

Hard Diagrams of Split Links

Corentin Lunel 

INRIA Université Côte d’Azur, Montpellier, France

Arnaud de Mesmay 

Laboratoire d’Informatique Gaspard Monge, 5 Boulevard Descartes, Champs-sur-Marne,
77454 Marne-la-Vallée Cedex 2, France

Jonathan Spreer 

School of Mathematics and Statistics, The University of Sydney, Australia

Abstract

Deformations of knots and links in ambient space can be studied combinatorially on their diagrams via local modifications called Reidemeister moves. While it is well-known that, in order to move between equivalent diagrams with Reidemeister moves, one sometimes needs to insert excess crossings, there are significant gaps between the best known lower and upper bounds on the required number of these added crossings. In this article, we study the problem of turning a diagram of a split link into a split diagram, and we show that there exist split links with diagrams requiring an arbitrarily large number of such additional crossings. More precisely, we provide a family of diagrams of split links, so that any sequence of Reidemeister moves transforming a diagram with c crossings into a split diagram requires going through a diagram with $\Omega(\sqrt{c})$ extra crossings. Our proof relies on the framework of bubble tangles, as introduced by the first two authors, and a technique of Chambers and Liokumovitch to turn homotopies into isotopies in the context of Riemannian geometry.

2012 ACM Subject Classification Mathematics of computing → Geometric topology; Theory of computation → Computational geometry

Keywords and phrases Knot theory, hard knot and link diagrams, Reidemeister moves, extra crossings, split links, bubble tangles, compression representativity

Digital Object Identifier 10.4230/LIPIcs.SoCG.2025.67

Related Version *Full Version*: <https://arxiv.org/abs/2412.03372>


Funding *Corentin Lunel*: Partially supported by the ANR project ANR-20-CE48-0007 (AlgoKnot). *Jonathan Spreer*: Supported by the Australian Research Council under the Discovery Project scheme (grant number DP220102588).

Acknowledgements We would like to thank Clément Maria for helpful discussions and Stephan Tillmann for his comments and suggestions on an earlier version of the paper. Moreover, the authors want to thank the anonymous referees for their helpful comments.


1 Introduction

The Reidemeister theorem [32] is a fundamental and powerful result in low-dimensional topology. It states that any two link diagrams represent equivalent links¹ if and only if they can be related by a sequence of planar isotopies and local moves, called **Reidemeister moves**, pictured in Figure 1. This theorem is at the heart of many theoretical results as well as computational applications. Indeed, many knot invariants, from the most basic ones such as tri-colourability [1, Section 1.5] to more advanced ones such as the Jones polynomial [20]

¹ A link is an embedding of a collection of circles into \mathbb{R}^3 and two links are considered equivalent if they are ambient isotopic. Link diagrams are projections of links into a plane as shown, for instance, in Figure 2. We refer to Section 2 for standard definitions in knot theory.




or Khovanov homology [3], can be shown to be invariants by demonstrating that they are not modified by Reidemeister moves. From a computational point of view, the Reidemeister theorem allows for a discretisation of the space of possible transformations to consider when testing for knot equivalence. This fact is at the root of a straightforward algorithm to study algorithmic problems in knot theory: at the level of diagrams, apply Reidemeister moves in a random or brute-force manner until a desired property is verified.



■ **Figure 1** The three Reidemeister moves RI, RII, and RIII.

A primary example of such a knot theory problem is recognising the **unknot**, that is, the unique knot admitting a diagram with no crossings. This is a first instance of the fundamental problem of knot theory of deciding whether two knots are equivalent. It turns out that some unknot diagrams, called **hard unknot diagrams** or **culprits** [6, 21], exhibit an unwanted behaviour for the above algorithm. Namely, given an initial diagram \mathcal{D} of the unknot, the largest number of crossings of a diagram in any sequence of Reidemeister moves from \mathcal{D} to the 0-crossing diagram must be larger than in \mathcal{D} . This means one first needs to *add* crossings before being able to reach the 0-crossing diagram. The existence of such diagrams implies that it is not possible to untangle an unknot diagram by only applying Reidemeister moves that do not increase the number of crossings. An example of such a hard unknot diagram, called the **Goeritz culprit**, is shown in Figure 2. In [6], it is shown that, when diagrams in \mathbb{S}^2 are considered, at least one extra crossing is required to untangle this unknot. We do not know of systematic techniques or methods to prove easily that a given diagram of the unknot is hard. Instead, all proofs known to the authors resort to an exhaustive search in the graph of Reidemeister moves, which quickly becomes infeasible. Recently, new techniques based on reinforcement learning have been applied to find large numbers of hard unknot diagrams [2], but the proofs of hardness still involved exhaustive enumerations.



■ **Figure 2** The Goeritz culprit: using Reidemeister moves in \mathbb{S}^2 , one must add at least one extra crossing to untangle this unknot diagram.

Following the notations of [6, 21] we denote the number of crossings in a diagram \mathcal{D} by $\text{cr}(\mathcal{D})$. For diagrams \mathcal{D}_1 and \mathcal{D}_2 of equivalent links, and a sequence R of r Reidemeister moves transforming \mathcal{D}_1 into \mathcal{D}_2 , we define $\text{Top}(\mathcal{D}_1, R) = \max_{i \in \{0, 1, \dots, r\}} \{\text{cr}(\mathcal{D}^i) - \text{cr}(\mathcal{D}_1)\}$ where \mathcal{D}^i , $0 \leq i \leq r$, is the diagram after performing the first i moves of the sequence. The minimal number of extra crossings to pass from \mathcal{D}_1 to \mathcal{D}_2 is denoted by $\text{Add}(\mathcal{D}_1, \mathcal{D}_2)$,

which, formally, is the minimum of $\text{Top}(\mathcal{D}_1, R)$ taken over all sequences of Reidemeister moves R transforming \mathcal{D}_1 into \mathcal{D}_2 . Hence, $\text{Add}(\mathcal{D}_1, \mathcal{D}_2)$ is a lower bound on the number of crossings we must add in any sequence of Reidemeister moves performed on \mathcal{D}_1 to reach \mathcal{D}_2 .

Let \mathcal{D} be a diagram of the unknot and $\hat{\mathcal{D}}$ be its 0-crossing diagram. Naturally, \mathcal{D} is a hard unknot diagram if and only if $\text{Add}(\mathcal{D}, \hat{\mathcal{D}})$ is positive. This measure of complexity is called m in [6] (see also the *recalcitrance* in [21]). Studying these complexity measures and hard unknot diagrams is trickier than one might think. In fact, one of the purposes of [6] is to confirm or invalidate claims about previously known hard unknot diagrams using an exhaustive computer search over all possible sequences of Reidemeister moves. Still, the “hardest” known diagrams of the unknot have only been verified to require at least three extra crossings before they can be untangled. In contrast, the following conjecture is folklore.

► **Conjecture 1.** *Let m be an integer and let $\hat{\mathcal{D}}$ be the 0-crossing diagram of the unknot. Then there exists a diagram of the unknot \mathcal{D} with n crossings such that any sequence of Reidemeister moves from \mathcal{D} to $\hat{\mathcal{D}}$ passes through a diagram with at least $n + m$ crossings.*

Note that a proof of Conjecture 1, formulated in terms of recalcitrance, is claimed in [21], but concerns about the correctness of this proof are raised in [6].

In this article, as a possible step towards a proof of Conjecture 1, we shift our focus to **split links**. A link L is said to be split if there exists a sphere disjoint from L separating at least two link components of L . If such a sphere exists, there exists a link diagram in which two sublinks of L are disjoint: they are separated by a circle in the plane. Such a diagram is called a **split diagram**. By capping off the aforementioned circle with one disc above and one disc below the plane of projection, we can verify that conversely a split diagram witnesses a split link. Determining if a link is split is known as the **splitting problem**.


Given a split link L with a diagram \mathcal{D}_1 , we study $\text{Add}(\mathcal{D}_1, \mathcal{D}_2)$ where \mathcal{D}_2 is a split diagram of L . If the minimum of $\text{Add}(\mathcal{D}_1, \mathcal{D}_2)$ over all split diagrams \mathcal{D}_2 of L , called the **crossing-complexity** and denoted by $\text{CC}(\mathcal{D}_1)$, is positive, we call \mathcal{D}_1 a **hard split link**.

Our results. We exhibit a family of link diagrams $\mathcal{D}(p, q)$ of split links $L(p, q)$ with two unlinked sublinks. The first sublink $M = M(p, q)$ is made of two linked torus knots $T_{p,q}$, and the second is an unknot U surrounding one of the torus knots (see Figure 3 for an illustration). For any split diagram $\mathcal{D}'(p, q)$ of $L(p, q)$, we prove Theorem 2, implying that $\text{Add}(\mathcal{D}(p, q), \mathcal{D}'(p, q)) = \Omega(\min(p, q))$. More precisely, we have the following main theorem.

► **Theorem 2.** *For all $n \geq 2$, any sequence of Reidemeister moves transforming diagram $\mathcal{D}(n, n+1)$ of the link $L(n, n+1)$ with $2n^2 + 2$ crossings into a split diagram passes through a diagram with at least $2n^2 + \frac{2}{3}n$ crossings. In particular, there exist hard split links of arbitrarily large crossing-complexity.*

This is to our knowledge the first construction of split links with super-constant crossing-complexity, and we cannot rely on exhaustive search methods to prove Theorem 2. Instead, we develop new techniques to lower-bound the crossing-complexity. It was proved by Dynnikov [12, Theorem 2] that $\text{CC}(D) = O(|D|^2)$, where $|D|$ denotes the number of crossings in a diagram D , so there is still a significant gap between this upper bound and the $\Omega(\sqrt{|D|})$ lower bound provided by our theorem.

► **Remark 3.** Theorem 2 is one of the many knot-theoretical statements that seems intuitively clear but is surprisingly delicate to prove: inspecting Figure 3, “clearly” the only way to split U from the two components of M is through an overlay of U on top of one of these



■ **Figure 3** The link diagram $\mathcal{D}(p, q)$, $(p, q) = (7, 8)$: two linked torus knots $T_{7,8}$ and an unknot U .

components, yielding the claimed increase in the number of crossings. In fact, the following idea to prove Theorem 2 may seem straightforward: if there is a sequence of Reidemeister moves splitting U from M without adding too many crossings, then we can find a continuous family of planes $(P_t)_{t \in [0,1]}$ in \mathbb{R}^3 sweeping one of the torus knots in M such that none of the planes P_t intersects M in too many points. Known results on *bridge position* [30, 31] and *thin position* [14, 29] of torus knots then imply that such a family cannot exist².

However, this proof idea has two issues. First, during a sequence of Reidemeister moves, the unknot U might come to intersect itself, making it difficult to define a continuous family of planes that mimics its intersections with M . Second, the Reidemeister moves might lead U to move back and forth, leading to a non-monotone behaviour of the planes P_t , which is not allowed in the aforementioned positions commonly used in knot theory. This leads us to rely on more advanced tools to prove Theorem 2.

Overview of the proof. The proof of Theorem 2 works by contradiction. That is, we start by assuming that there exists a sequence of Reidemeister moves R transforming $\mathcal{D}(n, n+1)$ into a split diagram and such that $\text{Top}(\mathcal{D}(n, n+1), R)$ remains small. This implies that, throughout performing R , the number of crossings involving U and M always remains small.

The main tool that we rely on is the framework of bubble tangles introduced by the first and second authors in [26]. More precisely, we use the evolution of U throughout R to define a collection of 2-dimensional spheres that continuously sweep \mathbb{S}^3 (a *sweep-out*) and that – according to our assumptions – all have a small number of intersections with M . This sweep-out resembles the sphere decompositions of [26], but presents two notable differences. On the one hand it is simpler: it is linear and features no double bubbles. On the other hand, it is not monotone: a sphere involved in this sweep-out may go back-and-forth, while this behaviour is not allowed in the sphere decompositions from [26]. Despite this last difference, the bubble tangles, which are obstructions to thin sphere decompositions developed in [26], are versatile enough to show that the existence of this sweep-out leads to a contradiction.

² This argument is the basis of the proof in [8] showing that torus knot diagrams have high tree-width.

As mentioned in Remark 3, building the sweep-out is not straightforward. Intuitively, one would like to lift each unknot U to a sphere by capping it off above and below the diagram. However, the projection of the unknot U may intersect itself during the sequence of Reidemeister moves, which complicates the process. We alleviate this problem using results and methods from Chambers and Liokumovich [7] to transform homotopies of curves on a Riemannian surface into isotopies of similar length. The connection is the following: we think of the projection of the link M as a discrete metric for curves in the projection plane. In this discrete model, the **length** of a curve is given by its number of intersections with M , similarly to the cross-metric model commonly used in computational topology of surfaces (see for example [11]). Our assumptions imply that there exists a homotopy of the projection of U where intermediate curves all have small length. The techniques of [7] show that this implies that there also exists an isotopy of the projection of U with the same bounds on lengths. Since such an isotopy must consist of simple curves, we can then lift it into a sweep-out of \mathbb{S}^3 with 2-spheres, which all have a controlled number of intersections with M . However, a subtlety is that we cannot use the results of [7] out of the box. Indeed, the discrete metric defined by the link M is not fixed but evolves with the Reidemeister moves in R . In Section 3 we explain why the proof techniques in [7] can be adapted to deal with this issue.

The tools of the first and second authors [26] are then used to find a contradiction between the existence of the sweep-out and the topological properties of our link M . Namely, $L(n, n+1)$ consists of two torus knots $T_{n, n+1}$, which are embedded on tori in a specific way (they both have high *compression-representativity*). Hence, they can be used to define two **bubble tangles** using [26, Theorem 1.2] (In a nutshell, a bubble tangle is a way of choosing a “small side” for each sphere of the sweep-out, where the intuition is that the small side should be easy to sweep, see Section 2 for a formal definition). We then prove that, since in the initial diagram U lies in-between the two tori, different small sides are chosen for the corresponding sphere by each of the two bubble tangles. But the existence of the sweep-out with a small number of intersections with M forces the small sides of both bubble tangles to agree, leading to a contradiction.

Related work. The splitting problem has been studied several times as a useful and easier problem for understanding the unknot recognition problem [12, 23]. In 1961, Haken used normal surface theory to show that it is decidable [15]. Later, it was determined to be in **NP** [16] and also in co-**NP** [24, Theorem 1.6].

Several decision problems related to the splitting problem have been studied as well. For instance, the problem of deciding whether changing at most k crossings can transform a link diagram into the diagram of a split link is known to be **NP-hard** [22], and Lackenby provided an algorithm to detect links which can be split using exactly one crossing change [25].

Another natural question is to ask for the minimal number of Reidemeister moves needed to split a diagram. An exponential bound for this number was first provided by [19], and this bound was later greatly improved by Lackenby in [23], where he provided a polynomial bound using a combination of normal surface theory [16] and Dynnikov’s work on grid diagrams [12]. There is a quadratic lower bound on the number of moves needed to untangle a specific unknot diagram in [17], and it was shown in [9, 10] that finding the shortest sequence of Reidemeister moves to untangle an unknot is **NP-hard**.

Organisation of this paper. After going through our setup in Section 2, we explain how to use the results of [7] in Section 3. This step is crucial for our definition of sweep-outs. Then, we exploit obstructions from [26] to prove Theorem 2 in Section 4.

2 Setup and definitions

Knots and links. Many concepts from this paper come from knot theory: while we strive to be as self-contained as possible, we refer to standard textbooks [5, 28] for an introduction to this topic and to Hatcher [18] for background on algebraic topology. Throughout this article, we work in the piecewise-linear (PL) category, which means that all the objects and functions that we consider are piecewise-linear with respect to a fixed polyhedral decomposition of the ambient space (generally \mathbb{S}^3). Two embeddings i_1 and i_2 in a topological space S are **(ambient) isotopic** if there exists a continuous family of homeomorphisms $h : S \times [0, 1] \rightarrow S$ such that $h(i_1, 0) = i_2$ and $h(\cdot, 1)$ is the identity. A **knot**, resp. a **link**, is an embedding of the circle \mathbb{S}^1 , resp. of a disjoint union of circles, into \mathbb{S}^3 . In the following, we introduce the main definitions for knots for simplicity, but they apply identically to links. Two knots K_1 and K_2 are considered to be equivalent if they are isotopic. Since every knot misses at least one point of \mathbb{S}^3 , via stereographic projection we can equivalently consider knots to be embedded in \mathbb{R}^3 , and we freely switch between these two perspectives. The **unknot** is, up to equivalence, the unique embedding of \mathbb{S}^1 in \mathbb{R}^3 with image a triangle. A **torus knot** $T_{p,q}$ is a knot embedded on a surface of an unknotted torus \mathbb{T} in \mathbb{S}^3 , for example a standard torus of revolution. It winds p times around the revolution axis, and q times around the core of the torus. We refer to Figure 3 for an illustration of two torus knots and an unknot.

A **knot diagram** \mathcal{D} is a planar four-regular graph, where each crossing is decorated to indicate which strands are above and below. From such a diagram, one can easily obtain the data of a knot $K \hookrightarrow \mathbb{R}^3$ and a linear **projection** map $p : \mathbb{R}^3 \rightarrow P \simeq \mathbb{R}^2$ so that $p(K) = \mathcal{D}$ (respecting the decorations), where P is a plane of \mathbb{R}^3 called the **projection plane**. The **crossing number** of a knot is the minimal number of crossings among all of its diagrams. The Reidemeister theorem [32] shows that two diagrams represent equivalent knots if and only if they can be connected by a sequence of planar isotopies and local moves called **Reidemeister moves** and pictured in Figure 1.

Finally, recall that a **homotopy** between two simple closed curves of a surface Σ , $\gamma_0 : \mathbb{S}^1 \hookrightarrow \Sigma$ and $\gamma_1 : \mathbb{S}^1 \hookrightarrow \Sigma$ is a continuous map $\gamma : \mathbb{S}^1 \times [0, 1] \rightarrow \Sigma$ such that $\gamma(\mathbb{S}^1, 0) = \gamma_0$ and $\gamma(\mathbb{S}^1, 1) = \gamma_1$. In particular, closed curves are allowed to self-intersect in a homotopy, while this is disallowed in an isotopy. Throughout this paper, to distinguish between links and their projections, we use calligraphic letters to designate diagrams, and capital letters to designate links so that a diagram of the link M is denoted by \mathcal{M} .

► **Remark 4.** We work with the standard setting of knot diagrams in \mathbb{R}^2 , but it is also possible to work with diagrams in \mathbb{S}^2 , which therefore allow for more Reidemeister moves (this is the perspective taken in [6]). Our results also hold in that setting, since a knot diagram in \mathbb{S}^2 lifts to a knot $K \hookrightarrow \mathbb{S}^2 \times [-\varepsilon, \varepsilon] \subseteq \mathbb{R}^3$, with the natural projection map $p : \mathbb{S}^2 \times [-\varepsilon, \varepsilon], (s, t) \mapsto (s, 0)$. The definitions of the spheres obtained from the diagrams in Section 4 can be directly adapted to this setting, and the rest of the proof is identical.

Our link diagrams. Throughout this article, we write $L(p, q)$ for the split link consisting of two linked torus knots $T_{p,q}$, denoted by $M(p, q)$ and shown in Figure 3, and an unlinked unknot component U . We consider two diagrams of $L(p, q)$. The first, denoted by $\mathcal{D}(p, q)$, is shown in Figure 3, and the second diagram $\mathcal{D}'(p, q)$ is any split link diagram of $L(p, q)$.


Let R be a sequence of Reidemeister moves turning $\mathcal{D}(p, q)$ into $\mathcal{D}'(p, q)$, such that $\text{Top}(\mathcal{D}(p, q), R) \leq k$, $k \geq 0$. Our goal is to prove that k cannot be smaller than a function depending only on p and q . Since the values of p and q are mostly fixed and have little influence on our arguments, we mostly omit the parameters (p, q) from L , M , \mathcal{D} , and \mathcal{D}' .

The following lemma directly follows from known results on torus knots.

► **Lemma 5.** *Let $n \geq 2$, and let \mathcal{D}' be a link diagram equivalent to $\mathcal{D}(n, n+1)$ by Reidemeister moves. Then $\text{cr}(\mathcal{D}') \geq 2n^2$, that is, \mathcal{D}' has at least $2n^2$ crossings.*

Proof. Let \mathcal{D}' be a diagram of $L = L(n, n+1)$. First note that, since $n \geq 2$, it follows from a theorem of Murasugi [27, Proposition 7.5] that each of the torus knot components $T_{n, n+1}$ of L has at least $(n+1) \times (n-1) = n^2 - 1$ crossings. Since the two torus knots are linked, they share at least 2 crossings in each diagram of L , and it follows that $\text{cr}(\mathcal{D}) \geq 2n^2$. ◀


From a sequence of Reidemeister moves to continuous operations. As detailed above, we work with a sequence of Reidemeister moves R turning the link diagram \mathcal{D} of L into the split diagram \mathcal{D}' . From this sequence of Reidemeister moves, we can obtain an ambient isotopy $\Phi_R : \mathbb{R}^3 \times [0, 1] \rightarrow \mathbb{R}^3$ and a projection $p : \mathbb{R}^3 \rightarrow \mathbb{R}^2$ so that the diagrams $p(\Phi_R(L, t))$ follow the evolution of \mathcal{D} under the moves R , and in particular the combinatorial types of the diagrams $p(\Phi_R(L, t))$ only change for a finite number of values of t , one for each Reidemeister move. The projection p is regular except at the **critical times** of R , which are times where the projection $p \circ \Phi_R(L, t)$ displays a tangency or a triple point, see Figure 4 for an illustration. For any non-critical time t , we write $L_t = \Phi_R(L, t)$ and we denote the diagram defined by $p(L_t) = p(\Phi_R(L, t))$ by \mathcal{D}_t . Naturally, we have $\mathcal{D}_0 = \mathcal{D}$ and $\mathcal{D}_1 = \mathcal{D}'$.



■ **Figure 4** Critical times corresponding to Reidemeister moves RII (left) and RIII (right).

Note that the definition of $\text{Top}(\mathcal{D}, R)$ naturally coincides with $\sup_{t \in [0, 1]} \{\text{cr}(\mathcal{D}_t) - \text{cr}(\mathcal{D}_0)\}$. Indeed, the diagram \mathcal{D}_t at critical times has fewer intersections than one of $\mathcal{D}_{t+\epsilon}$ or $\mathcal{D}_{t-\epsilon}$ for ϵ small enough. Furthermore, $\text{cr}(\mathcal{D}_t)$ is constant for all t between two critical times.

In Section 3 we consider the movements of U and of M under the Reidemeister moves separately. We use $\mathcal{M}_t = p(\Phi_R(M, t))$ as a shorthand for the diagram of the sublink $M \subset L$ at time t . For \mathcal{U} we consider the homotopy $\phi_U : \mathbb{S}^1 \times [0, 1] \rightarrow \mathbb{R}^2$ in the plane induced by the projection $p(\Phi_R(U, t))$. We denote the corresponding curves by $\mathcal{U}_t = \phi_U(\mathbb{S}^1, t)$ and emphasise that we consider these as immersions of closed curves in the plane. That is, we forget which strand is over which at each self-crossing of \mathcal{U}_t . See Figure 5 for a summary of this setup.



■ **Figure 5** Definition of our homotopies from the sequence of Reidemeister moves R .

3 From homotopies to isotopies

We work with the definitions of Section 2. We start with a sequence of Reidemeister moves R turning \mathcal{D} into \mathcal{D}' , and consider the induced homotopy ϕ_U of the link component \mathcal{U} in the plane of projection. The goal of this section is to use results from [7] in order to locally alter the images $\mathcal{U}_t = \phi_U(\mathbb{S}^1, t)$ into simple closed curves, and hence to obtain an isotopy taking \mathcal{U}_0 to \mathcal{U}_1 . After these modifications, the altered simple closed curves representing U still sweep over the remainder \mathcal{M}_t of the diagrams \mathcal{D}_t , but this sweep no longer necessarily corresponds to a sequence of Reidemeister moves on the original link L .

Nevertheless, in the next section we use the isotopy from \mathcal{U}_0 to \mathcal{U}_1 to define a family of 2-spheres in \mathbb{S}^3 sweeping through M . This setup then allows us to use bubble tangles to obstruct small values of $\text{Top}(\mathcal{D}, R)$ in the initial sequence of Reidemeister moves R .

The key points of this process are given by the following statement.

► **Proposition 6.** *Let R be a sequence of Reidemeister moves turning \mathcal{D} into the split diagram \mathcal{D}' such that for all $t \in [0, 1]$, $\text{cr}(\mathcal{D}_t) \leq m$ for some integer $m \geq 0$. Then there exists an ambient isotopy $\Phi' : \mathbb{S}^3 \times [0, 1] \rightarrow \mathbb{S}^3$ and an isotopy $h : \mathbb{S}^1 \times [0, 1] \rightarrow \mathbb{R}^2$ such that*

1. $h(\mathbb{S}^1, 0) = \mathcal{U}_0$ and $h(\mathbb{S}^1, 1) = \mathcal{U}_1$; and
2. *for all $t \in [0, 1]$, the total number of crossings in the overlay of $p(\Phi'(M, t))$ and $h(\mathbb{S}^1, t)$ in \mathbb{R}^2 is at most m .*

We emphasise that in the second item of Proposition 6, we only consider the sublink M in $\Phi'(M, t)$, and not the entire link. That is, the proposition provides an ambient isotopy for M in \mathbb{R}^3 and an isotopy for $p(U)$ in \mathbb{R}^2 , while preserving a bound on the total number of intersections of $p(\Phi'(M, t))$ and $h(\mathbb{S}^1, t)$ in the plane of projection. This proposition follows from the techniques of Chambers and Liokumovich developed in Chapter 2 of [7], the proof is detailed in the full version of the article for completeness. We first state one of their main results.

► **Definition 7** (Chambers, Liokumovich [7, Definition 1.3]). *Two curves α and β on a Riemannian 2-manifold M are ϵ -image equivalent, $\alpha \sim_\epsilon \beta$, if (i) there exists a finite collection of disjoint intervals $\bigsqcup_{i=1}^n I_i \subset \mathbb{S}^1$ such that $\|\alpha(\mathbb{S}^1 \setminus \bigsqcup I_i)\| + \|\beta(\mathbb{S}^1 \setminus \bigsqcup I_i)\| < \epsilon$, and (ii) there exists a permutation σ of $\{1, \dots, n\}$ and a map $f : \{1, \dots, n\} \rightarrow \{0, 1\}$, such that $\alpha|_{I_i} = (-1)^{f(i)} \beta|_{I_{\sigma(i)}}$ for all i . Here, $\|\alpha\|$ denotes the length of the curve α .*


The main result we use for the proof of Proposition 6 is given by the following statement.

► **Theorem 8** (Chambers, Liokumovich [7, Theorem 1.1']). *Suppose that γ is a smooth homotopy of closed curves on a Riemannian 2-manifold M and that γ_0 is a simple closed curve. Then, for every $\epsilon > 0$, there exists an isotopy $\bar{\gamma}$ such that $\bar{\gamma}_0 = \gamma_0$ and $\bar{\gamma}_1$ is ϵ -image equivalent to a small perturbation of γ_1 . Additionally, for every t , there exists a t' such that $\bar{\gamma}_t$ is ϵ -image equivalent to a small perturbation of $\gamma_{t'}$. If γ_1 is simple, or is a point, then this homotopy also ends at γ_1 , up to a change in orientation.*

Informally, we think of Definition 7 and Theorem 8 as follows. When a curve α self-intersects in the plane, each crossing point can be **resolved** in one of two ways by reconnecting the endpoints in a small ball around the crossing point, see Figure 6 on the right. When all the crossings have been resolved in some way and we obtain a simple closed curve α' , we say that α' is a **resolution** of α . The theorem states that if we have a homotopy γ on a surface

between two simple curves γ_0 and γ_1 , one can obtain an *isotopy* $\bar{\gamma}$ between γ_0 and γ_1 ³ where each intermediate curve $\bar{\gamma}_t$ is a resolution of some intermediate curve $\gamma_{t'}$. Note that, here, the times t and t' may not coincide.

If M is endowed with a metric (for example, a Riemannian one), the lengths of γ_t and of $\bar{\gamma}_{t'}$ differ by an arbitrarily small quantity. Hence, Theorem 8 immediately implies that, for all $\varepsilon > 0$, if there exists a homotopy between two simple curves γ_0 and γ_1 where each intermediate curve has length at most ℓ , then there also exists an isotopy between γ_0 and γ_1 where each intermediate curve has length at most $\ell + \varepsilon$.




■ **Figure 6** Left: A diagram of the unknot U . Middle: A projection \mathcal{U}_t of this unknot, where the crossing information has been forgotten. Right: A resolution of \mathcal{U}_t .

Theorem 8 can be applied as a black-box to prove Proposition 6 in the particular case where the sublink M stays invariant throughout the sequence of Reidemeister moves R , i.e., $\Phi_R(M, t) = \Phi_R(M, 0)$ for all $t \in [0, 1]$. Indeed, in this case, we can think of \mathcal{M}_t as a discrete metric measuring the length of a curve \mathcal{U}_t by its number of intersections with \mathcal{M}_t . More formally, we can take γ to be the homotopy ϕ_U between \mathcal{U}_0 and \mathcal{U}_1 , which are both simple closed curves by the definition of the link diagram \mathcal{D} . Applying Theorem 8 provides us with an isotopy h between \mathcal{U}_0 and \mathcal{U}_1 where all intermediate curves \mathcal{U}'_t are obtained from resolving intersections of some $\mathcal{U}_{t'}$. In particular, for all $t \in [0, 1]$, the number of intersections between \mathcal{U}'_t and \mathcal{M}_t is at most the number of intersections between $\mathcal{U}_{t'}$ and \mathcal{M}_t , which is at most k since $\mathcal{M}_t = \mathcal{M}_0$ does not depend on t .

A careful reading of [7] shows that the general case of Proposition 6, where the diagram \mathcal{M}_t of the sublink M evolves during the sequence of Reidemeister moves, can also be obtained using the same proof techniques: The basic idea of the proof of Theorem 8 is to first decompose the homotopy of γ into a sequence of local moves, and then replace each curve γ_t by one of its resolutions.

For the first step, they track discrete times where the self-intersection pattern (i.e., the homeomorphism type) of γ_t changes. By an argument similar to the proof of the Reidemeister theorem, one can assume that this only happens at **critical events**, when a curve γ_t undergoes a homotopy move, which is a transformation analogue to a Reidemeister move but without any crossing information. This step is formalised by their Proposition 2.1 and Lemma 2.2. Note that this step is unnecessary for us, since by construction the local transformations undergone by ϕ_U are projections of Reidemeister moves.

³ Or γ_1 with its orientation reversed. Since it is safe to disregard orientations for our purpose, we always consider curves up to orientation reversal.



■ **Figure 7** (A) Reidemeister III move. (B) a problematic resolution of \mathcal{U}_t . (C) The unknot \mathcal{U}_t (black) seen as a curve sweeping through \mathcal{M} (blue).

Between these critical events, any homotopy of a curve γ_t can be applied similarly to any of its resolutions, yielding an isotopy. Hence, the idea is to replace each of these Reidemeister moves by a resolution of the move, as explained by their Figure 2. However, doing so in a naive way runs into discontinuity issues, a basic example of which is detailed in [7, Example 2], which is associated to their Figures 3 and 4. Therefore, the authors provide a more intricate workaround: the key to the proof of Theorem 8 is to show how to choose the correct resolutions and connect their isotopies together. This is achieved by defining an auxiliary graph of resolutions (see their Figure 7), synthesising how they are connected by local isotopies. The precise definition for this graph follows their Figure 8. The proof is then finalised by finding a path through this graph by using the handshaking lemma [13].

In our case, the critical events are exactly the times t when \mathcal{U}_t undergoes a Reidemeister move. In-between these critical events, there are other Reidemeister moves involving either (i) both \mathcal{U}_t and \mathcal{M}_t , or (ii) only \mathcal{M}_t . Note that in case (i), the Reidemeister move only changes the relative position of \mathcal{U}_t and \mathcal{M}_t , and hence such a move can be treated as leaving the isotopy type of \mathcal{M}_t unchanged. Therefore the homotopy of \mathcal{U}_t can be used as an isotopy of any of its resolutions in this case. In case (ii), \mathcal{M}_t changes, but \mathcal{U}_t , considered up to isotopy, does not. Therefore, by applying the same Reidemeister moves to \mathcal{M}_t , any motion of \mathcal{U}_t between critical events can be applied to any of its resolutions while preserving the number of intersections with \mathcal{M}_t . These motions can then be connected using the same handshaking argument as in the proof of Theorem 8. Summarising, the proof technique directly adapts to the case of an evolving metric, as long as these evolutions are applied appropriately throughout the new sequence of isotopies.

► **Remark 9.** We emphasise that in this section, we apply the framework of Theorem 8 to the *projection* of U without crossing information, which is thus considered as a closed, non-embedded, curve in the projection plane \mathbb{R}^2 . Instead, it might be tempting to apply these techniques *directly at the level of Reidemeister moves*. This way one may hope to prove that in any sequence of Reidemeister moves splitting \mathcal{U} from \mathcal{M} , one can assume that \mathcal{U} remains simple while preserving a bound on the number of added crossings. However, this proof strategy fails because some resolutions applied to \mathcal{U} can block the application of RIII moves involving \mathcal{U} and \mathcal{M} . Such a case is pictured in Figure 7(B).

4 Sweep-outs, bubble tangles and proof of the lower bound

For the remainder of this article, we use Proposition 6 to assume that the homotopy $\phi_U : \mathbb{S}^1 \times [0, 1] \rightarrow \mathbb{R}^2$ sweeping U across M is an isotopy. In other words, let h be the isotopy of \mathcal{U} in \mathbb{R}^2 and Φ' be the ambient isotopy of M in \mathbb{S}^3 from Proposition 6, and for all $t \in [0, 1]$, we replace $\phi_U(\mathbb{S}^1, t)$ by $h(\mathbb{S}^1, t) = \mathcal{U}_t$ which is a simple closed curve.

4.1 A sweep-out of 2-spheres

For each t , we associate a 2-sphere S_t to the simple curve \mathcal{U}_t , by extending \mathcal{U}_t infinitely towards the direction of the projection p . Seen in \mathbb{S}^3 this produces a torus pinched at ∞ , and cutting the surface at ∞ yields S_t (see Figure 8). By construction, S_t intersects $M_t = \Phi'(M, t)$ only in the pre-images of $\mathcal{U}_t \cap \mathcal{M}_t$ by the projection p . Altogether, this produces a continuous family of spheres S_t sweeping a continuous family of links M_t . By applying the ambient isotopy Φ'^{-1} to S_t and M_t , we can assume that M_t is fixed. We slightly abuse notation and make this assumption, while still denoting the family of spheres by S_t .




Figure 8 Gluing two infinite annuli on U and cutting the resulting pinched torus at ∞ in \mathbb{S}^3 .

We now relate $\text{Top}(\mathcal{D}(n, n+1), R)$ to our continuous family of spheres $\{S_t\}_{t \in [0, 1]}$.

► **Lemma 10.** *We have $\sup_{t \in [0, 1]} |S_t \cap M_t| - 2 \leq \text{Top}(\mathcal{D}(n, n+1), R)$.*

Proof. For \mathcal{M} and \mathcal{M}' two components of a link diagram, we denote by $\text{cr}(\mathcal{M}, \mathcal{M}')$ the number of crossings involving strands of \mathcal{M} and \mathcal{M}' . If $\mathcal{M} = \mathcal{M}'$, then $\text{cr}(\mathcal{M}, \mathcal{M}')$ is the number of self-crossings of \mathcal{M} .

Since \mathcal{U}_t is simple for all times t , we have $\text{cr}(\mathcal{D}_t) = \text{cr}(\mathcal{M}_t, \mathcal{U}_t) + \text{cr}(\mathcal{M}_t, \mathcal{M}_t)$. By definition, \mathcal{D}_t is the diagram obtained from $\mathcal{D}(n, n+1)$ at time t of the sequence of Reidemeister moves R , and is hence equivalent to $\mathcal{D}(n, n+1)$. It follows from Lemma 5, that $\text{cr}(\mathcal{M}_t, \mathcal{M}_t) \geq 2n^2$ at all times t . Furthermore, $\text{cr}(\mathcal{M}_t, \mathcal{U}_t) = |S_t \cap M_t|$ by construction of S_t . By definition of $\mathcal{D}(n, n+1) = \mathcal{D}_0$, we have $\text{cr}(\mathcal{D}_0) = 2n^2 + 2$. Hence, $\text{cr}(\mathcal{D}_t) - \text{cr}(\mathcal{D}_0) \geq 2n^2 + |M_t \cap S_t| - 2n^2 - 2$.

Thus, $\sup_{t \in [0, 1]} |S_t \cap M_t| - 2 \leq \sup_{t \in [0, 1]} \{\text{cr}(\mathcal{D}_t) - \text{cr}(\mathcal{D}_0)\} = \text{Top}(\mathcal{D}(n, n+1), R)$. ◀

4.2 Bubble tangles

The aim for the remainder of this section is to use results from [26] to exhibit an obstruction to $|S_t \cap M_t|$ remaining small throughout the sweep-out defined in Section 4.1.


The key objects that we rely on are bubble tangles. Before introducing those, we need the following definitions. A **double bubble** consists of two spheres intersecting on a disc (see the left of Figure 9 for an illustration). The complement of a double bubble consists of

three balls, and we say that they are **induced** by the double bubble. A 3-ball in \mathbb{S}^3 is said to be **L -trivial**, if it intersects $L \subset \mathbb{S}^3$ in a single unknotted segment or in the empty set, see Figure 9 on the middle for an example and on the right for a non-example.

► **Definition 11** ([26], Definition 2.3). *Let $L \in \mathbb{S}^3$ be a link and let $k \in \mathbb{N}$. A **bubble tangle** \mathcal{T} of order $k \geq 2$, is a collection of closed balls in \mathbb{S}^3 such that:*

1. *For all balls $B \in \mathcal{T}$, we have $|\partial B \cap L| < k$.*
2. *For all 2-spheres $S \in \mathbb{S}^3$ transverse to L , if $|S \cap L| < k$, then exactly one of the two⁴ balls B_i , $i \in \{1, 2\}$ with boundary S is in \mathcal{T} .*
3. *For all triples of balls $\{B_1, B_2, B_3\}$, if $\{B_1, B_2, B_3\}$ induces a double bubble transverse to L , then $\{B_1, B_2, B_3\} \notin \mathcal{T}$.*
4. *For every closed ball B in \mathbb{S}^3 , if B is L -trivial and $|\partial B \cap L| < k$, then $B \in \mathcal{T}$.*

Informally, we think of a bubble tangle as a way to choose, for each sphere S having less than k intersections with L , one of the two balls that it bounds. We refer to this ball as the **small side** of S , with the intuition that this small side intersects L in a simpler pattern than the other side. The essential property of a bubble tangle is the requirement that we cannot cover all of \mathbb{S}^3 using three small sides of a triple of spheres arranged in a double bubble. One of the main results of [26] states that if a bubble tangle of L of order k exists, any sweep-outs of \mathbb{S}^3 by 2-spheres contains at least one sphere with at least k intersections with L . As discussed in the introduction, our sweep-outs differ from those in [26]. However, we show that despite their differences, a bubble tangle can also be used to preclude the existence of a sweep-out S_t having few intersections with M .




■ **Figure 9** Left: a double bubble. Middle: an L -trivial ball. Right: a ball that is not L -trivial.


We denote by M^1 and M^2 the two torus links composing M . By definition, each of them can be embedded on a standard torus \mathbb{T} as pictured in Figure 11. A simple curve on $\Sigma \subseteq \mathbb{S}^3$ is **compressible** if it is non-contractible on Σ and bounds a disk in $\mathbb{S}^3 \setminus \Sigma$. The **compression-representativity** of a knot K embedded on a surface Σ is the minimum number of intersections between K and a compressible curve in Σ . The compression-representativity of M_1 and M_2 is n , see [26, Section 5]. Note that, since M^1 and M^2 are linked, their underlying tori intersect, but this is of no consequence for us. By [26, Theorem 1.2], there exist two bubble tangles $\mathcal{T}^1, \mathcal{T}^2$ of order $\frac{2}{3}n$, where each \mathcal{T}^i is the bubble tangle defined by M^i while completely forgetting about the other torus knot. Thus, \mathcal{T}^i is a bubble tangle *induced* by the torus on which M^i is embedded. Such bubble tangles induced by a surface are called **compression bubble tangles**: given a link embedded on a surface Σ , the small sides of the bubble tangle are the balls B which are either disjoint from Σ or intersect Σ trivially, i.e., the inclusion $i : B \cap \Sigma \rightarrow \Sigma$ is π_1 -trivial, which means that it

⁴ By the PL Schoenflies theorem [4, Theorem XIV.1], a PL 2-sphere in \mathbb{S}^3 bounds exactly two balls.

induces a trivial map on fundamental groups. For example, the annulus and disc obtained by intersecting the torus and green ball of Figure 10 have a π_1 -trivial inclusion in the torus, but the annulus stemming from the intersection between the red ball and the torus does not.



■ **Figure 10** The green ball intersects the torus trivially while the red ball does not.



■ **Figure 11** Link components M^1 and M^2 embedded on tori, and sphere S_0 .


4.3 Proof of Theorem 2

In order to prove Theorem 2, we assume that $\text{Top}(\mathcal{D}(n, n+1), R) < k = \frac{2}{3}n - 2$. By virtue of Lemma 10, it follows that $|M_t^i \cap S_t| \leq |M_t \cap S_t| < \frac{2}{3}n$ for $i \in \{1, 2\}$. Now, changes to $|M_t \cap S_t|$ can only happen at critical times where the number of crossings involving \mathcal{U}_t in \mathcal{D}_t increases or decreases. In particular, we can ignore RIII moves between \mathcal{U}_t and \mathcal{M}_t . Since we only consider moves involving both \mathcal{U}_t and \mathcal{M}_t , only RII moves are relevant for our study. We denote the critical times of RII moves involving both \mathcal{U}_t and \mathcal{M}_t by $(t_j)_{1 \leq j \leq s} \subset [0, 1]$. These times are the only times where S_t is not transverse to M_t , but **finitely tangent**, meaning that M_t is tangent to S_t , but their number of intersections is still finite.

Recall that, up to applying ϕ'^{-1} , we assume that for all $t \in [0, 1]$ M_t is fixed; we denote it by M . By 2, for all times $t \in [0, 1] \setminus \{t_j\}_{1 \leq j \leq s}$, S_t has a small side corresponding to a 3-ball $B_t^i \in \mathcal{T}^i$. We study how these small sides evolve throughout the sweep-out. Special attention must be paid to tangencies at times $\{t_j\}_{1 \leq j \leq s}$, where small sides are not defined. For convenience, we denote by $\Theta : \mathbb{S}^3 \times [0, 1] \rightarrow \mathbb{S}^3$ an isotopy describing the evolution of S_t , i.e., such that $\Theta(S_0, t) = S_t$.

We start by defining the **agreement** between bubble tangles \mathcal{T}^1 and \mathcal{T}^2 as a map $a : [0, 1] \setminus \{t_j\}_{1 \leq j \leq s} \rightarrow \{0, 1\}$, such that $a(t) = 1$ if $B_t^1 = B_t^2$ and 0 otherwise. With this definition, and under the assumption that $\text{Top}(\mathcal{D}(n, n+1), R) < k$, we can infer that $a(0) = 0$ and $a(1) = 1$: First, note that $a(1) = 1$. Indeed, \mathcal{U}_1 is disjoint from \mathcal{M} so that S_1 is disjoint from M . It follows that one side of S_1 is an empty ball with respect to both M^1 and M^2 and thus $a(1) = 1$ by 4. Second, the small side of S_0 in \mathcal{T}^1 is, by 4, the ball not containing M^1 . For \mathcal{T}^2 , as it is illustrated by Figure 11, S_0 contains an M^2 -trivial ball and M^1 on its small side, and the remaining of M^2 on the other side. It follows that the small sides disagree at time 0, and hence $a(0) = 0$.

We want to show that a stays constant on its domain of definition. For this, we introduce the following definition: Two disjoint spheres $S, S' \subset \mathbb{S}^3$ are **braid-equivalent** with respect to some link $L \subset \mathbb{S}^3$, if L forms a braid in the product region between S and S' , or, equivalently, if the product region between S and S' is homeomorphic to $S_\ell \times [0, 1]$ where S_ℓ is the 2-sphere with ℓ holes, $\ell \geq 0$, see Figure 12 for an illustration.



■ **Figure 12** A cross-section view of three spheres and a link. The two outermost spheres are braid-equivalent, but the innermost one is not braid-equivalent to the others.

We have the following statement for two disjoint braid-equivalent spheres.

► **Lemma 12** ([26], Lemma 3.2). *Let \mathcal{T} be a bubble tangle and $S, S' \subset \mathbb{S}^3$ be two braid-equivalent spheres. Define $\mathbb{S}^3 \setminus S = \{B_1, B_2\}$ and $\mathbb{S}^3 \setminus S' = \{B'_1, B'_2\}$ such that $B_1 \subset B'_1$. If $B_1 \in \mathcal{T}$ then $B'_1 \in \mathcal{T}$.*


The compression bubble tangles furthermore have the following property.

► **Lemma 13.** *Let \mathcal{T} be a compression bubble tangle of L , induced by an embedding of L into some surface $\Sigma \subset \mathbb{S}^3$, and let k be the order of \mathcal{T} . Then, for two closed 3-balls $A, B \subset \mathbb{S}^3$, if $A \in \mathcal{T}$, $B \subset A$, and $|\partial B \cap L| < k$, then $B \in \mathcal{T}$. That is, \mathcal{T} is stable by inclusion up to 1.*

Proof. Let A and B be two closed balls of \mathbb{S}^3 such that $|\partial A \cap L|, |\partial B \cap L| < k$, and $A \in \mathcal{T}$. By definition, $A \cap \Sigma$ is π_1 -trivial. Since $B \subset A$, we have $B \cap \Sigma \subset A \cap \Sigma$ so that the inclusion morphisms i_* satisfy $\pi_1(B \cap \Sigma) \rightarrow \pi_1(A \cap \Sigma) \rightarrow 0$. Hence, $B \cap \Sigma$ is π_1 -trivial and $B \in \mathcal{T}$. ◀

To handle the back-and-forths of our sweep-out, we introduce a new equivalence relation on spheres. Two intersecting spheres $S, S' \subset \mathbb{S}^3$ are **intersection-equivalent** if there exists an isotopy between them which stays constant on their intersection $S \cap S'$, see Figure 13. Note that, by this definition, a sphere S is intersection-equivalent with itself.

The structure of the remainder of the proof is to show that the agreement a is constant on $[0, 1]$ by showing that it is locally constant via Lemmas 15 and 16.



■ **Figure 13** Left: two intersection-equivalent spheres. Right: two spheres that are not intersection-equivalent: the red sphere has an annulus component that cannot be mapped to a component of the blue one. Indeed, the components of the blue sphere are discs and a sphere with 4 punctures.

Since we work in the piecewise-linear setting, we have the following observation:

► **Lemma 14.** *Let $t \in [0, 1]$. There exists a neighbourhood $V \subset [0, 1]$ of t such that for all spheres S_v , $v \in V$, S_v and S_t are either disjoint or intersection-equivalent.*

The proofs of the two following lemmas, given in the full version of the article, are similar.

► **Lemma 15.** *Let $t \in [0, 1] \setminus \{t_j\}_{1 \leq j \leq s}$ be a non-critical time. There exists a neighbourhood $V \subset [0, 1] \setminus \{t_j\}_{1 \leq j \leq s}$ of t such that a is constant on V .*

Sketch of proof. The idea is to keep track of the small sides of two close spheres for each \mathcal{T}^i . Either the two spheres are disjoint and the result follows from Lemma 12 and Lemma 13, or they are intersection-equivalent and the proof proceeds similarly by applying these two lemmas in succession on the parts inside and outside of a small side.

► **Lemma 16.** *Let $t \in [0, 1]$. Then there exists a neighbourhood $V \subset [0, 1]$ of t such that a is constant on $V \setminus \{t\}$.*

► **Proposition 17.** *The agreement a is constant on $[0, 1] \setminus \{t_j\}_{1 \leq j \leq s}$ and hence $a(0) = a(1)$.*

Proof. We can cover $[0, 1]$ by open discs from Lemma 16 on which a is constant. Since $[0, 1]$ is compact, only finitely many of them are enough to cover it. On each connected component of $[0, 1] \setminus \{t_j\}_{1 \leq j \leq s}$ the agreement function is constant by the continuity of a (a is locally constant). Moreover, we know from Lemma 16 that for $u < t_j < v$ close enough we have $a(u) = a(v)$. Hence, a is constant on $[0, 1] \setminus \{t_j\}_{1 \leq j \leq s}$, and $a(0) = a(1)$. ◀

Proof of Theorem 2. The initial discussion of this subsection states that $a(0) = 0$ and $a(1) = 1$. This contradicts Proposition 17. Thus, our assumption that $\text{Top}(\mathcal{D}(n, n+1), R) < \frac{2}{3}n - 2$ does not hold. Therefore, during the sequence, at least one diagram has at least $2n^2 + \frac{2}{3}n$ crossings. ◀

In the full version of the article, we discuss variants of our construction for which our arguments also provide lower bounds on the crossing-complexity.

References

- 1 Colin C. Adams. *The knot book*. American Mathematical Society, 1994.
- 2 Taylor Applebaum, Sam Blackwell, Alex Davies, Thomas Edlich, András Juhász, Marc Lackenby, Nenad Tomašev, and Daniel Zheng. The unknotting number, hard unknot diagrams, and reinforcement learning. *arXiv preprint*, 2024. doi:10.48550/arXiv.2409.09032.
- 3 Dror Bar-Natan. On Khovanov’s categorification of the Jones polynomial. *Algebraic & Geometric Topology*, 2(1):337–370, 2002.
- 4 R.H. Bing. *The geometric topology of 3-manifolds*, volume 40. American Mathematical Society, 1983.
- 5 Gerhard Burde and Heiner Zieschang. *Knots*. Walter de gruyter, 2002.
- 6 Benjamin A. Burton, Hsien-Chih Chang, Maarten Löffler, Clément Maria, Arnaud de Mesmay, Saul Schleimer, Eric Sedgwick, and Jonathan Spreer. Hard diagrams of the unknot. *Experimental Mathematics*, 0(0):1–19, 2023.
- 7 Gregory Chambers and Yevgeny Liokumovich. Converting homotopies to isotopies and dividing homotopies in half in an effective way. *Geometric and Functional Analysis*, 24, November 2013.
- 8 Arnaud de Mesmay, Jessica Purcell, Saul Schleimer, and Eric Sedgwick. On the tree-width of knot diagrams. *Journal of Computational Geometry*, 10(1):164–180, 2019. doi:10.20382/JOCG.V10I1A6.
- 9 Arnaud de Mesmay, Yo’av Rieck, Eric Sedgwick, and Martin Tancer. The unbearable hardness of unknotting. In *Proceedings of the 35th International Symposium on Computational Geometry (SoCG 2019)*, volume 129 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 49:1–49:19. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPICS.SOCG.2019.49.
- 10 Arnaud de Mesmay, Yo’av Rieck, Eric Sedgwick, and Martin Tancer. The unbearable hardness of unknotting. *Advances in Mathematics*, 381:107648, 2021.
- 11 Éric Colin De Verdière and Jeff Erickson. Tightening nonsimple paths and cycles on surfaces. *SIAM Journal on Computing*, 39(8):3784–3813, 2010. doi:10.1137/090761653.
- 12 I. A. Dynnikov. Arc-presentations of links: Monotonic simplification. *Fundamenta Mathematicae*, 190(1):29–76, 2006.
- 13 Leonhard Euler. Solutio problematis ad geometriam situs pertinentis (in latin). *Commentarii Academiae Scientiarum Imperialis Petropolitanae*, 8:128–140, 1736.
- 14 David Gabai. Foliations and the topology of 3-manifolds. iii. *Journal of Differential Geometry*, 26(3):479–536, 1987.
- 15 Wolfgang Haken. Theorie der Normalflächen. *Acta Mathematica*, 105(3):245–375, 1961.
- 16 Joel Hass, Jeffrey C. Lagarias, and Nicholas Pippenger. The computational complexity of knot and link problems. *Journal of the ACM (JACM)*, 46(2):185–211, 1999. doi:10.1145/301970.301971.
- 17 Joel Hass and Tahl Nowik. Unknot diagrams requiring a quadratic number of Reidemeister moves to untangle. *Discrete & Computational Geometry*, 44:91–95, 2010. doi:10.1007/S00454-009-9156-4.
- 18 Allen Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge ; New York, 2002.

- 19 Chuichiro Hayashi. The number of Reidemeister moves for splitting a link. *Mathematische Annalen*, 332:239–252, 2005.
- 20 Louis H. Kauffman. The Jones polynomial, knots, diagrams, and categories. *Bulletin of the American Mathematical Society*, 2023.
- 21 Louis H. Kauffman and Sofia Lambropoulou. Hard unknots and collapsing tangles. *Introductory Lectures on Knot Theory*, pages 187–247, 2011.
- 22 Dale Koenig and Anastasiia Tsvietkova. Unlinking, splitting, and some other NP-hard problems in knot theory. *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1496–1507, 2021. doi:10.1137/1.9781611976465.90.
- 23 Marc Lackenby. A polynomial upper bound on Reidemeister moves. *Annals of Mathematics*, pages 491–564, 2015.
- 24 Marc Lackenby. The efficient certification of knottedness and Thurston norm. *Advances in Mathematics*, 387:107796, 2021.
- 25 Marc Lackenby. Links with splitting number one. *Geometriae Dedicata*, 214(1):319–351, 2021.
- 26 Corentin Lunel and Arnaud de Mesmay. A Structural Approach to Tree Decompositions of Knots and Spatial Graphs. In *Proceedings of the 39th International Symposium on Computational Geometry (SoCG 2023)*, volume 258 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 50:1–50:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPICS.SOCG.2023.50.
- 27 Kunio Murasugi. On the braid index of alternating links. *Transactions of the American Mathematical Society*, 326(1):237–260, 1991.
- 28 Dale Rolfsen. *Knots and links*. AMS Chelsea Pub, Providence, R.I, 2003. OCLC: ocm52901393.
- 29 Martin Scharlemann. Thin position in the theory of classical knots. In *Handbook of knot theory*, pages 429–459. Elsevier, 2005.
- 30 Horst Schubert. Über eine numerische Knoteninvariante. *Mathematische Zeitschrift*, 61(1):245–288, 1954.
- 31 Jennifer Schultens. Bridge numbers of torus knots. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 143(3), pages 621–625. Cambridge University Press, 2007.
- 32 Kurt van Reidemeister. Elementare Begründung der Knotentheorie (in german). *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, 5:24–32, 1927.