8

Intersection Patterns of Convex Sets

In Chapter 1 we covered three simple but basic theorems in the theory of convexity: Helly's, Radon's, and Carathéodory's. For each of them we present one closely related but more difficult theorem in the current chapter. These more advanced relatives are selected, among the vast number of variations on the Helly–Radon–Carathéodory theme, because of their wide applicability and also because of nice techniques and tricks appearing in their proofs.

The development started in this chapter continues in Chapters 9 and 10. One of the culminations of this route is the (p, q)-theorem of Alon and Kleitman, which we will prove in Section 10.5. The proof ingeniously combines many of the tools covered in these three chapters and illustrates their power.

Readers who do not like higher dimensions may want to consider dimensions 2 and 3 only. Even with this restriction, the results are still interesting and nontrivial.

8.1 The Fractional Helly Theorem

Helly's theorem says that if every at most $d+1$ sets of a finite family of convex sets in $\mathbf{R}^d$ intersect, then all the sets of the family intersect. What if not necessarily all, but a large fraction of $(d+1)$-tuples of sets, intersect? The following theorem states that then a large fraction of the sets must have a point in common.

8.1.1 Theorem (Fractional Helly theorem). *For every dimension $d \geq 1$ and every $\alpha > 0$ there exists a $\beta = \beta(d, \alpha) > 0$ with the following property. Let $F_1, \ldots, F_n$ be convex sets in $\mathbf{R}^d$, $n \geq d+1$, and suppose that for at least $\alpha\binom{n}{d+1}$ of the $(d+1)$-point index sets $I \subseteq \{1, 2, \ldots, n\}$, we have $\bigcap_{i \in I} F_i \neq \emptyset$. Then there exists a point contained in at least βn sets among the F_i.*

Although simple, this is a key result, and many of the subsequent developments rely on it.

The best possible value of β is $\beta = 1-(1-\alpha)^{1/(d+1)}$. We prove the weaker estimate $\beta \geq \frac{\alpha}{d+1}$.

Proof. For a subset $I \subseteq \{1,2,\ldots,n\}$, let us write F_I for the intersection $\bigcap_{i\in I} F_i$.

First we observe that it is enough to prove the theorem for the F_i closed and bounded (and even convex polytopes). Indeed, given some arbitrary $F_1,\ldots,F_n$, we choose a point $p_I \in F_I$ for every $(d+1)$-tuple I with $F_I \neq \emptyset$ and we define $F_i' = \mathrm{conv}\{p_I\colon F_I \neq \emptyset, i \in I\}$, which is a polytope contained in F_i. If the theorem holds for these F_i', then it also holds for the original F_i. In the rest of the proof we thus assume that the F_i, and hence also all the nonempty F_I, are compact.

Let $\leq_{\mathrm{lex}}$denote the *lexicographic ordering* of the points of $\mathbf{R}^d$ by their coordinate vectors. It is easy to show that any compact subset of $\mathbf{R}^d$ has a unique lexicographically minimum point (Exercise 1). We need the following consequence of Helly's theorem.

8.1.2 Lemma. *Let $I \subseteq \{1,2,\ldots,n\}$ be an index set with $F_I \neq \emptyset$, and let v be the (unique) lexicographically minimum point of F_I. Then there exists an at most d-element subset $J \subseteq I$ such that v is the lexicographically minimum point of F_J as well.*

In other words, the minimum of the intersection F_I is always enforced by some at most d "constraints" F_i, as is illustrated in the following drawing (note that the two constraints determining the minimum are not determined uniquely in the picture):

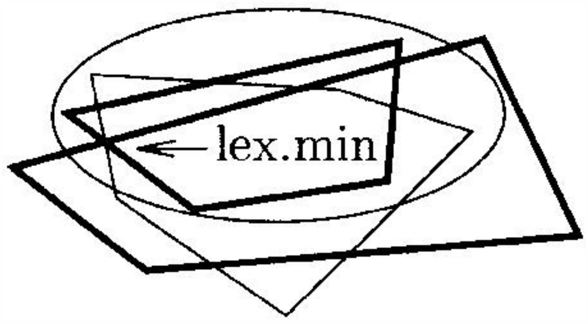

Proof. Let $C = \{x \in \mathbf{R}^d\colon x <_{lex} v\}$. It is easy to check that C is convex. Since v is the lexicographic minimum of F_I, we have $C \cap F_I = \emptyset$. So the family of convex sets consisting of C plus the sets F_i with $i \in I$ has an empty intersection. By Helly's theorem there are at most $d+1$ sets in this family whose intersection is empty as well. The set C must be one of them, since all the others contain v. The remaining at most d sets yield the desired index set J. □

Let us remark that instead of taking the lexicographically minimum point, one can consider a point minimizing a generic linear function. That formulation is perhaps more intuitive, but it appears slightly more complicated for rigorous presentation.

We can now finish the proof of the fractional Helly theorem. For each of the $\alpha\binom{n}{d+1}$ index sets I of cardinality $d+1$ with $F_I \neq \emptyset$, we fix a d-element set $J = J(I) \subset I$ such that F_J has the same lexicographic minimum as F_I.

The theorem follows by double counting. Since the number of distinct d-tuples J is at most $\binom{n}{d}$, one of them, call it J_0, appears as $J(I)$ for at least $\alpha\binom{n}{d+1}/\binom{n}{d} = \alpha\frac{n-d}{d+1}$ distinct I. Each such I has the form $J_0 \cup \{i\}$ for some $i \in \{1, 2, \ldots, n\}$. The lexicographic minimum of F_{J_0} is contained in at least $d + \alpha\,\frac{n-d}{d+1} > \alpha\,\frac{n}{d+1}$ sets among the F_i. Hence we may set $\beta = \frac{\alpha}{d+1}$. □

Bibliography and remarks. The fractional Helly theorem is due to Katchalski and Liu [KL79]. The quantitatively sharp version with $\beta = 1-(1-\alpha)^{1/(d+1)}$ was proved by Kalai [Kal84] (and the main result needed for it was proved independently by Eckhoff [Eck85], too). Actually, there is an exact result: If the maximum size of an intersecting subfamily in a family of n convex sets in $\mathbf{R}^d$ is m, then the smallest possible number of intersecting $(d+1)$-tuples is attained for the family consisting of $n - m + d$ hyperplanes in general position and $m - d$ copies of $\mathbf{R}^d$. But there are many other essentially different examples attaining the same bound.

These assertions are consequences of considerably more general results about the possible intersection patterns of convex sets in $\mathbf{R}^d$. For explaining some of them it is convenient to use the language of simplicial complexes. Let $\mathcal{F} = \{F_1, F_2, \ldots, F_n\}$ be a family of convex sets in $\mathbf{R}^d$. The *nerve* $\mathcal{N}(\mathcal{F})$ of $\mathcal{F}$ is the simplicial complex with vertex set $\{1, 2, \ldots, n\}$ whose simplices are all $I \subseteq \{1, 2, \ldots, n\}$ such that $\bigcap_{i \in I} F_i \neq \emptyset$. A simplicial complex obtainable as $\mathcal{N}(\mathcal{F})$ for some family of convex sets in $\mathbf{R}^d$ is called *d-representable*. A characterization of d-representable simplicial complexes for a given d is most likely out of reach. There are several useful *necessary* conditions for d-representability. One certainly worth mentioning is *d-collapsibility*, which means that a given simplicial complex $\mathcal{K}$ can be reduced to the void complex by a sequence of *elementary d-collapsings*, where an elementary d-collapsing consists in deleting a face $S \in \mathcal{K}$ of dimension at most $d-1$ that lies in a unique maximal face of $\mathcal{K}$ and all the faces of $\mathcal{K}$ containing S. The proof of the d-collapsibility of every d-representable complex (Wegner [Weg75]) uses an idea quite similar to the proof of the fractional Helly theorem.

While no characterization of d-representable complexes is known, the possible *f-vectors* of such complexes (where f_i is the number of i-dimensional simplices, which correspond to $(i+1)$-wise intersections here) are fully characterized by a conjecture of Eckhoff, which was proved by Kalai [Kal84], [Kal86] by an impressive combination of several methods. The same characterization applies to d-collapsible complexes as well (and even to the more general *d-Leray complexes*; these