Excision Made Easier

Rich Schwartz

November 9, 2021

The purpose of these notes is to give a swifter, lighter proof of the Excision Theorem. The proof follows what is in Hatcher's book up to a point (Steps 1 and 2) but then departs. What makes Hatcher's proof complicated is that he wants a canonically defined chain homotopy between the relevant objects. This requires a lot of algebraic sophistication. If we abandon the goal of getting a chain homotopy and just argue for a homology isomorphism, we can dispense with some of the algebraic complexity. In the last section of the notes I'll ruminate on the differences between the two proofs.

1 The Goal

Let X be a space. Let A^o be the interior of a subset $A \subset X$. Suppose we have sets $A, B \subset X$ such that such that $X = A^o \cup B^o$. The version of the Excision Theorem I will prove is

Theorem 1.1 *$H_n(A, A \cap B) \cong H_n(X, B)$ for all n.*

As usual, H_n denotes the nth singular homology group. This version is readily convertible to the other version treated in the book.

2 The First Isomorphism

Let $C_n(A+B) = C_n(A) \oplus C_n(B)$. In other words, $C_n(A+B)$ consists of those formal sums of singular simplices such that each image is either contained in A or in B. Note that the boundary operator ∂ maps $C_n(A+B)$ into $C_{n-1}(A+B)$ and respects chains in A and chains in B. This means that we can make the homology groups we need to make.

Define

$$C_n(A+B,B) = C_n(A+B)/C_n(B).$$

This group is just the set of formal sums of simplices which lie in A but not in B. Also define

$$C_n(A, A\cap B) = C_n(A)/C_n(A\cap B).$$

This group is also the set of formal sums of simplices which lie in A but not in B. Thus there is a canonical bijection

$$C_n(A+B,B) \leftrightarrow C_n(A, A\cap B).$$

This bijection commutes with ∂. Therefore, when we form the relative homology groups based on these chain complexes, we get

$$H_n(A+B,B) \cong H_n(A, A\cap B). \tag{1}$$

Really, these are the same sets.

There is an inclusion map

$$\iota : C_n(A+B,B) \to C_n(X,B). \tag{2}$$

This map commutes with ∂ and so induces a map of homology:

$$\iota^{\#} : H_n(A+B,B) \to H_n(X,B). \tag{3}$$

In view of Equation 1 it suffices to prove that $\iota^{\#}$ is an isomorphism. This is what we will do.

3 Barycentric Subdivision

3.1 Diameter

The diameter of a simplex σ is defined to be the maximal distance between a pair of points in σ. This is realized by a pair of vertices. To see this, suppose that $(v,w) \in \sigma$ maximize diameter. Take a huge sphere centered at v and start shrinking it. The first point it touches σ is at a vertex, by convexity. Hence $d(v,w) = d(v,w')$ where w' is the vertex. Here $d(\cdot)$ denotes distance. The same argument shows that $d(v,w') = d(v',w')$ where v' is the first vertex hit by a shrinking sphere centered at w'.

3.2 Geometric View

The *barycenter* of a simplex is the average of its vertices. The barycentric subdivision of a point is just that point. The barycentric subdivision of an interval is the union of the two intervals obtained by cutting it in half. In general, the barycentric subdivision of a simplex σ is define to be the cone over the barycentric subdivision of its boundary.

One nice property of barycentric subdivision is that it is natural under affine transformations. To get the barycentric subdivision of some simplex $A(\sigma)$ we can subdivide σ and then apply A. This gives the same answer as just subdividing $A(\sigma)$ directly. The reason for this invariance: The barycenter itself, being an average of the vertices, is affinely natural.

Lemma 3.1 *Let d be the diameter of the n-simplex σ. The maximum diameter of a simplex in the barycentric subdivision of σ is at most $\frac{n}{n+1} \times d$.*

Proof: Let σ' be a simplex in the subdivision which maximizes the diameter d'. Let d' be its diameter. This diameter is realized by a pair of vertices v, w. If one of these vertices is not the barycenter of σ then both vertices lie in a face of $\partial\sigma$. In this case, the result follows from induction on n. So, we may assume that v is the barycenter of σ. That is $d' = d(v, w)$.

Consider the line ℓ containing v and w. This line hits $\partial\sigma$ in some third point x. The points go w, v, x in order on ℓ. The diameter of σ is at least $d(w, x)$. So, it suffices to prove that

$$\frac{d(v,w)}{d(w,x)} \leq \frac{n}{n+1}.$$

But affine transformations preserve ratios of distances and so it suffices to establish this result when σ is the regular simplex. In this case an explicit calculation gives the result. ♠

Corollary 3.2 (Shrinking) *Let σ be any simplex and let $\epsilon > 0$ be given. Then there is some N such that all simplices in the Nth barycentric subdivision of σ all have diameter less than ϵ.*

Proof: The point is that we can make $(n/(n+1))^N$ as small as we like by choosing N large. ♠

3.3 Algebraic View

We fix some high dimensional Euclidean space whose dimension is larger than the dimensions of the groups we care about at any given time. Let L_n denote the set of formal sums of linear maps of the n-simplex into $\boldsymbol{R}^N$. Each linear simplex is described by its vertices.

Given any point $b \in \boldsymbol{R}^N$ and any linear simplex σ we let $b\sigma$ denote the linear simplex obtained by coning σ off to b. In other words, if $a_0, ..., a_n$ are the vertices of σ, then $b, a_0, ..., a_n$ are the vertices of $b\sigma$. We extend linearly to get a map from L_n into L_{n+1}.

Lemma 3.3 *$b\partial + \partial b = I$, the identity.*

Proof: It suffices to prove this for some simplex $\sigma = a_0, ..., a_n$. We have

$$b(\partial\sigma) = \sum_{i=0}^{n}(-1)^n ba_0, ...\widehat{a}_i...a_n.$$

$$\partial(b\sigma) = a_0...a_n + \sum_{i=0}^{n}(-1)^{i+1} ba_0...\widehat{a}_i...a_n.$$

Add these up and you just get $a_0...a_n$ ♠

It is convenient to set $L_{-1} = \boldsymbol{Z}$ and to let $\partial(\sum a_i\sigma_i) = \sum a_i$ be the augmentation map from L_0 to L_{-1}. The algebraic version of barycentric subdivision is a map $S : L_n \to L_n$ defined inductively. S is the identity on L_{-1} and L_0. In general, define

$$S(\lambda) = b_\lambda S(\partial\lambda). \tag{4}$$

Here b_λ is the barycenter of λ.

Lemma 3.4 *$S\partial = \partial S$.*

Proof: This works for $n = 0$ because S is the identity on L_{-1} and L_0. In general:

$$\partial(S\lambda) = \partial b_\lambda S(\partial\lambda) = (I - b_\lambda\partial)S(\partial\lambda) = S(\partial\lambda) + b_\lambda\partial S\partial\lambda.$$

Hence

$$\partial S\lambda + S\partial\lambda = b_\lambda\partial S\partial\lambda =^* b_\lambda(S\partial\partial\lambda) = 0.$$

The starred equality comes from the fact (by induction) that $S\partial = \partial S$ on L_{n-1}. ♠

3.4 The Homotopy

We are going to define a map $T : L_n \to L_{n+1}$. We define $T = 0$ on L_{-1}. We then inductively define

$$T(\lambda) = b_\lambda(\lambda - T(S(\lambda)). \tag{5}$$

To make the next proof work we formally define $\partial = 0$ on L_{-1}. This next lemma is in Hatcher's book as well. Actually, I think that Hatcher's derivation is a bit shorter.

Lemma 3.5 $T\partial + \partial T = I - S$.

Proof: For $n = -1$ we have $T = 0$ and $S = I$. So, this works for $n = -1$. In general let λ be a linear n-simplex and let $b = b_\lambda$.

$$\partial T\lambda = \partial b(\lambda - T(\partial\lambda)) =^0$$
$$(I - b\partial)(\lambda - T(\partial\lambda)) =$$
$$\lambda - T\partial\lambda - b\partial\lambda + b\partial T(\partial\lambda) =^1$$
$$\lambda - T\partial\lambda - b\partial\lambda + b((I - S) - T\partial)\partial\lambda =$$
$$\lambda - T\partial\lambda - b\partial\lambda + b\partial\lambda - bS\partial\lambda - bT\partial\partial\lambda =$$
$$\lambda - T\partial\lambda - bS\partial\lambda =^2 .$$
$$\lambda - T\partial\lambda - S\lambda.$$

Equality 0 comes from $b\partial + \partial b = I$. Equality 1 is induction. Equality 2 comes from the definition of S. The other equalities just amount to expanding things out. ♠

4 The Final Argument

The operator S induces a map from $C_n(X, B)$ to $C_n(X, B)$. We just push forward the operation using the maps involved in a given chain. That is, if we have a chain $\sum a_i f_j$ we just apply S to the original n-simplex and then sum over the restrictions of f_j to this simplex. Likewise T defines a map from $C_n(X,B)$ to $C_{n+1}(X,B)$. The same relations established above hold, namely

$$S\partial = \partial S, \qquad T\partial + \partial T = I - S.$$

What is more, these same maps preserve the chain complex $C(A+B,B)$. We can either think of S and T as being maps on $C(X,B)$ or on $C(A+B,B)$.

4.1 Surjectivity

Lemma 4.1 *If $c \in C_n(X,B)$ is a cycle and m is a positive integer, then $S^m(c)$ is a cycle in $C_n(X,B)$ and $[c] = [S^m(c)]$ in $H_n(X,B)$.*

Proof: By iteration, it suffices to prove this for $m = 1$. Again we note that we have $\partial S(c) = S(\partial c)$. But $\partial c \in B$ and S maps chains in A to chains in B. Hence $S(\partial c) = 0$ in $C_n(X,B)$. This proves that $S(c)$ is a cycle. Next, we have

$$c - S(c) = T(\partial c) + \partial T(c).$$

But $\partial c \in B$ and so $T(\partial c) \in B$ as well. Hence $c - S(c)$ differs by a boundary. Hence $[c] = [S(c)]$ in $H_n(X,B)$. ♠

Corollary 4.2 *$\iota^\#$ is a surjective map from $H_n(A+B,B)$ to $H_n(X,B)$.*

Proof: Let c be a cycle in $C_n(X,B)$. By compactness and the Shrinking Corollary, we can choose m so that $S^m(c)$ is a sum of simplices such that each one either has its image in A or in B. This uses the fact that the interiors of A and B make a covering of X. But $[c] = [S^m(c)]$ and $S^m(c)$ is in the image of $\iota^\#$. ♠

4.2 Injectivity

Lemma 4.3 *If $c \in C_n(A+B,B)$ is a cycle then $S^m(c)$ is a cycle in $C_n(A+B,B)$ and $[x] = [S^m(c)]$ in $H_n(A+B,B)$.*

Proof: This has exactly the same proof as Lemma 4.1, because both S and T act on $C(A+B,B)$. ♠

Corollary 4.4 *$\iota^\#$ is an injective map from $H_n(A+B,B)$ to $H_n(X,B)$.*

Proof: By linearity it suffices to show that the kernel of the map is trivial. Suppose that $c = \iota^\#(c')$ is 0 in $H^n(X,B)$. This means that $c = \partial d$ where $d \in C_{n+1}(X,B)$. In terms of chains, this means that

$$\partial d = c' + \beta,$$

where $\beta \in C_n(B)$. The point is that c and c' are the same chain, just considered in different groups.

By compactness we can find some integer m such that

$$S^m(d) \in C_{n+1}(A+B,B).$$

This means that $S^m(d) = \iota^\#(d')$. But then

$$\partial(d') = \partial S^m(d) = S^m \partial d = S^m c' + S^m \beta.$$

This shows that $[S^m(c')] = 0$ in $H_n(A+B,B)$. Lemma 4.3 now says that $[c'] = [S^m(c')] = 0$ in $H_n(A+B,B)$. ♠

The proof is done.

5 Discussion

The proof above avoids a lot of the algebraic complexity in Hatcher's proof. Why is this possible. Well, Hatcher wants a single chain homotopy that works for all n at the same time and for all chains at the same time. We are working with individual chains and so for each one we can pick some integer m which works. In Hatcher's case, the integer m varies with the simplex.

The varying m poses problems. To understand this, suppose that we have a chain involving 2 simplices. Maybe we need to take $m = 3$ on the first one and $m = 7$ on the second one. If we do this, then the faces common to both simplices no longer cancel out. The way to fix this problem would be to add "buffers" between unequally subdivided boundaries. The buffer would essentially be some finite sum of the operations S and T applied to the boundaries. I believe that this is exactly what Hatcher's final map ρ does.