Knot or not?

Unpicking tangled string is even harder than it looks, but the effort has paid astounding dividends, says Richard Elwes

AS MAGIC tricks go, it is nowhere near as spectacular as sawing your assistant in half or pulling a rabbit out of an empty hat. But for mathematicians, there's nothing in the magician's repertoire that trumps the amazing vanishing knot. The drums roll, the cymbals crash, and with a triumphant "Hey, presto!" the conjuror pulls an impossibly convoluted tangle of rope into a straight cord.

Anyone who has wrestled with intransigent shoelaces can tell you that the trick does not work for any old knot. The secret of success lies in good preparation: the magician ensures the desired outcome by carefully knotting the rope beforehand, following a well-established procedure. But what is it exactly about the conjuror's knot that means it can be pulled apart just like that, but the shoelace cannot?

Such questions turn out to have far-reaching significance. Molecules of DNA are often found in Byzantine tangles, and whether or not they can be untied seems to be a crucial factor in determining the likelihood of gene mutation, a driving force of evolution. The mechanical properties of the polymers that are ubiquitous in modern life also depend largely on how their long molecules get knotted. In physics, knots turn up unexpectedly in fundamental areas, from quantum computing to statistical mechanics.

The trouble is, finding a satisfactory answer to the problem of the conjuror's knot turns out to be decidedly tricky. The branch of mathematics that has developed as a result, known as knot theory, quickly runs up against a conundrum that has tied up some of the best mathematical brains of the past two centuries: given two randomly tangled bundles of string, how can you tell for sure whether they are different, or really the same knot in disguise?

Now we seem to be closing in on an answer, and a perfect, all-discerning way of describing knots might be within our grasp. What's more, the latest twists in the long and winding road to understanding knots reveal just how fundamental they are, and how they may even point the way to a new understanding of physical reality.

When mathematicians consider knots in a piece of string, they are careful – unlike the rest of us – to ensure that they first join the ends of the string to form a closed loop. This locks the knot into the string, allowing it to be pulled and twisted around with complete freedom without changing the loop's fundamental knottedness. Cutting and regluing while manipulating the string are, of course, strictly forbidden.

Following this definition, the most basic knot is just a simple loop. Not much of a knot, you might say, and indeed it is known as the "unknot". But trivial though it might seem, the humble unknot reveals the basic puzzle: it can take on many different convoluted guises, if you stretch and bend it enough. A tangled bundle of string might reduce to the unknot if you push and pull it in the right way – as the conjuror's knot does – or it might not.

Confronted with two knotted loops, the simplest way to check if they are the same thing in disguise may be to experiment: try and see if you can pull one into the shape of the other, with all its under and overcrossings in the same place. Catalogues of knots that have been laboriously compiled since the late 19th century provide a useful reference.

That's OK for fairly simple knots, but as the number of crossings in a knot increases, the number of different possible matches grows exponentially, making the task incredibly complicated. With just 12 crossings, more than 2000 distinct knots must be considered. That complexity tripped up even the pioneering knot cataloguers: some catalogues contained duplicate knots that were spotted only decades later (see "The Perko pair", page 34).

Knot theory thus reached a kind of impasse, where it remained for much of the 20th century. Brute-force trial and error could not provide an efficient way of telling knots apart and, despite the occasional advance, the search for a slick mathematical solution proceeded sluggishly.

When a breakthrough did come, it was from an unexpected direction. In 1984, Vaughan Jones, a New Zealander working at the University of Pennsylvania in Philadelphia, was investigating the maths underlying quantum mechanics when he began to notice similarities between his results and certain aspects of knot theory. That chance observation quickly led to the discovery of some easy algebraic manipulations that could be used to unpick a knot mathematically.

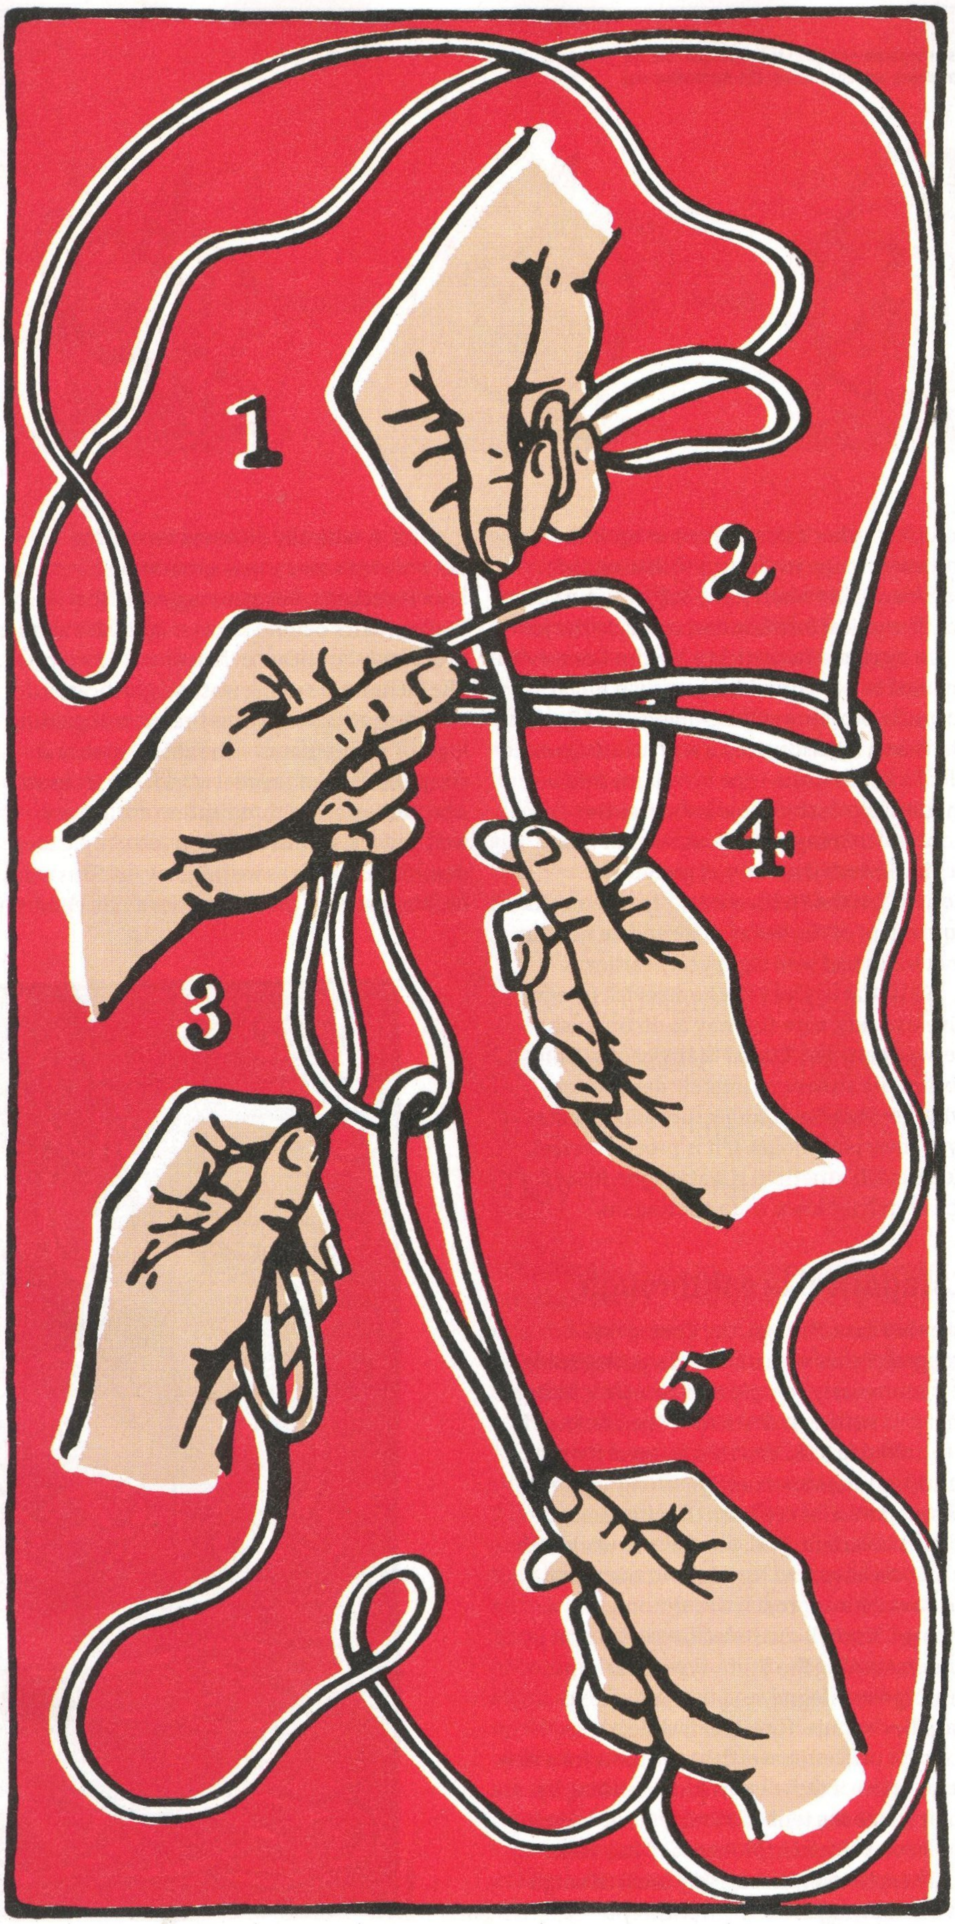

DAVID WARDLE

What Jones provided was an algebraic rule relating three knots that differed only at one crossing. Where the first knot had an undercrossing, the second knot would have an overcrossing and the third knot no crossing at all. A complex knot could then, in effect, be broken down into a succession of unknots by considering Jones's rule at each crossing in turn. The result was a series of mathematical expressions from which, with a little algebraic sleight of hand, a simple formula matching the original knot popped out.

How this mathematical jiggery-pokery related to physically unpicking a knot, and what information the resulting formula encoded, was far from clear. The main thing was that, however much you disguised a knot by pulling and twisting it before you started calculating, the same formula always emerged. In other words, Jones's formula was an "invariant" of a particular knot – something highly prized by mathematicians.

> "How can you tell if two knots are really different, or just the same thing in disguise?"

Being easy to calculate and pretty powerful, the Jones invariant has become an indispensable tool in several scientific disciplines where knots turn up. A case in point is biochemistry, where it is used to analyse how enzymes called topoisomerases and recombinases chop up and reassemble knotted loops of DNA during cell replication. Before Jones, trial and error and contrived, complex algebra were needed to model this process mathematically. The invariant provided an easy method to compare the DNA formations before and after, and so begin to work out what some of the enzymes were doing.

Useful as Jones's invariant proved in these other areas, it was not the full answer. For this to be so, the invariant would have to be a two-way street: not only should one knot always be described by the same mathematical formula, but one formula should always lead back to the same knot. This second condition is not fulfilled, as two unrelated knots can share the same Jones invariant.

Improvements on Jones's basic idea – dubbed "quantum invariants" because, like the original, they were based in the mathematics of quantum mechanics – arrived thick and fast. But none of these turned out to be a perfect, unique knot descriptor either. What's more, a fundamental enigma remained: where ▶

THE PERKO PAIR

The first and last knots in this sequence appeared for almost 100 years as separate entries in the catalogues started by Peter Guthrie Tait. They are one and the same, however, as the intervening sequence shows. The duplication was discovered in 1974 by amateur topologist Kenneth Perko

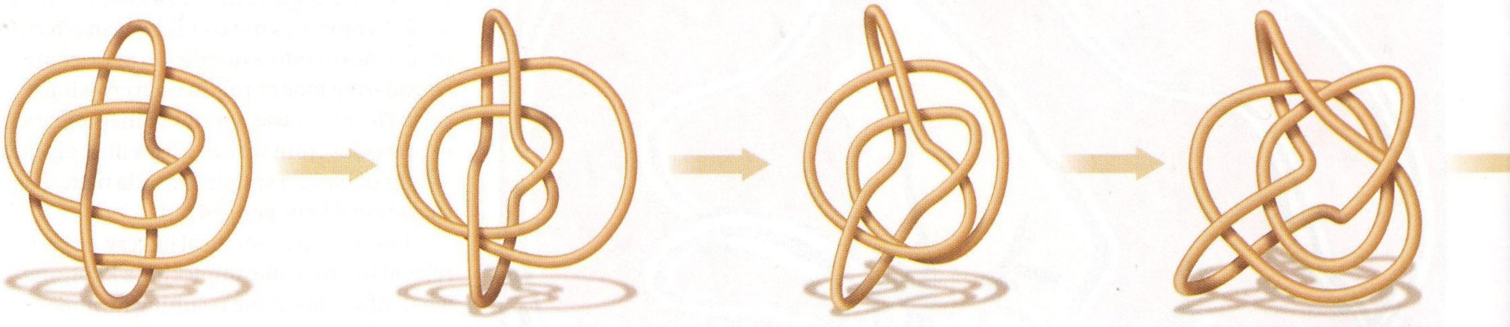

did these strange algebraic expressions come from, and what did they really say about knots?

Answering those questions demanded a radically different tack. In 1989, two Russians working independently, Victor Vassiliev of the Independent University of Moscow and Mikhail Goussarov at the Steklov Mathematical Institute, also in Moscow, considered what would happen if, instead of just passing under or over each other at a crossing, the strands of the knot were allowed to crash straight through each other. That might seem an odd way to attack the problem – after all, knots in the real world cannot do that kind of thing – but it paid dividends. Picking over these new knots revealed a dazzling array of what are known as "finite-type" invariants.

Individually, some finite-type invariants have found applications in problems of polymer physics, but for mathematicians their power lies in the sheer numbers of them that are available. In fact, the set describing any particular knot contains infinitely many finite-type invariants, and researchers have now shown that sequences of these invariants can be combined to build the Jones invariant and the other quantum invariants that tell real knots apart from one another.

Vassiliev's own observations went further. Of all the knots he investigated, he noted that no two were described by exactly the same set of finite-type invariants. That led him to a celebrated conjecture: if two knots really are different, then they will have at least one finite-type invariant that will be different. Equally, if the finite-type invariants describing two knots are the same, then the knots themselves must be the same.

To date, no knots have been found that do not obey Vassiliev's conjecture. Case closed, you might think: the finite-type invariants amount to a unique set of "fingerprints" belonging to one, and only one, knot. But mathematicians are not that easily satisfied. Quite apart from objecting that the conjecture has yet to be formally proved, they find the solution rather messy. Why put up with infinitely many invariants to distinguish between knots when there might be a single mathematical formula that could do the job?

In 1993, mathematician Maxim Kontsevich seemed to have provided just such a neat formula. Working at the University of Bonn in Germany, he discovered a way to amalgamate all of a knot's finite-type invariants into one mighty expression, now known as the Kontsevich integral. This was one of four achievements that in 1998 garnered Kontsevich the Fields medal – mathematics' equivalent of the Nobel prize and an honour that Jones had won eight years earlier.

So is that it? Does the Kontsevich integral give us a tidy, sure-fire way to distinguish between any two knots? Many people think so, but the answer ultimately depends on whether Vassiliev's conjecture about finite-type invariants holds. If it is proved correct, Kontsevich's method is watertight. If not, the search is blown wide open again.

Categorical breakthrough

In either case, though, the Kontsevich integral comes with a rather large downside: it is immensely complex. So complex, in fact, that just writing it down for a specific knot is a formidable task. Even the integral for the innocuous unknot is a frightening algebraic expression, a far cry from the simplicity of Jones's invariant. Still, you have to work with what you have, and throughout the 1990s Kontsevich's approach was the only show in town. Knot theorists' efforts went into understanding the finite-type invariants, taming the integral, and trying to prove Vassiliev's conjecture. Meanwhile, the concept proved its worth as the basis for a new generation of techniques for describing by their shape the tens of thousands of complex molecules stored in biological databases.

Then, in 1999, an exciting breakthrough came – once again from an entirely unexpected quarter. A radically different technique not only spawned a new generation of knot invariants, but also suggested that the mathematics behind knots might have a more profound significance than anyone suspected. The technique's name was categorification.

Categorification turns the normal guiding logic of mathematics – the abstraction and simplification of the real world – on its head. Abstraction and simplification are all very well, but the results are often too simple to describe as much as we might want. This can be illustrated by the difficulties small children

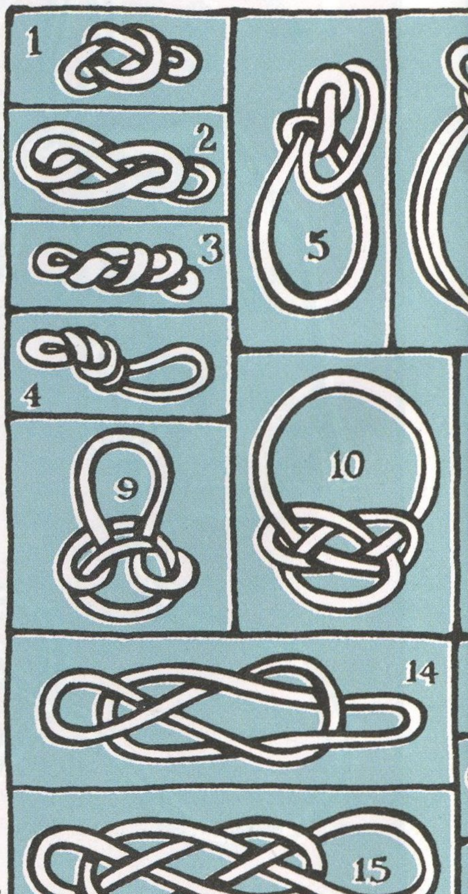

DAVID WARDLE

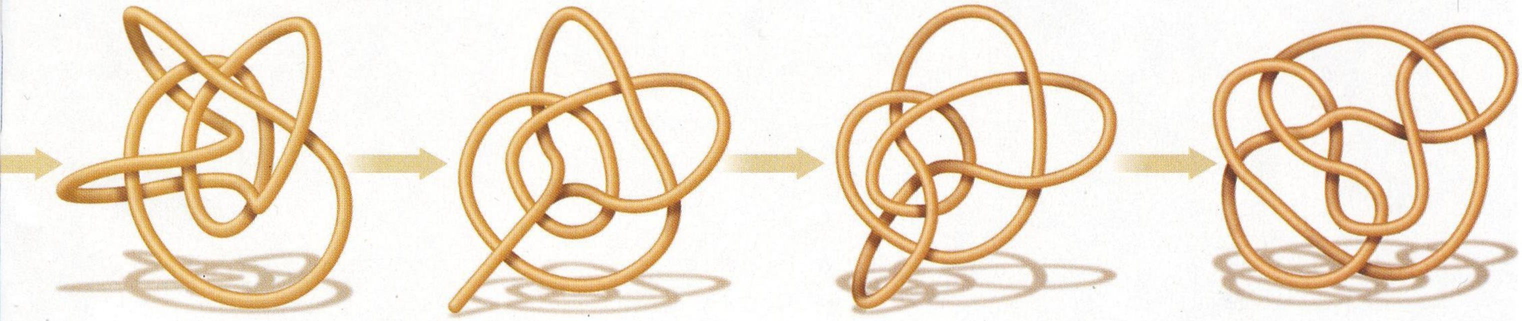

wrestle with when learning to count. Why do three apples and three oranges both reduce to the same "3", when apples and oranges are entirely different things?

The answer is that the mathematical structure that "3" inhabits – the system of ordinary numbers – is an abstraction shorn of any information about the things it represents. In this case, categorification involves replacing this system with a richer structure, a "category", in which the rigid equations that define the number system – "1+2=3", for example – are replaced with weaker statements comparing the sizes of different sets of objects. This category offers a real-world flexibility that numbers on their own do not: its sets can be different even if they have the same size. In this structure, the original number system appears as the category's "shadow", obtained by collapsing all sets of size 3 down to just one representative: the number 3.

Might what works for the number system also work for other mathematical objects? This was the philosophy adopted by Mikhail Khovanov, a mathematician at the University of California, Davis, when he revisited Jones's invariant in 1999. Instead of breaking it down into smaller finite-type expressions, he looked for some grander structure of which it was just the most visible shadow.

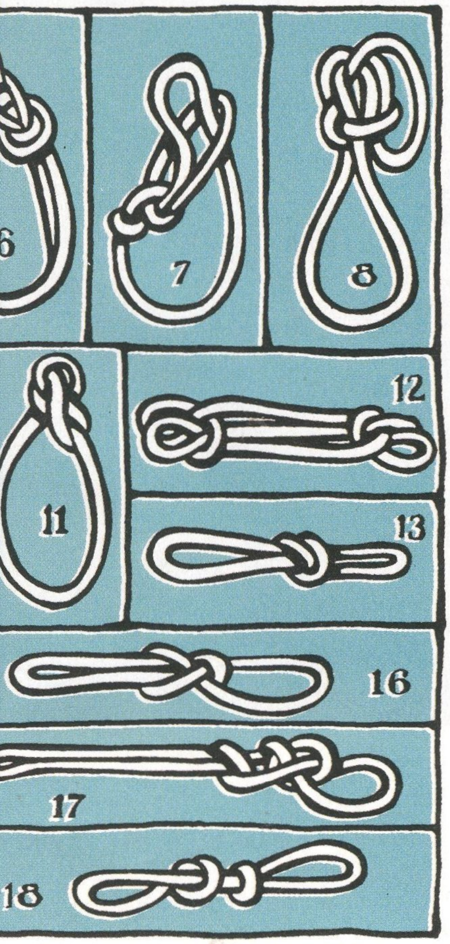

"The inconsistencies we see between relativity and quantum theory may be an illusion"

His search was a spectacular success. The overarching category that he found is conceptually difficult, but mathematically it is far less awkward than Kontsevich's integral, and it's a more reliable description of a knot than Jones's formula. Even better, thanks to an ingenious computer program written in 2006 by Dror Bar-Natan of the University of Toronto, Canada, it can now be computed efficiently for any knot, potentially opening up its use to researchers in other areas.

Khovanov's category is not perfect: there are still instances of knots that share the same category. So work continued and, in 2005, together with Lev Rozansky of the University of North Carolina, Chapel Hill, Khovanov unveiled a new invariant operating at an even higher level. Not only does this categorify many quantum invariants beyond Jones's, it also subsumes Khovanov's original category, as well as several other knot invariants discovered in the meantime.

The power of the Khovanov-Rozansky category is getting us close to the perfect knot description, though early indications are that we are not there yet. A still broader family of quantum invariants may yet need to be incorporated before we can say we have the loose ends of the knot problem tied up. Nevertheless, we seem to be edging towards the ultimate mathematical solution.

With early experiments in categorification yielding such riches, physicists and mathematicians are waking up to the idea that the approach might apply to more than just knots. Recalling the deep connection of knots with quantum theory that inspired Jones in the first place, some researchers think they have spotted a tantalising hint that whole chunks of mathematical physics are just shadows of larger categorical structures.

Striking analogies have already been found between the categorical descriptions of quantum mechanics and Einstein's relativity – the twin theoretical pillars of modern physics that have so far seemed fundamentally incompatible. And John Baez, Alexander Hoffnung and Christopher Rogers of the University of California, Riverside, have recently argued that string theory – a favoured starting point for a theory of everything – can be viewed as a categorification of particle physics (www.arxiv.org/abs/0808.0246).

If a categorification embracing both relativity and quantum theory can be found, then the inconsistencies that physicists see between these two theories might yet prove to be an illusion. Such a tying together of all of physics would indeed be a categorical triumph for the humble knot. ●

Richard Elwes is a mathematician and writer based in Leeds, UK