Heegaard Floer Homology

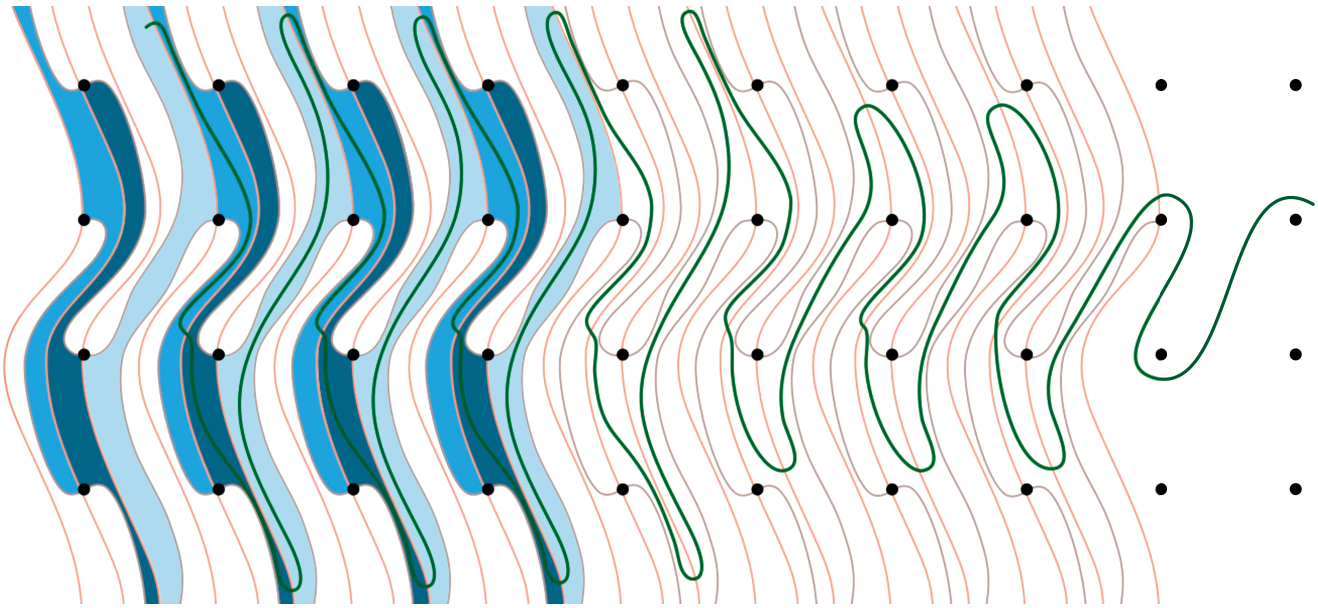

Joshua Evan Greene

Low-dimensional topology encompasses the study of manifolds in dimension four and lower. These are the shapes and dimensions closest to our observable experience in spacetime, and one of the great lessons in topology during the 20th century is that these dimensions exhibit unique phenomena that render them dramatically different from higher ones. Many outlooks and techniques have arisen in shaping the area, each with its own particular strengths: hyperbolic geometry, quantum algebra, foliations, geometric analysis,...

Heegaard Floer homology, defined at the turn of the 21st century by Peter Ozsváth and Zoltán Szabó, sits amongst these varied approaches. It consists of a powerful collection of invariants that fit into the framework of a $(3+1)$-dimensional topological quantum field theory. It represents the culmination of an effort to elucidate invariants from gauge theory using symplectic geometry. The resulting invariants have not only cast light on the structure of the earlier ones, they have opened the doors to many new ones, and they have led to tremendous breakthroughs on the problems of low-dimensional topology.

Joshua Evan Greene is a professor of mathematics at Boston College. His email address is `joshua.greene@bc.edu`.

DOI: https://doi.org/10.1090/noti2194

This is a story about Heegaard Floer homology: how it developed, how it has evolved, what it has taught us, and where it may lead.

Heegaard...

To set the stage, we return to the beginning of the 20th century, when Poincaré forged topology as an area of independent interest. In 1901, he famously and erroneously asserted that, as is the case for 1- and 2-manifolds, any 3-manifold with the same ordinary homology groups as the 3-dimensional sphere S^3 is, in fact, homeomorphic to S^3. By 1905, he had vanquished his earlier claim by way of an example that now bears his name: the Poincaré homology sphere P^3. We will describe this remarkable space in a moment and re-encounter it in many guises. Poincaré then revised his inaccurate assertion to ask whether any simply connected 3-manifold with the same homology groups as S^3 is, in fact, homeomorphic to S^3. We now know this to be the case following Perelman's resolution of W. Thurston's geometrization conjecture, and we will close with a folklore "proof" which is conditional on two conjectures in Heegaard Floer homology.

Heegaard diagrams. Poincaré described the homology sphere P^3 by way of a scheme developed by Heegaard in his 1898 dissertation, a variation on the age-old theme of decomposing a complicated object into simpler pieces. In this variation, the pieces are *handlebodies*: a genus-g

handlebody (colloquially, a *g-holed doughnut*) is the compact, oriented 3-manifold obtained by attaching g solid handles to a ball, as displayed in Figure 1.

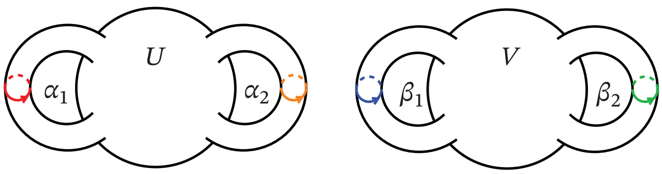

Figure 1. A pair of genus-2 handlebodies, the solid spaces enclosed by the surfaces shown, along with a complete collection of belt circles on each.

A *Heegaard diagram* consists of a closed, oriented surface $\Sigma = \Sigma_g$ of genus g, along with a pair of transverse multicurves $\alpha = \alpha_1 \cup \cdots \cup \alpha_g$ and $\beta = \beta_1 \cup \cdots \cup \beta_g$, each composed of g pairwise disjoint, homologically independent, simple closed curves. The data (Σ, α, β) encode a closed, oriented 3-manifold Y, as follows. We form a pair of genus-g handlebodies U and V with a complete collection of belt circles surrounding the 1-handles on each. We then glue U and V along their boundaries to Σ in such a way that the map $\partial U \to \Sigma$ preserves orientation, mapping belt circles onto the α curves, and the map $\partial V \to \Sigma$ reverses orientation, mapping belt circles onto the β curves. We thereby obtain a *Heegaard decomposition* $U \cup_\Sigma V$ of the resulting space Y. Every closed, oriented 3-manifold Y admits a Heegaard decomposition and so a Heegaard diagram.

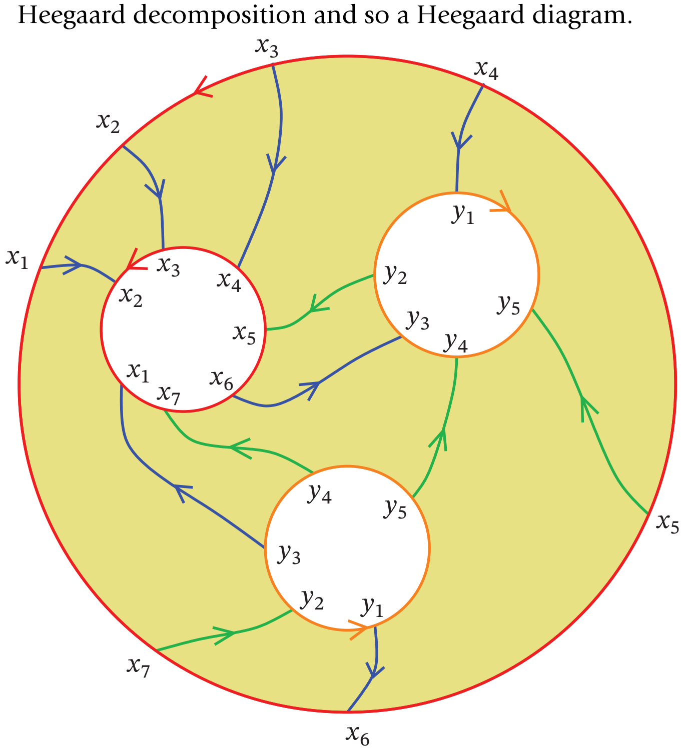

Figure 2. Poincaré's Heegaard diagram of P^3.

To understand Figure 2, identify the circles in pairs in such a way that like labels match. The result of the identification is a closed surface Σ_2 with a red simple closed curve α_1 and an orange simple closed curve α_2. The blue segments close up to a simple closed curve β_1, and the green ones to a simple closed curve β_2.

A Heegaard diagram leads to a presentation of the ordinary homology group of the space it presents. Each pair of curves α_i and β_j has a well-defined intersection number which we can calculate by orienting and making the signed count of intersection points between them. Collecting these values into a $g \times g$ matrix, we obtain a presentation matrix M for the first homology group of the manifold. For the Heegaard diagram at hand, we obtain the matrix $M = \left(\begin{smallmatrix} 3 & 2 \\ -2 & -1 \end{smallmatrix}\right)$ of determinant 1, confirming the fact that P^3 is a homology sphere. A similar procedure results in a presentation for the fundamental group $\pi_1(P^3)$. With a bit more effort, one finds that it has a 2-to-1 map onto the group of orientation-preserving isometries of the dodecahedron; hence P^3 is not simply connected.

A glimpse of the construction. We can now describe "half" of the construction of Heegaard Floer homology. Let Y denote a closed, oriented 3-manifold. Present Y by means of a Heegaard diagram $H = (\Sigma, \alpha, \beta)$.[1] Let G denote the set of g-tuples of points on Σ with the property that there is one point in the g-tuple on each curve in α and one point on each curve in β. In place of the intersection matrix M that we formed to calculate $H_1(Y)$, we make a new matrix whose (i, j) entry is a formal sum of the intersection points between α_i and β_j, signed by their intersection numbers. For instance, for the Heegaard diagram of P^3 of Figure 2, we obtain the matrix

$$\begin{pmatrix} x_1 + x_2 + x_3 + x_4 - x_6 & x_5 + x_7 \\ -y_1 - y_3 & y_2 - y_4 - y_5 \end{pmatrix}.$$

The monomials appearing in the formal expansion of the determinant of this matrix correspond precisely to the set G. There are 19 in the example. Let $CF(H)$ denote the free abelian group generated by the set G. This is the "Heegaard half" of Ozsváth and Szabó's construction: the underlying group of a chain complex whose homology computes the Heegaard Floer homology of Y. The group depends solely on the combinatorics of the curve systems, and it has a (mod 2) grading coming from the signs of the monomials in the determinant expansion.

... and Floer

To describe the "Floer half," the required differential, we must chart a course through a century of significant advances that developed since the time of Heegaard and Poincaré, passing through Morse theory, symplectic geometry, and gauge theory.

[1] *Strictly speaking, H should be* admissible, *a minor technicality.*

Morse theory and Morse homology. In the 1930s, Morse developed an influential approach to understanding smooth manifolds. Equip a smooth manifold M with a smooth, real-valued function $f : M \to \mathbb{R}$. The set of critical points $\mathrm{Crit}(f)$ consists of the points p at which the derivative of f vanishes. The function f is called a *Morse function* if its critical points are non-degenerate. Non-degeneracy means that in suitable coordinates around p, f takes the form $f(p) + x_1^2 + \cdots + x_k^2 - x_{k+1}^2 - \cdots - x_n^2$ for a well-defined value $0 \le k \le n$ called the *index* of f at p and denoted $\mathrm{ind}(p)$.

How many critical points must a Morse function f have on M? By placing a metric g on M and analyzing the gradient flow of a suitably generic Morse function f, one sees that M admits a cell decomposition with one k-cell for each index-k critical point of f. By forming the associated cellular chain complex, we obtain the *Morse inequality* $\#\mathrm{Crit}(f) \ge \beta(M)$, where $\beta(M)$ denotes the sum of the Betti numbers of M, addressing the stated question.

In the early 1980s, Witten reinterpreted the relationship between Morse theory and homology by swapping out cells for a physically-motivated count of "tunneling" effects. Each regular point of f lies along a gradient flow trajectory $u : \mathbb{R} \to M$, unique up to $\mathbb{R}$-translation. If M is closed, then $\mathrm{im}(u)$ limits to a pair of critical points, and we define the index of u to be the difference of the indices of its limit points. Assuming the Morse–Smale genericity conditions, there exist finitely many index-1 trajectories. Witten described the Morse complex associated with the triple (M, f, g): it is the chain complex freely generated by $\mathrm{Crit}(f)$ whose differential ∂ counts index-1 trajectories with sign. Thus, for $p \in \mathrm{Crit}(f)$, we have $\partial p = \sum_q \partial(p,q) \cdot q$, where q ranges over $\mathrm{Crit}(f)$ and $\partial(p,q)$ is a signed count of index-1 trajectories from p to q. The fact that $\partial^2 = 0$ involves examining flow-lines of index 2, which form a 1-dimensional space with a natural compactification by broken flow-lines. The resulting Morse homology groups return the ordinary homology of M. This complex is implicit in earlier work of Smale and Thom. However, Witten's interpretation stimulated the brilliant innovations of Floer to follow.

Symplectic geometry and Lagrangian Floer homology. Consider a smooth manifold L in which a particle moves subject to a conservative force. In how many positions could the particle begin at rest so that, after the force is applied for one unit of time, it returns to rest, possibly in a different location? In the 1980s, Arnold and Givental generalized this question and postulated an answer to it that became one of the driving themes in the nascent field of symplectic geometry. A *symplectic manifold* is a smooth manifold equipped with a closed, non-degenerate 2-form ω. For example, the cotangent bundle $M = T^*L$ is naturally a symplectic manifold within which L is *Lagrangian*, i.e., a submanifold of half the dimension and on which ω vanishes. Arnold and Givental conjectured that if M is an arbitrary symplectic manifold, L_0 is a special class of Lagrangian submanifold therein, and L_1 is a Hamiltonian displacement of L_0 in M, then $\#(L_0 \cap L_1) \ge \beta(L_0)$, in direct line with the Morse inequality.

In the mid-1980s, Floer attacked the Arnold–Givental conjecture by marrying the construction of the Morse complex with Gromov's newly developed theory of pseudoholomorphic curves. Floer's basic idea was to do Morse theory with the symplectic action functional defined on the space of paths from one Lagrangian L_0 to another L_1 (not necessarily displacements of one another). Remarkably, the critical points correspond to the intersection points $L_0 \cap L_1$, and, with a suitable metric in place, the gradient flow trajectories to so-called pseudoholomorphic Whitney disks between a pair of intersection points $x, y \in L_0 \cap L_1$.

A *Whitney disk* from x to y is a smooth embedding of a disk $\phi : D \to M$, $D = \{z \in \mathbb{C} \,|\, |z| \le 1\}$, as in Figure 3: namely, $\phi(-i) = x$; $\phi(i) = y$; $\phi(z) \in L_0$ for $z \in \partial D$, $\mathrm{Re}(z) \le 0$; and $\phi(z) \in L_1$ for $z \in \partial D$, $\mathrm{Re}(z) \ge 0$. Disks

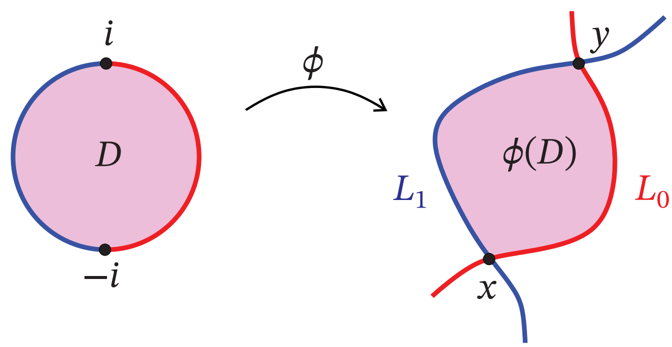

Figure 3. A Whitney disk ϕ from x to y.

like these were first exploited by Whitney in the course of proving his embedding theorem in the 1940s. Now suppose that M is equipped with an *almost-complex structure* j, meaning an automorphism of its tangent bundle TM that squares to -1. The prototype is the tangent bundle $TD^2 \approx D^2 \times \mathbb{C}$, which carries the almost-complex structure given by multiplication by i on each fiber. A Whitney disk is *pseudoholomorphic* with respect to j if it intertwines the two almost-complex structures: $d\phi \circ i = j \circ d\phi$. Much like a gradient trajectory of a Morse function, a Whitney disk ϕ possesses a *Maslov index* $\mu(\phi) \in \mathbb{Z}$, and disks fall into one-parameter families related by $\mathbb{R}$-translation (conformal automorphisms of the disk which fix $\pm i$). The Maslov index of ϕ calculates the *expected dimension* of the space of pseudoholomorphic representatives in its homotopy class.

Guided by the Morse–Witten complex, we form the abelian group $CF(L_0, L_1)$ freely generated by $L_0 \cap L_1$. Let $\mathcal{M}(L_0, L_1)$ denote the set of $\mathbb{R}$-translation classes of pseudoholomorphic disks of Maslov index one between a pair of intersection points in $L_0 \cap L_1$. We define an endomorphism $\partial : CF(L_0, L_1) \circlearrowleft$ by declaring its value on generators $x \in L_0 \cap L_1$ and extending linearly: $\partial x = \sum_y \partial(x, y) \cdot y$, where y ranges over $L_0 \cap L_1$ and $\partial(x, y)$ is a signed count of disks $\phi \in M(L_0, L_1)$ from x to y. Floer showed that, under favorable conditions, $\partial^2 = 0$, giving a chain complex and corresponding homology group $HF(L_0, L_1)$. Moreover, when L_1 is a displacement of L_0, we have $HF(L_0, L_1) \approx H^*(L_0)$. The Arnold–Givental conjecture then follows under suitable hypotheses.

The construction of Lagrangian Floer homology has since become the model for many related constructions, including Heegaard Floer homology. The latter arose as a result of parallel developments in gauge theory, to which we now turn.

Gauge theory and 4-manifolds. In the 1860s, Maxwell laid down the theory of electromagnetism, describing four fundamental equations that govern the relationship between the electrical and magnetic fields on spacetime. Over the ensuing decades, Maxwell's laws got recast succinctly as a condition on the curvature 2-form of a connection defined on a principal $U(1)$-bundle $P \to M$, where the 2-form corresponds to the electromagnetic field. The equations are invariant under alteration to the "gauge of measurement," i.e., the action of the *gauge group* of bundle automorphisms, an infinite-dimensional Lie group. In the 1950s, Yang and Mills described the strong nuclear force in similar terms. In their theory, the structure group $U(1)$ is replaced by the non-abelian Lie group $SU(2)$, with a condition on the connection resembling Maxwell's.[2] The Yang–Mills equations admit a special kind of solution called an *instanton*, as it is generically isolated in the space of connections in the case of classical spacetime: it appears "only for an instant." In dimension four, instantons satisfy a first-order PDE, the *anti-self-duality equations*. However, extracting solutions to them is difficult, owing to the equations' non-linearity. Through the 1970s, geometers had done so only in very special circumstances.

A dramatic paradigm shift occurred in the early 1980s, when Donaldson explained that the topology of the moduli space of instantons depends only on the smooth topology of M. As an application, he obtained a remarkable constraint on the *intersection lattice* of a smooth, oriented 4-manifold X. This is the free abelian group $H^2(X)/\text{Tors}$ equipped with the symmetric, bilinear form defined by $Q_X(\alpha, \beta) = (\alpha \cup \beta) \cap [X] \in \mathbb{Z}$. The name of the lattice is explained by the fact that when X is smooth, we may represent the Poincaré duals to α and β by a pair of transverse, smoothly embedded, closed, oriented surfaces, and then $Q_X(\alpha, \beta)$ equals their oriented intersection number.

[2] *We now understand the strong force in terms of $SU(3)$ gauge theory and the weak force in terms of $SU(2)$.*

For example, the intersection lattice of a connected sum of n copies of $\mathbb{C}P^2$ has an orthonormal basis consisting of the class of a complex line in each plane; thus, it is isometric to the *Euclidean lattice* I_n. Remarkably, Donaldson showed that no other definite lattices arise.

Theorem 1. *Let X denote a smooth, closed, simply connected 4-manifold. If Q_X is positive-definite, then the intersection lattice of X is isometric to I_n.*

Only months earlier, and using completely different tools, Freedman uncovered equally remarkable facts about *topological* 4-manifolds.

Theorem 2. *For every unimodular, integral lattice Λ, there exists a topological, closed, simply connected 4-manifold X with intersection lattice isometric to Λ.*

Theorems 1 and 2 express a sharp contrast between smooth and topological 4-manifolds. Perhaps most strikingly, they lead to the existence of different—"exotic"—smooth structures on $\mathbb{R}^4$, whereas $\mathbb{R}^n$ admits a unique one for all $n \neq 4$.

Instanton homology. Inspired by Donaldson's work, Floer introduced another version of infinite-dimensional Morse theory designed to count the number of instantons $I(X) \in \mathbb{Z}$ on a smooth, closed, oriented 4-manifold X. The idea is to cut X along a 3-manifold Y, so that $X = X_+ \cup_Y X_-$, similar to a Heegaard decomposition a dimension down. We assign a group $I(Y)$ that roughly measures how instantons look along a cylinder $Y \times \mathbb{R}$. The groups $I(Y)$ and $I(-Y)$ are dual to one another, and we obtain relative classes $I(X_\pm) \in I(\pm Y)$ that measure instantons on the two pieces. The pairing $\langle I(X_+), I(X_-)\rangle$ then recovers $I(X)$.

Floer defined the instanton homology $I(Y)$ by setting up a version of Morse theory with the Chern–Simons functional CS on the space of connections on an $SU(2)$-bundle $P \to Y$. For technical reasons, we assume that Y is a homology sphere. The critical points of CS are orbits of *flat* $SU(2)$-connections, those whose curvature vanishes. The bundle P is trivial, since Y is a 3-manifold, so flat connections are in 1-1 correspondence with representations $\pi_1(Y) \to SU(2)$ modulo conjugation. This provides a direct link between $I(Y)$ and the fundamental group. The trajectories of CS are finite-energy instantons on the cylinder $Y \times \mathbb{R}$. One of Floer's key technical advances was to make sense of indices in this infinite-dimensional setting in order to define the appropriate differential.

What Floer obtained was a prototype of a *topological quantum field theory* (TQFT), a notion first codified by Atiyah. Loosely speaking, the $(n+1)$-dimensional *cobordism category* has as its objects the closed, oriented n-manifolds. A morphism $X : Y_0 \to Y_1$ is a cobordism: a smooth, compact, oriented $(n+1)$-manifold with oriented boundary $\partial X = -Y_0 \sqcup Y_1$. An $(n+1)$-dimensional TQFT is a functor from the cobordism category to the category of $\mathbb{Z}$-modules and linear maps, along with some additional properties. In particular, instanton Floer homology assigns a map $F_X : I(Y_0) \to I(Y_1)$ to a smooth 4-dimensional cobordism $X : Y_0 \to Y_1$.

Casson's invariant. In between the definitions of Lagrangian and instanton Floer homology, Casson developed an invariant of homology spheres based on representations of the fundamental group into $SU(2)$. Begin with a genus-g Heegaard splitting of a homology sphere: $Y = U \cup_\Sigma V$. The space of representations from $\pi_1(\Sigma)$ to $SU(2)$, modulo conjugation, defines a variety $R(\Sigma)$ within which sits the pair of subvarieties $R(U)$ and $R(V)$. They intersect in a discrete set of points whose signed count is the Casson invariant $\lambda(Y) \in \mathbb{Z}$, independent of the choice of Heegaard splitting.

Taubes showed that, under a non-degeneracy assumption, Casson's invariant is the Euler characteristic of instanton Floer homology: $\lambda(Y) = \sum_k (-1)^k \operatorname{rk} I_k(Y)$. Moreover, the representation variety $R(\Sigma)$ carries a natural symplectic structure with respect to which $R(U)$ and $R(V)$ are Lagrangians. Putting a finer point on the matter, the Atiyah–Floer conjecture asserts that the Lagrangian Floer homology of the pair $R(U)$, $R(V)$ coincides with the instanton Floer homology $I_*(Y)$. While this conjecture remains open, it became a key inspiration in the construction of Heegaard Floer homology.

Seiberg–Witten theory. In Fall 1994, gauge theory sent another seismic wave through 4-manifold topology in the form of the Seiberg–Witten monopole equations. The equations are orders of magnitude easier to manipulate than the instanton equations, since moduli spaces of monopoles are compact, in contrast to the typically noncompact moduli spaces of instantons. Moreover, the abelian structure group $U(1)$ replaces the non-abelian one $SU(2)$, which simplifies analysis of singularities in the moduli space. Nevertheless, by physical considerations, Witten conjectured that they contain the same information as the instanton equations, an equivalence later confirmed by Feehan and Leness. This suggestion was initially met with astonishment, following a decade of hard-fought gains of Donaldson theory, but it quickly bore out new and spectacular results. For example, within one month of the equations' introduction, Kronheimer and Mrowka settled a celebrated conjecture of Thom.

Theorem 3. *A complex curve in $\mathbb{C}P^2$ has the least genus of any closed, oriented surface in its homology class.*

Following the advent of the monopole equations, it was apparent that there should exist a corresponding $(3+1)$-dimensional TQFT. In Kronheimer and Mrowka's description, the *monopole Floer homology* of a closed, oriented 3-manifold Y is modeled on the Morse homology of a manifold with boundary equipped with an S^1-action. The role of the Morse function is now played by the so-called Chern–Simons–Dirac functional. There are three related versions of the invariant, and each one is a module over the graded polynomial ring $\mathbb{Z}[U]$, the S^1-equivariant cohomology of a point. As powerful as they are, the monopole invariants are difficult to calculate except in very special circumstances.

The Construction and Structure of HF

Inspired by the Atiyah–Floer conjecture, we may ask: is there an instance of Lagrangian Floer homology that coincides with monopole Floer homology? This is the very question that Ozsváth and Szabó took up as the monopole theory developed. Their investigations led them to define a revolutionary set of invariants for objects of low-dimensional topology.

The construction. Once more, we present a closed, oriented 3-manifold Y by a genus-g Heegaard diagram (Σ, α, β). The parameter space of unordered g-tuples of points on Σ forms the g-fold *symmetric product* $\operatorname{Sym}^g(\Sigma)$. A choice of complex structure j on Σ induces a complex structure $\operatorname{Sym}^g(j)$ on $\operatorname{Sym}^g(\Sigma)$. This space plays the role of the ambient symplectic manifold in the construction. Its appearance is related to the fact that $\operatorname{Sym}^g(\Sigma)$ parametrizes solutions of the *vortex equations* on Σ, a dimensional reduction of the Seiberg–Witten equations.

The g-fold Cartesian product $\Sigma^{\times g}$ contains a pair of g-dimensional tori, $\alpha_1 \times \cdots \times \alpha_g$ and $\beta_1 \times \cdots \times \beta_g$. The projection map $\Sigma^{\times g} \to \operatorname{Sym}^g(\Sigma)$ is injective on these tori, since the α_i are pairwise disjoint from one another, as are the β_j. Therefore, we obtain a pair of g-dimensional tori $\mathbb{T}_\alpha, \mathbb{T}_\beta \subset \operatorname{Sym}^g(\Sigma)$, and they intersect transversally, since α and β do.

In this formulation, these tori are not Lagrangian, but Ozsváth and Szabó showed how to adapt the framework of Lagrangian Floer homology to them to extract an invariant of Y. Thus, we form the free abelian group $CF(\mathbb{T}_\alpha, \mathbb{T}_\beta)$ generated by the intersection points $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$, and we form a differential by counting disks in $M(\mathbb{T}_\alpha, \mathbb{T}_\beta)$. Observe that our description thus far recovers the "Heegaard half" glimpsed early on: the intersection points $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ are in bijective correspondence with the generating set G we had described earlier. Furthermore, the sign attached to each generator in the determinantal expansion is just its sign as an intersection point between the tori.

However, there is one problem with the invariant we have just produced: it simply returns the order of the ordinary homology group! Indeed, the development of Heegaard Floer homology briefly languished in this stage. It turns out that there is literally just one point missing in order for the theory to attain all its expected richness: we simply revise the construction by placing a basepoint z on

$\Sigma - \alpha - \beta$ and letting ∂ count only those disks in $M(\mathbb{T}_\alpha, \mathbb{T}_\beta)$ that avoid the submanifold $V_z = \{z\} \times \mathrm{Sym}^{g-1}(\Sigma) \subset \mathrm{Sym}^g(\Sigma)$. We thereby obtain a pair (CF, ∂) that relies on a pointed Heegaard diagram $H = (\Sigma, \alpha, \beta, z)$ presenting Y and a complex structure j on Σ. As usual in Lagrangian Floer theory, we may need to perturb $\mathrm{Sym}^g(j)$ to a nearby almost-complex structure in order to ensure that the moduli space $M(\mathbb{T}_\alpha, \mathbb{T}_\beta)$ is cut out transversely, by analogy to the Morse–Smale conditions. This done, we have the "fundamental theorem" of Heegaard Floer homology, due to Ozsváth and Szabó from 2001.

Theorem 4. *The map ∂ is a differential, and the chain homotopy type of (CF, ∂) is an invariant of the underlying 3-manifold Y.*

We denote the chain homotopy type by $CF(Y)$ and its homology $H_*(CF, \partial)$ by $HF(Y)$.[3]

Invariance and functoriality. The fact that the invariants do not depend on the analytic data used to construct them follows established lines in Floer homology. The independence of the choice of Heegaard diagram relies on a theorem of Reidemeister and Singer from the 1930s: any pair of Heegaard diagrams that present the same manifold are related by a sequence of moves, much like Reidemeister moves on knot diagrams. Any such move induces a quasi-isomorphism between the corresponding Floer chain complexes.

Originally, only the *isomorphism type* of $HF(Y)$ was established to be an invariant of Y. However, we would like to promote $HF(Y)$ to a concrete *group* in order to discuss maps between these invariants, such as those induced by cobordism, and to refer to specific elements thereof, like the contact invariant defined below. Juhász, D. Thurston, and Zemke showed how to do so once we equip Y with—perhaps unsurprisingly—a basepoint, which we assume is in place henceforth. This development resembles the passage, a century earlier, from the isomorphism invariance of simplicial homology to its full functoriality.

The Heegaard Floer 4-manifold invariant is defined in similar diagrammatic terms. A cobordism $X : Y_0 \to Y_1$ between a pair of connected 3-manifolds is presented by means of a Heegaard *triple* diagram $(\Sigma, \alpha, \beta, \gamma)$. The cobordism invariant counts pseudoholomorphic triangles between the triple of Lagrangian tori $\mathbb{T}_\alpha$, $\mathbb{T}_\beta$, $\mathbb{T}_\gamma \subset \mathrm{Sym}^g(\Sigma)$, yielding a map $F_X : HF(Y_0) \to HF(Y_1)$. The cobordism maps are essential for applications, for instance by furnishing a grading on the theory.

Refinements. The basepoint enables us to define richer algebraic structures. For instance, let $CF^-(H)$ denote the $\mathbb{Z}[U]$-module freely generated by $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ with the differential ∂ adjusted by weighting the count of each disk $\phi \in M(\mathbb{T}_\alpha, \mathbb{T}_\beta)$ by U^n, where $n = n_z(\phi) := \#\phi^{-1}(V_z) \geq 0$. Its homology $HF^-(Y)$ is an invariant that models one of the monopole invariants, and there are variations on the construction that recover the others. Indeed, the expected $\mathbb{Z}[U]$-module structure was one clue to the introduction of the basepoint.

The conjectured equivalence of the Heegaard Floer and monopole theories loomed over the area for well over a decade. At last, monumental efforts by two teams, Colin, Ghiggini, and Honda, and Kutluhan, Li, and Taubes, have established the analogue to the Atiyah–Floer conjecture in this setting.

Theorem 5. *The monopole Floer groups and corresponding Heegaard Floer groups are isomorphic for every closed, oriented 3-manifold.*

Structural properties. Essential to the theory is a bundle-theoretic object called a *spin*^{c} *structure*. We content ourselves to report a few essential facts about spin^c structures on a smooth manifold M in dimension 3 or 4: the set $\mathrm{Spin}^c(M)$ admits a free, transitive action by $H^2(M)$; there is a first Chern class map $c_1 : \mathrm{Spin}^c(M) \to H^2(M)$; and, when $\dim M = 4$, there is a restriction map $\mathrm{Spin}^c(M) \to \mathrm{Spin}^c(\partial M)$. For these reasons, the reader may safely conflate $\mathrm{Spin}^c(M)$ with the more familiar group $H^2(M)$.

We now focus on a closed, oriented 3-manifold Y for which $b_1(Y) = 0$. This is the case of interest in many applications, and the definition and structure of the invariant simplifies in this setting. The group $HF(Y)$ is finitely generated, and it decomposes according to a pair of gradings:

$$HF(Y) = \bigoplus_{s \in \mathrm{Spin}^c(Y)} HF(Y, s),$$

$$HF(Y, s) = \bigoplus_{m \in \mathbb{Q}} HF_m(Y, s).$$

The $\mathrm{Spin}^c(Y)$ grading originates at the chain level. The $\mathbb{Q}$-grading is called the *Maslov grading*. The support of the group $HF(Y, s)$ is contained in gradings $m \in d(Y, s) + \mathbb{Z} \subset \mathbb{Q}$, where $d(Y, s)$ denotes a distinguished grading called the *d-invariant*. The relative $\mathbb{Z}$-grading lifts the (mod 2) grading coming from the sign of an intersection point in a Heegaard diagram, and its enhancement to a $\mathbb{Q}$-grading uses properties of the maps induced by cobordism and the normalization $d(S^3, s) = 0$ for the unique $s \in \mathrm{Spin}^c(S^3)$.

The d-invariant alone is responsible for many spectacular applications of Heegaard Floer homology to low-dimensional topology. It was inspired by an analogous construction due to Frøyshov in the context of monopole Floer homology, and it possesses several useful properties. For one, $\mathrm{rk}\, HF_m(Y, s) \geq 1$ for $m = d(Y, s)$, leading to the inequality

$$\mathrm{rk}\, HF(Y) \geq |\mathrm{Spin}^c(Y)| = |H_1(Y)|. \tag{1}$$

[3] *These are the so-called "hat" versions of the invariant, usually denoted $\widehat{CF}$ and $\widehat{HF}$.*

For another, if Y bounds a smooth, positive-definite 4-manifold X, then $4d(Y,s) \leq c_1(t)^2 - b_2(X)$ for all $t \in \mathrm{Spin}^c(X)$, $t|Y = s$.

As an application, consider a smooth, closed, definite manifold X as in Theorem 1. Removing a small ball from X gives a positive-definite filling of S^3. Therefore, $0 = 4d(S^3,s) \leq c_1(t)^2 - b_2(X)$ for all $t \in \mathrm{Spin}^c(X)$. The classes $c_1(t) \in H^2(X)$ account for the set of so-called *characteristic elements* $\lambda \in H^2(X)$, i.e., those for which $Q_X(\lambda,\beta) \equiv Q_X(\beta,\beta) \pmod 2$ for all $\beta \in H^2(X)$. A theorem of Elkies from 1995 asserts that if $\operatorname{rk} L \leq \chi^2$ for every characteristic element χ in a positive-definite, unimodular lattice L, then L is isometric to the Euclidean lattice. Thus, we obtain a reproof of Theorem 1 without any assumption on π_1. This proof template is due to Kronheimer, who devised it in the setting of monopole Floer homology. In fact, Elkies proved his theorem in response to an inquiry of his! A reproof of Theorem 3 due to Strle issues along similar lines.

Examples. One of the triumphs of Heegaard Floer homology is its broad computability. We give a few sample calculations and motivate an important conjecture.

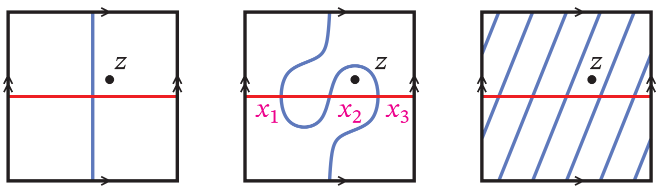

Figure 4. Two pointed genus-1 Heegaard diagrams of S^3 and one of a lens space.

1. The three-sphere admits the Heegaard diagram H shown on the left of Figure 4. Since it has a single intersection point, the differential d on $CF(H)$ vanishes, and we obtain $HF(S^3) \approx \mathbb{Z}_{(0)}$. The subscript denotes the Maslov grading, and we suppress the single spin^c grading. As a check of invariance, consider the pointed Heegaard diagram of S^3 shown in the middle of Figure 4. We see a single homotopy class of Whitney disk ϕ between the points of $\mathbb{T}_\alpha \cap \mathbb{T}_\beta = \alpha \cap \beta$ in $\mathrm{Sym}^1(\Sigma_1) = \Sigma_1$ for which $n_z(\phi) = 0$, namely one from x_2 to x_1. The Riemann mapping theorem guarantees a unique pseudoholomorphic representative of ϕ modulo $\mathbb{R}$-translation for any choice of complex structure on Σ. Hence $\partial x_1 = 0$, $\partial x_2 = \pm x_1$, and $\partial x_3 = 0$, leading to a consistent calculation of $HF(S^3)$.

2. A 3-dimensional lens space admits a genus-1 Heegaard diagram H, as shown on the right of Figure 4. For it, $CF(H)$ has one generator in each spin^c grading, ∂ vanishes, and we get equality in (1). The d-invariant in this case recovers the classical Reidemeister torsion, a complete invariant of the homeomorphism types of lens spaces.

3. For some complexity, let $p,q,r \in \mathbb{Z}$ denote a triple of pairwise coprime integers, and form the *Brieskorn homology sphere* $\Sigma(p,q,r) = \{(x,y,z) \in \mathbb{C}^3 \mid x^p + y^q + z^r = 0, |x|^2 + |y|^2 + |z|^2 = 1\}$. For instance, $-\Sigma(2,3,5)$ is another description of P^3. There is an algorithm for calculating the invariant of Brieskorn spheres, based on the existence of special definite fillings of them. From it, we obtain $HF(P^3) = \mathbb{Z}_{(2)}$, a calculation which would be hard to carry out from the Heegaard diagram of Figure 2, owing to the difficulty of analyzing disks in $\mathrm{Sym}^2(\Sigma_2)$. As another family of examples, $HF(\Sigma(2,3,6n+1)) \approx \mathbb{Z}_{(0)} \oplus \mathbb{Z}^n_{(-1)} \oplus \mathbb{Z}^n_{(-2)}$ for $n \geq 1$.

The three-sphere, lens spaces, and P^3 attain equality in (1). In fact, all manifolds with elliptic geometry do. By definition, an *L-space* is a manifold Y for which $\operatorname{rk} HF(Y) = |H_1(Y)|$, in homage to the lens spaces. From the point of view of Heegaard Floer homology, these spaces have the simplest possible invariant. We will see more examples to come, but the scarcity of small examples supports the following conjecture.

Conjecture 6. *The only irreducible L-space homology spheres are S^3 and $\pm P^3$.*

Knots

One of the novelties of the Heegaard Floer package is a version for knots and links defined independently by Ozsváth and Szabó and by Rasmussen. It developed as a way to describe the behavior of Heegaard Floer homology under Dehn surgery, and it strongly resembles a theory due to Khovanov described below.

The classical invariants of knots. Figures 5 and 6 display an assortment of knots. (The notation 6_1 comes from the knot tables.) An important infinite family are the *torus knots*, those that can be positioned on the genus-1 Heegaard surface in S^3. The (p,q)-torus knot $T_{p,q}$ admits the

Figure 5. András Stipsicz, John Morgan, Zoltán Szabó, and Peter Ozsváth (with multiplicity two). New York City, 2007.

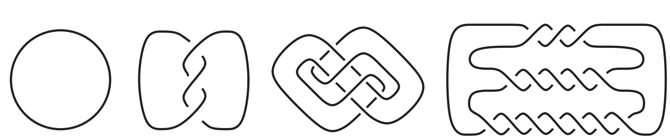

Figure 6. A few knots: the unknot ○, trefoil $T_{2,3}$, stevedore knot 6_1, and pretzel $P(-3,5,7)$.

succint description as the locus of the variety $\{z^p + w^q = 0\} \subset \mathbb{C}^2$ on the unit sphere $\{|z|^2 + |w|^2 = 1\}$ for coprime $p, q \in \mathbb{Z}$, echoing the definition of the Brieskorn spheres.

In the 1920s, Alexander introduced his famous polynomial invariant of knots and links. It is derived from the ordinary homology of the infinite cyclic cover of the knot complement $S^3 - K$. The invariant takes the form of a Laurent polynomial $\Delta(K)$, and it possesses several pleasant features. For one, it satisfies the *skein relation*[4] $\Delta(\nearrow\!\!\!\!\nwarrow) - \Delta(\nwarrow\!\!\!\!\nearrow) = (t^{1/2} - t^{-1/2})\Delta()()$. The pictures indicate a triple of oriented links with identical projections outside of a small window where they differ as shown. In fact, Δ is completely characterized by this relation and the value $\Delta(\bigcirc) = 1$. As examples, $\Delta(T_{p,q}) = (t^{pq} - 1)(t-1)/(t^p - 1)(t^q - 1)$, up to symmetrization, and $\Delta(P(-3,5,7)) = 1$.

Knots enjoy a special relationship with surfaces. A *Seifert surface* for a knot K is a compact, orientable surface $F \subset S^3$ with boundary $\partial F = K$.

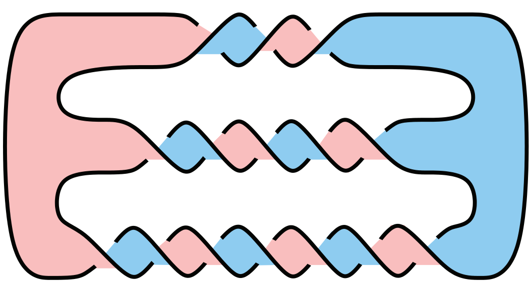

Figure 7. A genus-1 Seifert surface for $P(-3,5,7)$.

Seifert himself gave an algorithm for producing one from a knot diagram in the 1930s. Define the *genus* $g(K)$ of a knot K as the minimum genus of a Seifert surface for K. Thus, $g(K) = g(\bigcirc) = 0$ iff $K \simeq \bigcirc$. A particular case of interest occurs when K is *fibered*, a term coined by Seifert. This means that the knot complement fibers over the circle, and the closure of each fiber is a Seifert surface for K. For instance, the unknot is fibered, since each disk in the product fibration $S^3 - \bigcirc \approx S^1 \times \mathrm{int}(D^2)$ closes to a Seifert surface for $\bigcirc$. Examples of fibered knots abound and include all torus knots. Notably, the genus of a fibered knot equals the genus of its fiber surface.

A Seifert surface for K leads to a construction of the infinite cyclic cover of its complement and, in turn, to a direct relationship with its Alexander polynomial. For instance, $\deg \Delta(K) \le g(K)$, and $\Delta(K)$ is monic and attains equality in this bound if K is fibered. There is a related invariant, the knot signature $\sigma(K)$, which satisfies $|\sigma(K)|/2 \le g(K)$. However, Δ and σ do not detect the Seifert genus and fiberedness of every knot: for instance, $P(-3,5,7)$ has genus one and does not fiber, although $\Delta = 1$ and $\sigma = 0$ for it.

[4] *The dynamic terminology comes from Conway, picturing the knot as a length of yarn.*

Dehn surgery. In 1910, Dehn devised a general method called *surgery* for constructing 3-manifolds. As he writes, "These results have the particular value of yielding a very simple method for the construction of infinitely many Poincaré spaces [homology spheres], of which the discoverer [Poincaré] had constructed only one, in a complicated way." Dehn's construction begins with a knot $K \subset S^3$. A closed regular neighborhood of K is diffeomorphic to a solid torus. We excise its interior from S^3 to produce the knot exterior X_K, a compact manifold with torus boundary. We obtain a 3-manifold by regluing a solid torus $S^1 \times D^2$ to X_K along their boundaries in such a way that a curve $\{\theta\} \times \partial D^2$ glues to a curve that wraps p times longitudinally and q times meridionally around K. The diffeomorphism type of the result depends only on K and the slope p/q, and we denote it $K(p/q)$.

For instance, with the simplest non-trivial knot, the trefoil $T_{2,3}$, and the simplest non-trivial slope, $p/q = 1$, Dehn's construction recovers P^3. More generally, $T_{2,3}(1/n) \approx -\Sigma(2,3,6n-1)$.

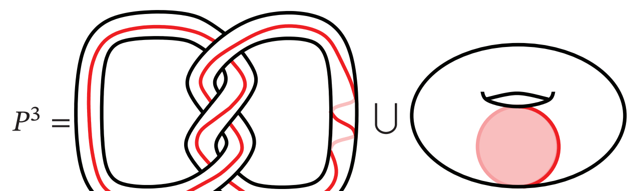

Figure 8. Dehn's description of the Poincaré homology sphere as $+1$-surgery along $T_{2,3}$.

As another example, for $p > 1$ and the unknot $\bigcirc$, we obtain the 3-dimensional *lens space* $L(p,q) \coloneqq \bigcirc(p/q)$. The unknot exterior $X_\bigcirc$ is diffeomorphic to a solid torus, so $L(p,q)$ admits a genus-1 Heegaard diagram; the right side of Figure 4 depicts one for $L(5,2)$.

Dehn surgery is popular, both for its elegance and its utility: in the 1960s, Lickorish and Wallace showed that every closed, oriented 3-manifold may be obtained by Dehn surgery along some link in S^3. It also undergirds many important results in knot theory. For example, in the late 1980s, Gordon and Luecke used combinatorial methods to establish the Dehn surgery characterization of S^3.

Theorem 7. *$K(p/q) \approx S^3$ iff $K \simeq \bigcirc$ or $p/q = 1/0$.*

As a consequence, knots in S^3 are determined by their complements: $X_{K_1} \approx X_{K_2}$ iff $K_1 \simeq K_2$.

The surgery exact triangle. Dehn surgery and Floer homology enjoy a special relationship dating back to the original work of Floer. Suppose that M is a compact, oriented manifold with torus boundary; $\gamma_1, \gamma_2, \gamma_3$ are a triple of simple closed curves on ∂M with the property that each pair meet in a single point; and $Y_i = M(\gamma_i)$ denotes the Dehn filling, $i = 1, 2, 3$.[5] Then their Floer homology groups fit into an exact triangle:

$$\begin{tikzcd} HF(Y_1) \arrow[rr] & & HF(Y_2) \arrow[dl] \\ & HF(Y_3) \arrow[ul] & \end{tikzcd}$$

As an application, suppose that Y_1 and Y_2 are L-spaces and $|H_1(Y_3)| = |H_1(Y_1)| + |H_1(Y_2)|$. The triangle then implies that Y_3 is an L-space, as well. For example, since S^3 and P^3 are L-spaces, iterated use of the triangle implies that $T_{2,3}(p/q)$ is an L-space for all $p/q \geq 1$. The maps appearing in the exact triangle are maps induced by cobordism.

Knot Floer homology. Every knot K in a 3-manifold Y is presented by a *doubly-pointed* Heegaard diagram $H = (\Sigma, \alpha, \beta, z, w)$. As before, the data (Σ, α, β) describe a Heegaard decomposition for Y. We join the basepoints z and w by an embedded arc in Σ disjoint from α and push its interior into the handlebody in which the α curves are belt circles; and we do the same with respect to β. The union of the two arcs gives the knot K. We define a chain

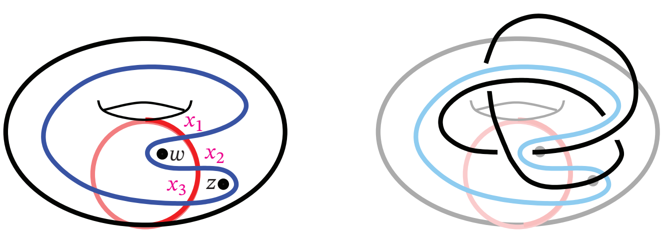

Figure 9. A doubly-pointed Heegaard diagram of $T_{2,3}$.

complex (CFK, ∂_0) on the same generating set as CF, where the differential ∂_0 counts disks $\phi \in M(\mathbb{T}_\alpha, \mathbb{T}_\beta)$ for which $n_z(\phi) = n_w(\phi) = 0$. As with the case of a closed 3-manifold, the homology group $H_*(CFK, \partial_0)$ is a knot invariant $HFK(K)$.[6]

Like the 3-manifold invariant, the knot Floer homology of a knot $K \subset S^3$ admits a bigrading:

$$HFK(K) = \bigoplus_{a \in \mathbb{Z}} HFK(K, a),$$
$$HFK(K, a) = \bigoplus_{m \in \mathbb{Z}} HFK_m(K, a).$$

The grading a is akin to the spin^c grading. It is called the *Alexander grading*, anticipating the relationship with Δ to come. The grading m is the Maslov grading.

In fact, the Alexander grading comes from a *filtration* on $CF(H)$. A related structure controls the invariant $HF(K(n))$ by an algebraic procedure called a *surgery formula*. This is the relationship that Ozsváth and Szabó and Rasmussen sought to capture, and it has proven to be a very rich vein to mine. By general homological algebra, $HFK(K)$ is the E_2 page in a *spectral sequence* which abuts to $HF(S^3) \approx \mathbb{Z}$. The spectral sequence produces a knot invariant $\tau(K) \in \mathbb{Z}$ whose power we will showcase below.

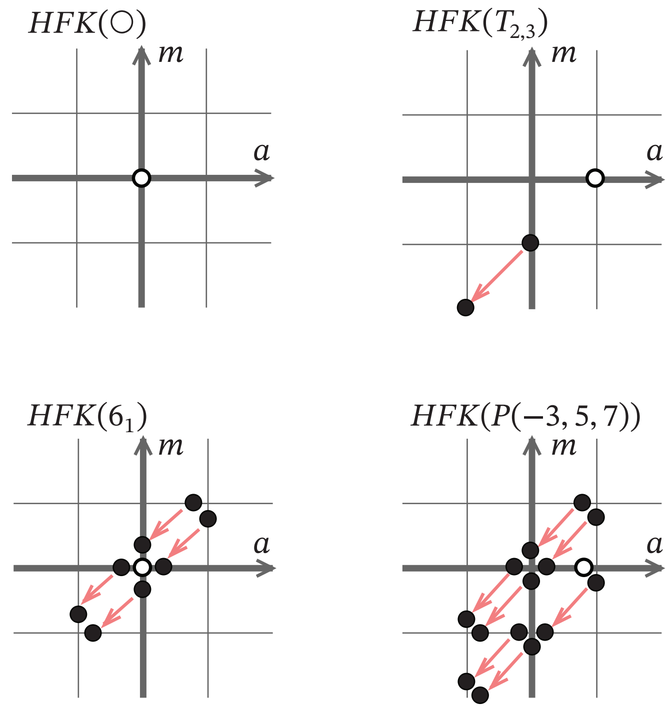

Figure 10. Sample knot Floer homology groups. A cluster of n dots nearby the lattice point (a, m) conveys that $HFK_m(K, a) \approx \mathbb{Z}^n$. The arrows represent the induced differential on $HFK(K)$ whose homology computes $HF(S^3) \approx \mathbb{Z}$. The highlighted dot represents the surviving generator, and its Alexander grading equals $\tau(K)$.

The following theorem collects several striking features of knot Floer homology, which strengthen the aforementioned properties of the Alexander polynomial.

Theorem 8. *Knot Floer homology*

(1) *categorifies the Alexander polynomial:*
$$\Delta(K) = \sum_{a,m \in \mathbb{Z}} (-1)^m \operatorname{rk} HFK_m(K, a) \cdot t^a;$$

(2) *detects the Seifert genus:*
$$g(K) = \max\{a : HFK(K, a) \neq 0\}; \text{ and}$$

(3) *detects fiberedness:*
$$K \text{ is fibered iff } \operatorname{rk} HFK(K, g(K)) = 1.$$

It is interesting to check the conclusions about $P(-3, 5, 7)$ implied by Theorem 8 and the calculation shown in Figure 10.

Part (1) of Theorem 8 is due to Ozsváth and Szabó and Rasmussen. Part (2) is due to Ozsváth and Szabó, utilizing

[5] *Gordon points out the medical-dental duality between Dehn surgery and filling.*

[6] *As before, we must equip K with a basepoint in order to obtain a concrete group.*

important results about foliations and contact and symplectic geometry. Since the unknot is the unique knot of genus 0, we obtain the immediate corollary that knot Floer homology detects the unknot.

Corollary 9. *$HFK(K) \approx HFK(\bigcirc)$ iff $K \simeq \bigcirc$.*

Part (3) is due to Ghiggini in the case $g(K) = 1$ and to Ni in the general case. The special case $g(K) = 1$ alone leads to a noteworthy corollary.

Corollary 10. *$K(p/q) \approx P^3$ iff $K \simeq T_{2,3}$ and $p/q = 1$.*

That is, the only way to obtain the Poincaré homology sphere by Dehn surgery along a knot in S^3 is Dehn's original description!

Juhász showed that both $HF(Y)$ and $HFK(K)$ are special cases of a more general invariant associated to a *sutured 3-manifold*. This is a structure developed by Gabai in the 1970s as a means of keeping track of decompositions of 3-manifolds along embedded surfaces. By studying the behavior of Floer homology for sutured manifolds under such decompositions, Juhász gave elegant re-proofs of parts (3) and (4) of Theorem 8, as well as a host of other applications. Moreover, Juhász's work has inspired similar constructions and results in both instanton and monopole Floer homology, reversing the course of influence between Heegaard Floer homology and these invariants.

Complements

Knot Floer homology and sutured Floer homology are two important variations on the construction of Heegaard Floer homology. Here we report on a few others and their implications.

Combinatorial Floer homology. In 2006, Sarkar and his collaborators made the pivotal observation that every knot and 3-manifold admits a special kind of Heegaard diagram whose affiliated Floer chain complex can be calculated in simple terms from the combinatorics of the diagram. Thus, it is possible to compute Heegaard Floer invariants without passage to symplectic geometry or gauge theory whatsoever, distilling them to their combinatorial essence!

Sarkar and Wang described an algorithm that takes as input a pointed Heegaard diagram for a 3-manifold Y and outputs another pointed Heegaard diagram H for Y with the property that every region apart from the one containing the basepoint z is either 2- or 4-sided. The region containing z will thus be a many-sided polygon. Such a Heegaard diagram is called a *nice diagram*. They proved that if H is a nice diagram and x and y are generators of $CF(H)$, then $\partial(x, y) = \pm 1$ iff x and y differ in one or two coordinates, and the intersection points in which they differ are connected by an embedded 2- or 4-gon on H that avoids z and the other points comprising x and y. Moreover, $\partial(x, y) = 0$ otherwise. Thus, ∂ simply counts 2- and 4-gons on H! For instance, using the Heegaard diagram of P^3 in Figure 2, the Sarkar–Wang algorithm produces a genus-2 nice diagram. As they report: "There are 335 generators and 505 differentials for this diagram. We leave the actual computation using this diagram to the patient reader."

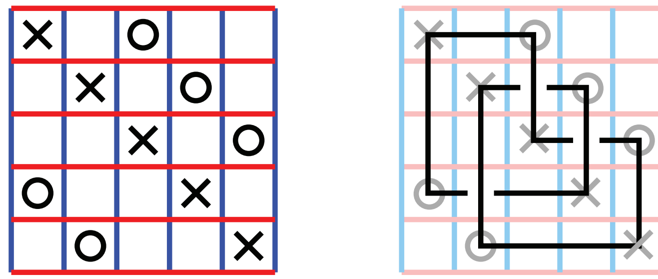

Figure 11. A grid diagram of $T_{2,3}$. Tic-tac-toe symbols replace z and w.

Manolescu, Ozsváth, and Sarkar developed a particularly satisfying method for producing nice diagrams adapted to knots in S^3. Take the standard genus-1 Heegaard splitting $(\Sigma, \alpha_1, \beta_1)$ of S^3, and let α and β denote the union of n parallel copies of α_1 and β_1, respectively. Place sets of basepoints $z = \{z_1, \dots, z_n\}$ and $w = \{w_1, \dots, w_n\}$ on $\Sigma - \alpha - \beta$ in such a way that each component of $\Sigma - \alpha$ and $\Sigma - \beta$ contains a single z_i and a single w_j. We thereby obtain a *multipointed* Heegaard diagram H of a link $L \subset S^3$. If we cut open Σ along α_1 and β_1 into a square, then we obtain a *grid diagram* of L, a representation studied in various guises, dating at least back to Brunn in the 1890s. The Floer chain complex $(CF(H), \partial_0)$ affiliated with a grid diagram of a link on an $n \times n$ grid thus contains $n!$ generators, and the differential d_0 counts embedded rectangles on $\Sigma - z - w$. The homology group $H_*(CF(H), \partial_0)$ calculates a stabilized version of knot Floer homology. Making use of grid diagrams, it is possible to define and prove the invariance of knot Floer homology without any analysis.

Bordered Floer homology and immersed curves. Bordered Floer homology, defined by Lipshitz, Ozsváth, and D. Thurston, is a cut-and-paste method for calculating the Heegaard Floer 3-manifold invariant. Thus, it represents a kind of $(2+1+1)$-dimensional TQFT. The idea is familiar: we cut a 3-manifold Y along a surface into two pieces and recover $CF(Y)$ as a pairing between relative invariants of the pieces. These invariants possess a deeper layer of algebraic structure, utilizing the A_∞-structure of pseudo-holomorphic polygon counts. There is a related version for knots due to Ozsváth–Szabó that recovers HFK; using it, we can now quickly calculate the invariant for knots with dozens of crossings.

Hanselman, Rasmussen, and Watson gave an elegant description of the bordered Floer invariant of a compact, oriented 3-manifold with torus boundary. In place of an algebraic object, the invariant takes the form of the regular homotopy class of an immersed multicurve in a punctured torus. Many topological applications issue from this description. For instance, the trio re-proved a theorem of Eftekhary in support of Conjecture 6: any L-space homology sphere is atoroidal.

Khovanov homology. In 1999, Khovanov defined an invariant of links in S^3 based on the ideas of TQFT and categorification. It developed out of Kauffman's state sum model for Jones's remarkable polynomial invariant of links defined in the early 1980s, which initiated the influx of quantum algebra into low-dimensional topology.

Like its close cousin, the Alexander polynomial, the Jones polynomial assigns a Laurent polynomial $V(K) \in \mathbb{Z}[q, q^{-1}]$ to a knot $K \subset S^3$. It satisfies the skein relation $q^{-1}V(\nearrow\!\!\!\!\nwarrow) - qV(\nwarrow\!\!\!\!\nearrow) = (q^{1/2} - q^{-1/2})V()\,()$, and its value on links is determined by this relation and the value $V(\bigcirc) = 1$. Kauffman's construction of $V(L)$ begins with a planar projection D of a link L with n double points. Each double point $\times$ can be smoothed in one of two ways: $)($ and $\asymp$. A choice of smoothing at each crossing is thus parametrized by a vertex v of the n-cube $[0,1]^n$ and results in a collection D_v of disjoint circles in the plane. In Kauffman's model, each complete resolution D_v contributes a Laurent polynomial in q that depends only on the vertex label and the number of circles in D_v. Their total sum then calculates $V(L)$.

Khovanov *categorified* this construction by applying a particular $(1+1)$-dimensional TQFT to the cube of resolutions. An edge $e = (v_1, v_2)$ of the n-cube specifies a pair of complete resolutions that coincide away from the neighborhood of a single crossing in D. They thereby cobound a *saddle cobordism* $S_e \subset \mathbb{R}^2 \times [0,1]$. The TQFT assigns to each complete resolution D_v a free abelian group and to each saddle cobordism S_e a corresponding map. The sum of these groups and maps gives a chain complex $(CKh(D), d)$ whose (co)homology is an invariant of the link $Kh(L)$. The invariant comes equipped with an integer bigrading $Kh(L) = \bigoplus_{i,j} Kh^{i,j}(L)$, and it *categorifies* the Jones polynomial: $V(L) = \sum_{i,j}(-1)^i \mathrm{rk}\, Kh^{i,j}(L) \cdot q^j$. Almost by definition, it also comes equipped with a *skein exact sequence*

$$\cdots \to Kh(\times) \to Kh()\,() \to Kh(\asymp) \to Kh(\times) \to \cdots$$

for an *unoriented skein triple* of links related by the indicated crossing and smoothings in some diagram.

Khovanov homology has inspired connections with Floer theory since its definition. Initially, it offered a clue to the structural properties of knot Floer homology. Since then, many direct links have been established, many with a similar form.

Theorem 11. *There exist spectral sequences with E^2 page isomorphic to $Kh(\overline{K})$ abutting to*

(1) *the Heegaard Floer homology of the branched double-cover, $\widehat{HF}(\Sigma(K))$;*
(2) *instanton knot Floer homology, $I^\natural(K)$; and*
(3) *knot Floer homology, $\widehat{HFK}(K)$.*

Here $\overline{K}$ denotes the mirror image of K.

Part (1) is due to Ozsváth and Szabó. It was the first explicit bridge between the Floer and Khovanov worlds. A key fact is that the branched double-covers of an unoriented skein triple of links fit into a triple of Dehn fillings to which the exact triangle applies. Iterating the triangle leads to a link surgery spectral sequence that establishes (1).

Part (2) is due to Kronheimer and Mrowka, who proved that their invariant $I^\natural$ detects the unknot. As a corollary, Khovanov homology does, too.

Corollary 12. $Kh(K) \approx Kh(\bigcirc) \iff K \simeq \bigcirc$.

Part (3) is due to Dowlin, confirming a conjecture of Rasmussen. The proof involves a cube of resolutions for $\widehat{HFK}$ based on singular knots. Combined with Corollary 9, it re-proves Corollary 12.

Remarkably, it remains unknown whether the Jones polynomial itself detects the unknot.

Topological Applications

Heegaard Floer homology has led to an impressive array of applications to low-dimensional topology. Here we showcase a few highlights.

Dehn surgery. When does Dehn surgery along a knot in S^3 yield a lens space, the simplest class of 3-manifold? Around 1990, Berge described the following elegant construction. Take a knot K that lies on a genus-2 Heegaard surface $\Sigma \subset S^3$ and that intersects a belt circle for each handlebody in a single transverse point of intersection. Then

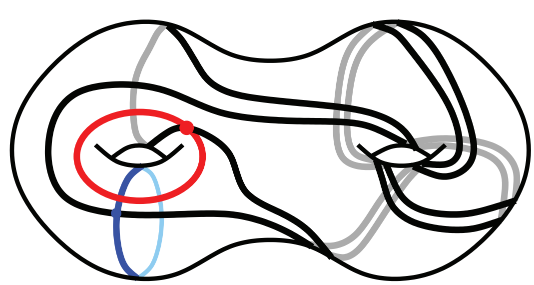

Figure 12. A knot in S^3 with a lens space surgery.

Dehn surgery along K, with slope given by a push-off of K on Σ, results in a lens space. For instance, all torus knots admit lens space surgeries, and Berge found many more examples using this construction. The Berge conjecture, from 1990, posits that these are all.

Conjecture 13. *Berge's construction accounts for all of the ways to obtain a lens space by Dehn surgery along a knot in S^3 with an integer slope.*[7]

Heegaard Floer homology tells us a lot about a knot with an integer surgery to a lens space, thanks in large part to the fact that lens spaces are L-spaces. We call a knot $K \subset S^3$ an *L-space knot* if some (positive) surgery along K yields an L-space.

Theorem 14. *The following conditions hold for an L-space knot $K \subset S^3$:*

(1) *K is fibered;*
(2) *$\tau(K) = \deg \Delta(K) = g(K)$; and*
(3) *the non-zero coefficients of $\Delta(K) = \sum_i c_i \cdot T^i$ are ± 1 and alternate in sign.*

The various parts make use of the surgery formula, the surgery sequence, and Theorem 8. One application is a proof of a conjecture of Gordon first proven in 2007 by Kronheimer, Mrowka, Ozsváth, and Szabó using monopole Floer homology.

Theorem 15. *$K(p/q) \approx L(p,q)$ iff $K \simeq \bigcirc$.*

We also recover an alternate proof of Theorem 7 and so the knot complement theorem.

A derivative of the Berge conjecture is the *realization problem*: which lens spaces $L(p,q)$ are realized as an integer surgery $K(p)$? Conjecturally, they are just the ones resulting from Berge's construction. Every knot surgery $K(p/q)$ bounds a smooth, definite 4-manifold whose intersection lattice $\Lambda(p/q)$ depends only on the fraction p/q. If $L(p,q) \approx K(p)$, then Theorem 1 implies that $\Lambda(p/q)$ embeds isometrically into the Euclidean lattice I_n, $n = \operatorname{rk} \Lambda(p/q) + 1$. This condition translates into a significant constraint on p/q, but it does not precisely identify the fractions p/q resulting from Berge's construction. Heegaard Floer homology offers more. By parlaying Theorem 1 together with properties of the d-invariant, we arrive at the following curious combinatorial condition: the orthogonal complement to $\Lambda(p/q) \subset I_n$ is spanned by a so-called *changemaker vector* $\sigma = (\sigma_1, \dots, \sigma_n)$. This means that, given coins with values $\sigma_1, \dots, \sigma_n \geq 0$, it is possible to make exact change from them in any amount up to their total value. Remarkably, this constraint pinpoints the lens spaces coming from Berge's construction, as the author showed in 2010.

Theorem 16. *The following conditions are equivalent: $L(p,q)$ is integer surgery along a knot $K \subset S^3$; $\Lambda(p/q)$ embeds in I_n orthogonal to a changemaker vector; and $L(p,q)$ is integer surgery along a Berge knot B. Moreover, $HFK(K) \approx HFK(B)$.*

[7]*Only torus knots have non-integral surgeries to lens spaces, following the celebrated cyclic surgery theorem of Culler, Gordon, Luecke, and Shalen, from 1987.*

The solution to the realization problem follows.

What further light might Heegaard Floer homology throw onto Conjecture 13? Call a knot $K' \subset L(p,q)$ *simple* if it is presented by placing a pair of basepoints on the standard genus-1 Heegaard diagram of $L(p,q)$ (cf. Figure 4) and *Floer simple* if it attains equality in the inequality $\operatorname{rk} HFK(K') \geq \operatorname{rk} HF(L(p,q)) = p$. Conjecturally, knot Floer homology detects simple knots.

Conjecture 17. *A knot in a lens space is Floer simple iff it is simple.*

This conjecture, due to Baker, Grigsby, and Hedden and to Rasmussen, generalizes Theorem 8 and, if true, would imply Conjecture 13.

Contact geometry. A *contact form* on a 3-manifold Y is a 1-form $\alpha \in \Omega^1(Y)$ with the property that $\alpha \wedge d\alpha$ equals a positive multiple of the volume form. A *contact structure* is the 2-plane field $\xi \coloneqq \ker(\alpha)$ of a contact form α. Contact structures often arise naturally on the boundaries of compact symplectic manifolds. They also relate to taut foliations, which are totally integrable 2-plane fields whose leaves are pierced by a common closed transversal curve. For instance, Eliashberg and W. Thurston showed how to perturb a taut foliation into a pair of contact structures $\xi_\pm$ with the property that both $(\pm Y, \xi_\pm)$ are convex boundaries of symplectic 4-manifolds.

Heegaard Floer homology assigns an element $c(\xi) \in HF(Y)$ to a contact structure ξ on Y. If $c(\xi) \neq 0$, then ξ is *tight*, meaning that there is no disk in Y whose boundary is tangent to ξ. Eliashberg showed that contact structures that are not tight are classified up to isotopy by the homotopy class of the 2-plane field, so the theory is geometrically much richer for tight ones. For example, the contact structure $\ker(dz - y\,dx)$ on $\mathbb{R}^3$ is tight, and it extends to a tight contact structure ξ_0 on S^3. In fact, ξ_0 is the unique tight contact structure on S^3, and $c(\xi_0)$ generates $HF(S^3)$.

The contact invariant has led to the classification of tight contact structures on large classes of 3-manifolds. It has also enabled us to distinguish which kinds of structures on 4-manifolds can fill a given contact structure: weak versus strong symplectic filling, Stein filling, etc. For instance, if ξ is Stein-fillable, then $c(\xi) \neq 0$. At present, there is no characterization of the vanishing of $c(\xi)$ in terms of contact topology alone. For the case of the contact structure ξ_+ coming from a taut foliation, Ozsváth and Szabó showed not only that $c(\xi_+)$ is non-zero, it certifies that $\operatorname{rk} HF(Y) > |H_1(Y)|$. Thus, an L-space does not admit a taut foliation.

There is a version of knot theory in the presence of a contact structure. A knot in a contact manifold (Y, ξ) is *Legendrian* if it is everywhere tangent to ξ. One reason for the importance of Legendrian knots is that they are used to carry out Dehn surgery in the setting of contact geometry. To a Legendrian knot $\mathcal{K}$ we can associate its smooth knot type K and a couple of numerical invariants which capture how the 2-plane field $\xi|K$ twists around it. They define invariants of Legendrian isotopy. Historically, it was challenging to distinguish pairs of Legendrian knots with these same classical invariants.

Knot Floer homology, particularly in its combinatorial form, gives an effective way to distinguish such pairs. Projecting a Legendrian knot $\mathcal{K}$ in $(\mathbb{R}^3, \xi_0)$ onto the xz-plane gives its *front projection*, a picture which closely resembles a grid diagram. We adapt a grid diagram H to the underlying knot K from a front projection of $\mathcal{K}$, and we construct a generator $x \in CFK(H)$ by selecting the intersection point in the top-right corner of each square containing a w basepoint. The element x is a cycle, and its homology class defines a Legendrian invariant $\lambda(\mathcal{K}) \in HFK(K)$ by which one can distinguish pairs of Legendrian knots with the same classical invariants.

Homology cobordism. Piecewise linear homology spheres in dimension n form a group Θ^n under the relation of homology cobordism: $Y_0 \sim Y_1$ iff there exists a (PL) cobordism $X : Y_0 \to Y_1$ such that the inclusion of Y_j into X induces an isomorphism on homology for $j = 0, 1$.[8] The abelian group operation on Θ^n is induced by connected sum, inversion is induced by orientation reversal, and the zero element is represented by the n-sphere. In the 1960s, Kervaire showed that $\Theta^n = 0$ for all $n \geq 4$. By contrast, in the 1950s, Rokhlin had already introduced a homomorphism $\rho : \Theta^3 \to \mathbb{Z}_2$ taking $[P^3]$ to the non-trivial element.

A major motivation for the study of Θ^3 stems from high-dimensional topology. In the 1970s, Galewski and Stern and Matumoto showed that the existence of a closed, non-triangulable, topological manifold in each dimension ≥ 5 is equivalent to the assertion that $\ker(\rho)$ contains all the 2-torsion in Θ^3. This would follow if there were a homomorphism $\varphi : \Theta^3 \to \mathbb{Z}$ that lifts ρ, meaning that $\varphi \equiv \rho \pmod 2$. This possibility, which is still open, served as one motivation for Casson, whose invariant λ lifts ρ but does not vanish on every null-bordant homology sphere.

In 2013, Manolescu exploited a Pin(2)-symmetry in the Seiberg–Witten equations to define an equivariant version of monopole Floer homology.[9] Using it, he extracted a well-defined mapping $\beta : \Theta^3 \to \mathbb{Z}$ which, while it does not add under connected sum, lifts ρ and negates under orientation reversal. These properties suffice to solve the triangulation problem.

Theorem 18. *All 2-torsion in Θ^3 is contained in* $\ker(\rho)$. *Consequently, in every dimension ≥ 5, there exist closed, topological manifolds that cannot be triangulated.*

Still, the structure of Θ^3 remains deeply mysterious: it is unknown whether it contains any torsion whatsoever or, for that matter, infinitely divisible elements. All progress on understanding Θ^3 since Rokhlin owes to advances in gauge theory and Floer homology. For instance, Furuta showed that Θ^3 contains a $\mathbb{Z}^\infty$ subgroup generated by Brieskorn spheres. The d-invariant of a homology sphere in its unique spin^c structure descends to a non-trivial homomorphism $d : \Theta^3 \to \mathbb{Z}$. It follows that Θ^3 contains a $\mathbb{Z}$ *summand*, a fact first proven by Frøyshov using instanton Floer homology. We now know quite a bit more, following work of Dai, Hom, Stoffregen, and Truong, from 2018.

Theorem 19. *Θ^3 contains a $\mathbb{Z}^\infty$ summand.*

The argument involves the *involutive Heegaard Floer homology* invariant due to Hendricks and Manolescu. The invariant is designed to recover a part of the Pin(2)-equivariant monopole group in the Heegaard Floer setting.[10] It is conceivable that $\Theta^3 \approx \mathbb{Z}^\infty$, which would have further implications for high-dimensional triangulations.

Slice genus. Knots in 3-manifolds probe the topology of 4-manifolds through the surfaces they bound and slice therein. Consider a smoothly embedded 2-sphere in S^4. Slice it by an equatorial 3-sphere, and assume that the intersection is transverse and connected. We thereby obtain a *slice knot* $K \subset S^3$ and a properly embedded *slice disk* $D \subset B^4$ to either side of the slice. As an example, reconsider the stevedore knot. Figure 13 displays a slice disk for it in a collar $S^3 \times I \subset B^4$ with three critical points.

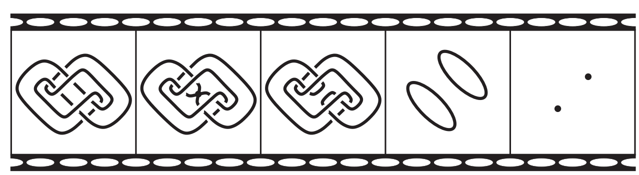

Figure 13. A movie reel of a ribbon disk filling 6_1.

This disk has the distinctive feature that the I-coordinate restricts to a Morse function on it with no local maxima. Such a disk is called a *ribbon disk*, and the famous *slice-ribbon problem* of Fox from the 1960s asks whether every slice knot bounds a ribbon disk.

Not every knot is slice. Define the *slice genus* $g_*(K)$ as the minimum value of $g(F)$ over all smoothly, properly embedded, oriented surfaces $F \subset B^4$ with $\partial F = K$. In general, we have the classical bounds $|\sigma(K)|/2 \leq g_*(K) \leq g(K)$. For instance, $g_*(T_{2,3}) = 1$, since $|\sigma(T_{2,3})|/2 = g(T_{2,3}) = 1$. However, the gap between $|\sigma|/2$ and g is typically quite large for torus knots. Nevertheless, in connection with the study of surface singularities, Milnor conjectured the following result.

[8] *The PL and smooth categories coincide in dimensions 3 and 4, so we elide the distinction there.*

[9] Pin(n) *is the Lie group double cover of $O(n)$. The terminology is wordplay deriving from the fact that* Spin(n) *is the Lie group double cover of $SO(n)$.*

[10] *For technical reasons to do with* naturality, *it is not (yet) possible to define* Pin(2)*-equivariant Heegaard Floer homology.*

Theorem 20. $g_*(T_{p,q}) = g(T_{p,q}) = (p-1)(q-1)/2$.

Kronheimer and Mrowka proved Theorem 20 in 1993 using instanton Floer homology. Subsequently, there have been very elegant re-proofs using knot homology theories. Ozsváth and Szabó showed that the τ-invariant bounds the slice genus: $|\tau(K)| \leq g_*(K)$. Using Theorem 14(2), Theorem 20 follows. Rasmussen gave a celebrated re-proof by defining a related invariant in Khovanov homology with similar properties to τ. The remarkable feature of Rasmussen's proof is that it avoids the use of any analysis. Sarkar gave another analysis-free argument by re-proving the bound involving τ using grid homology.

Concordance. Concordance bridges the topics of homology cobordism and slice genus. A *concordance* between a pair of oriented knots K_0, K_1 is a smooth, properly embedded annulus in the cylinder $S^3 \times [0,1]$ with oriented boundary $K_j \subset S^3 \times \{j\}$, $j = 0,1$. Concordance defines an equivalence relation on knots, and concordance classes form an abelian group $\mathcal{C}$ in the same way as Θ^3, with the unknot representing zero. Note that a knot represents zero in $\mathcal{C}$ iff it is slice. In addition, if two knots are concordant, then Dehn surgeries along them with the same slope are homology cobordant. Thus, obstructions to homology cobordism translate into obstructions to concordance.

As with Θ^3, the structure of $\mathcal{C}$ remains highly elusive, and it is an object of intense study. The knot signature descends to a homomorphism $\sigma : \mathcal{C} \to \mathbb{Z}$. Thus, $T_{2,3}$ represents an element of infinite order. An *amphichiral* knot, one that is isotopic to its own mirror image, represents an element of order dividing two in $\mathcal{C}$. The class of the figure-eight knot (not pictured) has order two for this reason. It is conceivable that $\mathcal{C} \approx \mathbb{Z}^\infty \oplus (\mathbb{Z}/2\mathbb{Z})^\infty$ and that amphichiral knots account for all of the torsion. However, we cannot rule out the possibility that $\mathcal{C}$ contains n-torsion, for any value $n > 2$, or infinitely divisible elements.

Concordance also exemplifies the distinction between the smooth and topological categories in dimension four. A knot K is *topologically slice* if it bounds a *locally flat* disk $D \subset B^4$: that is, for each point x interior to D, there exists a neighborhood $U \subset B^4$ of x and a continuous chart $\varphi : (\mathbb{R}^4, 0) \xrightarrow{\sim} (U, x)$ for which $\varphi^{-1}(D)$ is a coordinate 2-plane (mutatis mutandis for $x \in \partial D$). We define the topological slice genus g_*^{top} in a similar way. Classical invariants obstruct topological sliceness: for instance, $|\sigma(K)/2| \leq g_*^{top}(K) \leq g_*(K)$. The existence of a topologically slice knot that is not smoothly slice was proven following Theorems 1 and 2. In fact, the existence of such a knot can be used to manufacture an exotic smooth structure on $\mathbb{R}^4$. The pretzel $P(-3,5,7)$ is one example. A celebrated theorem of Freedman related to Theorem 2 asserts that a knot K is topologically slice if $\Delta(K) = 1$, as it is for this pretzel. However, it is not smoothly slice, since $\tau = 1$ for it (Figure 10).

We can probe matters further by restricting attention to the subgroup $\mathcal{C}_{TS} \subset \mathcal{C}$ of topologically slice knots modulo smooth concordance. The invariant τ defines a homomorphism $\mathcal{C}_{TS} \to \mathbb{Z}$ that is non-trivial (consider $P(-3,5,7)$), so $\mathcal{C}_{TS}$ contains a $\mathbb{Z}$ summand. In fact, the authors of Theorem 19 showed that it contains a $\mathbb{Z}^\infty$ summand by a similar argument.

Lastly, concordance leads to a conjectural partial order on the set of knots. A smooth concordance $A : K_0 \to K_1$ is *ribbon* if the restriction of I to A has no local maxima. Write $K_0 \lesssim K_1$ if there is a ribbon concordance from K_0 to K_1. In the 1980s, Gordon conjectured that if $K_0 \lesssim K_1$ and $K_1 \lesssim K_0$, then $K_0 \simeq K_1$. The following result of Zemke, from 2019, offers support.

Theorem 21. *A ribbon concordance $A : K_0 \to K_1$ induces an injection $F_A : HFK(K_0) \hookrightarrow HFK(K_1)$ of bigraded groups. Consequently, $K_0 \lesssim K_1 \implies g(K_0) \leq g(K_1)$. Moreover, if $K_0 \lesssim K_1$ and $K_1 \lesssim K_0$, then $HFK(K_0) \approx HFK(K_1)$.*

Unknotting number. What is the fewest number of crossing changes required to convert a diagram of a knot K into the unknot? In spite of its disarming simplicity, the *unknotting number* $u(K)$ has proven to be one of the most challenging knot invariants to estimate. For instance, it is unknown whether it is additive under connected sum. An unknotting sequence in g steps for a knot K leads to a smooth slice surface for K of genus g, whence $g_*(K) \leq u(K)$. The standard diagram of $T_{p,q}$ can be unknotted by $(p-1)(q-1)/2$ crossing changes. Combining these two facts with Theorem 20, we deduce that $u(T_{p,q}) = (p-1)(q-1)/2$.

In a different direction, a knot K is *alternating* if it admits a planar projection whose crossings alternate over and under as the diagram is traversed. Figure 6 displays two alternating diagrams and three alternating knots. Kohn conjectured the following result in 1990, which McCoy proved in 2014.

Theorem 22. *If an alternating knot K has unknotting number one, then any alternating diagram of K displays an unknotting crossing.*

The argument uses similar techniques to the solution of the realization problem, Theorem 16.

Challenges

For all their power and computability, there remains much more we would like to know about Heegaard Floer homology. What do its cycles signify? Does it admit a characterization by a simple set of axioms, like the Eilenberg–Steenrod axioms for ordinary homology? One major outstanding issue is to relate Heegaard Floer homology directly to the fundamental group π_1 and instanton Floer homology. Another is to relate it to the geometry of 3-manifolds, such as the existence of a taut foliation or a hyperbolic structure.

The *L-space conjecture* attempts to address the last two issues by reconciling three different notions of size for a 3-manifold, in terms of Heegaard Floer homology, geometry, and the fundamental group.

Conjecture 23. *The following conditions on a closed, orientable, irreducible* 3*-manifold* Y *with* $b_1(Y) = 0$ *are equivalent:*

(1) $\operatorname{rk} HF(Y) > |H_1(Y)|$ (*Y is* <u>*not*</u> *an L-space*);
(2) *Y admits a taut foliation; and*
(3) $\pi_1(Y) \neq 1$ *admits a total order invariant under left-multiplication.*

The geometric relationship (1) $\Longleftrightarrow$ (2) bears a passing resemblance to fiberedness detection, Theorem 8(3). Of the six implications posited by Conjecture 23, only the direction (2) $\Longrightarrow$ (1) noted above is known to hold in general. Building on the work of many researchers, Conjecture 23 holds for the class of *graph manifolds*, those whose geometric decomposition consists of non-hyperbolic pieces.

To come full circle, we close with a conditional "proof" of the Poincaré conjecture using Heegaard Floer homology. Suppose that Y is a simply connected 3-manifold. Thus, Y is a homology sphere, and by passing to a prime summand, we may assume that it is irreducible. A closed transversal to a taut foliation on a manifold has infinite order in its fundamental group. Hence Y does not admit a taut foliation, so Conjecture 23 (1) $\Longrightarrow$ (2) implies that it is an irreducible L-space homology sphere. Now Conjecture 6 implies that $Y \approx S^3$, as $\pi_1(P^3) \neq 1$. We invite the reader to supply the missing parts.

ACKNOWLEDGMENTS. The author thanks Siddhi Krishna, András Stipsicz, Ari Turner, and the referees for many valuable comments and corrections. The excellent references below helped shape this account and elaborate on many of the topics treated here. The author was supported by NSF CAREER Award DMS-1455132 and a Simons Fellowship.

References

[Ati88] Michael Atiyah, *Topological quantum field theories*, Inst. Hautes Études Sci. Publ. Math. **68** (1988), 175–186 (1989). MR1001453

[Bot88] Raoul Bott, *Morse theory indomitable*, Inst. Hautes Études Sci. Publ. Math. **68** (1988), 99–114 (1989). MR1001450

[Don02] S. K. Donaldson, *Floer homology groups in Yang-Mills theory*, Cambridge Tracts in Mathematics, vol. 147, Cambridge University Press, Cambridge, 2002. With the assistance of M. Furuta and D. Kotschick. MR1883043

[Gor99] C. McA. Gordon, *3-dimensional topology up to 1960*, History of topology, North-Holland, Amsterdam, 1999, pp. 449–489. MR1674921

[Hom17] Jennifer Hom, *A survey on Heegaard Floer homology and concordance*, J. Knot Theory Ramifications **26** (2017), no. 2, 1740015, 24. MR3604497

[Juh15] András Juhász, *A survey of Heegaard Floer homology*, New ideas in low dimensional topology, Ser. Knots Everything, vol. 56, World Sci. Publ., Hackensack, NJ, 2015, pp. 237–296. MR3381327

[KM07] Peter Kronheimer and Tomasz Mrowka, *Monopoles and three-manifolds*, New Mathematical Monographs, vol. 10, Cambridge University Press, Cambridge, 2007. MR2388043

[Lin16] Francesco Lin, *Lectures on monopole Floer homology*, Proceedings of the Gökova Geometry-Topology Conference 2015, Gökova Geometry/Topology Conference (GGT), Gökova, 2016, pp. 39–80. MR3526838

[Man] Ciprian Manolescu, *Homology cobordism and triangulations*, Proceedings of the International Congress of Mathematicians—Rio de Janeiro 2018. Vol. II. Invited lectures, World Sci. Publ., Hackensack, NJ, 2018, pp. 1175–1191. MR3966804

[Man15] Ciprian Manolescu, *Floer theory and its topological applications*, Jpn. J. Math. **10** (2015), no. 2, 105–133. MR3392529

[Mor14] John W. Morgan, *100 years of topology: work stimulated by Poincaré's approach to classifying manifolds*, The Poincaré conjecture, Clay Math. Proc., vol. 19, Amer. Math. Soc., Providence, RI, 2014, pp. 7–29. MR3308756

[OS04a] Peter Ozsváth and Zoltán Szabó, *Heegaard diagrams and holomorphic disks*, Different faces of geometry, Int. Math. Ser. (N. Y.), vol. 3, Kluwer/Plenum, New York, 2004, pp. 301–348. MR2102999

[OS04b] Peter Ozsváth and Zoltán Szabó, *Holomorphic disks and three-manifold invariants: properties and applications*, Ann. of Math. (2) **159** (2004), no. 3, 1159–1245. MR2113020

[OS04c] Peter Ozsváth and Zoltán Szabó, *Holomorphic disks and topological invariants for closed three-manifolds*, Ann. of Math. (2) **159** (2004), no. 3, 1027–1158. MR2113019

[OSS15] Peter S. Ozsváth, András I. Stipsicz, and Zoltán Szabó, *Grid homology for knots and links*, Mathematical Surveys and Monographs, vol. 208, American Mathematical Society, Providence, RI, 2015. MR3381987

[Ras05] Jacob Rasmussen, *Knot polynomials and knot homologies*, Geometry and topology of manifolds, Fields Inst. Commun., vol. 47, Amer. Math. Soc., Providence, RI, 2005, pp. 261–280. MR2189938

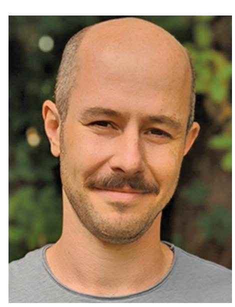

Joshua Evan Greene

Credits

Opening image is courtesy of Jonathan Hanselman and Liam Watson.

Figure 5 is courtesy of Jonathan Bloom.

All other figures are courtesy of the author.

Photo of the author is courtesy of Daphne Feng.