III.2 Convex Set Systems

A convenient way to construct large simplicial complexes is through specifying sets and recording their intersection patterns.

Nerves. Let F be a finite collection of sets. The *nerve* consists of all subcollections whose sets have a non-empty common intersection,

$$\mathrm{Nrv}\, F \quad = \quad \{X \subseteq F \mid \bigcap X \neq \emptyset\}.$$

The nerve is an abstract simplicial complex since $\bigcap X \neq \emptyset$ and $Y \subseteq X$ implies $\bigcap Y \neq \emptyset$. We can geometrically realize it in Euclidean space, so it makes sense to talk about the topology or the homotopy type of the nerve. We will sometimes do this without explicit construction of the geometric realization. As an example consider a collection of touching but otherwise not overlapping closed disks, as in Figure III.5. We have pairwise but no triplewise intersections. The nerve is therefore an abstract graph and it can be geometrically realized by connecting the centers of the kissing disks.

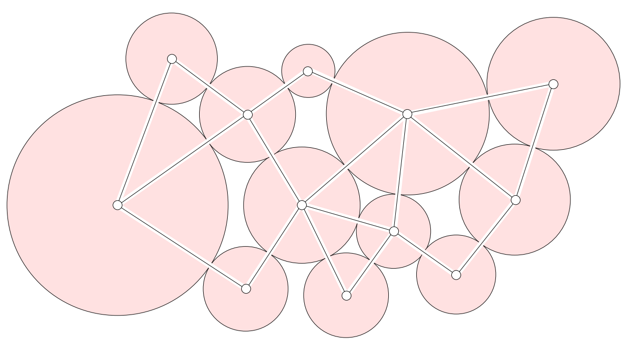

Figure III.5: A collection of twelve disks and its nerve drawn as a straight-line graph.

If the sets in the collection are convex then the nerve and the union have the same homotopy type.

NERVE THEOREM. Let F be a finite collection of closed, convex sets in Euclidean space. Then $\mathrm{Nrv}\, F \simeq \bigcup F$.

The requirement on the sets can be relaxed without sacrificing the conclusion. Specifically, if $\bigcup F$ is triangulable, all sets in F are closed, and all non-empty common intersections are contractible then $\mathrm{Nrv}\, F \simeq \bigcup F$.

Sets with common points. If the convex sets are in d-dimensional Euclidean space then their intersection patterns are restricted. For example, if three closed intervals intersect in pairs then they intersect as a triple.

Helly's Theorem. Let F be a finite collection of closed, convex sets in $\mathbb{R}^d$. Every $d+1$ of the sets have a non-empty common intersection iff $\bigcap F \neq \emptyset$.

Proof. We prove only the non-obvious direction, by induction over the dimension, d, and the number of sets, n. The implication is clearly true for $n = d+1$ and for $d = 1$. Now suppose we have a minimal counterexample consisting of $n > d+1$ closed, convex sets in $\mathbb{R}^d$, which we denote as $X_1, X_2, \ldots, X_n$. By minimality of the counterexample, the set $Y_n = \bigcap_{i=1}^{n-1} X_i$ is non-empty and disjoint from X_n. Because Y_n and X_n are both closed and convex we can find a $(d-1)$-dimensional plane h that separates and is disjoint from both sets, as in Figure III.6. Let F' be the collection of sets $Z_i = X_i \cap h$, for $1 \leq i \leq n-1$,

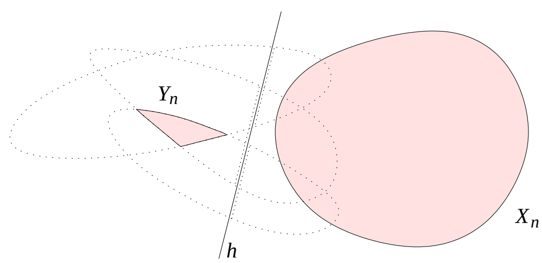

Figure III.6: The $(d-1)$-plane separates the n-th set from the common intersection of the first $n-1$ sets in F.

each a non-empty, closed, convex set in $\mathbb{R}^{d-1}$. By assumption, any d of the first $n-1$ sets X_i have a common intersection with X_n. It follows that the common intersection of the d sets contains points on both sides of h implying that any d of the sets Z_i have a non-empty common intersection. By minimality of the counterexample, this implies $\bigcap F' \neq \emptyset$. This intersection is

$$\bigcap F' = \bigcap_{i=1}^{n-1} (X_i \cap h) = Y_n \cap h.$$

But this contradicts the choice of h as a $(d-1)$-plane disjoint from Y_n. □

Similar to the Nerve Lemma, convexity is a convenient but unnecessarily strong requirement. The conclusion in Helly's Theorem still holds if the sets in F are closed and all their non-empty common intersections are contractible.

Čech complexes. A *unit ball* in d dimensions is a set $x + \mathbb{B}^d$; its center is the point $x \in \mathbb{R}^d$ and its radius is 1. Let F be a finite collection of unit balls in $\mathbb{R}^d$. Clearly, the balls have a non-empty intersection iff their centers lie inside a common unit ball. Indeed $y \in \bigcap F$ iff $\|x - y\| \leq 1$ for all centers x. An easy consequence of Helly's Theorem is therefore the following, historically earlier result.

JUNG'S THEOREM. Let S be a finite set of point in $\mathbb{R}^d$. Every $d + 1$ of the points lie in a common unit ball iff all points in S lie in a common unit ball.

For a non-negative radius r and for each point $x \in S$ we consider the ball $B_x(r) = x + r\mathbb{B}^d$. The *Čech complex* of S and r is isomorphic to the nerve of the set of this collection of balls,

$$\check{\text{C}}\text{ech}(r) = \{\sigma \subseteq S \mid \bigcap_{x \in \sigma} B_x(r) \neq \emptyset\}.$$

This complex does not necessarily have a geometric realization in $\mathbb{R}^d$ but it is fine as an abstract simplicial complex; see Figure III.7. For larger radius the

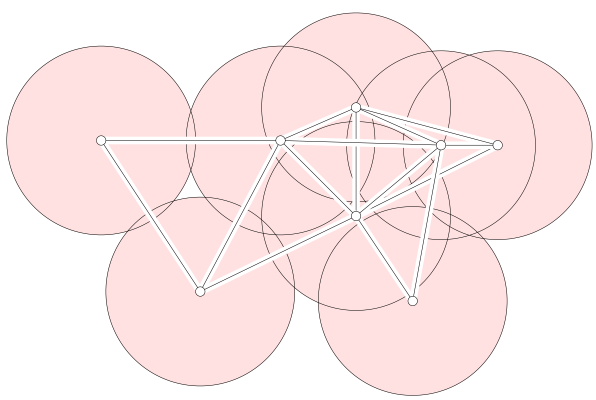

Figure III.7: Nine points with pairwise intersections among the disks indicated by straight edges connecting their centers. The Čech complex fills nine of the ten possible triangles as well as the two tetrahedra. The only difference between the Rips and the Čech complexes is the tenth triangle, which belongs only to the former.

disks are bigger and create more overlaps while retaining the ones for smaller radius. Hence $\check{\text{C}}\text{ech}(r_0) \subseteq \check{\text{C}}\text{ech}(r)$ whenever $r_0 \leq r$. If we continuously increase the radius, from 0 to ∞, we get a discrete family of nested Čech complexes. We will come back to this construction later.

Smallest enclosing balls. Let $\sigma \subseteq S$ be a subset of the given points. Deciding whether or not σ belongs to Čech(r) is equivalent to deciding whether or not σ fits inside a ball of radius r. Let the *miniball* of σ be the smallest closed ball that contains σ, which we note is unique. The radius of the miniball is smaller than or equal to r iff $\sigma \in$ Čech(r), so finding it solves our problem. Observe that the miniball is already determined by a subset of $k+1 \leq d+1$ of the points, which all lie on its boundary. If we know this subset then we can verify the miniball by testing that it indeed contains all the other points. In a situation in which we have many more points than dimensions the chance that a point belongs to this subset is small and discarding it is easy. This is the strategy of the miniball algorithm. It takes two disjoint subsets τ and υ of σ and returns the miniball that encloses τ and contains all points of υ on its boundary. To get the miniball of σ we call MINIBALL$(\sigma, \emptyset)$.

```
ball MiniBall(τ, υ)
  if τ = ∅ then return Ball(υ)
           else choose a random point u ∈ τ;
                B = MiniBall(τ − {u}, υ);
                if u ∉ B then
                   B = MiniBall(τ − {u}, υ ∪ {u})
                endif
  endif; return B.
```

When τ is empty we have a set υ of at most $d+1$ points, which we know all lie on the boundary. We can therefore compute the miniball directly, using the function BALL. To analyze the running time, we assume the dimension d is a constant and ask how often we execute the test "$u \notin B$". Let $t_j(n)$ be the expected number of such tests if we call MINIBALL for $n = \operatorname{card} \tau$ and $d+1-j = \operatorname{card} \upsilon$. Obviously, $t_j(0) = 0$, and it is reassuring that the constant amount of work needed to compute the ball for the at most $d+1$ points in υ is payed for by the test that initiated the call. Consider $n > 0$. We have one call with parameters $n-1$ and j, one test "$u \in B$", and one call with parameters $n-1$ and $j-1$. The probability that the second call indeed happens is only $\frac{j}{n}$. Hence

$$t_j(n) \quad \leq \quad t_j(n-1) + 1 + \frac{j}{n} t_{j-1}(n-1).$$

It follows that $t_0(n) \leq n$ and $t_1(n) \leq 2n$. More generally we have $t_j(n)$ bounded from above by some constant times n, where the constant is less than $(j+1)!$. In summary, for constant dimension the algorithm takes expected constant time per point.

Rips complexes. Instead of checking all subcollections for non-empty common intersections, we may just check pairs and add 2- and higher-dimensional simplices whenever we can. The *Rips complex* of S and r consists of all subsets of diameter at most $2r$,

$$\text{Rips}(r) \quad = \quad \{\sigma \subseteq S \mid \operatorname{diam} \sigma \leq 2r\}.$$

Clearly, the edges in the Rips complex are the same as in the Čech complex. Furthermore $\text{Čech}(r) \subseteq \text{Rips}(r)$ because the latter contains every simplex warranted by the given edges. We now prove $\text{Rips}(r) \subseteq \text{Čech}(\sqrt{2}r)$.

A simplex is *regular* if all its edges have the same length. A convenient representation for dimension d is the *standard d-simplex*, Δ^d, spanned by the endpoints of the unit coordinate vectors in $\mathbb{R}^{d+1}$; see Figure III.8. Each edge of

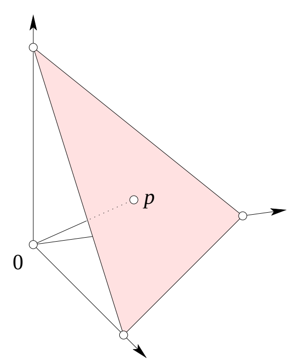

Figure III.8: The standard triangle.

Δ^d has length $\sqrt{2}$. By symmetry, the distance of the origin from the standard simplex is its distance from the barycenter, the point p whose $d+1$ coordinates are all equal to $\frac{1}{d+1}$. That distance is therefore $\|p\| = 1/\sqrt{d+1}$. The barycenter is also the center of the smallest d-sphere that passes through the vertices of Δ^d. Writing r_d for the radius of that sphere, we have $r_d^2 = 1 - \|p\|^2 = \frac{d}{d+1}$. For dimension 1 this is indeed half the length of the interval, and for dimension 2 it is the radius of the equilateral triangle. As the dimension goes to infinity, the radius grows and approaches 1 from below. Using induction on the dimension, it is not difficult to prove that every d-simplex of diameter at most $\sqrt{2}$ fits inside a ball of radius r_d. The d-simplex thus belongs to the Čech complex for radius r_d, implying $\text{Rips}(\sqrt{2}/2) \subseteq \text{Čech}(r_d)$ in $\mathbb{R}^d$. More generally, $\text{Rips}(r/\sqrt{2}) \subseteq \text{Čech}(rr_d) \subseteq \text{Čech}(r)$ since $r_d \leq 1$ for all d. This implies the claimed relationship between Rips and Čech complexes.

Bibliographic notes. The concept of nerve has been introduced in the early years of combinatorial topology [1]. The Nerve Theorem goes back to Leray [4]. It has a complicated literature, with version differing in the generality of the assumption and the strength of the conclusion. Helly's Theorem is even older, being proved first for convex sets and then from sets with contractible common intersections [2, 3]. The Čech complexes are inspired by the theory of Čech homology, from which they borrow their name. Algorithms for finding the smallest ball enclosing a finite set of points have been studied in computational geometry, culminating in the minidisk algorithm of Welzl which has versions that are efficient even for large sets in high dimensions [6]. The Rips complex appears in Vietoris [5] but receives its name from later work by Rips.

[1] P. S. ALEXANDROV. Über den allgemeinen Dimensionsbegriff und seine Beziehungen zur elementaren geometrischen Anschauung. *Math. Ann.* **98** (1928), 617–635.

[2] E. HELLY. Über Mengen konvexer Körper mit gemeinschaftlichen Punkten. *Jahresber. Deutsch. Math.-Verein.* **32** (1923), 175–176.

[3] E. HELLY. Über Systeme von abgeschlossenen Mengen mit gemeinschaftlichen Punkten. *Monatsh. Math. Physik* **37** (1930), 281–302.

[4] J. LERAY. Sur la forme des espaces topologiques et sur les points fixes des représentations. *J. Math. Pure Appl.* **24** (1945), 95–167.

[5] L. VIETORIS. Über den höheren Zusammenhang kompakter Räume and eine Klasse von zusammenhangstreuen Abbildungen. *Math. Ann.* **97** (1927), 454–472.

[6] E. WELZL. Smallest enclosing disks (balls and ellipsoids). In *New Results and New Trends in Computer Science*, H. A. Maurer (ed.), Springer-Verlag, Lecture Notes in Computer Science **555** (1991), 359–370.