for the blue points. So let r' be a red point and b' a blue point contained in H^-, and fix a green point g. Since $g \in H^+$ and $r', b' \in H^-$ it follows that the arcs gr' and gb' must cross the great circle H, and we define the points $r^* = gr' \cap H$ and $b^* = gb' \cap H$. It follows from the definition of the points r^* and b^* that the arc rb^* is contained in the triangle grb'. Similarly, the arc r^*b is contained in $gr'b$, and the arc r^*b^* is contained in $gr'b'$.

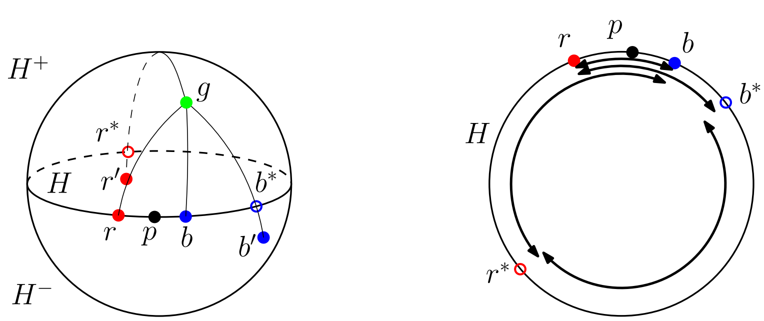

Figure 4: On the left: The arcs gr' and gb' intersect the great circle H, which gives us the points r^* and b^*. On the right: The arcs rb, rb^*, r^*b, r^*b^* cover H.

Now we make a crucial observation. We have defined four arcs on $H \cong S^1$:

$$rb,\ rb^*,\ r^*b,\ r^*b^*,$$

and these arcs form the edges of a cycle (S^1). Among these four arcs, p is contained *only* in the arc rb, which follows from the fact that p is a boundary point. This implies that the arcs rb, rb^*, r^*b, r^*b^* must cover H. (Excercise 3). Therefore the triangles $grb, grb^*, gr^*b, gr^*b^*$ will cover H^+, and we are done.

The proof of Lemma 2 for $d > 2$ is a straightforward generalization of the proof given above. The key is the crucial observation in the end, and it so happens that this observation has a natural extension which is stated as Claim 3, below.

For $d \geq 2$, let K be a finite collection of $(d-1)$-simplices on S^{d-1}. A point $p \in S^{d-1}$ will be called *generic with respect to K* if and only if p is not contained in any of the faces of the simplices of K, of dimension less than $d-1$. In other words, p is generic with respect to K if and only if for each simplex of $S \in K$, p is either in the relative interior of S, or disjoint from S. (We may omit 'with respect to K' when it is clear from the situation what K is). For a generic point p, let the *order* of p denote the number of $(d-1)$-simplices of K which contain p in their relative interiors.

Claim 3. *For $k \geq 2$, let $B = \{a_1, \ldots, a_k, b_1, \ldots, b_k\}$ be distinct points in general position on S^{k-1}. Let K denote the collection of the $(k-1)$-simplices formed by k-tuples of B with no repeated indices. Either the order of every generic point is even, or the order of every generic point is odd.*

Proof. One way to prove this follows by noticing that K is a continuous image of the boundary of the k-dimensional cross-polytope, X^k, and that $\partial X^k \cong S^{k-1}$. Thus K is defined by a continuous map $f : S^{k-1} \to S^{k-1}$, and the claim follows by considering the (Brouwer) degree of f. But here we also give a direct argument.

Let $K^{<(k-1)}$ be the union of faces of the simplices of K, of dimension less than $k-1$. Then $S^{k-1} \setminus K^{<(k-1)}$ is a collection of finitely many open cells, and any pair of generic points can be connected by a path in S^{k-1} that does not pass through any faces of dimension less than $k-2$. Thus it suffices to consider how the order changes as we pass through a face of dimension $k-2$. For any face F of dimension $k-2$, there are preciesly two points a_i and b_i (for some $1 \leq i \leq k$) such that $\text{conv}_S(F \cup a_i)$ and $\text{conv}_S(F \cup b_i)$ are $(k-1)$-simplices of K. Let H be the unique great $(k-2)$-sphere that contains F. If a_i and b_i are contained in the same open hemisphere bounded by H, then the order changes by ± 2 as we pass through F. If a_i and b_i are contained in opposite open hemispheres bounded by H, then the order stays the same as we pass through F. □

Proof of Lemma 2. Suppose U does not cover S^{d-1}. Since U is the union of finitely many simplices, U is closed and has a boundary, which is a subset of finitely many $(d-2)$-faces of simplices of K. Let p be a point of the boundary of U with the property that it is contained in the relative interior of a unique $(d-2)$-face. Clearly such a point must exist, so suppose p is contained in the relative interior of the unique $(d-2)$-face, $F = \text{conv}_S \{a_1, \ldots, a_{d-1}\}$, where $a_i \in A_i$.

There is a unique great $(d-2)$-sphere, H, which contains the points $a_1, \ldots, a_{d-1}$, which bounds disjoint open hemispheres H^+ and H^-. If there exists points $x^+ \in A_d \cap H^+$ and $x^- \in A_d \cap H^-$, then p belongs to the $(d-1)$-simplices spanned by $x^+ \cup F$ and $x^- \cup F$, which have disjoint relative interiors, share the common face F, and belong to U. This is impossible since p is a boundary point of U, so we may assume that $A_d \subset H^-$.

For every $1 \le i \le d-1$, we must have $A_i \cap H^+ \neq \emptyset$. If not, there exists an A_i such that $A_i \cup A_d \subset H \cup H^-$, which, by the general position assumption, means that $A_i \cup A_d$ is contained in some open hemisphere. Pick points $a \in A_d \subset H^-$ and $p_i \in A_i \cap H^+$, and let $b_i = H \cap \mathrm{conv}_S\{p_i, a\}$. It follows from the general position assumption that the set of points $J = \{a_1, \dots, a_{d-1}, b_1, \dots, b_{d-1}\}$ is in general position on H. Let K denote the set of $(d-2)$-simplices spanned by the $(d-1)$-tuples of J with no repeated indices. By our choice of p, it follows that p is a generic point in H with respect to K.

Let $T \neq F$ be a $(d-2)$-simplex of K. It follows from how we defined the points of J, that any point in the relative interior of T is contained in the relative interior of a $(d-1)$-simplex spanned by U, for instance,

$$x \in \mathrm{int}\,\mathrm{conv}_S\{a_1, a_2, b_3, b_4, \dots, b_{d-1}\} \subset \mathrm{int}\,\mathrm{conv}_S\{a, a_1, a_2, p_3, p_4, \dots, p_{d-1}\} \subset U.$$

This means that p is covered only once (in H) by the $(d-2)$-simplices of K. So by Claim 3, with $k = d-1$, the simplices of K must cover H, which implies that $H \subset U$. Therefore

$$H^- \subset \bigcup_{X \in K} \mathrm{conv}_S(a \cup X)$$

□

Most of the work is done now that we have established Lemma 2, and it will be quite easy to prove Theorem 1. Just remember that we are proving the *spherical* version of Theorem 1, so rather dealing with sets $A_1, \dots, A_{d+1}$ in $\mathbb{R}^d$ in general position with respect to the origin, we assume that the sets are in general position on S^{d-1}.

Proof of Theorem 1. The sets $A_1, \dots, A_d$ satisfy the conditions of Lemma 2, and they define the set U. If there exists a point $a \in A_{d+1}$ such that $-a \cap U \neq \emptyset$, then $-a$ is contained in some $(d-1)$-simplex, S of U, which means that the simplex spanned by $a \cup S$ contains the origin. On the other hand, if $-A_{d+1} \cap U = \emptyset$ then U cannot cover S^{d-1} so by Lemma 2 there is some $1 \le i \le d$ and an open hemisphere H such that $A_i \cup A_{d+1} \subset H \subset U$, which is a contradiction. □