Chapter 5

Sequences

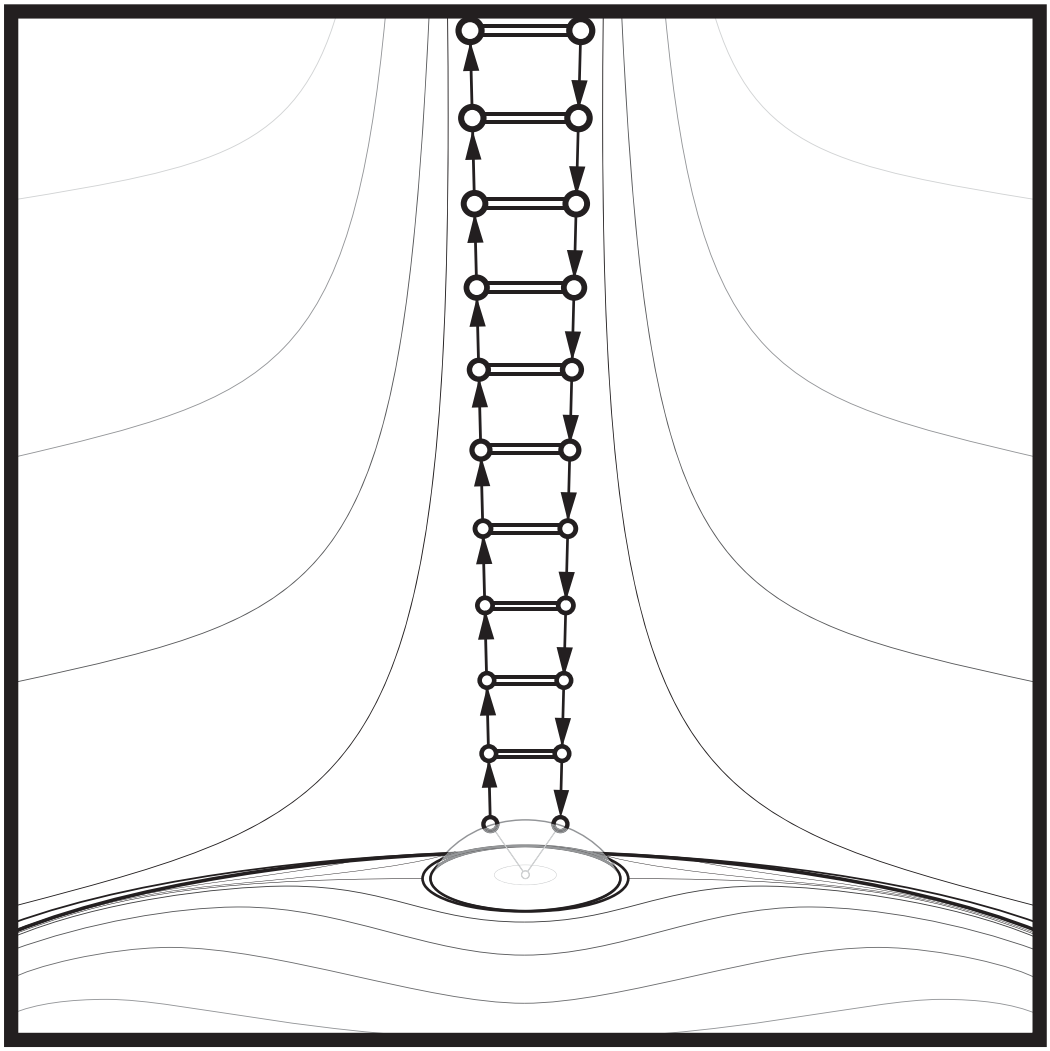

Homology takes as its input a chain complex – a hierarchical assembly line of parts – and returns its global features. The elemental tools for analyzing homology are likewise linelike devices: **exact sequences**. When outfitted with functoriality, these sequences assemble into the homological engines of inference.

5.1 Homotopy invariance

Homology begins by replacing topological spaces with complexes of algebraic objects (vector spaces, abelian groups, or modules, depending on one's preferences). Other topological notions – continuous functions, homeomorphisms, homotopies, *etc.* – also have analogues at the level of chain complexes. Generalizing from the pattern set in §4.10, define a **chain map** to be any graded homomorphism $\varphi_\bullet\colon \mathcal{C} \to \mathcal{C}'$ between chain complexes that respects the grading and commutes with the boundary maps. This is best expressed in the form of a commutative diagram as per (4.10):

$$\begin{array}{ccccccccc} \cdots & \longrightarrow & C_{n+1} & \xrightarrow{\partial} & C_n & \xrightarrow{\partial} & C_{n-1} & \xrightarrow{\partial} & \cdots \\ & & \downarrow{\scriptstyle \varphi_\bullet} & & \downarrow{\scriptstyle \varphi_\bullet} & & \downarrow{\scriptstyle \varphi_\bullet} & & \\ \cdots & \longrightarrow & C'_{n+1} & \xrightarrow{\partial'} & C'_n & \xrightarrow{\partial'} & C'_{n-1} & \xrightarrow{\partial'} & \cdots \end{array} \tag{5.1}$$

Commutativity means that homomorphisms are path-independent in the diagram: $\varphi_\bullet \circ \partial = \partial' \circ \varphi_\bullet$. Chain maps are the analogues of continuous maps, since, via respect for the boundary operators, neighbors are sent to neighbors. The appropriate generalization of a homeomorphism to chain complexes is therefore an invertible chain map – one which is an isomorphism for all $C_n \to C'_n$. Clearly, a homeomorphism $f\colon X \to Y$ induces a chain map $f_\bullet\colon \mathcal{C}_X^{\text{sing}} \to \mathcal{C}_Y^{\text{sing}}$ which is an isomorphism. As such, $H_\bullet^{\text{sing}}(X) \cong H_\bullet^{\text{sing}}(Y)$.

The extension to homotopy is more subtle. Recall that $f, g\colon X \to Y$ are homotopic if there is a map $F\colon X \times [0,1] \to Y$ which restricts to f on $X \times \{0\}$ and to g on $X \times \{1\}$. A **chain homotopy** between chain maps $\varphi_\bullet, \psi_\bullet\colon \mathcal{C} \to \mathcal{C}'$ is a homomorphism $F\colon \mathcal{C} \to \mathcal{C}'$ sending n-chains to $(n+1)$-chains so that $\partial' F - F\partial = \varphi_\bullet - \psi_\bullet$:

$$\begin{array}{ccccccccc} \cdots & \longrightarrow & C_{n+1} & \xrightarrow{\partial} & C_n & \xrightarrow{\partial} & C_{n-1} & \xrightarrow{\partial} & \cdots \\ & \swarrow^{F} & \downarrow{\scriptstyle \psi_\bullet}\ \downarrow{\scriptstyle \varphi_\bullet} & \swarrow^{F} & \downarrow{\scriptstyle \psi_\bullet}\ \downarrow{\scriptstyle \varphi_\bullet} & \swarrow^{F} & \downarrow{\scriptstyle \psi_\bullet}\ \downarrow{\scriptstyle \varphi_\bullet} & \swarrow^{F} & \\ \cdots & \longrightarrow & C'_{n+1} & \xrightarrow{\partial'} & C'_n & \xrightarrow{\partial'} & C'_{n-1} & \xrightarrow{\partial'} & \cdots \end{array} \tag{5.2}$$

One calls F a map of **degree** $+1$, indicating the upshift in the grading.[1] Note the morphological resemblance to homotopy of maps: a chain homotopy maps each n-chain to a $n+1$-chain, the algebraic analogue of a 1-parameter family. The difference between the ends of the homotopy, $\partial' F - F\partial$, gives the difference between the chain maps.

Lemma 5.1. *Chain homotopic maps induce the same homomorphisms on homology.*

[1] The overuse of the term *degree* in graphs, maps of spheres, and chain complexes is unfortunate.

Proof. Consider $[\alpha] \in H_\bullet(\mathcal{C})$. Assuming $\varphi_\bullet$ and $\psi_\bullet$ are chain homotopic maps from $\mathcal{C}$ to $\mathcal{C}'$,

$$\begin{aligned} H(\varphi_\bullet)[\alpha] - H(\psi_\bullet)[\alpha] &= [(\varphi_\bullet - \psi_\bullet)\alpha] \\ &= [(\partial' F + F\partial)\alpha] \\ &= [\partial'(F\alpha)] + [F(\partial\alpha)] = 0, \end{aligned}$$

since α is a cycle and $\partial'(F\alpha)$ is a boundary. ⊙

The following theorem is proved by constructing an explicit chain homotopy [176]:

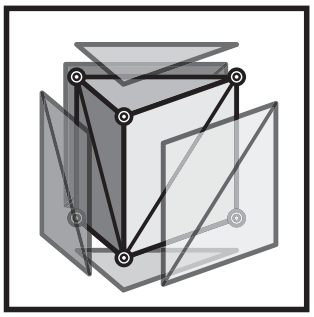

Theorem 5.2 (Homotopy Invariance of Homology). *Homotopic maps $f, g\colon X \to Y$ induce chain homotopic maps $f_\bullet, g_\bullet$ from $\mathcal{C}_X^{\text{sing}}$ to $\mathcal{C}_Y^{\text{sing}}$.*

The idea behind the proof is simple. For each singular n-simplex σ, one considers the F-image of $\sigma \times [0,1]$ as a family of singular n-simplices parameterized by the homotopy. This prism is then triangulated into singular $(n+1)$-simplices that encode the homotopy. It is shown that this chain map $\mathcal{P}$ (called a **prism operator**) is a chain homotopy.

Corollary 5.3. *Singular homology is a homotopy invariant of topological spaces.*

5.2 Exact sequences

In the analogy between topological spaces and algebraic complexes, there is a special class of complexes that are elementary building blocks. A complex $\mathcal{C} = (C_\bullet, \partial)$ is **exact** when its homology vanishes: $\ker \partial_n = \operatorname{im} \partial_{n+1}$ for all n. Exact sequences are as useful and primal as the nullhomologous spaces they mirror.

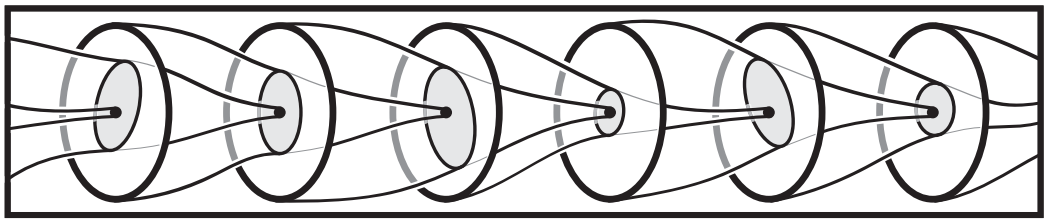

Example 5.4 (Exact sequence) ◎

The following simple examples of exact sequences help build intuition:

1. Two groups are isomorphic, $\mathbf{G} \cong \mathbf{H}$, iff there is an exact sequence of the form:

$$0 \longrightarrow \mathbf{G} \longrightarrow \mathbf{H} \longrightarrow 0$$

2. The 1st Isomorphism Theorem says that for a homomorphism φ of $\mathbf{G}$, the following sequence is exact:

$$0 \longrightarrow \ker \varphi \longrightarrow \mathbf{G} \xrightarrow{\varphi} \operatorname{im} \varphi \longrightarrow 0$$

   Such a 5-term sequence framed by zeroes is called a **short exact sequence**. In any such short exact sequence, the second map is injective; the penultimate, surjective.
3. More generally, the kernel and cokernel of a homomorphism $\varphi\colon \mathbf{G} \to \mathbf{H}$ fit into an exact sequence:

$$0 \longrightarrow \ker \varphi \longrightarrow \mathbf{G} \xrightarrow{\varphi} \mathbf{H} \longrightarrow \operatorname{coker} \varphi \longrightarrow 0$$

4. Consider $C = C^\infty(\mathbb{R}^3)$, the vector space of differentiable functions and $\mathcal{X} = \mathcal{X}(\mathbb{R}^3)$, the vector space of C^∞ vector fields on $\mathbb{R}^3$. These fit together into an exact sequence,

$$0 \longrightarrow \mathbb{R} \longrightarrow C \xrightarrow{\nabla} \mathcal{X} \xrightarrow{\nabla\times} \mathcal{X} \xrightarrow{\nabla\cdot} C \longrightarrow 0\ , \tag{5.3}$$

   where ∇ is the gradient differential operator from vector calculus, and the initial $\mathbb{R}$ term in the sequence represents the constant functions on $\mathbb{R}^3$. The exactness of this sequence encodes the fact that *curl-of-grad* and *div-of-curl* vanish, as well as the fact that, on $\mathbb{R}^3$, all curl-free fields are gradients and all div-free fields are curls. This one exact sequence compactly encodes many of the relations of vector calculus.

◎

The most important examples of exact sequences are those relating homologies of various spaces and subspaces. The critical technical tool for the generation of such weaves an exact thread through a loom of chain complexes.

Lemma 5.5 (Snake Lemma). *If $\mathcal{A} = (A_\bullet, \partial)$, $\mathcal{B} = (B_\bullet, \partial)$, and $\mathcal{C} = (C_\bullet, \partial)$ form a short exact sequence of chain complexes,*

$$0 \longrightarrow A_\bullet \xrightarrow{i_\bullet} B_\bullet \xrightarrow{j_\bullet} C_\bullet \longrightarrow 0\ ,$$

then this induces the **long exact sequence***:*

$$\longrightarrow H_n(\mathcal{A}) \xrightarrow{H(i)} H_n(\mathcal{B}) \xrightarrow{H(j)} H_n(\mathcal{C}) \xrightarrow{\delta} H_{n-1}(\mathcal{A}) \xrightarrow{H(i)} \ . \tag{5.4}$$

Moreover, the long exact sequence is **natural***: a commutative diagram of short exact sequences and chain maps*

$$\begin{array}{ccccccccc}
0 & \longrightarrow & A_\bullet & \longrightarrow & B_\bullet & \longrightarrow & C_\bullet & \longrightarrow & 0 \\
 & & \downarrow{\scriptstyle f_\bullet} & & \downarrow{\scriptstyle g_\bullet} & & \downarrow{\scriptstyle h_\bullet} & & \\
0 & \longrightarrow & \tilde{A}_\bullet & \longrightarrow & \tilde{B}_\bullet & \longrightarrow & \tilde{C}_\bullet & \longrightarrow & 0
\end{array}$$

induces a commutative diagram of long exact sequences

$$\begin{array}{ccccccccc} \longrightarrow & H_n(\mathcal{A}) & \longrightarrow & H_n(\mathcal{B}) & \longrightarrow & H_n(\mathcal{C}) & \xrightarrow{\delta} & H_{n-1}(\mathcal{A}) & \longrightarrow \\ & \downarrow{\scriptstyle H(f)} & & \downarrow{\scriptstyle H(g)} & & \downarrow{\scriptstyle H(h)} & & \downarrow{\scriptstyle H(f)} & \\ \longrightarrow & H_n(\tilde{\mathcal{A}}) & \longrightarrow & H_n(\tilde{\mathcal{B}}) & \longrightarrow & H_n(\tilde{\mathcal{C}}) & \xrightarrow{\delta} & H_{n-1}(\tilde{\mathcal{A}}) & \longrightarrow \end{array} \tag{5.5}$$

An *exact sequence of chain complexes* means that there is a short exact sequence in each grading, and these short exact sequences fit into a commutative diagram with respect to the boundary operators. The induced **connecting homomorphism** $\delta\colon H_n(\mathcal{C}) \to H_{n-1}(\mathcal{A})$ comes from the boundary map in $\mathcal{C}$ as follows:

1. Fix $[\gamma] \in H_n(\mathcal{C})$; thus, $\gamma \in C_n$.
2. By exactness, $\gamma = j(\beta)$ for some $\beta \in B_n$.
3. By commutativity, $j(\partial\beta) = \partial(j\beta) = \partial\gamma = 0$.
4. By exactness, $\partial\beta = i\alpha$ for some $\alpha \in A_{n-1}$.
5. Set $\delta[\gamma] := [\alpha] \in H_{n-1}(\mathcal{A})$.

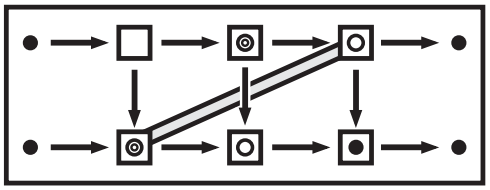

Every topologist, no matter how wedded to geometric intuition, must possess a thorough understanding of the connecting homomorphism. This is perhaps best grasped via animation; a static illustration is a poor substitute, but nevertheless conveys the critical shift in grading that exactness and weaving enacts. The reader should demonstrate that δ is well-defined and that the resulting long exact sequence is indeed exact. Doing so solidifies the invaluable technique of diagrammatic argument.

5.3 Pairs and Mayer-Vietoris

Of the many exact sequences that a topologist must master, two are central: the long exact sequence of a pair, and the Mayer-Vietoris sequence. The former unwraps the relative homology of §4.7. Given $A \subset X$ (a subset in the singular case, or a subcomplex in the cellular), the following sequence of chain complexes is exact:

$$0 \longrightarrow C_\bullet(A) \xrightarrow{i_\bullet} C_\bullet(X) \xrightarrow{j_\bullet} C_\bullet(X, A) \longrightarrow 0\ ,$$

where $i\colon A \hookrightarrow X$ is an inclusion and $j\colon (X, \varnothing) \hookrightarrow (X, A)$ is an inclusion of pairs. This yields the **long exact sequence of the pair** (X, A):

$$\longrightarrow H_n(A) \xrightarrow{H(i)} H_n(X) \xrightarrow{H(j)} H_n(X, A) \xrightarrow{\delta} H_{n-1}(A) \longrightarrow\ , \tag{5.6}$$

The connecting homomorphism δ takes a relative homology class $[\alpha] \in H_n(X, A)$ to the homology class $[\partial\alpha] \in H_{n-1}(A)$. This is an excellent sequence for decomposing

homology of quotient spaces in terms of the homological analogues of images and kernels of a projection.

Example 5.6 (Spheres) ◎

Computing the homology of the sphere $\mathbb{S}^k$ is a simple application of the long exact sequence of the pair $(\mathbb{D}^k, \partial\mathbb{D}^k)$, made simpler still by using reduced homology:

$$\longrightarrow \tilde{H}_n(\mathbb{D}^k) \xrightarrow{H(j)} H_n(\mathbb{D}^k, \partial\mathbb{D}^k) \xrightarrow{\delta} \tilde{H}_{n-1}(\partial\mathbb{D}^k) \xrightarrow{H(i)} \tilde{H}_{n-1}(\mathbb{D}^k) \longrightarrow .$$

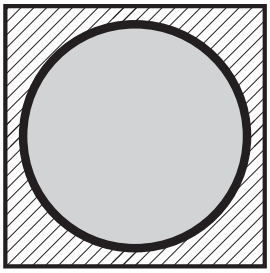

By definition, $\partial\mathbb{D}^k \cong \mathbb{S}^{k-1}$. By excision, $H_\bullet(\mathbb{D}^k, \partial\mathbb{D}^k) \cong \tilde{H}_\bullet(\mathbb{D}^k/\partial\mathbb{D}^k) \cong \tilde{H}_\bullet(\mathbb{S}^k)$. For all n, $\tilde{H}_n(\mathbb{D}^k) = 0$. As the first and last terms in the sequence above vanish, exactness yields the recurrence relation

$$\tilde{H}_n(\mathbb{S}^k) \cong \tilde{H}_{n-1}(\mathbb{S}^{k-1}),$$

for all n. Beginning with the explicit and trivial computation of $\tilde{H}_n(\mathbb{S}^0) \cong 0$ for $n > 0$ and $\cong \mathbb{Z}$ for $n = 0$, one inducts via the above to show that $\tilde{H}_n(\mathbb{S}^k) \cong \mathbb{Z}$ for $n = k$, and $\cong 0$ else. ◎

The second key sequence is derived from a decomposition of X into subsets (or subcomplexes) A and B. In the singular setting, one requires $X = \mathrm{int}(A) \cup \mathrm{int}(B)$; in the cellular case, subcomplexes suffice. Consider the short exact sequence

$$0 \longrightarrow C_\bullet(A \cap B) \xrightarrow{\phi_\bullet} C_\bullet(A) \oplus C_\bullet(B) \xrightarrow{\psi_\bullet} C_\bullet(A + B) \longrightarrow 0,$$

with chain maps $\phi_\bullet\colon c \mapsto (c, -c)$, and $\psi_\bullet\colon (a, b) \mapsto a + b$. The term on the right, $C_\bullet(A + B)$, consists of those chains which can be expressed as a sum of chains on A and chains on B. In cellular homology with A, B subcomplexes, $C_\bullet(A+B) \cong C_\bullet(X)$; in the singular category, one shows (via the techniques of Čech homology) that $H_\bullet(A+B) \cong H_\bullet(X)$. In both settings, the resulting long exact sequence yields the **Mayer-Vietoris sequence**:

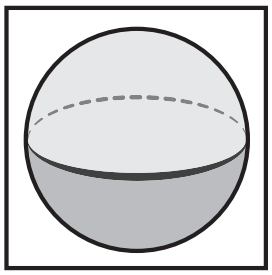

$$\longrightarrow H_n(A \cap B) \xrightarrow{H(\phi)} H_n(A) \oplus H_n(B) \xrightarrow{H(\psi)} H_n(X) \xrightarrow{\delta} H_{n-1}(A \cap B) \longrightarrow$$

The connecting homomorphism decomposes a cycle in X into a sum of chains in A and B, then takes the boundary of one of these chains in $A \cap B$. This sequence captures the additivity of homology, *cf.* Lemma 3.8.

Example 5.7 (Spheres, redux) ◎

The computation of $H_\bullet(\mathbb{S}^k)$ can be carried out via Mayer-Vietoris as follows. Let A and B be upper and lower hemispheres, homeomorphic to $\mathbb{D}^k$, intersecting at an equatorial $\mathbb{S}^{k-1}$.

$$\longrightarrow \tilde{H}_n(\mathbb{D}^k)\oplus\tilde{H}_n(\mathbb{D}^k) \xrightarrow{H(\psi)} \tilde{H}_n(\mathbb{S}^k) \xrightarrow{\delta} \tilde{H}_{n-1}(\mathbb{S}^{k-1}) \xrightarrow{H(\phi)} \tilde{H}_{n-1}(\mathbb{D}^k)\oplus\tilde{H}_{n-1}(\mathbb{D}^k) \longrightarrow$$

As $\tilde{H}_\bullet(\mathbb{D}^k) \cong 0$, one obtains by exactness that $\delta\colon \tilde{H}_n(\mathbb{S}^k) \cong \tilde{H}_{n-1}(\mathbb{S}^{k-1})$ for all n and all k. Thus, again, via $\tilde{H}_\bullet\mathbb{S}^0$, one has immediately that $\tilde{H}_n(\mathbb{S}^k) \cong 0$ unless $n = k$, where it equals $\mathbb{Z}$. ◎

There are many other exact sequences, only a few of which this chapter will unfurl. For future use, note the existence of relative versions of the two sequences above. The long exact sequence of a **triple** (X, A, B), where $B \subset A \subset X$, is derived from the short exact sequence of chains:

$$0 \longrightarrow C_\bullet(A, B) \xrightarrow{i_\bullet} C_\bullet(X, B) \xrightarrow{j_\bullet} C_\bullet(X, A) \longrightarrow 0\ ,$$

where, as before, $i_\bullet$ and $j_\bullet$ are induced by inclusions on pairs. In the case of Mayer-Vietoris, if one decomposes the pair (X, Y) for $Y \subset X$ as $X = A \cup B$ and $Y = C \cup D$, with $C \subset A$ and $D \subset B$, then the following is exact:

$$0 \longrightarrow C_\bullet(A \cap B, C \cap D) \xrightarrow{\phi_\bullet} C_\bullet(A, C) \oplus C_\bullet(B, D) \xrightarrow{\psi_\bullet} C_\bullet(A + B, C + D) \longrightarrow 0,$$

As before, when X lies in the union of interiors of A and B (likewise with Y in the union of C and D), then the penultimate term in the sequence becomes $H_\bullet(X, Y)$ when passing to homology. Section 5.12 will put this relative tool to use.

5.4 Equivalence of homology theories

This chapter slowly builds the argument that even elementary homological algebra is a powerful upgrade to the basic linear algebra so useful in applied mathematics. Exactness and commutativity are two such simple tools: much more is available. Consider the following diagrammatic lemma from homological algebra.

Lemma 5.8 (The 5-Lemma). *Given a commutative diagram of abelian groups of the form*

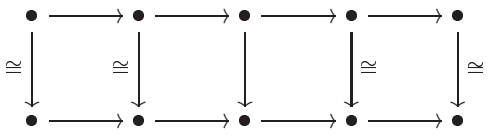

(5.7)

whose top and bottom rows are exact, and whose four outer *vertical maps are isomorphisms; the middle vertical map is an isomorphism as well.*

This lemma is extremely useful: one example suffices. This text handles competing homology theories glibly, with the justification that they usually agree.[2] This almost always can be shown using the 5-Lemma and induction.

[2]Exotic spaces on which they disagree are not of primary importance in most applications.

Theorem 5.9. *On simplicial complexes, singular and simplicial homology are isomorphic.*

Proof. Choose a filtration $X_i \subset X_{i+1} \subset \cdots$ for X that adds one simplex per step and induct on this sequence order. On the 0-skeleton of X, the isomorphism is clear. Assume that σ is the k-simplex which when added to X_i yields X_{i+1}. This σ is glued to X_i along its boundary $\partial\sigma$, homeomorphic to the sphere $\mathbb{S}^{k-1}$. The Mayer-Vietoris sequences for (X_i, σ) in simplicial and singular homology fit together in a commutative diagram,

$$\begin{array}{ccccccccc}
H_n^{\text{simp}}\partial\sigma & \longrightarrow & H_n^{\text{simp}}X_i \oplus H_n^{\text{simp}}\sigma & \longrightarrow & H_n^{\text{simp}}X_{i+1} & \longrightarrow & H_{n-1}^{\text{simp}}\partial\sigma & \longrightarrow & H_{n-1}^{\text{simp}}X_i \oplus H_{n-1}^{\text{simp}}\sigma \\
\downarrow\cong & & \downarrow\cong & & \downarrow & & \downarrow\cong & & \downarrow\cong \\
H_n^{\text{sing}}\partial\sigma & \longrightarrow & H_n^{\text{sing}}X_i \oplus H_n^{\text{sing}}\sigma & \longrightarrow & H_n^{\text{sing}}X_{i+1} & \longrightarrow & H_{n-1}^{\text{sing}}\partial\sigma & \longrightarrow & H_{n-1}^{\text{sing}}X_i \oplus H_{n-1}^{\text{sing}}\sigma
\end{array}$$

with vertical arrows induced from the map interpreting a simplicial chain as a singular chain. By induction and previous computations of the homologies of balls and spheres, four of the five vertical maps are isomorphisms. The 5-Lemma completes the induction step and the proof. ⊙

The same result holds for other homology theories, such as Čech and cellular, assuming the appropriate defining structures (covers, cell structures, *etc.*) exist. The proof of Theorem 5.9 works for cellular homology of a regular cell complex. In the general case, more care concerning the definitions is required. In all these cases, the isomorphism is *natural*, meaning that, as in the case of the long exact sequences of §5.2, a map $f : X \to Y$ between spaces induce the same homomorphisms on homology, independent of which theory (singular, cellular, *etc.*) is used, so long as these homologies are well-defined. This sameness is expressed as a commutative diagram.

$$\begin{array}{ccc}
H_\bullet^{\text{cell}}X & \xrightarrow{H(f)} & H_\bullet^{\text{cell}}Y \\
\downarrow\cong & & \downarrow\cong \\
H_\bullet^{\text{sing}}X & \xrightarrow[H(f)]{} & H_\bullet^{\text{sing}}Y
\end{array}$$

5.5 Cellular homology, redux

The treatment of cellular homology in Chapter 4 was incomplete, especially in regards to defining the incidence number $[\sigma : \tau]$ for a pair of cells $\sigma \trianglelefteq \tau$, and showing that the boundary operator in Equation (4.6) yields a chain complex. Moreover, the boundary operator was only defined for *regular* cell complexes, and a formal definition of induced orientation was never given at all outside of the simplicial setting. These issues can now be rectified.

Let X be a finite-dimensional cell complex. Choose an orientation for each cell τ: recall from Example 4.18 that this is a choice of generator for $H_n(\tau, \partial\tau; \mathbb{Z})$ for $n = \dim \tau$. To define the induced orientation on the boundary of the cell (without using one's right hand), use the connecting homomorphism $\delta : H_n(\tau, \partial\tau) \to H_{n-1}(\partial\tau)$ from the long exact sequence on the pair $(\tau, \partial\tau)$. Recall from §4.3 that the incidence

number $[\sigma\colon \tau]$ of a face $\sigma \trianglelefteq \tau$ records (for a *regular* cell complex) whether the orientations on $\partial\tau$ and σ agree $(+1)$ or disagree (-1). This can now be both interpreted and extended to the general cellular setting as a degree:

$$[\sigma\colon \tau] := \deg\left(H_n(\tau, \partial\tau) \xrightarrow{\delta} H_{n-1}(\partial\tau) \xrightarrow{H(q)} H_{n-1}(\sigma/\partial\sigma) \right), \tag{5.8}$$

where for τ an n-cell, $q\colon \partial\tau \to \sigma/\partial\sigma$ is the map that quotients out the complement of σ in $X^{(n-1)}$. In short, *incidence number is the degree of the attaching map on the boundary*. For a non-regular cell complex, this number may be nonzero: witness the cell structure on $\mathbb{P}^2$ that has one cell in each dimension $0, 1, 2$, and an attaching map from the 2-cell to the 1-skeleton of degree 2.

What remains is to show that the incidence numbers cancel to give the cellular boundary operator as per Equation (4.7). This requires a deeper look at what the quotient map in (5.8) is. Consider $X^{(n)} \subset X$ the n-skeleton comprised of all cells of dimension less than or equal to n. To analyse how the n-cells are glued onto the $(n-1)$-skeleton, one focuses on the long exact sequence on the pair $(X^{(n)}, X^{(n-1)})$ in *singular* homology near grading n:

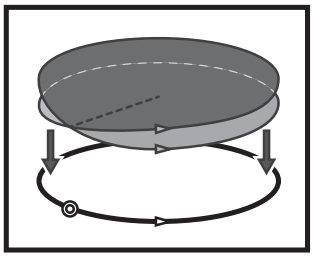

$$\longrightarrow H_n X^{(n)} \xrightarrow{H(j)} H_n(X^{(n)}, X^{(n-1)}) \xrightarrow{\delta} H_{n-1} X^{(n-1)} \xrightarrow{H(i)} H_{n-1} X^{(n)} \longrightarrow .$$

The critical observation is this: $H_n(X^{(n)}, X^{(n-1)}) \cong H_n(X^{(n)}/X^{(n-1)}) \cong C_n^{\text{cell}}(X)$, since the quotient is a wedge of spheres of dimension n, one for each n-cell of X. This allows for a definition of the cellular chain complex via weaving together the long exact sequences of incident skeletal pairs:

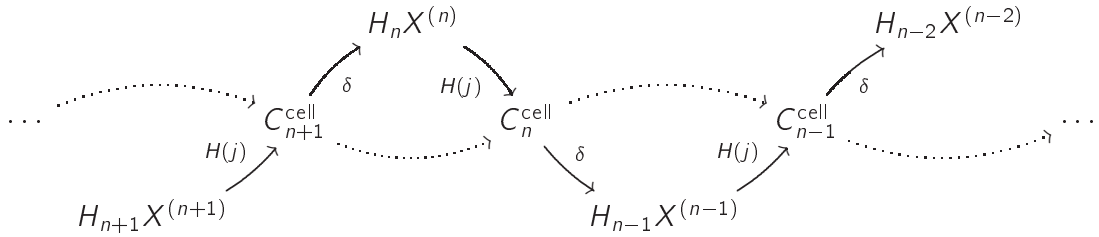

One defines the cellular boundary operator ∂^{cell} from this diagram as the composition $\partial^{\text{cell}} := H(j) \circ \delta$. This gives an immediate proof that $\partial^2 = 0$ in cellular homology, since $\partial^2 = H(j) \circ (\delta \circ H(j)) \circ \delta$, and the middle two terms vanish due to exactness. That this definition of cellular homology agrees with that of §4.3 follows from a close examination of ∂^{cell}, which takes $X^{(n)}/X^{(n-1)}$ to $X^{(n-1)}/X^{(n-2)}$ by means of δ and quotients. On an n-cell τ, the n-sphere $\tau/\partial\tau$ is sent to the wedge of $(n-1)$-spheres $\sigma/\partial\sigma$ for each $\sigma \trianglelefteq \tau$. This map, coming from δ, yields precisely the incidence number $[\sigma\colon \tau]$ on each face. Compared to juggling the combinatorics of incidence number cancellation, this exploitation of exactness is incisive.

5.6 Coverage in sensor networks

Sensors – devices which return data tied to a location – are ubiquitous. The problem of collating distributed pieces of sensor data over a communications network is an engineering challenge for which the tools of topology and homological algebra seem strangely fitting.

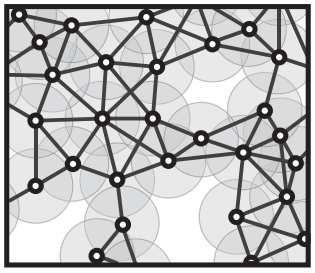

One simple-to-state problem is that of **coverage**. Fix a domain $\mathcal{D} \subset \mathbb{R}^2$ and consider a finite collection $\mathcal{Q}$ of sensors nodes in the plane with two tasks: they (1) sense a neighborhood of their locale in $\mathcal{D}$; and (2) communicate with other sensors. Both of these actions are assumed to be local in the sense that individual nodes cannot extract sensing data from or communicate data over all of $\mathcal{D}$. The problem of coverage, or more precisely, **blanket coverage**, is the question of whether there are holes in the sensor network – are there any regions in $\mathcal{D}$ which are not sensed? Other important coverage problems include **barrier coverage**, in which one wants to determine whether a sensor network separates $\mathcal{D}$ or surrounds a critical region, and **sweeping coverage**, the time-dependent problem familiar to users of robotic vacuum sweepers.

When the coordinates of the sensors are known, computational geometry suffices to determine coverage. For non-localized sensors, the following simple application of homology gives effective criteria. Specific assumptions are kept to a minimum, for clarity and ease of proofs:

1. Sensors are modeled as a finite collection of nodes $\mathcal{Q} \subset \mathbb{R}^2$.
2. Each sensor is assumed to have a unique identity which it broadcasts; nearby neighbors detect the transmission and establish a communication link.
3. Communication is symmetric and generates a communications graph, X, on $\mathcal{Q}$ with corresponding flag complex $\mathsf{F} = \mathsf{F}(X)$.
4. Sensor coverage regions are based on communication proximity: for any subset of sensors $S \subset \mathcal{Q}$ which pairwise communicate, the union of coverage regions of S contains the convex hull of S in $\mathbb{R}^2$.
5. One fixes a 'fence' cycle $\mathsf{C} \subset X$ whose image in $\mathbb{R}^2$ is a simple closed curve bounding a domain $\mathcal{D} \subset \mathbb{R}^2$ of interest.

One wants to know whether $\mathcal{D}$ is contained in the coverage region of the network. The critical assumption is the fourth, connecting the communications and sensing of the network by means of the flag complex F. It is satisfied by systems with radially-symmetric communications networks (or *unit disc graphs*) and radially symmetric sensing regions with the proper ratio between sensing and communication [87, 88], but asymmetric systems are permissible in this framework.

Theorem 5.10 ([87]). *Given $\mathcal{Q}$, X, F, C, and $\mathcal{D}$ as above, then all of $\mathcal{D}$ is contained in the sensor-covered region if, equivalently:*

1. $[\mathsf{C}] = 0 \in H_1(\mathsf{F})$.

2. *There exists* $[\alpha] \in H_2(\mathsf{F}, \mathsf{C})$ *with* $\partial\alpha = \mathsf{C}$.

Proof. Equivalence of the two conditions comes from the long exact sequence of the pair (F, C) induced by the inclusion $i\colon \mathsf{C} \hookrightarrow \mathsf{F}$,

$$\cdots \longrightarrow H_2(\mathsf{F}, \mathsf{C}) \xrightarrow{\delta} H_1(\mathsf{C}) \xrightarrow{H(i)} H_1(\mathsf{F}) \longrightarrow \cdots,$$

since $\ker H(i) = \operatorname{im} \delta$ by exactness. Consider (from §2.2) the shadow map $\mathcal{S}\colon \mathsf{F} \to \mathbb{R}^2$ which sends vertices of the flag complex F to the physical sensor locations $\mathcal{Q} \subset \mathcal{D}$ and which sends a k-simplex of F to the (potentially singular) k-simplex given by the convex hull of the vertices implicated. By construction, $\mathcal{S}$ takes the pair (F, C) to $(\mathbb{R}^2, \partial\mathcal{D})$, acting as a homeomorphism on the second terms of the pair. This map induces the following diagram on the long exact sequences of the pairs:

$$\begin{array}{ccccccc} \cdots & \longrightarrow & H_2(\mathsf{F}, \mathsf{C}) & \xrightarrow{\delta} & H_1(\mathsf{C}) & \longrightarrow & \cdots \\ & & \downarrow{\scriptstyle H(\mathcal{S})} & & \cong \downarrow{\scriptstyle H(\mathcal{S})} & & \\ \cdots & \longrightarrow & H_2(\mathbb{R}^2, \partial\mathcal{D}) & \xrightarrow{\delta} & H_1(\partial\mathcal{D}) & \longrightarrow & \cdots \end{array} \tag{5.9}$$

By assumption, $\partial\alpha = \mathsf{C} \neq 0$; hence, $H(\mathcal{S})\delta[\alpha] = H(\mathcal{S})[\partial\alpha] \neq 0$. Naturality implies that the diagram is commutative: $\delta H(\mathcal{S}) = H(\mathcal{S})\delta$. Commutativity implies that $\delta H(\mathcal{S})[\alpha] \neq 0$, and thus $H(\mathcal{S})[\alpha] \neq 0$. If the sensors do not cover some point $p \in \mathcal{D}$, then p does not lie in the image of $\mathcal{S}$; thus, the map $\mathcal{S}\colon \mathsf{F} \to \mathbb{R}^2$ is a composition of maps $\mathsf{F} \to \mathbb{R}^2{-}p \hookrightarrow \mathbb{R}^2$. Diagram (5.9) is restructured as:

$$\begin{array}{ccccc} & & H_2(\mathsf{F}, \mathsf{C}) & \xrightarrow{\delta} & H_1(\mathsf{C}) \\ & \swarrow & \downarrow{\scriptstyle H(\mathcal{S})} & & \cong \downarrow{\scriptstyle H(\mathcal{S})} \\ H_2(\mathbb{R}^2{-}p, \partial\mathcal{D}) & & H_2(\mathbb{R}^2, \partial\mathcal{D}) & \xrightarrow{\delta} & H_1(\partial\mathcal{D}) \\ & \searrow & & & \end{array} \tag{5.10}$$

However, $H_2(\mathbb{R}^2{-}p, \partial\mathcal{D}) = 0$, as the long exact sequence of the pair $(\mathbb{R}^2{-}p, \partial\mathcal{D})$ reveals:

$$\longrightarrow H_2(\mathbb{R}^2{-}p) \longrightarrow H_2(\mathbb{R}^2{-}p, \partial\mathcal{D}) \xrightarrow{\delta} H_1(\partial\mathcal{D}) \xrightarrow{H(i)} H_1(\mathbb{R}^2{-}p) \longrightarrow$$

The first term in this sequence is zero since $\mathbb{R}^2{-}p \simeq \mathbb{S}^1$, which has vanishing H_2. The two last terms are $H_1(\partial\mathcal{D}) \cong H_1(\mathbb{R}^2{-}p)$, and, moreover, $H(i)\colon H_1(\partial\mathcal{D}) \to H_1(\mathbb{R}^2{-}p)$ gives the winding number (§4.12) of $\partial\mathcal{D}$ about $p \in \mathbb{R}^2$. Since p lies in the interior of $\mathcal{D}$, the winding number is ± 1, and $H(i)$ is an isomorphism. Exactness then implies that $H_2(\mathbb{R}^2{-}p, \partial\mathcal{D}) = 0$. Commutativity of (5.10) completes the proof. ⊙

The assumption on sensor coverage specifies that certain regions are guaranteed to be covered while passing no information about lack of coverage elsewhere. As such,

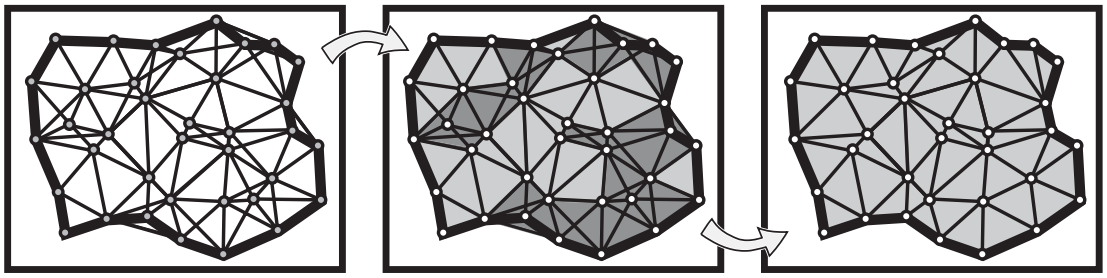

the homological coverage criterion cannot be if-and-only-if. It, like the assumptions on which it is built, is a conservative criterion. However, when the homological coverage criterion fails, choosing a basis for $H_1(\mathsf{F})$ which is *sparse* (implicating few nodes and edges) gives information about where the coverage holes may reside. Finally, *any* relative cycle $\alpha \in Z_2(\mathsf{F}, \mathsf{C})$ with $\partial\alpha = \mathsf{C}$ suffices to cover $\mathcal{D}$: only those nodes implicated in α are required to be actively sensing/communicating. This allows one to conserve power or establish a sleep-wake cycle by homological means.

5.7 Degree and computation

The crucial step in the proof of Theorem 5.10 used a winding number (degree) whose computation was, fortunately, obvious. More difficult degree computations are often possible by means of local formulae. For example, given a point $p \in \mathbb{R}^2$ and a closed curve $\gamma\colon \mathbb{S}^1 \to \mathbb{R}^2 - p$, the winding number of γ about p is easily computed as follows. Draw a ray from p and perturb it to intersect the image of γ in a finite set of points $Q = \{q_i\}_1^K$. At each intersection point q_i, the curve *kisses* or *crosses* the ray; either the curve traverses (left-to-right or right-to-left, since both are oriented) or it osculates, touching the ray and immediately turning back. Each action contributes a *local* degree: ± 1 if crossing and 0 if kissing. The net winding number of γ about p is the sum of these local contributions to degree.

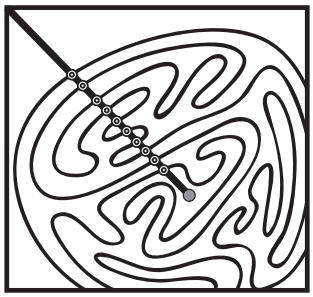

This simple example of a local computation rewards rumination. What happens if, instead of a curve in the plane, one needs to know whether a cycle in an ad hoc non-localized network surrounds a node (as in, *e.g.*, a network of security cameras [151])? To know whether a curve surrounds a point in the plane, it suffices to know the local behavior of the curve at a (small) finite number of points. What the curve does elsewhere is irrelevant. This has the pattern and stamp of topology. This intuitively simple procedure has a rigorous footing in transversality [169]; better still is the use of local homology. Assume $f\colon \mathbb{S}^n \to \mathbb{S}^n$ and $q \in f^{-1}(p)$ is an isolated point in the inverse image. Define the **local degree** of f at q to be

$$\deg(f; q) := \deg\, H(f)\colon H_n(U, U-q) \to H_n(V, V-p), \tag{5.11}$$

for U and V sufficiently small neighborhoods of q and p satisfying $f(U) \subset V$. The local degree is an integer since $H_n(\mathbb{S}^n, \mathbb{S}^n - p) \cong \mathbb{Z}$ for any point $p \in \mathbb{S}^n$. The validity of local computation can be shown using basic tools:

Proposition 5.11. *Assume that for $f : \mathbb{S}^n \to \mathbb{S}^n$, p has discrete inverse image $f^{-1}(p) = Q = \{q_i\}_i$. Then*

$$\deg f = \sum_i \deg(f, q_i). \tag{5.12}$$

Proof. From Lemma 5.5, the long exact sequences of the pairs $(\mathbb{S}^n, \mathbb{S}^n - Q)$ and $(\mathbb{S}^n, \mathbb{S}^n - p)$ form a commutative diagram:

$$\begin{array}{ccccccc}
H_n(\mathbb{S}^n - Q) & \longrightarrow & H_n(\mathbb{S}^n) & \xrightarrow{H(j)} & H_n(\mathbb{S}^n, \mathbb{S}^n - Q) & \xrightarrow{\delta} & H_{n-1}(\mathbb{S}^n - Q) \\
\downarrow{\scriptstyle H(f)} & & \downarrow{\scriptstyle H(f)} & & \downarrow{\scriptstyle H(f)} & & \downarrow{\scriptstyle H(f)} \\
H_n(\mathbb{S}^n - p) & \longrightarrow & H_n(\mathbb{S}^n) & \xrightarrow[H(j)]{\cong} & H_n(\mathbb{S}^n, \mathbb{S}^n - p) & \xrightarrow{\delta} & H_{n-1}(\mathbb{S}^n - p)
\end{array}$$

The first terms in both rows and the last term in the bottom row vanish. By exactness, the lower map $H(j)$ is an isomorphism. Thus, by commutativity

$$\deg f = \deg \left(H_n(\mathbb{S}^n) \xrightarrow{H(j)} H_n(\mathbb{S}^n, \mathbb{S}^n - Q) \xrightarrow{H(f)} H_n(\mathbb{S}^n, \mathbb{S}^n - p) \right).$$

Choose a small neighborhood V of p so that $f^{-1}(V) = \sqcup_i U_i$ is a disjoint collection of neighborhoods of the q_i. It follows that

$$H_n(\mathbb{S}^n, \mathbb{S}^n - Q) \cong^{(1)} H_n\left(f^{-1}(V), f^{-1}(V) - Q\right) \cong^{(2)} \bigoplus_i H_n(U_i, U_i - q_i),$$

where the isomorphisms come from (1) excision, and (2) additivity. The local degree at q_i, $\deg(f, q_i)$, is by definition the degree of the induced map $H(f)\colon H_n(U_i, q_i) \to H_n(V, p)$. Thus, by additivity,

$$\deg f = \deg \left(H_n(\mathbb{S}^n) \to \bigoplus_i H_n(U_i, U_i - q_i) \to H_n(V, V - p) \right) = \sum_i \deg(f, q_i).$$

⊙

Corollary 5.12. *If $f : \mathbb{S}^n \to \mathbb{S}^n$ is not surjective, then* $\deg(f) = 0$.

Proof. If not surjective, apply (5.12) to an empty inverse image. ⊙

The long exact sequence of the pair also permits computation of degree in slightly different settings. For example, consider the case of a map of the form $f : (\mathbb{D}^n, \partial\mathbb{D}^n) \to (\mathbb{D}^n, \partial\mathbb{D}^n)$ that maps a closed disc to itself, restricting to a map

on the boundary sphere. It is sensible to speak of the degree of f by using relative homology: $\deg(f) := \deg\left(H_n(\mathbb{D}^n, \partial\mathbb{D}^n) \to H_n(\mathbb{D}^n, \partial\mathbb{D}^n)\right)$. This is well-defined since $H_n(\mathbb{D}^n, \partial\mathbb{D}^n) \cong \mathbb{Z}$. Moreover, it is easily computed in terms of what f does either on the boundary or on the interior, as follows. There is both an induced map $\tilde{f}\colon \mathbb{S}^n \to \mathbb{S}^n$ on the quotient sphere, given by collapsing the boundary to a point, and a restriction map $\overline{f}\colon \mathbb{S}^{n-1} \to \mathbb{S}^{n-1}$ given by restriction to the boundary.

Lemma 5.13. *For the above,* $\deg(f) = \deg(\tilde{f}) = \deg(\overline{f})$.

Proof. That $\deg(f) = \deg(\tilde{f})$ follows from naturality. The equivalence of this to $\deg(\overline{f})$ comes from the long exact sequence applied to $f\colon (\mathbb{D}^n, \partial\mathbb{D}^n) \to (\mathbb{D}^n, \partial\mathbb{D}^n)$:

$$\begin{array}{ccccccccc} 0 = H_n(\mathbb{D}^n) & \longrightarrow & H_n(\mathbb{D}^n, \partial\mathbb{D}^n) & \xrightarrow{\delta} & H_{n-1}(\partial\mathbb{D}^n) & \longrightarrow & H_{n-1}(\mathbb{D}^n) = 0\,, \\ & & \downarrow{\scriptstyle H(f)} & & \downarrow{\scriptstyle H(\overline{f})} & & \\ 0 = H_n(\mathbb{D}^n) & \longrightarrow & H_n(\mathbb{D}^n, \partial\mathbb{D}^n) & \xrightarrow{\delta} & H_{n-1}(\partial\mathbb{D}^n) & \longrightarrow & H_{n-1}(\mathbb{D}^n) = 0 \end{array}$$

By exactness, the connecting homomorphisms δ on top and bottom are isomorphisms. Thus, by commutativity, $\deg(f) = \deg(\overline{f})$. ⊙

Example 5.14 (Colorings) ◎

There are numerous results in combinatorics that are inherently topological, several of which involve **colorings**. The following is a classical example implicating degree. Consider a 2-simplex T with vertices labeled by $\{0, 1, 2\}$. Let T' be a subdivision of T. Label all the new vertices of T' using any label $\{0, 1, 2\}$ subject to the following *boundary condition*: on $\partial T'$, each vertex must *not* be labeled by the label of T on the vertex opposite that edge. Hence, on the portion of the boundary of T' connecting the $\{0\}$ and $\{1\}$ vertices of T, the label $\{2\}$ must not be used.

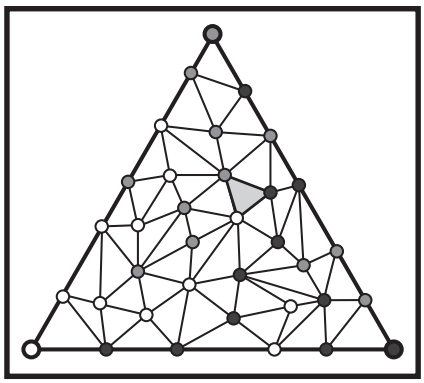

Lemma 5.15 (Sperner's Lemma). *Any* $\{0, 1, 2\}$*-labeling of the vertices of* T' *obeying the boundary condition must possess at least one triangle with all three labels at vertices.*

Proof. Consider the piecewise-linear simplicial map $f\colon T' \to T$ that sends each simplex in T' to the unique simplex in T determined by the convex hull of the labels in T'. Note that the conclusion of the theorem holds if and only if f is surjective. The boundary coloring condition implies that $f\colon \partial T' \to \partial T$ with degree $+1$. Thus, by Corollary 5.12 and Lemma 5.13, f is surjective. ⊙

From this proof, one sees clearly that the boundary condition can be relaxed to allow for any coloring of the vertices on the boundary that imparts a nonzero degree.

Furthermore, it is clear that the proof holds for a subdivided n-simplex with $n+1$ distinct colors. With the appropriate homological restrictions, other cellular spaces are likewise admissible: everything hinges on homology. Sperner's Lemma can be viewed as a discretized version of the Brouwer fixed point theorem of §4.13; as such, it is useful in discrete versions of fair-division and consensus problems [283]. ◎

Sperner's Lemma can also be restated as saying that any $(n+1)$-coloring of the vertices of a simplicial $\mathbb{S}^n$ must have an *even* number of $(n+1)$-colored simplices. In this setting, the proof above reduces to an examination of a simplicial map $\mathbb{S}^n \to \Delta^n$. The topological features of such maps are covered by another classical theorem from algebraic topology implicating degrees and exact sequences.

5.8 Borsuk-Ulam theorems

There are a number of results which fall under the names Borsuk-Ulam; all orbit about spheres and the antipodal map $a\colon \mathbb{S}^n \to \mathbb{S}^n$ that sends $x \mapsto -x$. The following is the key step in these results, written in the language of degree theory.

Theorem 5.16 (Borsuk-Ulam Theorem). *Odd maps of $\mathbb{S}^n$ have odd degree; even maps of $\mathbb{S}^n$ have even degree.*

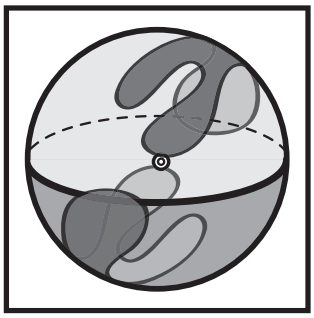

Proof. The proof in the (harder) odd case is sketched, the key to which is a long exact sequence to track antipodes. Assume that f is odd, so that $f \circ a = a \circ f$. Recall from Example 1.2, the quotient space $\mathbb{S}^n/a$ consisting of equivalence classes of antipodal points is one definition of the real projective space $\mathbb{P}^n$. Let $\pi\colon \mathbb{S}^n \to \mathbb{P}^n$ denote the quotient projection map; this map is a 2-to-1 local homeomorphism. Let $\overline{f}\colon \mathbb{P}^2 \to \mathbb{P}^2$ be the induced map. There is a commutative diagram of short exact sequences of chain complexes in $\mathbb{F}_2$ coefficients,

$$\begin{array}{ccccccccc}
0 & \longrightarrow & C_\bullet(\mathbb{P}^n) & \xrightarrow{\tau} & C_\bullet(\mathbb{S}^n) & \xrightarrow{\pi_\bullet} & C_\bullet(\mathbb{P}^n) & \longrightarrow & 0\ , \\
 & & \downarrow{\scriptstyle \overline{f}_\bullet} & & \downarrow{\scriptstyle f_\bullet} & & \downarrow{\scriptstyle \overline{f}_\bullet} & & \\
0 & \longrightarrow & C_\bullet(\mathbb{P}^n) & \xrightarrow{\tau} & C_\bullet(\mathbb{S}^n) & \xrightarrow{\pi_\bullet} & C_\bullet(\mathbb{P}^n) & \longrightarrow & 0
\end{array}$$

where τ is the **transfer map**, lifting a chain in $\mathbb{P}^n$ via a to an antipodal pair of chains in $\mathbb{S}^n$. The $\mathbb{F}_2$ coefficients ensures that the sequence is exact. The corresponding long exact sequences yield:

$$\begin{array}{ccccccccccc}
0 & \longrightarrow & H_n(\mathbb{P}^n) & \xrightarrow{H(\tau)} & H_n(\mathbb{S}^n) & \xrightarrow{H(\pi)} & H_n(\mathbb{P}^n) & \xrightarrow{\delta} & H_{n-1}(\mathbb{P}^n) & \longrightarrow & 0\ , \\
 & & \downarrow{\scriptstyle H(\overline{f})} & & \downarrow{\scriptstyle H(f)} & & \downarrow{\scriptstyle H(\overline{f})} & & \downarrow{\scriptstyle H(\overline{f})} & & \\
0 & \longrightarrow & H_n(\mathbb{P}^n) & \xrightarrow{H(\tau)} & H_n(\mathbb{S}^n) & \xrightarrow{H(\pi)} & H_n(\mathbb{P}^n) & \xrightarrow{\delta} & H_{n-1}(\mathbb{P}^n) & \longrightarrow & 0
\end{array}$$

By exactness and knowledge of $H_\bullet(\mathbb{P}^n)$, one argues that in the above diagram, δ and $H(\tau)$ are isomorphisms, while $H(\pi) = 0$. By inducting on the dimension n and using the commutative square penultimate to the right, one shows that $H(\overline{f})$ is an isomorphism. By commutativity (or the 5-Lemma), this implies that $H(f)\colon H_n(\mathbb{S}^n;\mathbb{F}_2) \to H_n(\mathbb{S}^n;\mathbb{F}_2)$ is an isomorphism. This, in turn, is the mod 2 reduction of $H(f)$ from $H_n(\mathbb{S}^n;\mathbb{Z}) \to H_n(\mathbb{S}^n;\mathbb{Z})$; *i.e.*, $\deg(f) \bmod 2 = 1$. ⊙

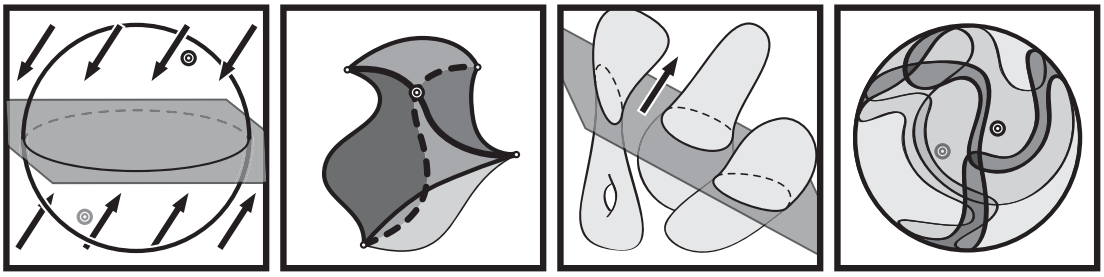

There are a number of famous corollaries of Theorem 5.16:

1. **Borsuk-Ulam:** Any map $\mathbb{S}^n \to \mathbb{R}^n$ must identify some pair of antipodal points.
2. **Radon:** Any map $\Delta^{n+1} \to \mathbb{R}^n$ of an $n+1$-simplex has a pair of disjoint closed faces (simplices of $\partial\Delta^{n+1}$) whose images in $\mathbb{R}^n$ intersect.
3. **Stone-Tukey:** Given a collection of n Lebesgue-measurable bodies in $\mathbb{R}^n$, there is a hyperplane which bisects evenly the volume of each body.
4. **Lusternik-Schnirelmann:** Any cover $\mathcal{U} = \{U_i\}_0^n$ of $\mathbb{S}^n$ by $n+1$ open sets must have at least one element U_j containing an antipodal pair of points.

Several of these theorems have physical interpretations. For example, it is common to express the first corollary above as saying that (assuming meteorological continuity) some antipodal pair of points on the earth have the same temperature and barometric pressure. Applications more relevant to this text lie in economics and *fair division* problems; for applications to combinatorics, see [219].

5.9 Euler characteristic

Sequences solve the mystery of the topological invariance of the Euler characteristic. In what follows, use field coefficients. The subtle step, following the theme of §5.1, is to lift the notion of Euler characteristic from a cell complex to an arbitrary (finite, finite-dimensional) sequence $C_\bullet$ of vector spaces:

$$\chi(C_\bullet) := \sum_k (-1)^k \dim\, C_k. \tag{5.13}$$

The alternating sum is a binary variant of exactness. A short exact sequence of vector spaces $0 \to A \to B \to C \to 0$ has $\chi = 0$, since $C \cong B/A$. By applying this to individual rows of a short exact sequence of (finite, finite-dimensional) chain complexes,

$$0 \longrightarrow A_\bullet \longrightarrow B_\bullet \longrightarrow C_\bullet \longrightarrow 0$$

one sees that χ of this sequence also vanishes: $\chi(A_\bullet) - \chi(B_\bullet) + \chi(C_\bullet) = 0$. This, then, provides the key to understanding why the Euler characteristic is a homological invariant. The following lemma is the homological version of the Rank-Nullity Theorem (Lemma 1.10):

Lemma 5.17. *The Euler characteristic of a chain complex $C_\bullet$ and its homology $H_\bullet$ are identical, when both are defined.*

Proof. From the definitions of homology and chain complexes, one has two short exact sequences of chain complexes, arranged like so:

$$\begin{array}{ccccccccc} 0 & \longrightarrow & \boxed{B_\bullet} & \longrightarrow & \boxed{Z_\bullet} & \longrightarrow & \boxed{H_\bullet} & \longrightarrow & 0\,. \\ 0 & \longrightarrow & Z_\bullet & \longrightarrow & C_\bullet & \longrightarrow & B^-_\bullet & \longrightarrow & 0 \\ & & \boxed{+} & & \boxed{-} & & \boxed{+} & & \end{array}$$

Here, $B^-_\bullet$ is the shifted boundary complex: $B^-_k := B_{k-1}$. By exactness, the Euler characteristic of the sum of these two sequences is zero. The Z terms cancel. The B terms cancel, since $\chi(B^-_\bullet) = -\chi(B_\bullet)$. This leaves $\chi(H_\bullet) - \chi(C_\bullet) = 0$. ⊙

Corollary 5.18. *For a compact cell complex X with subcomplexes A and B,*

$$\begin{aligned} \chi(X) &= \sum_k (-1)^k \dim\, H_k(X) \\ \chi(A \cup B) &= \chi(A) + \chi(B) - \chi(A \cap B) \\ \chi(X - A) &= \chi(X) - \chi(A) \end{aligned}$$

Furthermore, χ is a homotopy invariant among this class of spaces.

These results follow from applications of Lemma 5.17 to (1) the chain complex for cellular homology; (2) the Mayer-Vietoris sequence; and (3) the long exact sequence of the pair (X, A) respectively, the last requiring a little excisive effort to relate $C_\bullet(X, A)$ to $X - A$.

5.10 Lefschetz index

There is a generalization of Euler characteristic from spaces to self-maps. For any chain map $\varphi_\bullet : C_\bullet \to C_\bullet$ on a finite-dimensional chain complex $\mathcal{C}$ over a field $\mathbb{F}$, define the **Lefschetz index** as the graded alternating sum of the traces of chain maps $\tau(\varphi_\bullet) := \sum_k (-1)^k \mathrm{trace}\,(\varphi_\bullet : C_k \to C_k)$.

The analogue of Lemma 5.17 holds: the alternating sum of traces of $H(f)$ equals the alternating sum of traces of $f_\bullet$ via the same argument. This index, like the Euler characteristic which it mimics, is intimately connected to the question of fixed points, not of vector fields, but of self-maps in general. For a self-map $f : X \to X$ of a space

X, one defines its Lefschetz index as, equivalently:

$$\tau_f := \tau(f_\bullet) = \sum_k (-1)^k \mathsf{trace}\,(H(f)\colon H_k X \to H_k X)\,. \tag{5.14}$$

Theorem 5.19 (Lefschetz Fixed Point Theorem). *For X a finite (thus, compact) cell complex, any map $f : X \to X$ must have a fixed point if $\tau_f \neq 0$.*

Proof. The technical portion of the proof (omitted) is to show that f may be approximated by a cellular map (also labeled f) for a suitably subdivided cell structure on X in such a way that the approximation also has empty fixed point set. By compactness and the lack of fixed points, this subdivided cellular map sends each n-cell to a *different* cell. Thus, the trace of the chain map $f_\bullet : C_\bullet \to C_\bullet$ vanishes and $\tau_f = 0$. The approximation step does not change the action on homology or (therefore) Lefschetz index. ⊙

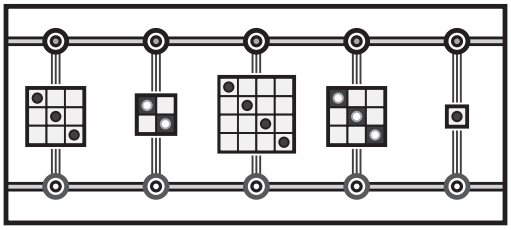

The theorem can be extended greatly. Assuming that the image $f(X)$ is contained in a compact subset of X, the theorem holds for any X homeomorphic to a neighborhood retract in some Euclidean space. [3] Moreover, the theorem holds not only for maps but for multivalued maps $F : X \rightrightarrows X$ with the appropriate modifications: see §5.11 and §7.7. Finally, the computation of the Lefschetz index can be reduced to a sum over the fixed point set of f: see §7.7. This hints at a different proof of Theorem 5.19, since the sum over the empty set vanishes.

The Lefschetz theorem provides a simple proof of Theorem 3.3, that any vector field on a manifold M with $\chi \neq 0$ has at least one fixed point. If a nonvanishing vector field existed, the time-ϵ map of the flow (for ϵ sufficiently small) would be a map of M without fixed points that is a small perturbation of the identity map. The Lefschetz index of this map is $\tau(\mathsf{Id}) = \chi(M)$, since the identity map on $H_\bullet$ has trace equal to the dimension of $H_\bullet$. Indeed, it follows that for any Euclidean neighborhood retract X with $\chi \neq 0$, any map $f : X \to X$ homotopic to the identity has a fixed point.

5.11 Nash equilibria

Among the many applications of fixed point theorems, few are as celebrated as that of the existence of **Nash equilibria** in multiplayer games. The following is a terse rendition.[4] Consider a collection of N players, each of whom can choose from a finite set A_i, of *pure* strategies in, say, a game or an economy. The **payoff function**, $f : \prod_i A_i \to \mathbb{R}^N$, records the numerical return to each of the N players upon executing

[3] Specifically, $Y \subset \mathbb{R}^n$ is a neighborhood retract if there is a retraction $r\colon U \to Y$ for U an open neighborhood. Any X homeomorphic to such a Y is called an **ENR**, or Euclidean neighborhood retract.

[4] Nash' 1950 article – less than a page – is terser still.

the strategies chosen in each A_i. Naturally, each player wishes to maximize payoff. A **Nash equilibrium** for this game is a choice of strategies x^* such that for each i, $f_i(x^*)$ is maximal with respect to varying all inputs except x_i^* – that is, each player has chosen an optimal strategy with respect to fixing everyone else's [known] strategies. The existence of a Nash equilibrium means, in principle, that *everyone is content*. Such an equilibrium may or may not exist.

The insightful step is to allow for *mixed* or probabilistic strategies. Each player chooses from his set of strategies A_i according to some fixed probability distribution, playing repeatedly. This has the effect of generating a **strategy space**, X_i, for each player, corresponding to the space of probability distributions on A_i: each X_i is a simplex. The payoff function extends by linearity to a continuous map $f : \prod_i X_i \to \mathbb{R}^N$, recording the *expected* return to each of the N players. Nash showed that for mixed strategies, an equilibrium always exists [239]. His initial proof used the Brouwer fixed point theorem (Theorem 4.26) in a manner not unlike that of Example 4.27; he quickly converged to a simpler proof via a better fixed point theorem for multivalued functions. The *Kakutani fixed point theorem* [187] states that for $F\colon \mathbb{D}^n \rightrightarrows \mathbb{D}^n$ a multivalued map whose graph is closed in $\mathbb{D}^n \times \mathbb{D}^n$ and has $F(x)$ convex for each x, then there exists a fixed point – some x satisfying $x \in F(x)$.

Nash's proof follows. For mixed strategies $x, x' \in \prod_i X_i$, one says that x' *counters* x if, for each i, the i^{th} strategy of x' is f-optimal with respect to all the not-i strategies of x: that is, the player-i payoff,

$$f_i(\ldots, x_{i-1}, x_i', x_{i+1}, \ldots),$$

is maximal with respect to varying only x_i'. Consider the multivalued map $F\colon \prod_i X_i \rightrightarrows \prod_i X_i$ which sends x to the set of countering strategies $\{x'\}$. This satisfies the criteria for the Kakutani theorem, since the domain is homeomorphic to some $\mathbb{D}^n$, the images of F are convex, and the graph of F is closed (via continuity of f). The existence of a self-countering strategy – a Nash equilibrium by definition – follows via Kakutani.

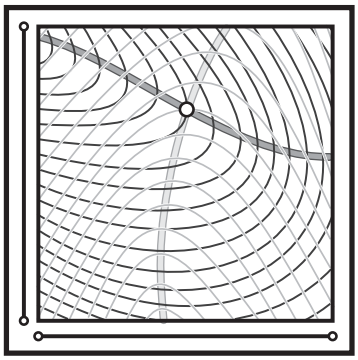

The classical Nash theorem has been extended in numerous ways. One such extension is to more general strategy spaces X_i. It is possible to handle this setting with the Lefschetz Theorem [289], relaxing the conditions on the strategy space and on images of F. Let $X = \prod_i X_i$ be a product of reasonable spaces such as cell complexes.[5] Assume $F\colon X \rightrightarrows X$ a multivalued map whose graph is closed in $X \times X$ and whose images $F(x)$ are acyclic (have $\tilde{H}_\bullet = 0$). It follows that the two factor projection maps $p_1, p_2\colon X \times X \to X$ can be used to define a Lefschetz index for F via $\tau_f = \tau(H(p_2 \circ p_1^{-1}))$, since $H(p_1)$ is an isomorphism. A multi-valued extension of the Lefschetz Theorem [106, 164] states that F must have a fixed point when $\tau_f \neq 0$. This implies the existence of the classical Nash equilibrium (in that case, $\tau_f = 1$), but it also allows for strategy spaces which are noncontractible, if the conditions on τ_f are met.

[5] The technical conditions can be relaxed to an ANR – absolute neighborhood retract – and can be made to work even in infinite-dimensional settings.

The lesson here – not for the last time – is that a qualitative theorem which requests convexity or piecewise-linearity can be persuaded to relax given the appropriate invocation of homology.

5.12 The game of Hex

Nash equilibria arise frequently in the context of mathematical games, itself a fruitful field of topological intricacies. The **game of Hex** is a classical game of topological type. The traditional version of the game is played on a rhombus-shaped board with a uniform hexagonal tiling. The two players (traditionally *black* and *white*, representing colors of markers used in the game) alternate laying down one marker of their color on an open hex cell on the board, filling that cell.

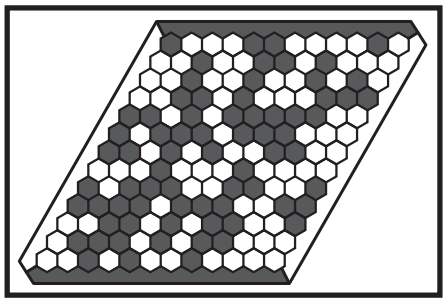

The player goals are, respectively, to build a connection from the top-to-bottom (black) or left-to-right (white) of the board, each player trying to win while blocking the other. Note that the corner pieces of the board border two sides. It is a classical result that this game always ends with one and only one winner. While it is obvious that this is a topological result, a clear topological proof is a worthy exercise. The Brouwer fixed point theorem was initially used by Nash and then Gale [138] to prove existence of a winner (Gale gave also a converse proof of the fixed point theorem via the Hex Theorem). The following is a different approach that may help the beginner learn to work with diagrams and exact sequences. It has the virtue of permitting very general (though not arbitrary) playing boards.

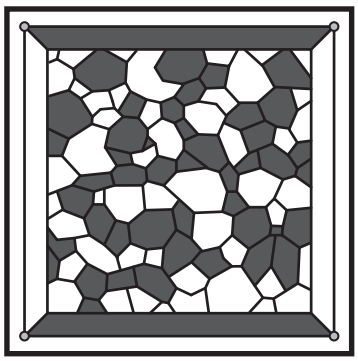

Let $D = [0,1]^2$ be a Euclidean square, outfitted with a definable cell decomposition satisfying the following: (1) there are a disjoint pair of (black) 2-cells, B_0, containing the top and bottom edges, and another (white) disjoint pair, W_0, containing the left and right edges; (2) each vertex in the cell structure has degree exactly three – there are three edges that terminate at each vertex. Note that the four corners of the square satisfy this condition. Players alternate choosing 2-cells which are then colored white-or-black depending on the player. Since the cell decomposition is finite, each game must terminate, and it suffices to consider what happens when all 2-cells are thus colored. Let B denote the union of the closure of each black 2-cell (including boundary cells B_0), and let W be the corresponding white subset. The classical Hex Theorem translated into this setting states that *either* there is a path in B connecting the two components of B_0 *or* a path in W connecting the two components of W_0, but *not both*.

Theorem 5.20 (Hex Theorem). *This game of Hex has one and only one winner.*

Proof. Clearly, $B \cup W$ covers D and $S = B_0 \cup W_0$ deformation retracts to ∂D. The pairs (B, B_0), (W, W_0), and (D, S) fit together via the relative Mayer-Vietoris sequence and the long exact sequences of pairs (§5.3) into a commutative diagram:

$$\begin{array}{ccccccc}
H_2(D,S) & \longrightarrow & H_1(B\cap W, B_0\cap W_0) & \longrightarrow & H_1(B,B_0)\oplus H_1(W,W_0) & \longrightarrow & H_1(D,S) \\
\uparrow & & \uparrow & & \uparrow \quad \uparrow & & \uparrow \\
H_2(D) & \longrightarrow & H_1(B\cap W) & \longrightarrow & H_1(B)\oplus H_1(W) & \longrightarrow & H_1(D) \\
 & & \uparrow & & \uparrow \quad \uparrow & & \\
 & & 0 = H_1(B_0\cap W_0) & & 0 = H_1(B_0)\oplus H_1(W_0) = 0 & &
\end{array}$$

Horizontal rows are Mayer-Vietoris; vertical columns are sequences of pairs; field $\mathbb{F}$ coefficients are used. Only the relevant portions are displayed. The entire right column is zero, as is $H_2(D)$ in the lower-left entry. All vertical maps are injective by exactness. Note the lone copy of $\mathbb{F}$ in the upper-left corner, due to $H_2(D, S)$: this *is* the source of the unique solution to the game. To explain: the goal is to prove the existence of either a path in B connecting B_0 or a path in W connecting W_0. Such a path is *precisely* a relative homology class in one of the factors of $H_1(B, B_0) \oplus H_1(W, W_0)$ which is *not* in the image of the absolute homology $H_1(B) \oplus H_1(W)$ (necessarily injective, by exactness of columns). From exactness of the middle row, $H_1(B \cap W) \cong H_1(B) \oplus H_1(W)$.

The crucial observation is this: $B \cap W$ is a 1-manifold with boundary, since each point of $B \cap W$ is either (1) along a 1-cell of D; (2) at a degree-three vertex of D with two-out-of-three cells of one color and the third another; or (3) one of the four corners of D. The *only* boundary points of $B \cap W$ are these four points. Therefore, $B \cap W$ consists of N circles and exactly two compact intervals, with endpoints the four corners of D. The homology $H_1(B \cap W) \cong \mathbb{F}^N$ injects into $H_1(B \cap W, B_0 \cap W_0)$, which must have dimension $N + 2$, since the two intervals become relative cycles. By reducing the above diagram to dimensions of the homologies and adding one more (vanishing) term to the left of the top row, one obtains:

$$\begin{array}{ccccccccc}
0 & \longrightarrow & 1 & \longrightarrow & N+2 & \longrightarrow & M & \longrightarrow & 0\,, \\
 & & \uparrow & & \uparrow & & \uparrow\uparrow & & \uparrow \\
 & & 0 & \longrightarrow & N & \longrightarrow & N & \xrightarrow{\cong} & 0
\end{array}$$

with vertical maps injective. The top row is short-exact and implies that $M+1 = N+2$; thus $M = N + 1$, *i.e.*,

$$\dim\left(H_1(B, B_0) \oplus H_1(W, W_0)\right) = \dim\left(H_1(B) \oplus H_1(W)\right) + 1.$$

There exists exactly one relative homology class for *either* $H_1(B, B_0)$ *or* $H_1(W, W_0)$ which is not an absolute homology class: the winner. ⊙

5.13 Barcodes and persistent homology

The capstone applications of this chapter comprise a short survey of the exciting work being done in **topological data analysis** using sequences and homologies. The motivation is that, for a parameterized family of spaces (*e.g.*, Vietoris-Rips complexes) modeling a point-cloud data set, qualitative features which persist over a larger parameter range have greater statistical significance. This branch of applied topology is advancing very rapidly; see [105, 306] for initial primary works, [55, 103, 144] for surveys, and [104, 305] for texts. The work described here spans contributions of Carlsson, de Silva, Edelsbrunner, Harer, Zomorodian, and many others, viewed from a representation-theoretic aspect.

Consider a sequence of spaces X_i with maps:

$$X_0 \longrightarrow X_1 \longrightarrow \cdots \longrightarrow X_{N-1} \longrightarrow X_N. \tag{5.15}$$

The sequence may be finite, as shown, or infinite. This is motivated by a sequence of Vietoris-Rips complexes of a set of data points with an increasing sequence of radii $(\epsilon_i)_{i=0}^N$, in which case the maps are inclusions. This topological sequence is converted to an algebraic sequence by passing to homology and invoking functoriality:

$$H_k(X_0) \longrightarrow H_k(X_1) \longrightarrow \cdots \longrightarrow H_k(X_{N-1}) \longrightarrow H_k(X_N). \tag{5.16}$$

The induced homomorphisms on homology encode local topological changes in the X_i: the question is, what are the *global* changes? A homology class in $H_\bullet(X_i)$ is said to *persist* if its image in $H_\bullet(X_{i+1})$ is also nonzero; otherwise it is said to *die*. A homology class in $H_\bullet(X_j)$ is said to be *born* when it is not in the image of $H_\bullet(X_{j-1})$. This is most easily seen in the context of H_0, classifying connected components of a space. In the context of increasing Vietoris-Rips complexes of a point cloud this sequence of homologies in grading zero gives information about *clustering* of points.

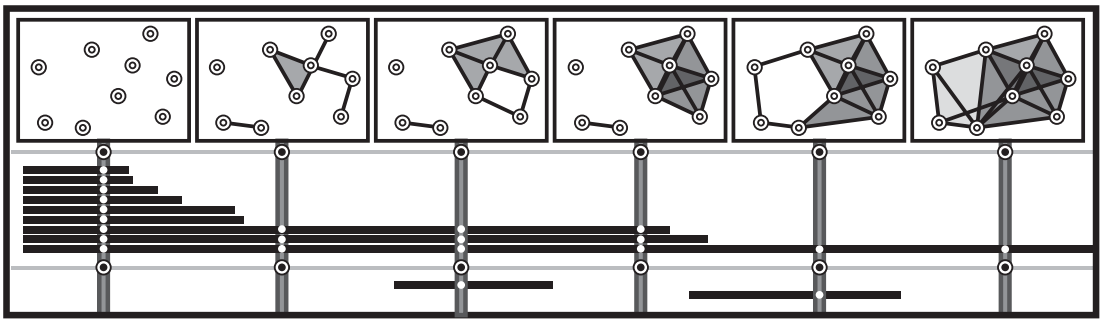

This language of birth, death, and everything-in-between is convenient, but suspiciously informal for mathematics. The injection of a little representation theory yields a principled approach that fulfills intuition in a manner consistent with the themes of this chapter. Consider the generalization of (5.16) to an arbitrary $\mathbb{Z}$-graded sequence of $\mathbb{F}$-vector spaces and linear transformations:

$$V_\bullet = \quad \cdots \longrightarrow V_{i-1} \longrightarrow V_i \longrightarrow V_{i+1} \longrightarrow \cdots. \tag{5.17}$$

Two such sequences $V_\bullet$ and $V'_\bullet$ are said to be *isomorphic* if there are isomorphisms $V_k \cong V'_k$ which commute with the linear transformations in $V_\bullet$ and $V'_\bullet$ as in Equation (5.1). The simplest such sequence is an **interval indecomposable** of the form

$$I_\bullet = \quad \cdots \longrightarrow 0 \longrightarrow 0 \longrightarrow \mathbb{F} \xrightarrow{\mathrm{Id}} \mathbb{F} \xrightarrow{\mathrm{Id}} \cdots \xrightarrow{\mathrm{Id}} \mathbb{F} \longrightarrow 0 \longrightarrow 0 \longrightarrow \cdots$$

where the *length* of the interval equals the number of nonzero terms minus 1, so that an interval of length zero consists of $0 \to \mathbb{F} \to 0$ alone. Infinite or bi-infinite intervals are also included. Intervals – indeed, arbitrary linear sequences – can be formally added by taking the direct sum, $\oplus$, term-by-term and map-by-map. Interval indecomposables are precisely *indecomposable* with respect to $\oplus$ and cannot be expressed as a sum of simpler sequences.

The following result is a simple version of a deeper decomposition theorem from representation theory:

Theorem 5.21 (Structure Theorem for finite linear sequences). *Any linear sequence of finite-dimensional vector spaces and linear transformations decomposes as a direct sum of interval indecomposables, unique up to reordering.*

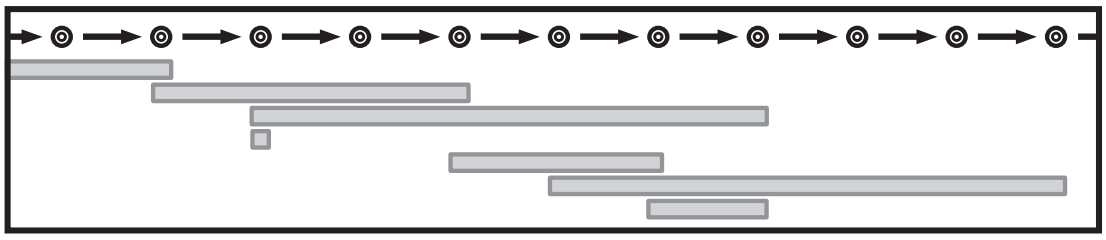

This prompts a graphical language for interpreting basic linear algebra. For example, any linear transformation $\mathbb{F}^n \xrightarrow{A} \mathbb{F}^m$ extends to a linear sequence with two nonzero entries. According to the Structure Theorem, such a sequence may be decomposed in precisely $\min(m, n) + 1$ different ways, depending on the number of interval indecomposables of length 1. The student of linear algebra knows this number of indecomposable intervals of length 1 under the guise of $\mathsf{rank}(A)$.

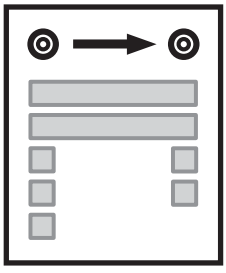

Applying the Structure Theorem to a sequence of homologies (as in Equation (5.16)) in field coefficients yields a pictograph that is called a homology **barcode**. The phenomena of homology class *birth*, *persistence*, and *death* corresponds precisely to the *beginning*, *middle*, and *end* of an interval indecomposable. The barcode provides a simple descriptor for topological evolution: the shorter an interval, the more ephemeral the hole; long bars indicate robust topological features with respect to the parameter. This is salient in the context of point clouds $\mathcal{Q}$ and Vietoris-Rips complexes $\mathsf{VR}_\epsilon(\mathcal{Q})$ using an increasing sequence $\{\epsilon_i\}$ as parameter. For ϵ too small or too large, the homology of $\mathsf{VR}_\epsilon(\mathcal{Q})$ is unhelpful. Instead of trying to choose an *optimal* ϵ, choose them *all*: the barcode reveals significant features.

There is no need to restrict to the case of metric-based simplicial complexes or spaces at all. One can begin with a **persistence complex**: a sequence of chain complexes $\mathcal{P} = (\mathcal{C}_i)$, $i \in \mathbb{Z}$, together with chain maps $x\colon \mathcal{C}_i \longrightarrow \mathcal{C}_{i+1}$. (For notational simplicity, suppress the index subscript on the chain maps x.) Note that each $\mathcal{C}_i = (C_{\bullet,i}, \partial)$ is itself a complex. The **persistent homology** of a persistence complex $\mathcal{P}$ on the interval $[i,j]$, denoted $H_\bullet(\mathcal{P}[i,j])$, is defined to be the image of the homomorphism $H(x^{j-i})\colon H_\bullet(\mathcal{C}_i) \to H_\bullet(\mathcal{C}_j)$ induced by x^{j-i}. The homology barcode is an infographic of persistent homology: $\dim H_k(\mathcal{P}[i,j])$ is equal to the number of intervals in the barcode of $H_k(\mathcal{P})$ containing the parameter interval $[i,j]$.

5.14 The space of natural images

The number, scope, and impact of examples of persistent homology and barcodes are too many to encapsulate: to date, example applications include computer vision [120, 173], Gaussian random fields [4], genetic markers [93], hypothesis testing [43], materials science [199, 213], molecular compression [139], sensor coverage [88], signal processing [107], and much more. One of the first examples of discovering subtle topological structure in a high-dimensional data set came from an examination of *natural images*. A collection of 4167 digital photographs of random outdoor scenes was assembled in the late 1990s by van Hateren and van der Schaaf [294, 236]. Mumford, Lee, and Pederson [237] sampled this data by choosing at random 5000 three-pixel by three-pixel squares within each digital image and retaining the top 20% of these with respect to contrast. The full data set consisted of roughly 8,000,000 vectors in $\mathbb{R}^9$ whose components represent grey-scale intensities. By normalizing with respect to mean intensity and high-contrast images (those away from the origin), and by utilizing a certain norm for contrast, the data set $\mathcal{M}$ was fit on a topological seven-sphere $\mathbb{S}^7 \subset \mathbb{R}^8$.

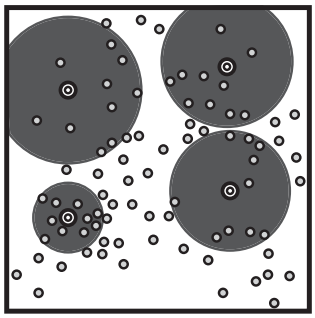

A cursory visualization reveals points distributed seemingly densely over the entire $\mathbb{S}^7$, prompting judicious use of density filtrations. A **codensity** function is used in [59] as follows. Fix a positive integer $k > 0$. For any point x_α in the data set, define $\delta_k(x_\alpha)$ as the distance in $\mathbb{R}^n$ from x_α to k^{th} nearest neighbor of x_α in the data set. For a fixed value of k, δ_k is a positive distribution over the point cloud which measures the radius of the ball needed to enclose k neighbors. Values of δ_k are thus inversely related to the point cloud density. The larger a value of k used, the more averaging occurs among neighbors, blurring finer variations. Denote by $\mathcal{M}[k,T]$ the subset of $\mathcal{M}$ in the upper T-percent of density as measured by δ_k. This is a two-parameter subset of the point cloud which, for reasonable values of k and T, represents an appropriate core.

The first interesting persistent homology computation on this data set occurs at the level of H_1. Taking a density threshold of $T = 25$ at neighbor parameter $k = 300$, with 5000 points sampled at random from $\mathcal{M}[k,T]$, computing the barcode for the first homology H_1 reveals a unique persistent generator. This indicates that the data

set is diffused about a primary circle in the 7-sphere. The structure of the barcode is robust with respect to the random sampling of the points in $\mathcal{M}[k, T]$. In practice, witness complexes from §2.3 provide small enough spaces for computations to be done quickly.

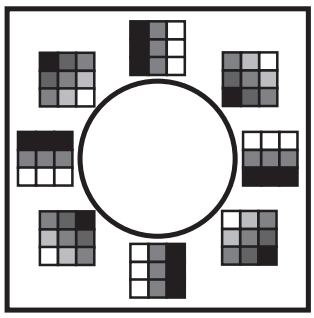

A closer examination of the data corresponding to this primary circle reveals a pattern of 3-by-3 patches with one light region and one dark region separated by a linear transition. This curve between light and dark is linear and appears in a circular family parameterized by the angle of the transition line. As seen from the barcode, this generator is dominant at the threshold and codensity parameters chosen. An examination of the barcodes for the first homology group H_1 of the data set filtered by codensity parameter $k = 15$ and threshold $T = 25$ reveals a different persistent H_1. The reduction in k leads to less averaging and more localized density sensitivity.

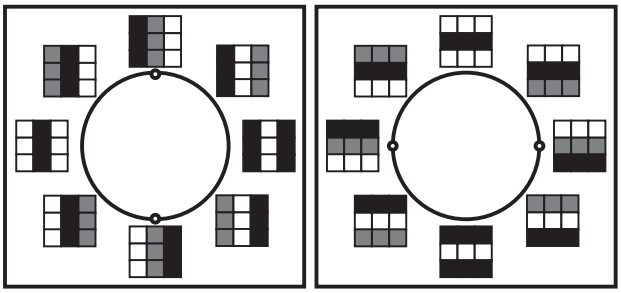

The barcode reveals that the persistent H_1 of samples from $\mathcal{M}[k, T]$ has dimension five. This does not connote the presence of five disjoint circles in the data set. Rather, by focusing on the generators and computing the barcode for H_0, it is observed that, besides the primary circle from the high-k H_1 computation, there are two secondary circles which come into view at the lower density parameter. A close examination of these three circles reveals that each intersects the primary circle twice, yet the two secondary circles are disjoint. As noted in [59], each secondary circle regulates images with three contrasting regions and interpolates between these states and the primary circle. The difference between the two secondary circles lies in their bias for horizontal and vertical stratification respectively.

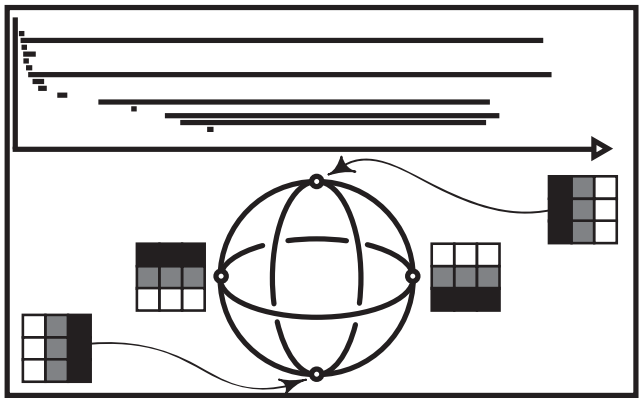

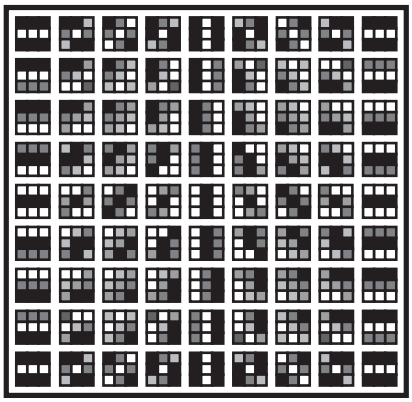

The barcodes for the second persistent homology H_2 are more volatile with respect to changes in density and thresholding. This is not surprising: the lowest order terms in any series expansion are always most easily perceived. However, there is indication of a persistent H_2 generator (in $\mathbb{F}_2$ coefficients) at certain settings of k and T. Combined with the basis of H_1 generators, one obtains predictive insight to the structure of the space of high-contrast patches. At certain density thresholds, the H_2 barcode suggests a two-dimensional completion of the low-k persistent H_1 basis into a Klein bottle K^2. The primary and secondary circles appear with the appropriate intersection properties. A comparison of homology computations in $\mathbb{F}_2$ and $\mathbb{F}_3$ coefficients resolves the ambiguity that $H_2(K^2;\mathbb{F}_2) \cong H_2(\mathbb{T}^2;\mathbb{F}_2)$ and verifies that the persistent surface found is K^2 and not $\mathbb{T}^2$.

5.15 Zigzag persistence

The icon of persistence is the monotone sequence, $\cdots \longrightarrow \bullet \longrightarrow \bullet \longrightarrow \bullet \longrightarrow \cdots$, where arrows connote maps of spaces or chains or the induced homomorphisms on homology. However, other non-monotone sequences are possible and relevant to data management [56]. For example, sequences of the form $\cdots \longrightarrow \bullet \longleftarrow \bullet \longrightarrow \bullet \longleftarrow \cdots$, arise in consistency tests for sampling point clouds as follows.

Given a large set of nodes $\mathcal{Q}$, as in §5.14, it may be infeasible to construct the full Vietoris-Rips complexes and compute persistent homology. A small sample of points is taken (perhaps at random) and a witness complex is constructed as in §2.3: this is indeed the method used for the natural images example in the previous section. To check for accuracy of the sampling, one could compare the homologies resulting from a pair of samples. However, it is not the dimensions of the homology that matter, but the correspondence. Assume that several samples of the data all detect the presence of a single hole. Is it the same hole? Or are there many holes in the true data set, each sampling detecting a different one? Given an ordered sequence of small subsamples, build simplicial complexes X_i based on them. One means of performing comparisons is to build an alternating sequence of spaces and inclusions:

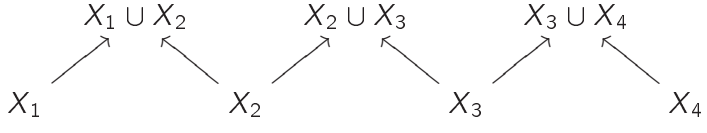

An increase in the complexity of a diagram prompts a concomitant increase in the sophistication of algebra brought to bear. Note that, although not monotone, this sequence is *linear*. A little deeper into representation theory (the classification theorem of Gabriel [136] for quivers of type A_n) reveals that the Structure

Theorem classifies indecomposables of *any* linear sequence,

$$\cdots \longleftrightarrow \bullet \longleftrightarrow \bullet \longleftrightarrow \bullet \longleftrightarrow \cdots,$$

where each $\bullet$ is an $\mathbb{F}$-vector space and each $\longleftrightarrow$ is a linear transformation going either to the left or to the right. An indecomposable interval, as before, is a sequence of copies of $\mathbb{F}$ with identity maps, in this setting respecting the directions of arrows in the original sequence. In classical persistence, all arrows go to the right. In consistency tests, the arrows alternate. The resulting **zigzag persistence** and associated barcodes exist, are computable [61], and are extremely useful. In the context of consistency checks, a long bar in the barcode means that a homological feature is sampled consistently over the sequence.

The moral of these latter sections and of this chapter is that the homology of a sequence is worth more than a sequence of homologies.

Notes

1. The Snake Lemma and 5-Lemma are two of many wonderfully useful general results in homological algebra: for others, see [142, 193, 300]. The subject of homological algebra does not exactly make for colorful reading, but [142], combined with the mantra that complexes are algebraic representations of spaces, is a good place to start.
2. This chapter on exact sequences is just the beginning of a diagrammatic calculus. The next step is to build a **double complex**, a bi-graded 2-d array of chains $\mathcal{C}_{\bullet,\bullet}$ with horizontal ∂ and vertical ∂' chain maps satisfying $\partial\partial' + \partial'\partial = 0$. Such techniques lead quickly to a **spectral sequence**, a structure reminiscent of a book whose pages are double complexes, outfitted with homomorphisms that *turn the pages*. Such structures, though notationally intricate, are quite powerful.
3. The Radon Theorem (in the corollaries of Theorem 5.16) is usually stated in terms of convex hulls of points in Euclidean space, and its proof is often by means of convex geometry. The deeper meaning of the Radon Theorem is, like that of the Helly Theorem in §6.6, topological in nature. All the Borsuk-Ulam type results are greatly generalizable, with the *Colored Tverberg Theorem* being one of the most general [41, 219]. It is to be suspected that such generalizations of the Borsuk-Ulam Theorem are useful – perhaps to economics most readily [33, 180].
4. One should not underestimate the utility of local degree computations as in §5.7; the ability to infer global features from a small number of local measurements is greatly desirable in the sciences.
5. The Cellular Approximation Theorem casually alluded to in the proof of Theorem 5.19 is deep and significant: any map between cell complexes can be homotoped to a cellular map relative to a subcomplex on which the map is already cellular [176].
6. Every proof of the Hex Theorem the author has seen uses two different strategies for the existence and uniqueness of the winner. The (new) proof presented here – though not the simplest or most direct – is pleasant in that a single diagram yields both.
7. The extension of the coverage criterion to higher-dimensional networks is possible, but less clean than the 2-d case presented here. The difficulty resides in specifying the boundary of the domain in question. For regions in the plane, it is not difficult to imagine choosing a cycle or even establishing a *fence* of sensors. In $\mathbb{R}^n$, one must specify a triangulated $n-1$ cycle: this seems awkward to the author, but some applications may permit this condition, in which case a simple modification of the existing proofs

suffices. A persistent homology approach is given in [88].

8. Why was the Čech complex of the sensor cover not used to obtain a coverage criterion? Determining the depths of overlaps of coverage sets requires explicit distances, hence, coordinates; *cf.* the use of Vietoris-Rips complexes versus Čech complexes in point cloud data.
9. As stated, the criterion for computing homological coverage in sensor networks is *centralized*, in the sense that nodes must upload connectivity data to a central computer. More desirable is a *decentralized* or *distributed* computation, performable by nodes communicating with neighbors. Algorithms for decentralized computation of homology have just recently emerged [30, 97, 209, 210, 267].
10. The available perspectives on persistent homology for an author to choose from are daunting, and the treatment in this chapter is necessarily elementary. See the books by Edelsbrunner-Harer [104] and Zomorodian [305] (and, eventually, the book by Blumberg, Carlsson, and Vejdemo-Johansson) for more details from several perspectives. In particular, there are different conventions for representing persistence information beyond the barcodes of §5.13.
11. Since the core idea of persistent homology is little more that iterated functoriality, it is difficult to declare when it was discovered. Early formulations of the notion of persistent homology appear independently in the work of Frosini and Ferri [53, 132, 133], the thesis of Robins [251], and the paper of Delfinado and Edelsbrunner [91]. The subsequent history is one of simultaneous crystallization of theorems, algorithms, and applications about this notion of persistence.
12. From the classification theorem of Gabriel [136], it is shown that not only linear sequences, but certain Dynkin diagrams (of types A, D, and E) have a nice structure theorem classifying indecoposables. Other results from representation theory are poised to impact persistence (such as the Auslander-Reiten quiver [110]). This is especially salient in the context of multi-dimensional persistence – an algebraically challenging scenario [62].
13. Besides being a clear and convenient descriptor for topological data analysis, barcodes possess a very useful **stability** in the context of point-cloud data – nearby point-clouds return barcodes which are close in a certain (*interleaving*) distance: see §7.2 for an example and §10.6 for details. It is this stability that makes barcodes useful in describing noisy point-cloud approximations.