Simplicial Complexes: Second Lecture

4 Nov, 2010

1 Overview

Today we have two main goals:

- Prove that every continuous map between triangulable spaces can be approximated by a simplicial map. To do this, we will introduce the idea of barycentric subdivision.
- Discuss various ways to triangulate a point cloud.

2 Simplicial Approximations

Suppose that K and L are simplicial complexes. Recall that a *vertex map* between these complexes is a function $\phi : Vert(K) \rightarrow Vert(L)$ such that the vertices of a simplex in K map to the vertices of a simplex in L. Given such a ϕ, we can create a *simplicial map* $f : |K| \rightarrow |L|$ by linearly extending ϕ over each simplex.

On the other hand, suppose we have an arbitrary continuous map $g : |K| \rightarrow |L|$. There is no reason to assume that g would be simplicial. On the other hand, we can hope to approximate g by a function f which is itself simplicial and is not "too far away" from f in some sense. That's the goal today, and we start by defining it rigorously. A simplicial map $f : |K| \rightarrow |L|$ is a *simplicial approximation* of g if, for every vertex $u \in K$, $g(St_K(u)) \subseteq St_L(f(u))$; in other words, if g maps points "near" v to points "near" $f(v)$, where points are considered "near" if they live in a common simplex. If the simplices in K are reasonably small, it seems likely that we can do this. Our goal now is to make this happen by repeatedly subdividing K.

2.1 Barycentric Subdivision

Another simplicial complex K' is a *subdivision* of K if $|K'| = |K|$ and every simplex in K is the union of simplices in K'.

One way to subdivide K is to "star" from an arbitrary point $x \in |K|$, a procedure which we now describe:

- Find the simplex $\sigma \in K$ such that $x \in int(\sigma)$.
- Remove the star of σ.
- Cone the point x over the boundary of the closed star of σ.

We obtain $sd(K)$, the *barycentric subdivision* of K by starring from the barycenter of each simplex in K, starting from the top-dimensional simplices and ending with the edges.

We can of course repeat this as many times as we like. Let $sd^j(K) = sd(sd^{j-1}(K))$ denote the jth barycentric subdivision of K. Intuitively, repeated subdivision should make the resulting simplices very small. We define the *mesh* of a simplicial complex to be the largest diameter of any simplex; in this case, this is of course just the length of the longest edge.

MESH LEMMA: Let K be a d-dimensional simplicial complex. Then $Mesh(sd(K)) \leq \frac{d}{d+1} Mesh(K)$.

2.2 The Simplicial Approximation Theorem

We again let $g : |K| \to |L|$ be a continuous but not necessarily simplicial map. We say that g satisfies the *star condition* if, for every vertex $u \in K$, there exists some vertex $v \in L$ such that $g(St_K(u)) \subseteq St_L(v)$.

If g satisfies the star condition, then it has a simplicial approximation, as we now show. First we construct a map $\phi : Vert(K) \to Vert(L)$ by mapping each vertex $u \in K$ to some vertex $v = \phi(u) \in L$ which satisfies the condition above (if there's more than one, we pick one). We claim that ϕ is in fact a vertex map. To see this, let $u_0, u_1, \ldots, u_k$ be the vertices of a simplex $\sigma \in K$ and choose some point $x \in int(\sigma)$. Then $x \in \bigcap_i st(u_i)$ and hence $g(x) \in \bigcap_i g(st_K(u_i)) \subseteq \bigcap_i st_L(\phi(u_i))$. Hence the stars of $\phi(u_0), \ldots, \phi(u_k)$ have nonempty mutual intersection, and thus these vertices span a simplex in L, as required. Letting f be the induced simplicial map, we see immediately that f is a simplicial approximation of g.

We are now ready to prove the big theorem for this lecture:

SIMPL. APPROX. THEOREM: Let K and L be simplicial complexes. If $g : |K| \to |L|$ is a continuous function, then there is a sufficiently large integer j such that g has a simplicial approximation $f : |sd^j(K)| \to |L|$.

proof: We cover $|K|$ by the open sets $g^{-1}(st_L(v)$, over all vertices $v \in L$. Since $|K|$ is compact, there exists a small positive number λ such that every set of diameter less than λ is contained entirely within one of these open sets (this is intuitively obvious and is formally called the Lebesgue Number Lemma). Appealing to the Mesh Lemma, we now choose j big enough that every simplex in $sd^j(K)$ has diameter less than $\frac{\lambda}{2}$, and consider the map $g : |sd^j(K)| \to |L|$. We choose an arbitrary vertex $u \in sd^j(K)$ and note that the set $st_{sd^j(K)}(u)$ must have diameter less than λ, and thus must lie entirely within one of the open sets $g^{-1}(st_L(v))$. In other words, g satisfies the star condition, and thus, by the construction above, has a simplicial approximation. $\square$

We close the lecture by noting an important fact: if f is a simplicial approximation of a map g, then f must also be homotopic to g.

3 In-Class Exercises

A Let K and L be the following two 1-dimensional simplicial complexes geometrically realized in $\mathbb{R}^2$. K has vertices $a_0 = (0,0)$ and $a_1 = (1,0)$, along with the edge (a_0, a_1), while L has vertices $b_0 = (0,0)$, $b_1 = (0, 0.5)$, and $b_2 = (0,1)$, along with the edges between (b_0, b_1) and (b_1, b_2). Define $g : |K| \to |L|$ by the formula $g(x, 0) = (0, x^2)$.

(a) Show that g does *not* satisfy the star condition.

(b) Find a large enough j such that $g : |sd^j(K)| \to |L|$ satisfies the star condition. Then find a simplicial approximation for this map.

B Suppose K, L, M are simplicial complexes. Suppose that $f_1 : |K| \to |L|$ is a simplicial approximation of $g_1 : |K| \to |L|$, and that $f_2 : |L| \to |M|$ is a simplicial approximation of $g_2 : |L| \to |M|$. Prove that $f_2 \circ f_1$ is a simplicial approximation of $g_2 \circ g_1$.

4 Point Cloud Triangulations and the Nerve Lemma

We now discuss a variety of ways to triangulate a collection of points. In fact, we will construct, in several different ways, a nested family of simplicial complexes from a given point cloud; later these families will be very important in the computation of persistent homology.

4.1 The Nerve Lemma

Let F be a finite collection of sets. We define the *nerve* of F to be the abstract simplicial complex given by all subcollections of F whose member have non-empty common intersection:

$$Nrv(F) = \{X \subseteq F \mid \cap X \neq \varnothing\}.$$

Note that the nerve is indeed a simplicial complex since $\cap X \neq \varnothing$ and $Y \subseteq X$ implies $\cap Y \neq \varnothing$. If we need to, we can geometrically realize the nerve in some Euclidean space, but we often just reason about it abstractly.

The main use of the nerve is the following. Suppose we want to represent some topological space $\mathbb{X}$ in a combinatorial fashion. We cover X by some collection of sets F and then take the nerve. Hopefully this nerve will faithfully represent the original space, and in certain situations this is guaranteed. Recall that a space is called *contractible* if it is homotopically equivalent to a point.

NERVE LEMMA: Let F be a finite collection of closed sets such that every intersection between its members is either empty or contractible. Then $Nrv(F)$ has the same homotopy type as $\cup F$.

Note that if F consists of convex sets in Euclidean space, then the hypothesis of the Nerve Lemma will be satisfied.

4.2 Cech Complexes

One example of a nerve is the following. Let P be a finite set of points in $\mathbb{R}^d$ (or indeed any metric space). For each fixed $r \geq 0$ and each $x \in S$, we define $B_x(r)$ to be the closed ball of radius r centered at x. We then define the Cech complex of S and r to be the nerve of the collection of sets $B_x(r)$, as x ranges over S. Put another way,

$$Cech_S(r) = \{\sigma \subseteq S \mid \bigcap_{x\in\sigma} B_x(r) \neq \varnothing\}.$$

We note three facts about Cech complexes before moving on:

- Since the sets $B_x(r)$ are all convex, the Nerve Lemma applies, and hence $Cech_S(r)$ has the same homotopy type as the union of r-balls around the points in S. This latter set can be thought of as a "thickening" by r of the point set S. It will play an important role in later lectures, and we denote it by S_r.

- Given two radii $r < r'$, we obviously have the inclusion $Cech_S(r) \subseteq Cech_S(r')$. Hence, if we let r from 0 to ∞, we produce a nested family of simplicial complexes, starting with a set of n vertices and ending with an enormous n-simplex, where $n = |S|$.

- A set of vertices $\sigma \subseteq S$ forms a simplex in $Cech_S(r)$ iff the set can be enclosed within a ball of radius r (do you see why this is true?). Hence deciding membership in the Cech complex is equivalent to solving a standard problem in computational geometry. It's also not very nice in high dimensions.

- The Cech complex is massive at large r, both in size and in dimension. As we will see later, most of the information it provides is also redundant, in that we can come up with a much smaller complex with the same homotopy type if we are little smarter.

4.3 Vietoris-Rips Complexes

As stated above, it is often a little nasty to compute the Cech complex. A much more convenient object is the following. Given a point cloud S and a fixed number $r \geq 0$, we define the *Vietoris-Rips* complex of S and r to be:

$$Rips_S(r) = \{\sigma \subseteq S \mid B_x(r) \cap B_y(r) \neq \varnothing, \forall x, y \in \sigma\}.$$

In other words, $Rips_S(r)$ consists of all subsets of S whose diameter is no greater than $2r$. From the definitions, it is trivial to see that $Cech_S(r) \subseteq Rips_S(r)$. In fact, the two subcomplexes share the same 1-skeleton (vertices and edges) and the Rips complex is the largest simplicial complex that can be formed from the 1-skeleton of the Cech. In other words, the Rips complex will in general be even larger than the Cech. However, it's also clearly easier to compute, since we need only measure pairwise distance between points. Unfortunately, we have no guarantees that the Rips complex will give us the homotopy type of any particular space. Of course, for $r < r'$, we again have the inclusion $Rips_S(r) \subseteq Rips_S(r')$.

4.4 Delaunay Triangulation and Alpha-Shapes

The Cech and the Rips complex both suffer from a common problem: the number of simplices becomes massive, especially for large r. We now give a construction which drastically limits the number of simplices, as well as reducing the dimension to that of the ambient space for points in general position.

Given a finite point set $S \subseteq \mathbb{R}^d$, we define the *Voronoi cell* of a point $p \in S$ to be:

$$V_p = \{x \in \mathbb{R}^d \mid d(x,p) \leq d(x,q), \forall q \in S\}.$$

We note V_p is a convex polyhedron; indeed, it is the intersection of the half-spaces of points at least as close to p as to q, taken over all $q \in S$. Furthermore, any two such Voronoi cells are either disjoint or meet in a common portion of their boundary. The collection of all Voronoi cells is called the Voronoi diagram of S; we note that it covers the entire ambient space $\mathbb{R}^d$.

We then define the *Delaunay triangulation* of S to be (isomorphic to) the nerve of the collection of Voronoi cells; more precisely,

$$Del(S) = \{\sigma \subseteq S \mid \bigcap_{p \in \sigma} V_p \neq \varnothing\}.$$

We note that a set of vertices $\sigma \subseteq S$ forms a simplex in $Del(S)$ iff these vertices all lie on a common $(d-1)$-sphere in $\mathbb{R}^d$. Assuming general position, we do in fact get a simplicial complex.

Alpha Complexes We again let S be a finite set of points in $\mathbb{R}^d$ and fix some radius r. As seen above, the complex $Cech_S(r)$ has the same homotopy type as the union of r-balls S_r, but requires far too many simplices for large r. We now define a much smaller complex, $Alpha_S(r)$, which is geometrically realizable in $\mathbb{R}^d$, and gives the correct homotopy type.

First, for each $p \in S$, we intersect the r-ball around p with its Voronoi region, to form $R_p(r) = B_p(r) \cap V_p$. These sets are convex (why?) and their union still equals S_r. We then define the *Alpha complex* of S and r to be the nerve of the collection of these sets:

$$Alpha_S(r) = \{\sigma \subseteq S \mid \bigcap_{p \in \sigma} R_p(r) \neq \varnothing\}.$$

By the Nerve Lemma, $Alpha_S(r)$ has the same homotopy type as S_r. On the other hand, it is certainly much smaller, both in cardinality and dimension, than $Cech_S(r)$; for example, $Alpha_S(r)$ will always be a subcomplex of $Del(S)$, and thus can have dimension no larger than the ambient dimension. As usual, we have inclusions $Alpha_S(r) \subseteq Alpha_S(r')$, for all $r < r'$, and we note that $Alpha_S(\infty) = Del(S)$.

5 In-Class Exercises

A Given a point cloud S and a radius r, prove that $Cech_S(r) \subseteq Rips_S(2r)$.

B Let S be a set of three points in the plane which form an acute triangle (all angles below $\frac{\pi}{2}$), and let T be a set of three points in the plane which form an obtuse triangle (one angle above $\frac{\pi}{2}$).

(a) Draw the Voronoi diagrams and Delaunay triangulations of S and T.

(b) Draw the family of Alpha-complexes $Alpha_S(r)$ and $Alpha_T(r)$, for all radii r (note that there's only a finite number of radii at which these complexes change!). Which family contains a member that is homeomorphic to a circle?