Chapter 8

Homotopy

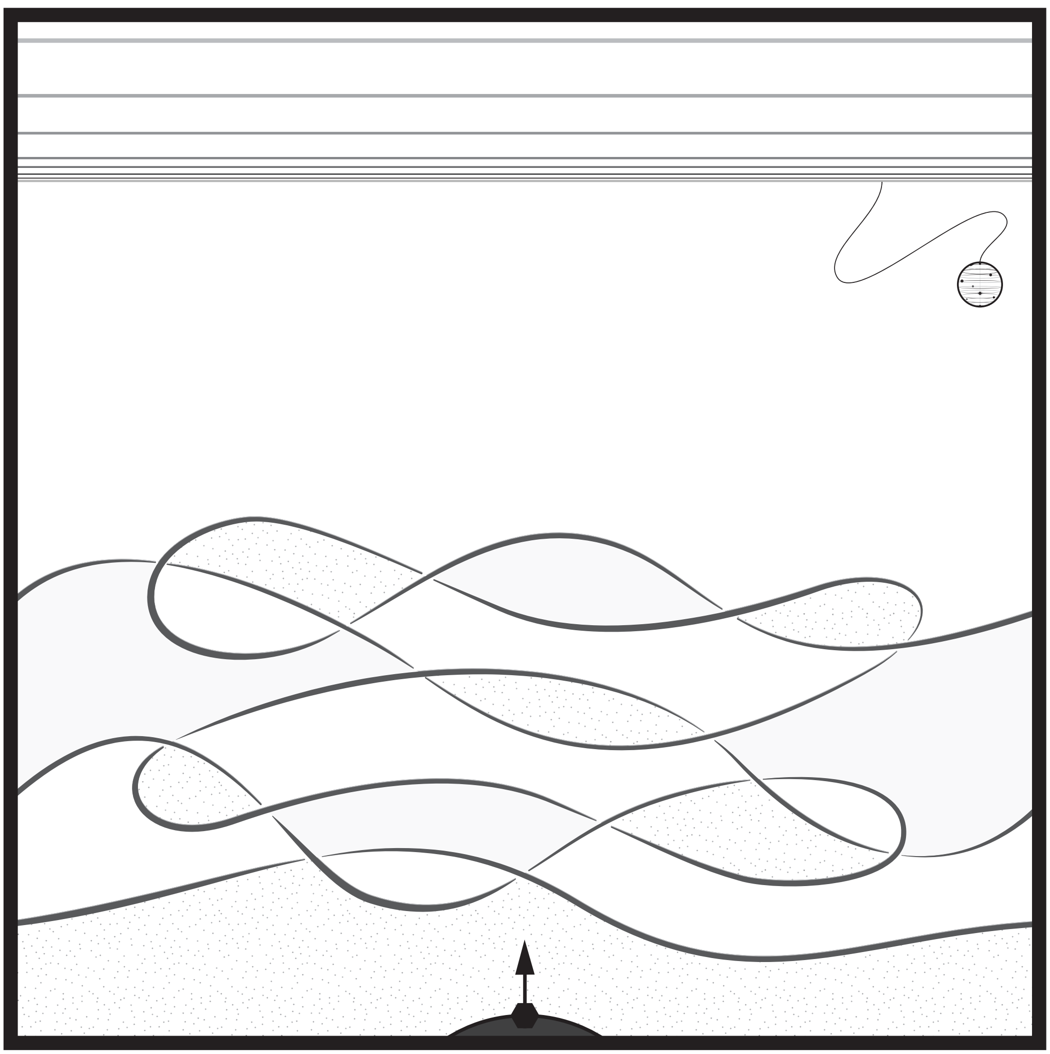

Deformation is the root operation in topology. Homotopy is the primal deformation, leading to homotopy equivalence and then homotopy theory. As compared to co/homology theory, homotopy theory is intuitive, winsome, and largely immune to computational methods. The intuition of homotopy combines with the practicality of co/homology to forge a more complete picture of algebraic topology.

8.1 Group fundamentals

Homological methods are almost entirely comprehensible (and computable) via linear algebra, usually over the reals or over $\mathbb{F}_2$. Homology with $\mathbb{Z}$ coefficients is a bit more subtle, but even the novice is so familiar with this ring that no detailed explanations are required for either intuition or computation. The general setting for homology is best managed using **R**-modules, and this structure has been both alluded to and exploited. In homotopy theory, it is no longer possible to avoid the use of general groups, though, in most every case, it will suffice to work with finitely presented groups described grammatically in terms of generators and relations: see Appendix A.2 for the appropriate keywords.

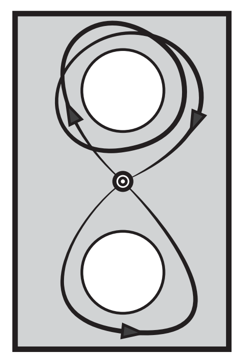

Let X be a space and $x_0 \in X$ be a designated **basepoint**. A *loop based at* x_0 is defined to be a map $\alpha\colon [0,1] \to X$ with $\alpha(0) = \alpha(1) = x_0$. The **fundamental group** $\pi_1(X, x_0)$ is defined on the set of homotopy classes of loops at x_0. That is, two loops α and β are equivalent if there is a homotopy of loops at x_0, $F_t\colon [0,1] \to X$, deforming $F_0 = \alpha$ to $F_1 = \beta$. Note that the basepoint is kept fixed throughout the homotopy, and it is this that permits a group operation given by concatenation of loops as follows:

$$\alpha \bullet \beta\colon [0,1] \to X \quad : \quad t \mapsto \begin{cases} \alpha(2t) & : \ 0 \leq t \leq \frac{1}{2} \\ \beta(2t-1) & : \ \frac{1}{2} \leq t \leq 1 \end{cases} .$$

It is to be checked that this extends to a well-defined associative operation of homotopy classes of loops: $[\alpha]\bullet[\beta] := [\alpha\bullet\beta]$. The **trivial loop** is the constant map $e\colon [0,1] \to \{x_0\}$. A loop is **contractible** if it is homotopic to the trivial loop. The inverse of a loop α is $\alpha^{-1}(t) := \alpha(1-t)$ and provides a true inverse on homotopy classes: $[\alpha][\alpha^{-1}] = [e] = [\alpha^{-1}][\alpha]$.

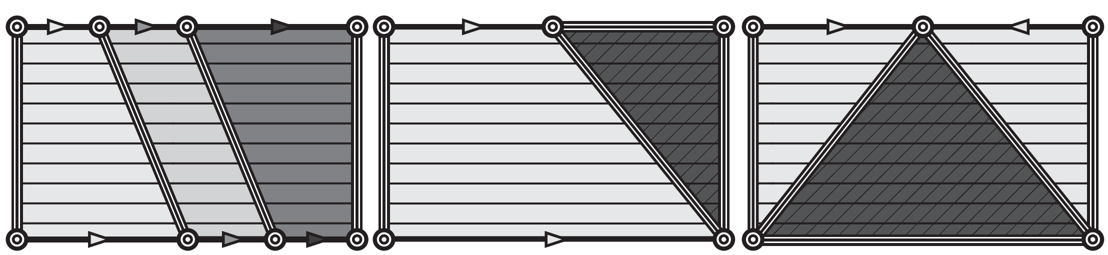

The basepoint is largely irrelevant, in that for X a path-connected space, changing the basepoint from x_0 to x_1 by means of a path $\gamma\colon [0,1] \to X$ leads to an isomorphism $\pi_1(X,x_0) \cong \pi_1(X,x_1)$ via $[\alpha] \mapsto [\gamma \bullet \alpha \bullet \gamma^{-1}]$. The notation $\pi_1(X)$ or π_1 will therefore be used when X is path-connected.

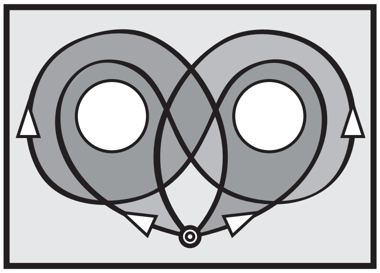

While H_1 comports with a linear-algebraic sensibility, π_1 insists upon the full algebraic regalia of a group: in general π_1 is *not* abelian. This is a frustrating and wonderful fact. Wonderful, in that π_1 yields information not captured by homology. Frustrating, in that *any* (finitely presented) group can arise as π_1 of a space – even such simple spaces as finite 2-dimensional cell complexes or compact smooth 4-manifolds. Determining such facts as whether a loop is contractible, or whether two given loops are homotopic, leads to provably uncomputable problems over finitely presented groups. Any computational homotopy questions must be limited to spaces from a suitably inoffensive class. This does not make the theory inapplicable, but it does dampen one's hopes. On the other hand, π_1 is made for working with homotopy theory. Compare the following to the task of proving why simplicial homology or Euler characteristic is a homotopy invariant.

Lemma 8.1. *Fundamental group π_1 is a homotopy invariant of spaces.*

Proof. Let $f\colon X \to Y$ be a map. There is an **induced homomorphism**,

$$\pi(f)\colon \pi_1(X,x_0) \to \pi_1(Y,f(x_0)),$$

that sends the homotopy class of a loop $\alpha\colon [0,1] \to X$ to the homotopy class of the loop $f \circ \alpha$ in Y. One observes that the induced homomorphism is, as in the case of co/homology, **functorial**, respecting identities and composition. For $f_t : X \to Y$ a homotopy of maps, the loops $f_t \circ \alpha$ are all homotopic. Thus, a homotopy-equivalence induces an isomorphism on π_1. ⊙

Example 8.2 (Examples of π_1) ◎

Certain simple spaces have abelian π_1: for example, $\pi_1(\mathbb{S}^1) \cong \mathbb{Z}$, since any loop α to $\mathbb{S}^1$ has a well-defined degree that fixes its homotopy and homology class (*cf.* the Hopf Theorem of §4.12). The annulus and Möbius strip are homotopic to a circle and thus have the same π_1. The 2-sphere $\mathbb{S}^2$ has $\pi_1(\mathbb{S}^2) \cong 1$, since any loop can be homotoped to one which is not onto $\mathbb{S}^2$ and thus factors through a punctured sphere $\mathbb{S}^2 - \star$, which is contractible. In like manner it is shown that spheres $\mathbb{S}^n$ for $n > 1$ are all **simply connected**, meaning that they are connected and have $\pi_1 \cong 1$. One predicts an accord between π_1 and H_1: the torus has $\pi_1(\mathbb{T}^2) \cong \mathbb{Z}^2$, and the real projective plane has $\pi_1(\mathbb{P}^2) \cong \mathbb{Z}_2$. However, nonabelian fundamental groups abound. The plane $\mathbb{R}^2$ with N points removed has π_1 a free group on N generators. Furthermore, every compact closed surface of genus $g > 1$ has a nonabelian fundamental group. These **surface groups** are beautiful, but hyperbolically complex. For example, a compact oriented genus g surface S_g has fundamental group presented as:

$$\pi_1(S_g) \cong \langle x_1, y_1, x_2, y_2, \ldots, x_g, y_g \;:\; x_1 y_1 x_1^{-1} y_1^{-1} x_2 y_2 x_2^{-1} y_2^{-1} \cdots x_g y_g x_g^{-1} y_g^{-1} = 1 \rangle \quad \odot$$

Lemma 8.3. *For X and Y pointed path-connected spaces:*

1. $\pi_1(X \times Y) \cong \pi_1(X) \times \pi_1(Y)$*, the Cartesian product of groups*
2. $\pi_1(X \vee Y) \cong \pi_1(X) * \pi_1(Y)$*, the free product of groups*

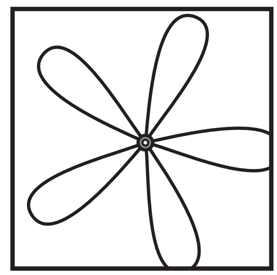

Any finite connected graph is homotopic (by collapsing out a spanning tree) to $\bigvee_1^N \mathbb{S}^1$, a wedge of N circles, which has as π_1 the free product on N elements: the group of all words on N symbols and their inverses, with no relations. Notice how this differs from H_1, which is the abelianization to $\mathbb{Z}^N$ that forgets the order in which one traverses loops. In contrast, the N-torus $\mathbb{T}^N$ has enough "room" to reverse the order of loops and has $\pi_1(\mathbb{T}^N) \cong \mathbb{Z}^N$. First homology H_1 cannot distinguish graphs from tori; π_1 can.

Computing π_1 in general is "easy" in that there is an almost-mechanical procedure for assembling π_1 from pieces, much like the Mayer-Vietoris sequence does for homology. Instead of using exact sequences, a more explicit statement using presentations is preferable.

Theorem 8.4 (Van Kampen Theorem). *Let $U \overset{\iota_U}{\hookleftarrow} U \cap V \overset{\iota_V}{\hookrightarrow} V$ be open and path-connected with finitely-presented fundamental groups:*

$$\pi_1(U) \cong \langle u_i \;:\; r_j = 1 \rangle_{i,j}$$
$$\pi_1(V) \cong \langle v_k \;:\; s_\ell = 1 \rangle_{k,\ell}$$
$$\pi_1(U \cap V) \cong \langle w_m \;:\; t_n = 1 \rangle_{m,n}$$

Then the union $U \cup V$ has fundamental group with presentation:

$$\pi_1(U \cup V) \cong \langle u_i, v_k \;:\; r_j = 1,\; s_\ell = 1,\; \pi(\iota_U)(w_m) = \pi(\iota_V)(w_m) \rangle_{i,j,k,\ell,m}. \qquad (8.1)$$

In other words, one takes the union of the generators and relations of U and V and declares new relations identifying generators $\pi_1(U \cap V)$ as mapped via $\iota_U : U \cap V \hookrightarrow U$ with those mapped via $\iota_V : U \cap V \hookrightarrow V$. The need for U and V open can be relaxed if, *e.g.*, they are subcomplexes of a cell structure. The construction permits induction, and stronger versions can be stated [221]. The difficulties of comparing presentations should not be underestimated: this theorem, though constructive, is not a panacea.

8.2 Covering spaces

Fundamental groups and induced homomorphisms display their power in classification theorems, the best example of which is for covering spaces. For the remainder of this section, all spaces will be assumed path-connected. A **cover** of a space X is a (covering) space $\tilde{X}$ and a map $p\colon \tilde{X} \to X$ which is a **local homeomorphism**. This means that to each $x \in X$, there is a small neighborhood $U \subset X$ of x with the property that $p^{-1}(U)$ is a disjoint union of homeomorphic copies of U, projected by p. The **fibers** $p^{-1}(x)$ are all discrete and have the same cardinality. Covers of X have the same local topology but (often) different global topology.

Example 8.5 (Circles) ⊚
The canonical example of a covering is the map $\mathbb{R} \to \mathbb{S}^1$ given by $t \mapsto e^{2\pi i t} \in \mathbb{S}^1 \subset \mathbb{C}$. However, there are other covers of $\mathbb{S}^1$ – the maps $e^{2\pi i t} \mapsto e^{2\pi n i t}$ for any $n \neq 0 \in \mathbb{Z}$ give an $|n|$-fold covering $\mathbb{S}^1 \to \mathbb{S}^1$. ⊚

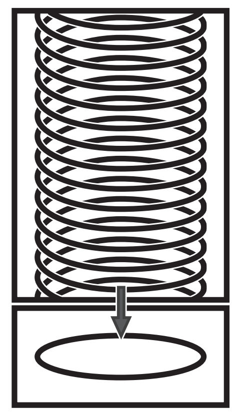

Two covers $p_1\colon \tilde{X}_1 \to X$ and $p_1\colon \tilde{X}_2 \to X$ are said to be *equivalent* if there is a homeomorphism $f\colon \tilde{X}_1 \to \tilde{X}_2$ such that $p_1 = p_2 \circ f$. For example, $e^{2\pi i t} \mapsto e^{\pm 2\pi n i t}$ are equivalent covers from $\mathbb{S}^1$ to $\mathbb{S}^1$ via the antipodal map. A **universal cover** is a cover $p : \tilde{X} \to X$ with $\tilde{X}$ simply connected. For all reasonable (connected and semi-locally simply connected) spaces X, a universal cover exists and is unique up to equivalence. Hence, if X is simply connected, it is its own universal cover, as is the sphere $\mathbb{S}^n$ for $n > 1$ – it has no nontrivial covers. On the other hand, $\mathbb{S}^n$ for $n > 1$ *is* the universal cover of lots of interesting quotient spaces, such as 3-dimensional **lens spaces** with finite π_1. For example, $\mathbb{S}^n \to \mathbb{P}^n$ is a double cover (the fiber has cardinality 2).

The problem of distinguishing between a space and a cover is salient in robot navigation and mapping. Assume a robot that moves about in an unknown environment and can use primitive vision/sensing to patch together explored local neighborhoods into a rough map of the environment based on (random or deterministic) exploration. While local patching is possible, global recurrence is more problematic (*cf.* being lost in the woods – *"Have we been here before?"*). It is a persistent problem to determine whether or not the robot has accurately mapped the region or one of its covering spaces.

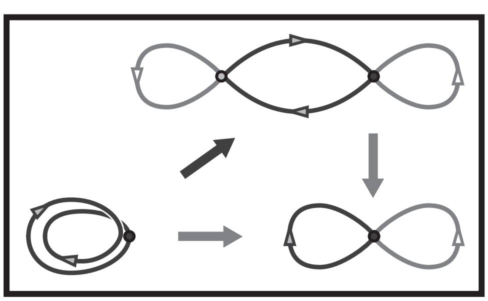

This partially motivates the question of classifying and distinguishing different covers of a fixed space X: this is a neat-and-tidy theory that mirrors the fundamental group perfectly. The reader is encouraged to try to classify all the different covers of $\mathbb{S}^1 \vee \mathbb{S}^1$. This is not an easy exercise, either in the setting of finite or infinite covers. However, the general theory *is* elementary (as elementary as is possible within the whirl of homotopy theory). This depends crucially on a single concept: a *lift*. A **lift** of a map $f\colon Y \to X$ to a cover $p\colon \tilde{X} \to X$ is a map $\tilde{f}\colon Y \to \tilde{X}$ such that $f = p \circ \tilde{f}$. The reader should augment this definition with some examples from the simplest types of covers: covers of $\mathbb{S}^1$, $\mathbb{T}^2$, and $\mathbb{S}^1 \vee \mathbb{S}^1$.

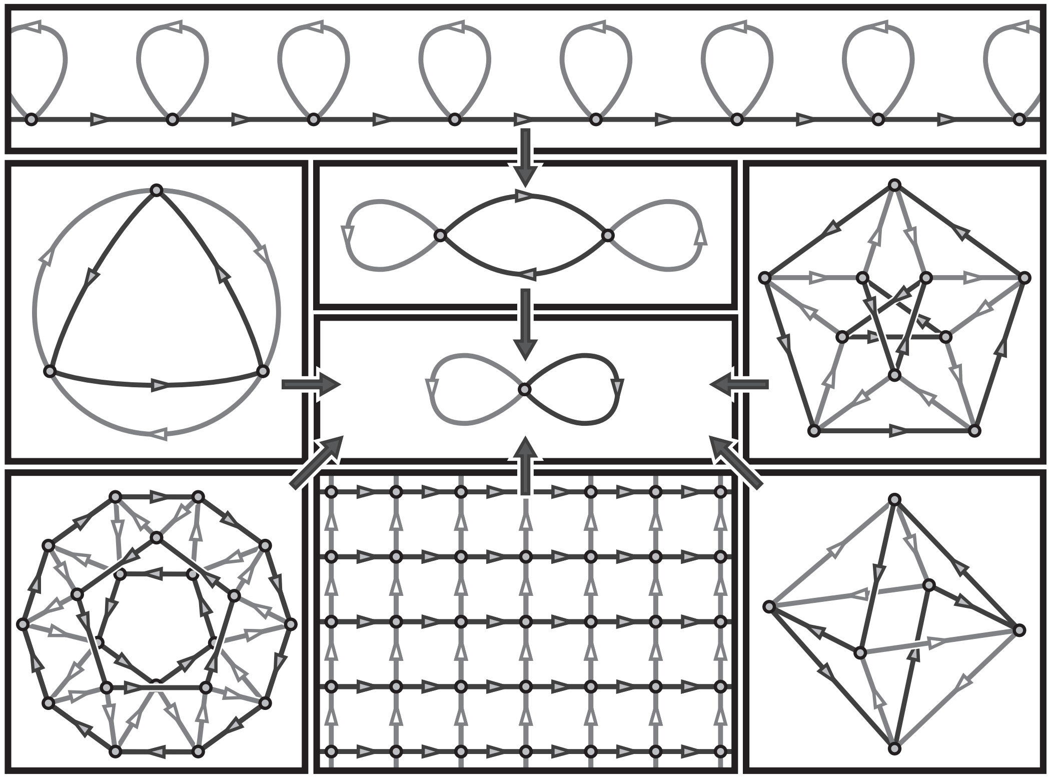

The following encapsulates the main results of covering space theory, with an emphasis on lifts:

Theorem 8.6 (Covering Space Theory). *Let X and Y be path-connected, locally path-connected, and locally simply connected spaces*[1]*, and let $p\colon (\tilde{X}, \tilde{x}_0) \to (X, x_0)$ be a cover.*

1. **[Lifting criterion]**: *A map $f : (Y, y_0) \to (X, x_0)$ lifts if and only if $\pi(f) < \pi_1(p)$: that is, nontrivial loops in the image of f are also nontrivial in the image of p.*
2. **[Homotopy lifting]**: *Any homotopy $f_t\colon (Y, y_0) \to (X, x_0)$ with an initial lift $\tilde{f}_0\colon (Y, y_0) \to (\tilde{X}, \tilde{x}_0)$ lifts uniquely to a homotopy $\tilde{f}_t\colon (Y, y_0) \to (\tilde{X}, \tilde{x}_0)$.*
3. **[Classification]**: *Covers of (X, x_0) up to covering space equivalence are in bijective correspondence with subgroups of $\pi_1(X, x_0)$.*

These results are sharp and tightly connected. Together, they provide a complete understanding of covering spaces (topological objects) in terms of subgroups of π_1 (algebraic objects). Of course, the correspondence is bidirectional, and as often as algebra enlightens topology, topology returns the favor: the best proof that every subgroup of a free group is free is a simple application of covering space theory [176, 218].

Example 8.7 (Euler angles) ⦾

[1] These can be relaxed with care [176].

As per §4.11, the orientation of an object (*e.g.*, an airplane or a *wiimote*) in $\mathbb{R}^3$ is as an element of SO_3 and can be written as an orthogonal matrix with determinant $+1$ or as a point in $\mathbb{P}^3$. It is more common in applications, however, to use angles to describe the object's orientation – in aviation, *e.g.*, one uses *roll*, *pitch*, and *yaw*. These angles – the *Euler angles* or any other choice of three cyclic variables – implicitly define a map $\mathbb{T}^3 \to \mathbb{P}^3$. Since the covers of $\mathbb{P}^3$ are classified by subgroups of $\pi_1(\mathbb{P}^3) \cong \mathbb{Z}_2$, there is only the trivial ($\mathbb{P}^3$) and universal ($\mathbb{S}^3$) cover. As $\mathbb{T}^3$ cannot be a cover, there is *no* good coordinate system that is everywhere a local homeomorphism to $\mathbb{P}^3$: the coordinates cannot have full rank at all image points. Here, π_1 acts both as a means of classification and obstruction. Note, however, that $\mathbb{S}^3$ is a perfectly good cover, meaning that one can faithfully use **quaternions** (the group structure on $\mathbb{S}^3$) without experiencing the same degeneracies. ⦾

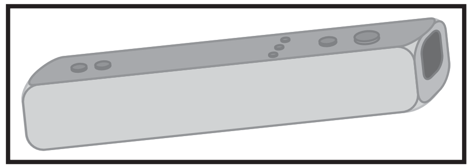

8.3 Knot theory

The fundamental group is well-suited to the theory of knots and links, a beautiful subject for visual topology [260]. Recall from Example 4.24 that a **knot** is an embedding of $\mathbb{S}^1$ into $\mathbb{S}^3$. Two knots are said to be equivalent (or of the same *knot type*) if there is an **ambient isotopy** – a homotopy of homeomorphisms – of $\mathbb{S}^3$ carrying one knot to the other. This fits with the intuition of deforming the strands without cutting or pulling a knot so tight as to cause it to vanish. The **unknot** is a knot equivalent to a standard $\mathbb{S}^1 \subset \mathbb{R}^2 \subset \mathbb{R}^3$.

The topological type of the complement $\mathbb{S}^3 - K$ of a knot is, clearly, an invariant, since an ambient isotopy drags the complement of one homeomorphically to that of the other. Thus, any algebraic-topological invariant provides a potential means of discriminating knot types. One is at first tempted to use homology; however, this is insufficient to the task. Every knot complement in $\mathbb{S}^3$ has the homology type of the circle, since, by Alexander duality, $\tilde{H}_k(\mathbb{S}^3 - K) \cong \tilde{H}^{2-k}(\mathbb{S}^1)$. The fundamental group is a much stronger, though not a complete, invariant.

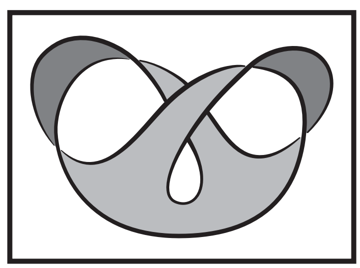

Example 8.8 (Genus) ⦾

The question of knot equivalence has led to a dizzying array of invariants, drawn on tools ranging from combinatorial trickery, covering spaces, Euler characteristic, geometry, Morse theory, Floer theory, and more. Among the simplest of invariants is the **genus** of a knot. A **Seifert surface** of a knot $K \subset \mathbb{S}^3$ is a punctured orientable surface $S \subset \mathbb{S}^3 - K$ embedded in the complement that *spans* the knot ($K = \partial S$). Such an S is homeomorphic to a punctured orientable surface of genus g. The *minimal* such g is defined to be the genus of the knot. This is by definition an invariant, since an ambient isotopy of $\mathbb{S}^3$ deforms a spanning surface along with the knot. Genus is by no means a complete invariant, since many distinct knot types have equal genus. However, genus is *unknot detecting* in the sense the genus of K is zero if and only if K is the unknot. ⊚

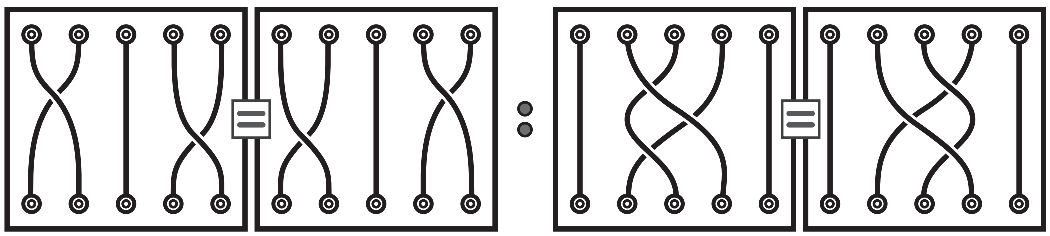

Example 8.9 (Braids) ⊚

Recall from Example 1.8 that one can describe periodic motions of robots (labeled or unlabeled) using loops in a configuration space ($\mathcal{C}^n(\mathbb{R}^2)$ or $\mathcal{UC}^n(\mathbb{R}^2)$ respectively). The description given is Chapter 1 was necessarily *ad hoc*. A better language is now available. The **braid group** on n strands is defined to be $B_n = \pi_1(\mathcal{UC}^n(\mathbb{R}^2))$; the **pure braid group** on n strands is $P_n = \pi_1(\mathcal{C}^n(\mathbb{R}^2))$. These are both, naturally, groups. The identity element is the constant loop; that is, *nobody moves*. Composition in the braid group is concatenation of braids: *first this, then that*. The inverse of a braid reverses motion. The braid group B_n has a clean presentation whose generators σ_i consist of crossing the i^{th} strand over the $(i+1)^{\text{st}}$:

$$B_n \cong \left\langle \sigma_1, \ldots, \sigma_{n-1} \; : \; \begin{array}{c} \sigma_i\sigma_j = \sigma_j\sigma_i \; : \; |i-j| > 1 \\ \sigma_i\sigma_{i+1}\sigma_i = \sigma_{i+1}\sigma_i\sigma_{i+1} \end{array} \right\rangle$$

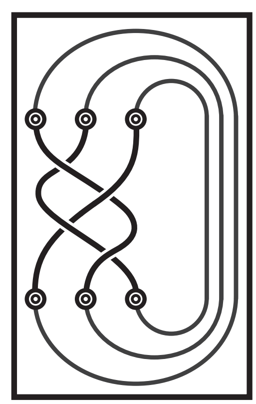

Braids provide not only an efficient language to describe robot motions but also algebraic descriptions of knots and links. A **closed** n**-braid** is the link obtained from a braid in B_n by connecting the points on the *bottom* of the braid to those on the *top* via n strands in the simplest possible manner. Equivalent braids give rise to isotopic closures. A theorem of Alexander [7] confirms the suspicion that every link can be represented as the closure of some braid. The smallest n for which a braid in B_n can be closed to form a given knot is a topological invariant called the **braid index** of the knot.

The π_1-based definition of braid groups extends to braids on any domain. One can consider, *e.g.*, braids on surfaces other than $\mathbb{R}^2$. A great deal of interesting structure is to be found in braid groups of graphs [118, 145]. ⊚

Example 8.10 (DNA and enzyme actions) ⊚

Protein chains, power cords, and DNA strands can coil into conformations one is tempted to call *Gordian*; however, a typical such chain is not a loop, and thus cannot be knotted. That has not prevented knot theorists from investigating knotting and linking phenomena in DNA, which *sometimes* comes in circular substrate molecules – loops. With a combination of electron micrography and gel electrophoresis, it is possible to sort out collections of cyclic DNA strands by knot type. This makes it possible to use knotted DNA as a test bed for determining the action of certain enzymes that aid in recombination. Since knotting, linking, and writhing of the chain prevents a simple *parallel* replication from separating from the parent chain, there must be some agents that aid in disassembly and reassembly of the chain. These are enzymes (*recombinase, topoisomerase, etc.*) whose actions are localized, vital, and largely hidden. By applying selective enzymes and analyzing changes to global knot type, the functionality of these enzymes can be inferred, quantified, and characterized rigorously [285]. ⊙

Example 8.11 (Flowlines) ⊙

Three-dimensional flows exhibit all kinds of knotting. Some of these flows are physical: smoke-rings and other types of vortices (as in, *e.g.*, superconducting fluids) provide beautiful dynamic examples of embedded loops in fluids, some of which are knotted/linked. Numerous authors [14, 16, 131, 232] have contributed to understanding lower bounds on the energy of a perfect fluid flow by means of knotting and linking of the flowlines (*cf.* helicity in Example 6.25), with parallel investigations in magentohydrodynamics: the text [16] is a good resource for this body of ideas.

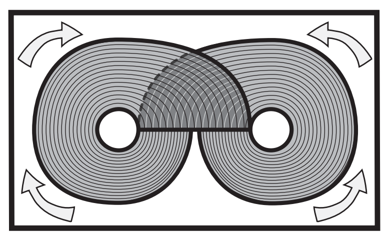

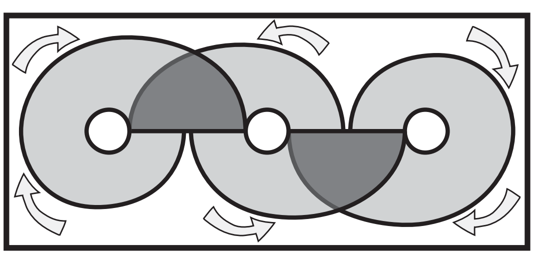

In general, any vector field on $\mathbb{S}^3$ or $\mathbb{R}^3$ may have periodic orbits, each of which is (by uniqueness of solutions to ODEs) an embedded loop – a knot. Together, these form a link of periodic orbits, which may or may not be a finite link. The basic question *"Which link types are possible?"*, even in the context of a sufficiently smooth or tame nonsingular vector field, is delicate. The empty link is possible, thanks to the solution to the **Seifert conjecture** [202]. Any finite link is possible (a simple exercise). For flows exhibiting chaotic dynamics (and thus infinitely many periodic orbits), Birman and Williams [39, 40] showed how to collapse sufficiently hyperbolic invariant sets onto a **template** – an embedded branched surface $\mathcal{T} \subset \mathbb{S}^3$ with a semiflow[2] – in a manner that preserves all knot and link data of periodic orbits. Their seminal work on the geometric **Lorenz attractor**

[2] A semiflow is an action of $\mathbb{R}^+$ instead of $\mathbb{R}$. One can flow forward in time uniquely, but not backward.

showed that although infinitely many knots types exist as periodic orbits in this flow, only certain types of knots and links can arise [39]. In contrast, there exist **universal templates** which contain *all* knots and (finite) links as periodic orbits of the semiflow [143]. These can arise in a number of interesting physical settings – explicit ODEs on $\mathbb{R}^3$ possessing *all* knots and links as solutions [147].

⊙

8.4 Higher homotopy groups

The notation $\pi_1(X)$ for fundamental group foreshadows the higher homotopy groups. The notation does not predict the resulting surprises.

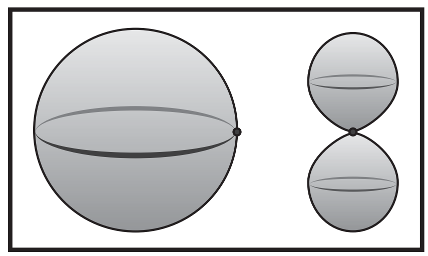

Fix a space X and a basepoint $x_0 \in X$. The **homotopy group** π_n of X at x_0 measures the number of ways to map a sphere $\mathbb{S}^n$ with a fixed basepoint s_0 into X up to (basepoint-preserving) homotopy. That is, π_n consists of homotopy classes of maps $f\colon (\mathbb{S}^n, s_0) \to (X, x_0)$. Note that this reduces to the loop-based definition in the case $n = 1$. For $n = 0$, $\pi_0(X, x_0)$ is a set whose cardinality measures the number of path-connected components of X. For all other $n > 0$, $\pi_n(X, x_0)$ has the structure of a group under the following multiplication operation. Given $f, g\colon (\mathbb{S}^n, s_0) \to (X, x_0)$, define the product $f \bullet g\colon (\mathbb{S}^n, s_0) \to (X, x_0)$ by sending an equator of $\mathbb{S}^n$ (containing s_0) to x_0, and mapping via f on the *upper* hemisphere and g on the *lower*. It is clear that the identity element is represented by the map $\mathbb{S}^n \mapsto x_0$. It is an easy exercise to show that, for connected spaces X, the homotopy groups are isomorphic for different choices of basepoint x_0: as in the case of π_1, basepoints are often suppressed unless explicitly needed.

The definition of homotopy groups is more elementary than that of co/homology. Indeed, as many algebraic topology courses begin with homotopy groups and only later turn to co/homology, the experienced reader may be frustrated at this late-in-time treatment of so fundamental a species. There is good reason to beware homotopy groups as an admixture of the divine and the devilish. A good example of a homotopy group computation is that of a sphere. One begins simply enough: $\pi_k(\mathbb{S}^n)$ is trivial for $0 \leq k < n$, and $\pi_n(\mathbb{S}^n) \cong \mathbb{Z}$. This compares favorably with the homology of spheres: $\pi_k(\mathbb{S}_n) \cong H_k(\mathbb{S}^n; \mathbb{Z})$ for all $1 \leq k \leq n$ (see Theorem 8.14 to come). But what of the higher homotopy groups? These *seem* even simpler than the sometimes problematic π_1. Any topologist who can't prove the following with a picture, isn't:

Proposition 8.12. *For $k > 1$, π_k is abelian.*

The computation of $\pi_3(\mathbb{S}^2)$ is the first hint at the mysterious nature of higher homotopy groups. Surprisingly unlike homology, $\pi_k(X)$ does not necessarily vanish for $k > \dim X$. One is misled by the case of $\mathbb{S}^1$, whose universal cover is contractible.

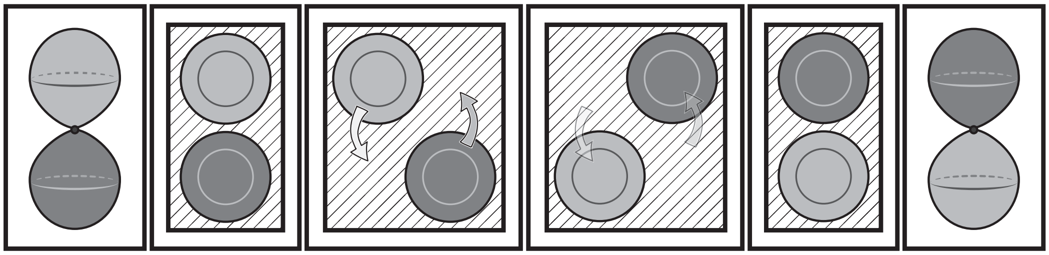

In §8.10, it will be shown that $\pi_3(\mathbb{S}^2) \cong \mathbb{Z}$. Higher homotopy groups do not offer much in the way of bubbly optimism for computational topology: perhaps the largest unsolved problem in algebraic topology is the computation of $\pi_k(\mathbb{S}^n)$ for large values of $k > n$. For example, $\pi_k(\mathbb{S}^2)$ as a function of $k > 0$ is:

$$0\,,\ \mathbb{Z}\,,\ \mathbb{Z}\,,\ \mathbb{Z}_2\,,\ \mathbb{Z}_2\,,\ \mathbb{Z}_{12}\,,\ \mathbb{Z}_2\,,\ \mathbb{Z}_2\,,\ \mathbb{Z}_3\,,\ \mathbb{Z}_{15}\,,\ \mathbb{Z}_2\,,\ \mathbb{Z}_2^2,\ \mathbb{Z}_2\times\mathbb{Z}_{12}\,,\ \mathbb{Z}_{84}\times\mathbb{Z}_2^2,\ \mathbb{Z}_2^2,\ \ldots$$

There is good reason to believe this problem will not be readily solved, not because of a lack of pattern in existing data on $\pi_\bullet(\mathbb{S}^n)$, but, as with much of Mathematics, because of a wild abundance of pattern.

Some computations *are* possible. Like π_1, higher homotopy groups are functorial: a map $f\colon X \to Y$ induces homomorphisms on π_n for all $n > 0$, and homotopic maps yield the same homomorphism (modulo basepoint considerations). Some maps preserve (higher) homotopy groups without necessarily being a homotopy equivalence of spaces. The following is a direct consequence of the lifting criterion of Theorem 8.6:

Corollary 8.13. *Covers induce isomorphisms on π_n for all $n > 1$.*

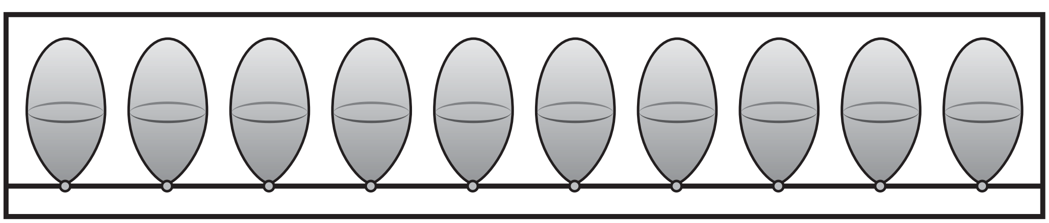

This leads to some interesting computations.

1. All graphs have $\pi_n \cong 0$ for $n > 1$.
2. All surfaces except $\mathbb{S}^2$ and $\mathbb{P}^2$ have $\pi_n \cong 0$ for $n > 1$.
3. $\mathbb{S}^2$ and $\mathbb{P}^2$ have all π_n isomorphic except for $n = 1$.
4. $\pi_2(\mathbb{S}^2 \vee \mathbb{S}^1) \cong \mathbb{Z}^\infty$, since the universal cover has an infinite number of nonhomotopic copies of $\mathbb{S}^2$.

8.5 Biaxial nematic liquid crystals

Recall from Example 4.25, nematic liquid crystals in $\mathbb{R}^2$ and $\mathbb{R}^3$ are composed of molecules whose idealized form is that of an axisymmetric rod. The corresponding singularities in the crystals are completely described by degree theory using homology (in $\mathbb{Z}$ coefficients for the 2-d case and $\mathbb{F}_2$ for 3-d).

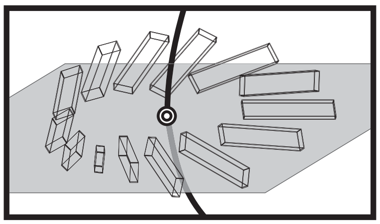

Let the reader note that in cases where the degree is computed from the director field ξ by means of a loop, then the homological degree of $\xi\colon \mathbb{S}^1 \to \mathbb{P}^1$ or $\xi\colon \mathbb{S}^1 \to \mathbb{P}^2$ classifies the singularity type: switching to the fundamental group π_1 returns no new information. However, in the case of a point-defect in $\mathbb{R}^3$, the relevant map of a surrounding sphere gives $\xi\colon \mathbb{S}^2 \to \mathbb{P}^2$. In homology, this degree is $\mathbb{F}_2$-valued, but in homotopy, one has $\pi(\xi)\colon \pi_2(\mathbb{S}^2) \to \pi_2(\mathbb{P}^2)$ which, being a map from $\mathbb{Z} \to \mathbb{Z}$, reveals a finer invariant.

The advantages of homotopy groups become more pronounced in more general liquid crystal structures. One important (though only more recently investigated) class of liquid crystals are the 3-d **biaxial nematics**, whose molecules are not axisymmetric, but rather have the form of a rectangular prism [6]. Recall that for the axisymmetric (nematic) case in $\mathbb{R}^3$, the director field takes values in the quotient of the rotation group SO_3 by the group of symmetries of an axisymmetric rod: $\mathbb{S}^1$. This quotient is clearly $\mathbb{P}^2$. However, in the biaxial setting, the director field takes values in the quotient of SO_3 by D_2, the symmetry group of a rectangle in the plane.

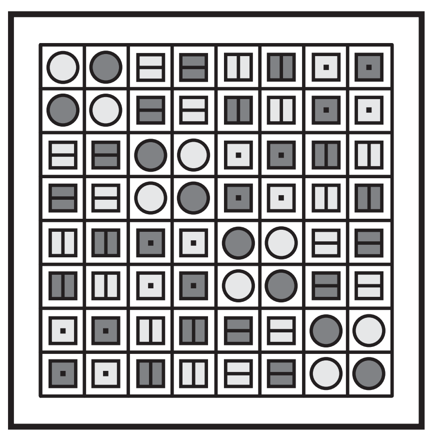

With a bit of work, this space can be shown to be homeomorphic to the quotient of $\mathbb{S}^3$ by Q_8, the **unit quaternions**. This is the (unique) non-abelian group $Q_8 = \{\pm 1, \pm i, \pm j, \pm k\}$ of order eight. This group, familiar from physics and 3-d vector calculus, expresses noncommutativity in the relations $ij = k = -ji$ (and permutations thereof). Because the action of Q_8 on $\mathbb{S}^3$ is regular, the quotient $\mathbb{S}^3/Q_8$ is a cover; since $\mathbb{S}^3$ is simply connected, the quotient has fundamental group $\pi_1 \cong Q_8$. It is intriguing that this director field has noncommutative fundamental group: it provides a much richer dictionary for curves of defect singularities than homology can [6]. Note also that because $\pi_2(\mathbb{S}^3/Q_8) \cong \pi_2(\mathbb{S}^3) = 0$, there are no point-like singularities up to homotopy.

There are many other interesting materials whose internal structure reveals a director field expressible as a quotient of the group SO_3 or SO_2 by some subgroup. These include not only various types of liquid crystals, but also metallic glasses, ferromagnets, and superfluid helium [29], all of which can exhibit disclinations and defects of various types. The best way to categorize these is via induced maps on homotopy groups.

8.6 Homology and homotopy

The relationships between $\pi_\bullet$, $H_\bullet$, and $H^\bullet$, are too many and too deep to encapsulate. One begins with the elementary observation that, while $\pi_1(\mathbb{S}^1 \vee \mathbb{S}^1)$ is a free group $\mathbb{Z} * \mathbb{Z}$, $H_1(\mathbb{S}^1 \vee \mathbb{S}^1)$ is the free-abelian group $\mathbb{Z} \oplus \mathbb{Z}$. The following is one of the few simple results that binds homotopy and homology groups together. It is the natural extension of the pattern seen with spheres.

Theorem 8.14 (Hurewicz Theorem). *There are homomorphisms*

$$\mathrm{Hur}_n \colon \pi_n(X) \to H_n(X; \mathbb{Z}),$$

which, for $n = 1$, is abelianization. If $n > 1$ and $\pi_k(X)$ is trivial for $0 \leq k \leq n-1$, then Hur_n is an isomorphism and Hur_{n+1} is surjective.

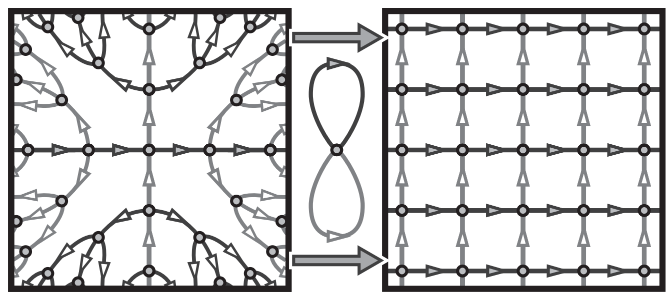

When a mapping between spaces is involved, the Hurewicz theorem pairs well with an extremely powerful result for proving homotopy equivalence of spaces.

Theorem 8.15 (Whitehead Theorem). *If $f \colon X \to Y$ is a map of cell complexes with $\pi(f) : \pi_n(X) \to \pi_n(Y)$ an isomorphism for all n, then f is a homotopy equivalence.*

Since homotopy equivalences are difficult to construct by hand, it is helpful to have implicit tools. One must be careful not to misread the result as saying that spaces with isomorphic homotopy groups are homotopic: it is the mapping and the induced homomorphisms that carry the theorem.

Example 8.16 (Eilenberg-MacLane spaces) ⊚

In homotopy theory, there are numerous approaches for decomposing spaces. One type of building block is an **Eilenberg-MacLane space**. Denoted, $K(\mathbf{G}, n)$, this is a (connected) space, unique up to homotopy type, whose homotopy groups are trivial, with the lone exception that $\pi_n(K(\mathbf{G}, n)) \cong \mathbf{G}$. For example, $\mathbb{S}^1$ is a $K(\mathbb{Z}, 1)$ since the circle has contractible universal cover and all higher homotopy groups vanish. It is not so easy to find Eilenberg-MacLane spaces within the class of finite cell complexes: existence results, though constructive, yield infinite-dimensional spaces as examples. The easiest-to-find finite-dimensional Eilenberg-MacLane spaces are of type $K(\mathbf{G}, 1)$. These include:

1. All knot complements in $\mathbb{S}^3$;
2. Configuration spaces $\mathcal{C}^n(\mathbb{R}^2)$ of points in the plane;
3. All state complexes (§2.11); hence, configuration spaces of graphs.

These spaces serve as the bridge to a surprising relationship between cohomology and homotopy. It is a theorem that the cohomology of a cellular space X is expressible in terms of homotopy classes of maps of X into Eilenberg-MacLane spaces; specifically, $H^n(X;\mathbf{G}) \cong [X, K(\mathbf{G}, n)]$, where $[X, Y]$ denotes a group of basepoint-preserving homotopy classes of maps $X \to Y$ (where the group structure is not entirely obvious: see [176, §4.3]). This is the first hint of the depth of the relationship between co/homology groups and homotopy groups. One simple example of this is that $\pi_1(X) \cong [\mathbb{S}^1, X]$ while $H^1(X;\mathbb{Z}) \cong [X, \mathbb{S}^1]$, thus revealing a type of duality linking π_1 and H^1; *cf.* §6.13. ⊙

8.7 Topological social choice

All of the applications to Economics in this text have thus far relied upon homological tools. Homotopy theory has something to contribute. Economists have long considered the problems associated with social choice and preferences. The following is a topological version of a classical social choice problem. Consider a set of preferences that is topologized as a space, X; examples include preferred prices, budget allocation ratios, or relative rankings of politicians. Given a population of n agents, each with a fixed preference, the state of that population's preferences is an n-tuple of points $\xi \in X^n$. The conversion of individual (local) preferences into a single (global) choice is via a **social choice** map $\Xi\colon X^n \to X$. To reflect reasonable conditions, such a map is required to satisfy the following properties:

1. **Continuity:** Ξ is continuous, so that small shifts in local preferences have small impact on the aggregate preference;
2. **Unanimity:** Ξ is the identity on the grand diagonal in X^n, so that a unanimous vote is accepted; and
3. **Anonymity:** Ξ is invariant under the action of a permutation on the factors of X^n.

The question of existence of a choice map sounds suspiciously like that of existence of an equilibrium in price- or game-theory. Here, instead of universal existence, there is a near-universal non-existence. The following theorem provides the basis for a nonexistence result.

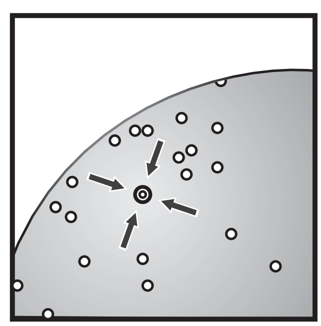

Theorem 8.17 ([102]). *If X admits a social choice map for some $n > 1$, then for each $k > 0$, $\pi_k(X)$ is abelian and uniquely divisible by n.*

Corollary 8.18. *If X is homotopic to a cell complex with finitely-generated $H_\bullet(X;\mathbb{Z})$ and has a social choice map for some $n > 1$, then X must be contractible.*

Proof. *(assuming Theorem 8.17)* As π_1 is abelian, it is isomorphic to H_1 via the Hurewicz theorem. Finiteness and divisibility imply both are zero. By Hurewicz again, $\pi_2 = H_2 = 0$. Induct to show that $\pi_k = 0$ for all higher k. Since X is (homotopic to) a cell complex and the trivial map $X \to \star$ induces isomorphisms on homotopy groups, the Whitehead theorem implies that X is contractible. ⊙

The reader familiar with Arrow's Impossibility Theorem for voting will note the similarities: the Arrow theorem is in the case where X is a finite set of rankings and the anonymity is not enforced (allowing for dictatorial outcomes) [22]. It is interesting to note that there are non-contractible spaces which do admit social choice maps. They are necessarily infinite-dimensional and algebraically subtle: Weinberger [301] shows that the infinite-dimensional real projective space $X = \mathbb{P}^\infty$ admits a social choice map for any n odd, but never for $n > 0$ even. It would also be interesting to connect this work with certain difficulties associated with managing swarms of mobile robots by generating a **consensus** in subspaces of the individual robot configuration spaces (*e.g.*, bearing or pose) [287].

8.8 Bundles

In linear algebra, a surjective linear transformation of vector spaces is characterized by the kernel and the image. Surjective maps $f: X \to Y$ between spaces are potentially wilder. The nicest type of nonlinear surjection has a homogeneity not unlike the linear case: the fiber [kernel] and the base [image] tell all, locally.

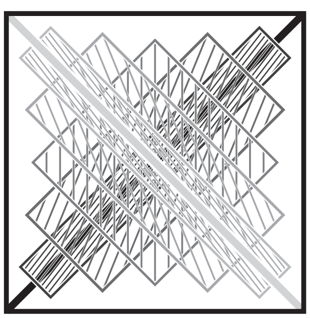

A **(fiber) bundle** is a space E together with a projection map $p: E \to B$ to a **base space** having **fibers** $p^{-1}(b)$ all homeomorphic to some fixed F, so that, on sufficiently small open sets $U \subset B$, $p^{-1}(U) \cong U \times F$. One thinks of the **total space** E of the fiber bundle as a family of F parameterized by B. In the case where $E = F \times B$, one says the bundle is **trivial**. Trivial bundles are all alike: every nontrivial bundle is nontrivial in its own way, bound up in the topology of the base and fiber. Simple examples include:

1. Covering spaces $p: \tilde{X} \to X$ are fiber bundles in which $E = \tilde{X}$ is the cover, $B = X$ the base, and the fiber F is discrete.
2. A Klein bottle is a nontrivial bundle over $\mathbb{S}^1$ with fiber $\mathbb{S}^1$. The difference between this and the trivial bundle $\mathbb{T}^2$ lies in a flip of the fiber.
3. The configuration space $\mathcal{C}^n(M)$ of n points on a connected manifold M is a fiber bundle over base M with fiber $\mathcal{C}^{n-1}(M - \star)$ for $\star$ a point.
4. The 3-sphere $\mathbb{S}^3$ is a nontrivial bundle over $\mathbb{S}^2$ with fiber $\mathbb{S}^1$. This elegant structure is called the **Hopf fibration** and is important in integrable Hamiltonian dynamics [45] and more. Each pair of fibers in $\mathbb{S}^3$ has linking number 1.
5. The **unit tangent bundle**, UT_*M, of a manifold M is the collection of unit tangent spheres in T_*M, expressed as a bundle over M with fiber $F \cong \mathbb{S}^{\dim M - 1}$. This bundle yields the Hopf fibration for $M = \mathbb{S}^2$.

A **vector bundle** is a bundle $p: E \to B$ whose fiber F is a vector space such that addition and scalar multiplication on the fibers extend to continuous maps on all of

E. Examples of vector bundles include tangent and cotangent bundles of a manifold, but others also exist.

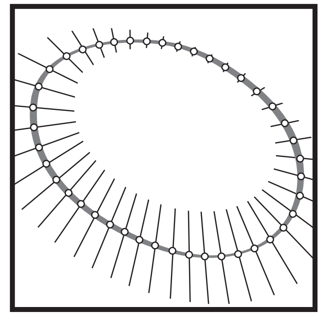

For example, there are, up to a natural equivalence, exactly two vector bundles over $\mathbb{S}^1$ with 1-dimensional fibers: one (the trivial $\mathbb{R}$-bundle) is equivalent to $T_*\mathbb{S}^1 \cong \mathbb{S}^1 \times \mathbb{R}^1$; the other (non-trivial) one is the **twisted bundle** with E homeomorphic to a Möbius band without boundary. Note that both these examples are homotopic to the base $\mathbb{S}^1$. Many of the constructs of this text take on richer meanings in the context of the cohomology of vector bundles [226]. For example, the Euler characteristic of a oriented connected compact n-manifold M lifts to an **Euler class** – a particular cohomology class $e \in H^n(M;\mathbb{Z}) \cong \mathbb{Z}$. This is a generalization of Euler characteristic in that the pairing of e with the fundamental class, the generator of $[M] \in H_n(M;\mathbb{Z})$, yields $e([M]) = \chi(M)$.

Example 8.19 (Fibered knots and magnetic fields) ⊙

A knot $K \subset \mathbb{S}^3$ is said to be **fibered** if the complement is a bundle over $\mathbb{S}^1$ with fibers of $p\colon \mathbb{S}^3 - K \to \mathbb{S}^1$ the soap-film-like Seifert surfaces of Example 8.8. The unknot is fibered (with fibers homeomorphic to a disc). Trefoil knots, as well as the classic figure-8 knot, are also fibered, though in general fibered knots are not common.

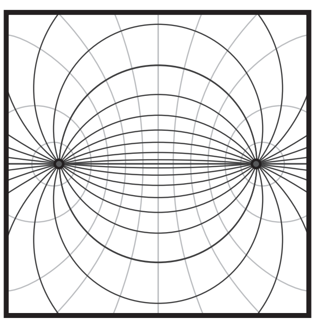

Fibered knots may be related to electromagnetics. A wire bent into a knot K in $\mathbb{S}^3$ emits, upon passing a current through it, an induced magnetic field on the complement. The magnetic vector field *coils* about the wire. For a fibered knot K, the magnetic field is transverse to the fiber bundle near K – the magnetic flowlines cross the fibers. One guesses [40] that the magnetic field is everywhere transverse to the fibers; this seems a reasonable conjecture for *relaxed* embeddings of K. *If* this is true, then any (fiber-transverse) magnetic field induced by a current through a figure-8 knot [143] (and many other fibered knots [148]) possesses closed field lines spanning all possible knot and link types. ⊙

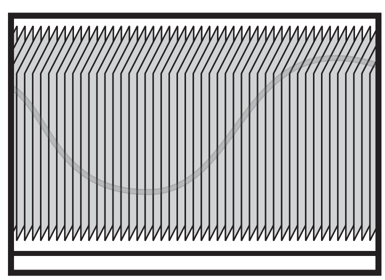

Nontrivial bundles are globally so: by definition, every bundle is locally trivial. What is the measure of nontriviality of a bundle? Euler characteristic is not a good invariant of the bundle structure, since $\chi(E) = \chi(B)\chi(F)$ for *all* bundles E over B with fiber F, whether trivial or not. Homology and cohomology of E can sometimes distinguish between bundles (*cf.* $\mathbb{T}^2$ and K^2), but not always; *e.g.*, a vector bundle is homotopic to its base. The question of triviality is related to that of the existence of sections. Recall from §1.4 that a vector field on a manifold is a section of the tangent bundle. For a general bundle $p\colon E \to B$, a **section** is a map $s\colon B \to E$ satisfying $p \circ s = \mathsf{Id}_B$. Not all bundles have sections: for example, the existence of a non-vanishing vector field on a manifold M is equivalent to the existence of a section of the unit tangent bundle UT_*M, and this does not exist if $\chi(M) \neq 0$. Trivial bundles always have sections. The degree to which sections fail to exist provides a measure of complexity relevant to applications.

8.9 Topological complexity of path planning

Recall from §1.5 the importance of configuration spaces in motion-planning problems for robotics: a configuration space converts the problem of physical motion planning to topological path planning. The insight of Farber [112] is to consider motion planning as a problem parameterized by the start and end configurations, and to crystalize this parametrization in terms of fiber bundles and invariants thereof. Given a path-connected configuration space X, consider the **path space** $\mathcal{P}(X)$, the space of all maps $\gamma\colon [0,1] \to X$ with the usual (compact-open) topology. There is a projection mapping $p\colon \mathcal{P}(X) \to X^2$ taking a path γ to its ordered endpoints $(\gamma(0), \gamma(1))$. With respect to this projection, the path space $\mathcal{P}(X)$ is a fiber bundle (assuming X is locally homogeneous, as in the case of a manifold – else it is a *fibration*: see §8.10). A **path planner** is a section of the path bundle: a continuous map $s\colon X^2 \to \mathcal{P}(X)$ satisfying $p \circ s = \mathsf{Id}$. The following result echoes the obstruction to the inverse kinematic map in §4.11.

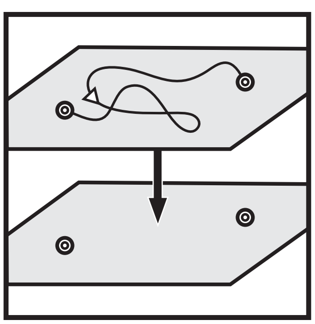

Lemma 8.20 (Farber [112]). *The bundle $p\colon \mathcal{P}(X) \to X^2$ possesses a section if and only if X is contractible.*

Proof. If $s\colon X^2 \to \mathcal{P}(X)$ is a section, then to each pair $(x_0, x_1) \in X^2$ is associated a path $s(x_0, x_1)\colon [0,1] \to X$ connecting x_0 to x_1, continuous in these parameters. Fix $x_1 \in X$ a basepoint, and consider $s(\cdot, x_1)\colon X \times [0,1] \to X$, which is a homotopy from $\mathsf{Id}\colon X \to X$ to the constant map $X \to \{x_1\}$. The argument is reversible: a deformation to a fixed basepoint can be unwound to yield a path planner (with all paths passing through the basepoint). ⊙

This initially discouraging result says that there are, in general, no stable motion-planning algorithms; fixing a motion-plan and varying the endpoints can lead to a discontinuous change in the plan, much in the same way that a portable GPS trip planner exhibits instabilities with respect to small changes in point-of-origin. However, similar negative results for the kinematics of a robot arm (§1.5) have not prevented the ubiquitous use of coupled rotation joints: lack of a continuous section simply means that the problem is more complicated and may exhibit instabilities. To what degree?

Bundles prompt a parametric version of the LS category of §7.9, yielding a notion of complexity relevant to the path-planning problem. The **sectional category**, secat, of a fiber bundle is the minimal number of open sets U_α covering B such that $p^{-1}(U_\alpha) \cong U_\alpha \times F$. This records the minimal number of trivial bundles needed to cover E. As with LScat, secat is difficult to compute, but can be bound by algebraic-topological invariants in the manner of Theorem 7.27. In the context of path-planning, sectional category is the degree of instability of the problem. After the initial question of obstruction – *can* or *cannot do* – there remains the issue of complexity. If one were to solve the stable path-planning problem piece-wise, how many pieces would be required? Farber defines the (reduced) **topological complexity** of the motion planning problem on X as:

$$\mathsf{TC}(X) := \mathsf{secat}\,(p\colon \mathcal{P}(X) \to X \times X) - 1. \tag{8.2}$$

Example 8.21 (Topological complexity) ⊚

The following examples of TC computations are surveyed in [113]:

1. **Spheres:** $\mathsf{TC}(\mathbb{S}^n)$ equals 1 for n odd; 2 for n even.
2. **Surfaces:** For a closed orientable surface S_g of genus g, $\mathsf{TC}(S_g) = 2$ for $g \leq 1$; $\mathsf{TC}(S_g) = 4$ for $g > 1$.
3. **Rotations:** $\mathsf{TC}(\mathsf{SO}_3) = 3$.
4. **Graphs:** For X a graph, $\mathsf{TC}(X)$ equals 0 (if X is a tree), 1 (if $X \simeq \mathbb{S}^1$), or 2 (else).
5. **Projective space:** $\mathsf{TC}(\mathbb{P}^n)$ is known only for $n < 24$. In general, it is *very* hard to compute [116]. ⊚

Theorem 7.27 can be used to give bounds on TC: upper bounds are regulated by dimension, and lower bounds are regulated by cohomology: for X a reasonably tame space, the topological complexity satisfies $\mathsf{cup}(X \times X) - 1 \leq \mathsf{TC}(X) \leq 2\dim X$.

Example 8.22 (Configuration spaces) ⊚

Those examples of TC that come closest to being relevant to robotics are for configuration spaces of points $\mathcal{C}^n(Y)$ or $\mathcal{UC}^n(Y)$. In particular, for a fixed space Y the computation of $\lim_{n\to\infty} \mathsf{TC}(\mathcal{C}^n(Y))$ gives a measure of asymptotic difficulty of collision avoidance. This is in general difficult to compute, as can be guessed from previous examples. It has been shown [115] that for labeled configuration spaces,

$$\mathsf{TC}(\mathcal{C}^n(Y)) = \begin{cases} 2n-2 & : \quad Y = \mathbb{R}^{2m+1} \\ 2n-3 & : \quad Y = \mathbb{R}^{2m} \\ 2\#V^{\text{ess}}(Y) & : \quad Y = \text{ a tree} \end{cases} . \tag{8.3}$$

Here, in the case of Y a tree, $\#V^{\text{ess}}$ stands for the number of **essential** vertices – vertices of degree strictly greater than two. These results are notable: (1) the lack of dependence on m for $\mathsf{TC}(\mathcal{C}^n(\mathbb{R}^m))$ means that the ambient space is largely irrelevant to collision-avoidance complexity; (2) the lack of dependence on n for $\mathsf{TC}(\mathcal{C}^n(Y))$ for Y a tree means that the critical lack of room on a tree collapses the complexity of collision avoidance to what happens at the essential vertices. ⊚

8.10 Fibrations

There is a far-reaching generalization of fiber bundles befitting homotopy theory. In spirit, a fibration is a surjective map $p\colon E \to B$ with the property that (for B path-connected) all fibers $p^{-1}(b)$ are homotopy equivalent, as opposed to homeomorphic. The proper definition does not reference the fibers at all but is founded in behavior of homotopies in manner not unlike covering spaces.

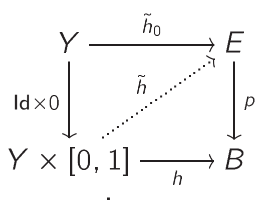

A [Hurewicz] **fibration** is an onto map $p\colon E \to B$ with the following **homotopy lifting property**: for any homotopy h_t of a space Y in B and a lift $\tilde{h}_0\colon Y \to E$ of h_0 to E, there is a lifted homotopy $\tilde{h}_t\colon Y \to E$. In other words, in the appropriate commutative diagram, the dotted lift exists. At first glance, this definition seems obtuse: where is the fiber, F, and why does it have a well-defined homotopy type?

Lemma 8.23. *Fibers of a fibration over a path-connected base have constant homotopy type.*

Proof. Pick a basepoint $b_0 \in B$ and let $F = p^{-1}(b_0)$. Let $\beta\colon [0,1] \to B$ be a path in B from b_0 to b_1. Define $h\colon F \times [0,1] \to B$ via β on $[0,1]$ and via collapse-to-a-point on F. This has a lift $\tilde{h}_0 = \mathsf{Id}\colon F \times \{0\} \to p^{-1}(b_0)$. Thus, by homotopy lifting, $\tilde{h}_t = p^{-1}(\beta(t))$ is a homotopy from F into $p^{-1}(b_1)$. Reverse the path and repeat to show a homotopy equivalence between fibers. ⊙

The idea of defining a fibration not in terms of explicit topological features of F, but rather in terms of implicit response of p to homotopy-lifting, is deep and presages the use of homotopy testing as a means to define and extend other notions. Indeed, this is the pattern for defining **cofibrations** – a dual notion that characterizes maps $\iota\colon B \to E$ in terms of possessing the **homotopy extension property** – see [176] for details. These ingredients – fibrations generalizing projection and cofibrations generalizing inclusion – form the basis for the abstraction of homotopy theory to **model categories** [246, 101]. Upon first pass, the reader should note merely that such generalizations exist and flow from the use of commutative diagrams.

One simple example suffices to demonstrate the power of this approach. The reader may wonder why exact sequences have not made an appearance in the context of homotopy groups: they *are* very important, but come with concomitant subtlety (*cf.* the difference between the Mayer-Vietoris and Van Kampen theorems). The definition of a fibration allows for inference of fiber behavior via the long exact sequence associated to a fibration. Given $p\colon E \to B$ with fiber homotopy type $[F]$, the following sequence is exact:

$$\cdots \longrightarrow \pi_n(F) \xrightarrow{\pi(\iota)} \pi_n(E) \xrightarrow{\pi(p)} \pi_n(B) \xrightarrow{\delta} \pi_{n-1}(F) \xrightarrow{\pi(\iota)} \cdots \qquad (8.4)$$

The maps are constructed as follows. Let $b_0 \in B$ be a basepoint, with $F = p^{-1}(b_0)$. The inclusion $\iota : F \hookrightarrow E$ sends a basepoint $f_0 \in F$ to $e_0 \in E$. The maps $\pi(\iota)$ and $\pi(p)$ are clear: it is the connecting homomorphism δ that is subtle. Consider the diagram whose top row is inclusion of the boundary $(n-1)$-sphere followed by collapse of same. The rightmost vertical arrow represents a class $[\alpha] \in \pi_n(B)$. Commutativity and the homotopy lifting property are used to generate the dotted vertical arrows and to show that the induced class $\delta([\alpha]) \in \pi_{n-1}(F)$ is well-defined.

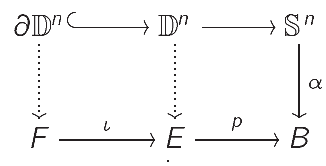

As a simple example, this sequence yields a direct proof that covers induce isomorphisms on π_n for $n > 1$: since the fiber is discrete, every third term of the sequence vanishes, yielding isomorphisms of the remaining pairs. More subtle is the example of the Hopf fibration $p : \mathbb{S}^3 \to \mathbb{S}^2$ with fiber $F = \mathbb{S}^1$; this induces the long exact sequence:

$$\cdots \longrightarrow \pi_n(\mathbb{S}^1) \longrightarrow \pi_n(\mathbb{S}^3) \longrightarrow \pi_n(\mathbb{S}^2) \longrightarrow \pi_{n-1}(\mathbb{S}^1) \longrightarrow \cdots$$

which, since $\pi_n(\mathbb{S}^1) = 0$ for all $n > 1$, implies that $\pi_n(\mathbb{S}^3) \cong \pi_n(\mathbb{S}^2)$ for all $n > 2$. Thus $\pi_3(\mathbb{S}^2) \cong \pi_3(\mathbb{S}^3) \cong \mathbb{Z}$.

8.11 Homotopy type theory

This text has largely avoided the (many and fruitful) applications of topology in the areas of logic and computer science, in part because of the significant overhead of definitions and formal structures required for a proper exposition. Some of the most recent activity in these domains is, however, too compelling not to limn. Let the reader beware that what follows is a severe redaction of a highly intricate and rapidly advancing research program.

Computer programmers use **types** to distinguish different classes of data: the reader may have seen λ-*calculus* or other systems that use formal rule systems to define functions from base or constant types. If this is not familiar, the reader may recall the difficulties stemming from *Russell's paradox* in set theory[3] and the resolution through distinguishing different types of objects (sets, classes, universes, *etc.*).

In (intensional) type theory, there are collections of *terms*, to which can be associated a *type* within a *universe* of types; these can range from variables to logical operators to functional types and more. There is a loose correspondence between type-theoretic constructs, logical constructs, and set-theoretic interpretations thereof. The novel ingredient in the recently-christened *homotopy type theory* is an injection of homotopy-theoretic perspectives. Beginning with the observation that a type is *something like* a space, as opposed to a set, and a term of a certain type is *something like* a point in said space, a homotopy-theoretic sensibility can inform and expand type theory. With the adjunction of one key axiom – the **univalence axiom**[4] of Voevodsky – one can construct a well-defined correspondence of the classical interpretations to the homotopy-theoretic constructs touched upon in this chapter. It is worth reproducing the table of correspondences from [291], which serves as a *Rosetta stone* for the

[3] *"The set of all sets that do not contain themselves..."* caused no small amount of trouble.

[4] The univalence axiom is about a universe $\mathcal{U}$ of types; it states that types which are formally equivalent in $\mathcal{U}$ are identical.

type-theoretic, logical, set-theoretic, and homotopy-theoretic views:

TYPES	LOGIC	SETS	HOMOTOPY
A	proposition	set	space
$a : A$	proof	element	point
$B(x)$	predicate	family of sets	fibration
$b(x) : B(x)$	conditional proof	family of elements	section
$0, 1$	$\bot, \top$	$\varnothing, \{\varnothing\}$	$\varnothing, \star$
$A + B$	$A \vee B$	disjoint union	coproduct of spaces
$A \times B$	$A \wedge B$	set of pairs	product of spaces
$A \to B$	$A \Rightarrow B$	set of functions	function space
$\sum_{(x:A)} B(x)$	$\exists_{x:A} B(x)$	disjoint sum	total space
$\prod_{(x:A)} B(x)$	$\forall_{x:A} B(x)$	product	space of sections
Id_x	equality $=$	$\{(x,x) \colon x \in A\}$	path space $\mathcal{P}(A)$

This table is meant to inspire rather than to define: proper definitions are more involved and require the categorical language of Chapter 10 that is the focus of the remainder of this text.

Notes

1. Knot theory began in earnest with the work of Kelvin and Tait as a problem in fluid dynamics. It was conjectured that atoms were knotted vortex tubes in the æther, and that a classification of knot types would reproduce and refine the periodic table. Poincaré's initial work on algebraic topology was likewise motivated by the desire to understand the dynamics of fluids. It is remarkable how much the field of topology owes to fluid dynamics. (The author, too, came to topology via fluids and dynamics.)
2. The Whitehead theorem has a computationally-friendly corollary: a map between *simply-connected* CW complexes that induces isomorphisms on all *homology* groups is a homotopy equivalence. Simply-connectivity is the key to this result.
3. Homotopy and homology are fundamentally entwined via configuration spaces. For a connected CW complex X, consider the unlabeled singular configuration space X^n/S_n of n unlabeled not-necessarily-distinct points on X. The **Dold-Thom theorem** states that in the limit as $n \to \infty$, the resulting configuration space has π_n isomorphic to $H_n(X;\mathbb{Z})$. This is very deep and *seems* to presage a topological version of statistical physics.
4. Corollary 8.18 on topological social choice was discovered by Weinberger, [301] see also Chichilnisky et al. [66], and Baryshnikov [22]; however, in the process of publication, it was realized that the result is implicit in the 1954 paper of Eckmann, who was not motivated by social choice at all, but rather by problems of generalized means and homotopy theory.
5. Section 8.8 hints at the theory of **characteristic classes**. Given a vector bundle $\pi\colon E \to M$, a characteristic class is an element of $H^\bullet(E)$ carrying data about the bundle. Among the more interesting characteristic classes besides the Euler class are the **Stiefel-Whitney**, **Chern**, **Pontryagin**, and **Thom** classes [226]. The application of homotopy groups to classifying defects in liquid crystals in §8.5 is just the beginning of a number of exciting instances of algebraic topology in condensed matter physics, recent examples of which use characteristic classes to explain experimentally observed

phenomena. This text has skipped most of the applications of topology to physics and fields, not because of lack of interest, but because of the requisite depth: bundles and characteristic classes are the starting point for modern approaches to quantum field theory.

6. It is hard to overstate the importance of fibrations and cofibrations within homotopy theory. For a proper treatment, see, *e.g.*, the excellent text of May [221].
7. The antecedent to the work described in §8.9 is the insightful work of Blum, Shub, and Smale on topological complexity of computations [42, 278]. This work defines a topological complexity for computing roots of a complex polynomial of degree k in terms of the section category of the bundle $\mathcal{C}^k(\mathbb{C}) \to \mathcal{UC}^k(\mathbb{C})$: see also, [17, 85, 295].
8. The cohomology computations hinted at in §8.9 go much deeper than explained in this text. See, *e.g.*, [112, 113, 114] for relations to cohomology operations, Steenrod squares, and the like. The topological complexity of the real projective spaces $\mathbb{P}^n$ is equal to its *immersion dimension* except for $n = 1, 3, 7$ [116] – this is a notoriously subtle and difficult to compute quantity.
9. The critique that **TC** does nothing to help with realistic motion-planning problems is perfectly true and perfectly ignorant of the illumination an obstruction theory brings.
10. The table in §8.11 is reproduced from [291] with some slight notational changes (legally, under their *Creative Commons* license). The reader is encouraged to see the source for more and better explanations.