Volume 20
Number 1
7 January 2024
Pages 1–268

Soft Matter

rsc.li/soft-matter-journal

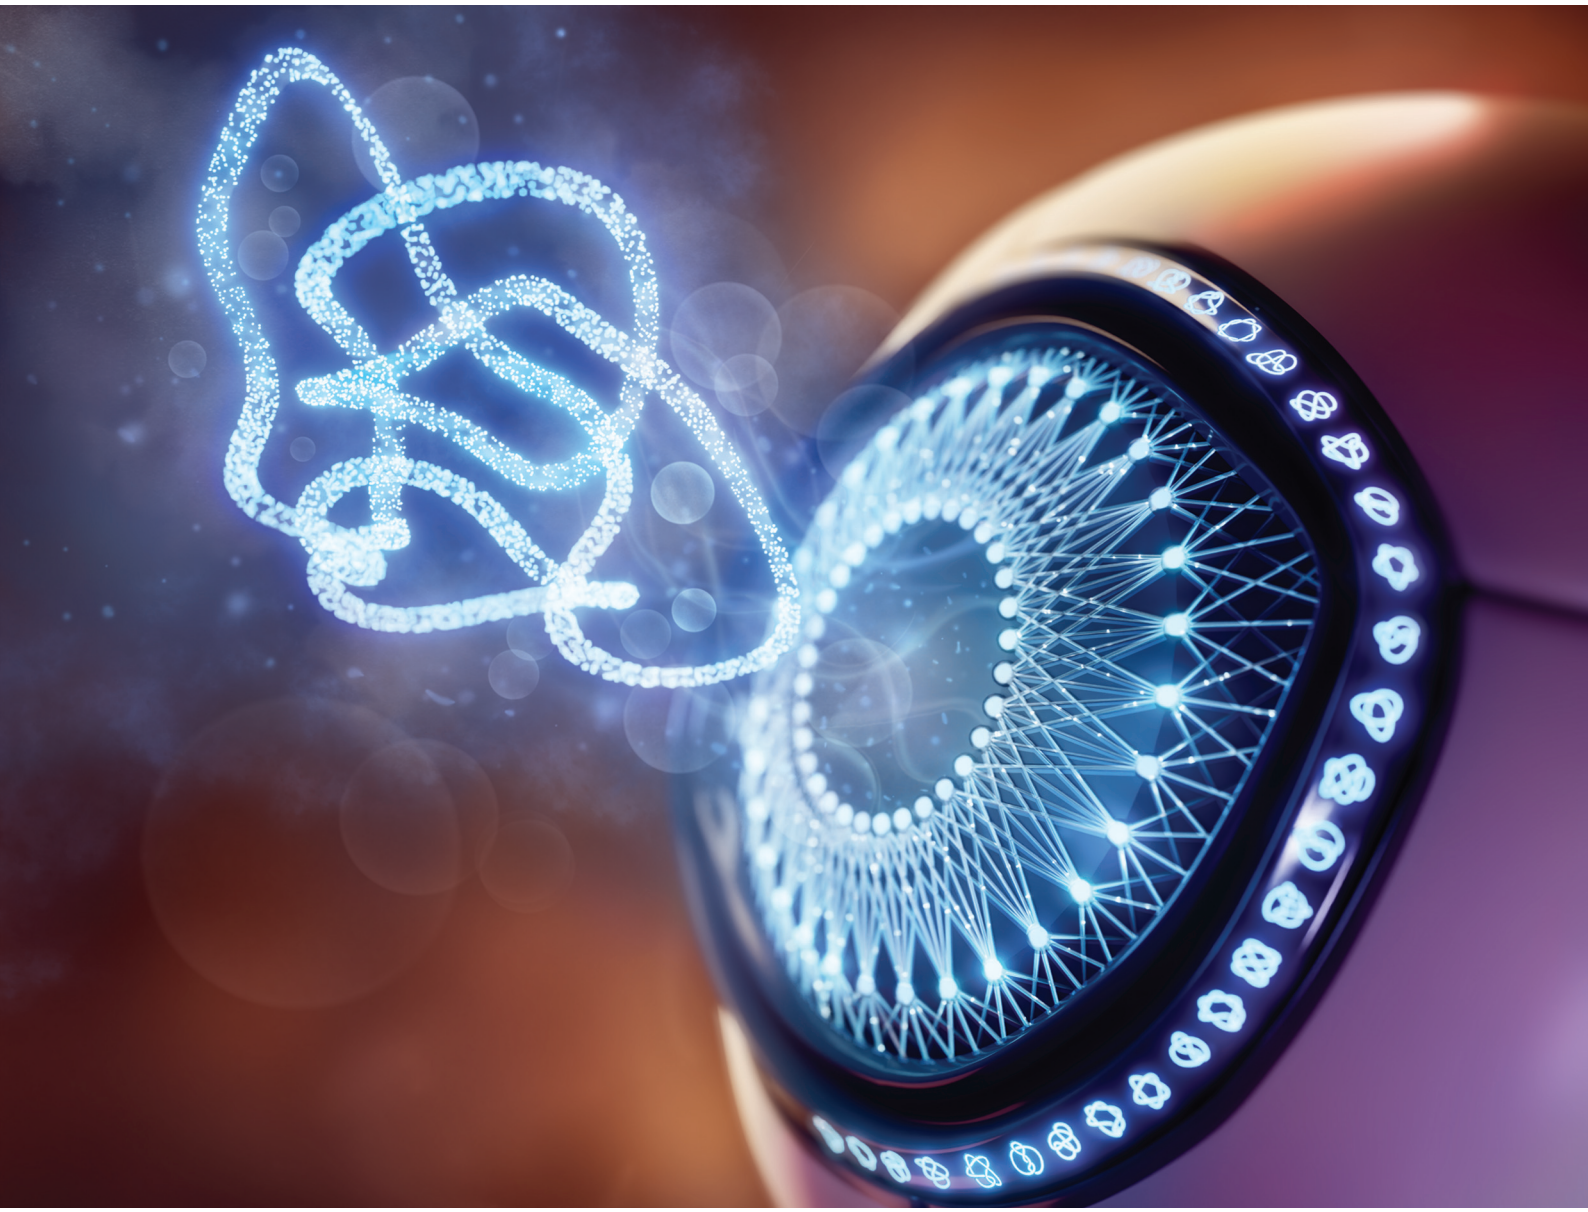

ISSN 1744-6848

ROYAL SOCIETY OF CHEMISTRY

PAPER
Davide Michieletto *et al.*
Geometric learning of knot topology

Soft Matter

PAPER

View Article Online
View Journal | View Issue

Check for updates

Cite this: *Soft Matter*, 2024, **20**, 71

Geometric learning of knot topology†

Joseph Lahoud Sleiman,‡[a] Filippo Conforto,‡[a] Yair Augusto Gutierrez Fosado[a] and Davide Michieletto*[ab]

Received 8th September 2023,
Accepted 18th October 2023

DOI: 10.1039/d3sm01199b

rsc.li/soft-matter-journal

Knots are deeply entangled with every branch of science. One of the biggest open challenges in knot theory is to formalise a knot invariant that can unambiguously and efficiently distinguish any two knotted curves. Additionally, the conjecture that the geometrical embedding of a curve encodes information on its underlying topology is, albeit physically intuitive, far from proven. Here we attempt to tackle both these outstanding challenges by proposing a neural network (NN) approach that takes as input a geometric representation of a knotted curve and tries to make predictions of the curve's topology. Intriguingly, we discover that NNs trained with a so-called geometrical "local writhe" representation of a knot can distinguish curves that share one or many topological invariants and knot polynomials, such as mutant and composite knots, and can thus classify knotted curves more precisely than some knot polynomials. Additionally, we also show that our approach can be scaled up to classify all prime knots up to 10-crossings with more than 95% accuracy. Finally, we show that our NNs can also be trained to solve knot localisation problems on open and closed curves. Our main discovery is that the pattern of "local writhe" is a potentially unique geometric signature of the underlying topology of a curve. We hope that our results will suggest new methods for quantifying generic entanglements in soft matter and even inform new topological invariants.

Open Access Article. Published on 20 October 2023. Downloaded on 5/28/2026 7:00:45 PM.
This article is licensed under a Creative Commons Attribution 3.0 Unported Licence.

1 Introduction

Knots are fascinating objects that have captured the attention of humans for centuries. From Incas' knotted Quipus,[1] and Lord Kelvin's theory of elements as knotted ether,[2] to sailors and climbers whose lives often rely on the strength of knotted rope, knots are deeply intertwined with history and art and often carry mystical meaning. The human obsession with knots brought Peter Guthrie Tait to compile the first knot tabulation of up to 10 crossings by hand;[3] currently, more than one million unique knots up to 16 crossings have been tabulated using computer programs.[4]

To rigorously prove that the early tabulated knots did not contain duplicates, so-called topological invariants and knot polynomials were developed, the first of which was the Alexander polynomial,[1,5,6] followed more recently by the Jones and HOMFLY polynomials.[1,7] Knot polynomials are mathematical constructs that can be computed on knot diagrams and are invariant under smooth deformations of the curve, *i.e.* deformations that preserve the curve topology. However, there are knots that share many topological invariants and cannot even be distinguished by knot polynomials. Famously, the 11-crossing Conway knot has the same Alexander polynomial as the unknot and shares the same Jones polynomial of its mutant, the Kinoshita–Terasaka (KT) knot.[1] More generally, all mutants of a knot have the same HOMFLY polynomials and the same hyperbolic volume,[1] while some composite knots share the same homeomorphic complements.[8–10]

Alongside the development of topological invariants, several attempts were made to identify a relationship between a specific geometric embedding of a knot and its underlying topology.[11] We note that this relationship is different from the one sought between so-called geometric and algebraic invariants,[12,13] *e.g.* between the hyperbolic volume of a knot and its Jones polynomial.[14] Perhaps one of the most rigorous results in this direction is the Fáry–Milnor theorem, stating that the total absolute curvature of non-trivially knotted curves must be greater than 4π.[15] Unfortunately, this result only imposes a weak constraint on the topology of the underlying curve, as an unknot can itself have large total curvature due to, for example, deformations of its contour. In parallel, a large body of work on so-called "ideal knots" was carried out with the aim of finding geometric features that could reflect the underlying knot topology. One impressive result in this context is that different DNA knots display a spatial separation when run on a gel electrophoresis that is linearly proportional to the so-called average crossing number;[16,17] this result entails that there is an intimate relationship between the physical shapes assumed by knots and their underlying topology. Another result that inspired our work is that the

[a] *School of Physics and Astronomy, University of Edinburgh, Peter Guthrie Tait Road, Edinburgh, EH9 3FD, UK. E-mail: davide.michieletto@ed.ac.uk*
[b] *MRC Human Genetics Unit, Institute of Genetics and Cancer, University of Edinburgh, Edinburgh EH4 2XU, UK*

† Electronic supplementary information (ESI) available. See DOI: https://doi.org/10.1039/d3sm01199b

‡ Joint first author.

Open Access Article. Published on 20 October 2023. Downloaded on 5/28/2026 7:00:45 PM. This article is licensed under a Creative Commons Attribution 3.0 Unported Licence.

total so-called "writhe" (see below) of an ideal knot is the same (up to a constant that is only a function of the curve length) as that of a non-ideal, thermally agitated curve with the same topology.[18] Though this suggests that "writhe" may be a good measure that is invariant under thermal fluctuations, there is no one-to-one relationship between the global writhe of a knot and its underlying topology; for instance, the global writhe of the 4_1 knot is 0, the same as the unknot.[11]

Thus, the problem of determining a curve topology based only on the geometric information of its segments (without using any projection or algebraic invariant) is an open challenge in knot theory that has ramification in many fields, for instance polymer physics, biophysics and fluid dynamics. In this paper, we propose to address this open challenge by using the power of artificial intelligence, and in particular deep learning, in recognising and classifying patterns in certain knot geometric features. Our main discovery is that by using a quantity we dub "local writhe", even simple machine learning (ML) algorithms can identify the topology of knotted curves undergoing thermal fluctuations with very high accuracy. We argue that this is an example of geometric learning, whereby the only quantity we pass to the ML algorithm is a quantity that can be computed from the Cartesian positions of a curve's segments, without the need to compute algebraic invariants such as Alexander or Jones polynomials. Our method can even distinguish 11-crossing knots that are otherwise impossible to distinguish using standard invariants (the Conway and KT knots). Finally, we show how this algorithm can be scaled to classify all 250 prime knots up to 10 crossings with 95% accuracy, and can even be employed to solve knot localisation problems. Overall, we argue that local writhe is an excellent feature – determined purely by the 3D positions of a curve segments – that results in patterns easily identifiable by ML algorithms. We argue that our results will be applied to other classification problems such as threading[19,20] and entanglements[21,22], and also prompt knot theorists to employ local writhe to define new geometric knot invariants.

2 Results

Two recent papers by Vandans[23] and Braghetto[24] have shown that machine learning is a promising tool to solve knot classification problems. They mostly considered the Cartesian position of the monomers, or adjacent monomer distances and dihedrals to classify the 5 simplest knots. In this work, we set out to test the use of a different type of geometric feature that our group recently utilised to identify essential crossings of a knot and plectoneme-like double folding of ring polymers.[25,26] More specifically, we focused on a generalisation of the Gauss linking integral applied to a single closed curve, often associated with its writhe[27] and average crossing number.[25,28] This choice is inspired by the intuition that writhe captures the geometrical entanglement of a curve with itself, and we thus define a generalised local segment-to-segment (StS) writhe as

$$\omega_{\mathrm{StS}}(x,y)=\frac{(\boldsymbol{t}(x)\times\boldsymbol{t}(y))\cdot(\boldsymbol{r}(x)-\boldsymbol{r}(y))}{|\boldsymbol{r}(x)-\boldsymbol{r}(y)|^3}, \qquad (1)$$

where $\boldsymbol{r}(x)$ and $\boldsymbol{t}(x)$ are the 3D position of, and the tangent at, segment x, respectively. Intuitively, eqn (1) captures the magnitude and the chirality of the entanglement between segment x and segment y (Fig. 1(A) and (B)). The quantity $\omega_{\mathrm{StA}}(x)=\oint_\gamma \omega_{\mathrm{StS}}(x,y)\mathrm{d}y$ is the local segment-to-all (StA) writhe and characterises how geometrically entangled segment x is with respect to the whole closed curve γ. In practice, the calculation of StS and StA writhe are conducted on discrete segments, taking a finite "window" with length $l_w = 10\sigma$ to smooth out short length fluctuations (see ESI† for details).

The StA writhe, $\omega_{\mathrm{StA}}(x)$, is a 1D geometrical representation of a knot that we hypothesise may display some patterns that are topology-dependent (Fig. 1(A)–(C)). Since complex pattern recognition is a task that naturally lends itself to being addressed using a machine learning approach, we thus asked ourselves if a neural network (NN) trained to recognise patterns within $\omega_{\mathrm{StA}}(x)$ was able to solve ambiguous knot classification problems. To do this, we built feed forward and recurrent (long-short term memory, LSTM) neural networks (FFNN and RNN, respectively) and trained them using 10^5 statistically uncorrelated and pre-labelled conformations for each knot. To generate these conformations, we initialised a bead-spring polymer with known topology, $N = 100$ beads, and persistence length $l_p = 10\sigma$ (other lengths and l_p are reported in the ESI†) using KnotPlot (knotplot.com) and subsequently evolved the polymer configurations in LAMMPS[29] *via* Langevin dynamics in an implicit solvent and fixed temperature, using a Kremer–Grest model[30] to preserve polymer topology (see Methods and ESI† for more details). The code to generate these conformations are available open access at **https://git.ecdf.ed.ac.uk/taplab/mlknotsproject**. We confirmed that the topology was conserved either by computing their Alexander determinant *via* KymoKnot (**https://kymoknot.sissa.it**)[31] or, when ambiguous, visually.

The NNs were built with an input layer that was determined according to the input representation being studied, *e.g.*, the Cartesian (*XYZ*) coordinate representation used 3 neurons (one for each dimension) per polymer bead. Other local input features, such as StA writhe, used one neuron per bead, while the StS writhe feature requires $N \times N$ input neurons. The optimal number of hidden layers, hidden units, learning rate and batch size were determined *via* an automated hyperparameter tuning method conducted on the Cartesian representation (KerasTuner[32]). Unless otherwise stated, our NNs contained 4 hidden layers, with around 4×10^5 trainable parameters. The output layer consisted of C output neurons, corresponding to the C knot types being classified, each implemented with a softmax activation function in order to return the probability that a given input is a certain knot type. We took the sparse categorical cross-entropy as the loss function, as the most appropriate for individual class probabilities and integer target labels, *i.e.* our knot types (Fig. 1(D)).

2.1 NNs trained with StA writhe yield more accurate knot classification than Cartesian features

We first tackle a 5-knot classification problem with the 5 simplest knots, which can be satisfactorily solved using NNs

 This journal is © The Royal Society of Chemistry 2024

Open Access Article. Published on 20 October 2023. Downloaded on 5/28/2026 7:00:45 PM.
This article is licensed under a Creative Commons Attribution 3.0 Unported Licence.

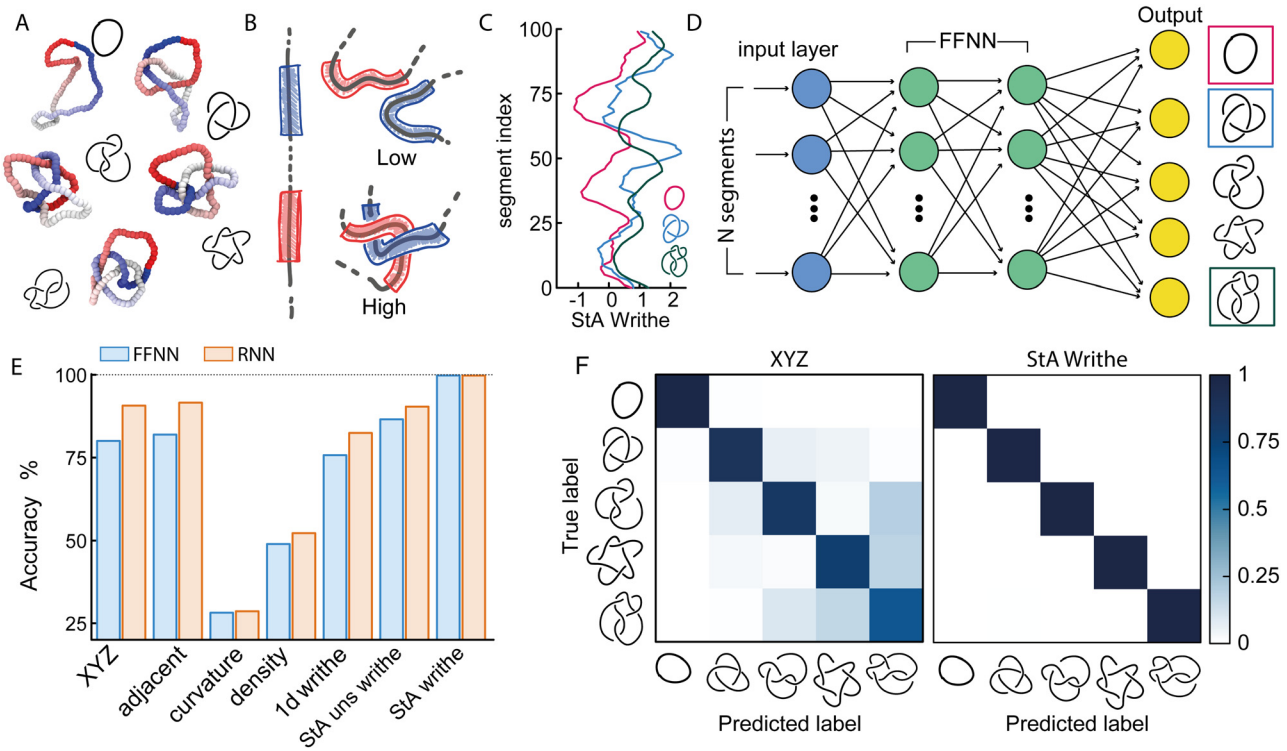

Fig. 1 (A) Examples of equilibrium knotted polymer conformations colour coded to indicate the knot contour (from red, to white, to blue). In this figure we consider the 5 simplest knots: 0_1, 3_1, 4_1, 5_1 and 5_2. (B) A graphical representation of StS writhe $\omega_{\mathrm{StS}}(x,y)$ showing an instance of small and large writhe between two segments. (C) Examples of patterns for $\omega_{\mathrm{StA}}(x)$ for three different knots. (D) A graphical representation of the (feed-forward) network. The input layer contains N (or $3N$) neurons corresponding to the size of the input feature representation, and the output layer yields a probability for each knot class. (E) Accuracy score, tested on unseen polymer conformations for different input features. The StA writhe classifies the 5-simplest knots with 99.9% accuracy irrespective of the network architecture. (F) Confusion matrices obtained by training the network with *XYZ* and StA writhe input features.

trained on center-of-mass-corrected Cartesian coordinates (*XYZ*) or adjacent bead input features.[23,24] In line with these previous works, we find that our NNs can accurately predict the topology of unseen conformations (80.1% accuracy with a FFNN and 86% accuracy with a recurrent NN architecture, Fig. 1(E)). These values are lower than the ones reported in ref. 23 since we use a smaller training dataset and simpler NNs. We then trained the same NNs using a range of other geometric features, such as local curvature, density and 1D writhe[26] (see ESI† for details), and found that most of them performed more poorly, or at best equally, with respect to the *XYZ* representation (Fig. 1(E)). A similar outcome was also obtained in ref. 24 In stark contrast, models trained using $\omega_{\mathrm{StA}}(x)$ outperformed all other models and were found to achieve 99.9% accuracy, irrespective of the FFNN or RNN architectures (we also tested random forest algorithms, see ESI†). Additionally, the networks reached the early stopping criterion in about 50% fewer epochs or less, compared to those trained using the *XYZ* representation (see ESI†). When plotted as a confusion matrix, the results clearly indicate that the *XYZ* input feature struggles to classify knots with a similar number of crossings, *e.g.* the 5_1 and 5_2 knots. In contrast, our local 3D writhe (StA) feature generated a near-perfect confusion matrix (Fig. 1(F)).

We found that these results are generally robust for different choices of dataset splitting, persistence length ($l_p = 1 - 10\ \sigma$), window length chosen to perform the StA calculation, and length of the chains (see ESI†). Nevertheless, they do display a significant reduction in accuracy when tested on knots generated using a different method (for instance freely jointed chains), and also when the window length for the StA writhe calculation is comparable to the full contour of the chain. In this case, the StA writhe is constant and equals the global writhe of the knot, which is not unique for different knots.[11] This is also in agreement with principal component analysis (PCA, see ESI†) of the StA-trained NNs, where we see that different knots are clearly separable in the reduced 2D PCA space, yet the 0_1 and 4_1 cluster together due to the fact that they share the global writhe (zero), which is related to the mean value of $\omega_{\mathrm{StA}}(x)$ along the contour.

2.2 NNs trained with StA writhe can distinguish knots with identical knot polynomials

Given that our NNs can distinguish knots with the same minimal number of crossings, *i.e.*, the 5-crossings knots, we asked ourselves if they could also solve more complicated problems where knots shared algebraic knot polynomials. To this end, we first considered three knots with identical Alexander polynomials: the square, granny, and 8_{20} knots (see Fig. 2(A)). The first two knots are 6-crossings knots consisting of trefoil composites with different chirality (hence they

Open Access Article. Published on 20 October 2023. Downloaded on 5/28/2026 7:00:45 PM.
This article is licensed under a Creative Commons Attribution 3.0 Unported Licence.

are homeomorphic knot complements), whereas the latter is an 8-crossings knot. Once again, we trained our FFNN using the $\omega_{StA}(x)$ profiles (Fig. 2(B)) and obtained a striking accuracy of 99.98%, compared with 91.8% obtained by training with COM-shifted *XYZ* coordinates (Fig. 2(C)).

We then asked ourselves if our NN could distinguish knots sharing multiple knot polynomials. As mentioned above, mutant knots share the same hyperbolic volume and several knot polynomials, including HOMFLY. We therefore performed simulations of the Conway (K11n34) knot and one of its mutants, the Kinoshita–Terasaka (KT, K11n42) knot. These 11-crossings knots have a number of identical knot invariants, sharing the same Jones, Alexander, and Conway polynomials.[1] Intriguingly, the latter two are also shared with the unknot. Thus, we generated 10^5 statistically uncorrelated conformations of $N = 200$ beads long polymers with the Conway, KT, and unknot topologies (Fig. 2(D)), and trained our FFNN to classify them either using a COM-subtracted *XYZ* or $\omega_{StA}(x)$ (Fig. 2(E)) representations. When tested on unseen conformations, we found that while the *XYZ*-trained NN could not distinguish the Conway and KT knots, both were accurately distinguished from the unknot (Fig. 2(F)). In marked contrast, we discovered that the StA-trained NN perfectly disentangles the three knots with 99.6% accuracy (Fig. 2(F)). We therefore conclude that the StA-trained NN has the ability to convert StA patterns into a topological knot classification, even for knots sharing multiple knot polynomials, such as mutants and composites. In turn, we argue that the StA writhe is a geometric quantity computed on a particular 3D embedding of a curve that carries high-density information about its underlying topology. Importantly, we stress that to classify these knots, the network does not compute any knot polynomial, as other standard software do.

Somewhat unsatisfactorily, we cannot fully pinpoint why StA writhe is so powerful at identifying different topologies. We hypothesise that the 1D patterns generated by StA writhe-specifically the sequence, sign and amplitudes assumed by consecutive maxima and minima - contain information on the relative orientation and magnitude of consecutive entanglements. As mentioned above, the average value of $\omega_{StA}(x)$ is related to the global writhe of the knot, which itself contains non-unambiguous information about its topology. Thus, we argue that the NNs can extract additional information from the full $\omega_{StA}(x)$ patterns, related to the chirality of individual entanglements and render the information unique. This hypothesis is also supported by the fact that the unsigned StA writhe (which cannot distinguish chirality) yields, in general, a lower accuracy (see Fig. 1(E)). We thus hypothesise that the information encoded in the pattern of the StA writhe may be related to the underlying knot's Dowker code. These hypotheses will be tested in more detail in future works.

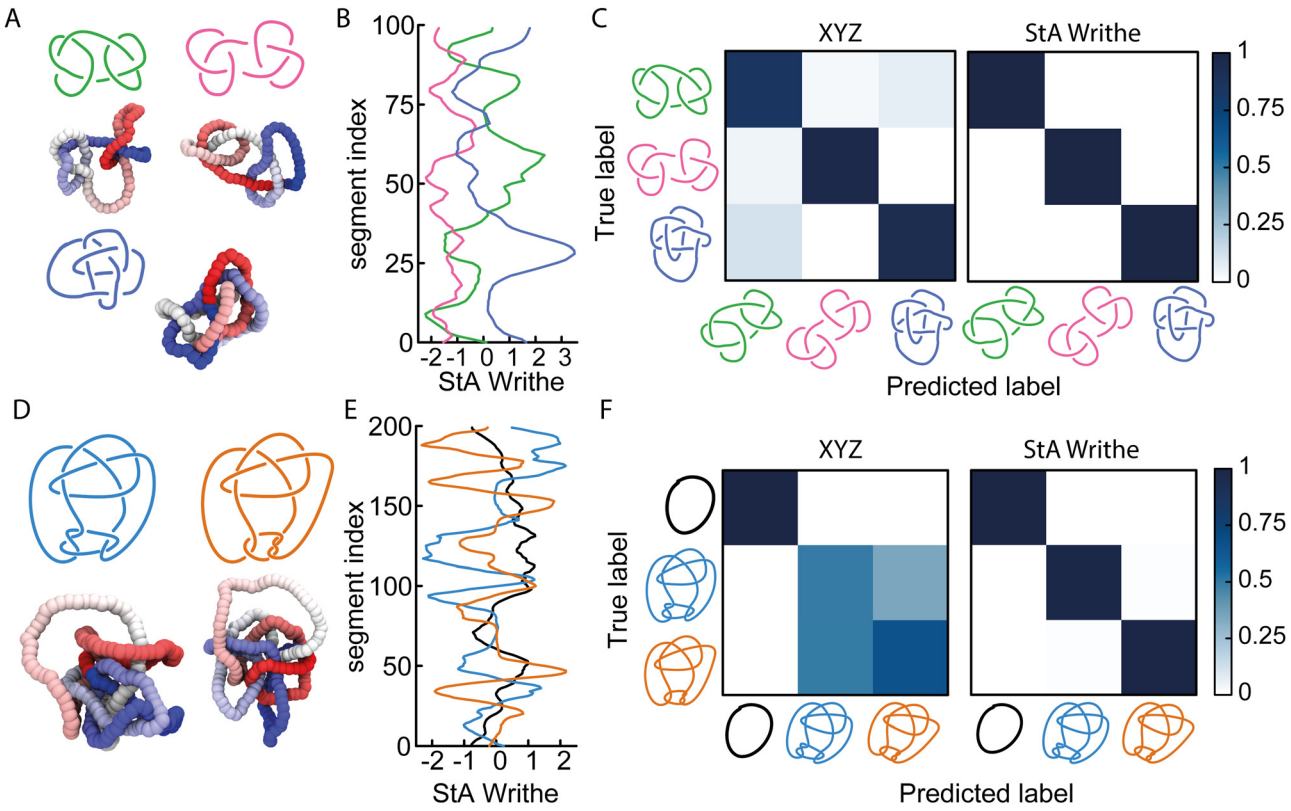

Fig. 2 (A) Snapshots of three knots with identical Alexander polynomial: square ($3_1^l\#3_1^r$), granny ($3_1^l\#3_1^l$) and 8_{20} knots. (B) Examples of StA writhe patterns from the three knots. (C) Confusion matrices obtained from a 3-class classification problem, training a FFNN with *XYZ* (91.7% accuracy) or StA writhe (99.9% accuracy) features. (D) Snapshots of Conway (blue) and KT (orange) knots. (E) Examples of StA writhe patterns, including the one from the unknot (black). (F) Confusion matrices obtained from a 3-class classification problem, training a FFNN with *XYZ* (67% accuracy) and StA writhe (99.6% accuracy).

 This journal is © The Royal Society of Chemistry 2024

Open Access Article. Published on 20 October 2023. Downloaded on 5/28/2026 7:00:45 PM.
This article is licensed under a Creative Commons Attribution 3.0 Unported Licence.

2.3 StS writhe outperforms StA writhe on knots with more than 7 crossings

To understand to what extent StA-trained NNs can be used to classify knotted curves, we trained our NNs on increasingly complex classification problems, and generated conformations of all prime knots up to 10-crossings. Among these 250 prime knots, there are over 30 that share the same Alexander polynomial (see ESI† for a table), making them challenging to classify using standard tools (for instance KymoKnot). We first noticed that *XYZ*-trained NNs rapidly declined in accuracy when we included knots with 6 or more crossings (Fig. 3(A) and (B)). In contrast, the confusion matrices from StA-trained NNs retained relatively high accuracies. However, we noticed that the knots 5_1 and 7_2 created some confusion even in the StA-trained NNs, causing a drop in accuracy to 98% (Fig. 3(C) and (D)). We argue that this was due to the fact that $\omega_{StA}(x)$ of the two knot types displayed similar patterns. For instance, we show two knot instances that yield particularly similar $\omega_{StA}(x)$ patterns in Fig. 3(C). Thus, to further distinguish these (and potentially other knots with similar $\omega_{StA}(x)$ curves) we decided to consider our original proposition of using the local StS writhe (eqn (1)); two examples of $\omega_{StS}(x,y)$ maps are reported in Fig. 3(E), for the same 5_1 and 7_2 knots configurations used to compute $\omega_{StA}(x)$ in Fig. 3(C). Interestingly, the $\omega_{StS}(x,y)$ maps appear very different, despite generating very similar StA curves when integrated along y and around the polymer contour. This is because a given segment x may itself have a certain sequence of negative and positive entanglements with other segments y. Once integrated along the contour in the y direction, different sequences may lead to similar overall values. Motivated by this, we trained our FFNNs using the StS writhe representation of the knots, and discovered we could restore a very high (99.8%) accuracy for the case of a database containing all knots up to 7-crossings (Fig. 3(F)). More specifically, the confusion between 5_1 and 7_2 knots is now resolved thanks to the StS writhe. Ultimately, the StS-trained NNs produced the most accurate models, achieving 95% for a 250-class classification task, including all prime knots up to 10 crossings. In comparison, the *XYZ*-trained and StA-trained NNs achieved 17% and 72% on the same problem, respectively (Fig. 3(G)).

Based on these results, we argue that the StS writhe is therefore the most scalable and precise geometric feature to utilise for knot classification problems. Most importantly, we would like to stress that the impressive accuracy demonstrated for a 250-class problem was achieved with a simple feed forward NN with 4 layers (around 3600k for the StS writhe and 400k parameters for StA writhe). A natural extension going forward will be to employ more complex architectures, and in particular convolutional NNs, to classify the 2D StS writhe maps.

2.4 StA-trained NNs can also solve knot *localisation* problems

In the final part of this paper we turn our attention to the knot localisation problem, *i.e.* determining the shortest knotted arc along the polymer contour. This task is challenging and

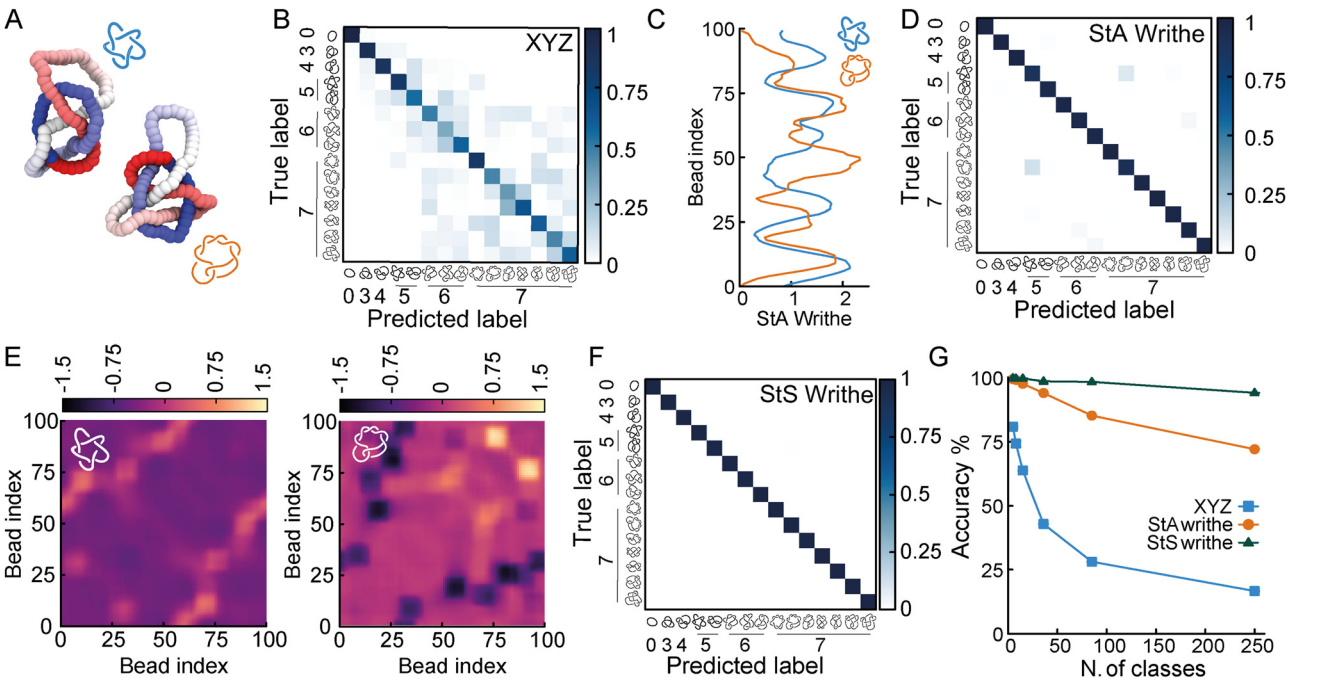

Fig. 3 (A) Two example conformations of 5_1 (blue) and 7_2 (orange) knots. (B) The *XYZ*-trained NN on a 15-class classification problem yields 63.8% accuracy and a rather non-diagonal confusion matrix. (C) Examples of $\omega_{StA}(x)$ curves for the two knots, displaying a degree of similarity between the pattern of maxima and minima. (D) The $\omega_{StA}(x)$-trained NN achieves 98% accuracy and the confusion matrix shows that 5_1 and 7_2 are the knots that are most confused with each other. (E) Examples of the $\omega_{StS}(x,y)$ geometric feature for the two knots corresponding to the $\omega_{StA}(x)$ profiles shown in (C). (F) Confusion matrix for a StS-trained FFNN to classify all knots up to 7 crossings, achieving 99.8% accuracy (see SI for more confusion matrices for more complex problems). (G) Accuracy as a function of the number of knot classes being distinguished, up to 10-crossing (250) prime knots.

Open Access Article. Published on 20 October 2023. Downloaded on 5/28/2026 7:00:45 PM.
This article is licensed under a Creative Commons Attribution 3.0 Unported Licence.

particularly important for open curves, such as linear polymers, DNA, and proteins,[33–38] which may contain entanglements and knots. In this context, identifying the shortest portion of a polymer that is knotted is akin to being able to identify entanglements in chain melts.

We first tackled this problem using the same FFNN architecture as in the knot classification task, but the accuracies generated were very low. We hypothesised that this was due to the fact that FFNNs do not preserve the sequential information along the polymer. For this reason, we consider a long-short term memory (LSTM) model, also known as a recurrent NN (RNN). More specifically, we employed a sequence-to-sequence LSTM, with an output layer corresponding to a binary sequence of $N = 100$ neurons, equivalent in dimension to the length of the input polymer. Each output neuron is passed through a sigmoid function, which converts the output into a probability between 0 and 1 representing the likelihood that a given monomer is within the knotted segment of the polymer conformation. The true output labels were generated using KymoKnot,[31] which employs a minimally-interfering closure algorithm followed by a standard Alexander determinant calculation to identify the start and end monomers of the knot. This data was then transformed into a vector of 100 bits, *i.e.* a value of 0 or 1, corresponding to whether a certain monomer was part of the knotted arc.

Unlike normal multi-class classification problems where the classes are mutually exclusive, here we consider a multi-label classification task, with mutually non-exclusive class labels (multiple classes per prediction).[39] To quantify the error in a multi-label classification task, we use the binary cross-entropy (BCE) function, suited to an output layer of sigmoid functions, given by

$$\mathrm{BCE} = \frac{1}{N}\sum_{i=1}^{N} y_i \log(\hat{y}_i) + (1 - y_i)\log(1 - \hat{y}_i) \qquad (2)$$

where y_i is the ith element in the true output vector, $\mathbf{y}$, $\hat{y}_i$ is the ith element in the predicted output vector, $\hat{\mathbf{y}}$, and N is the dimension of the output label, corresponding to the length of the polymer in our knot localisation task. This error is then used to optimise the model weights.

Finally, to determine the accuracy of the model, we converted the probabilities generated by the sigmoid function y_{prob} into binary values using a Heaviside step function ($y_{\mathrm{pred}} = \Theta(y_{\mathrm{prob}} - 0.5)$), and compared the result to the true binary value obtained using KymoKnot. The final accuracy is given by the binary accuracy, *i.e.* Accuracy = correct/total.

Overall, we find that the StA-trained RNNs perform extremely well, reaching above 90% accuracy in localising any knot that we tested: the 5 simplest knot types, 0_1, 3_1, 4_1, 5_1 and 5_2 (Fig. 4). We argue that this excellent performance relies on the effectiveness of RNNs in handling multi-scale sequential data and tracking multi-scale correlations along the polymer. This capability likely plays a major role in allowing the network to recognise that nearby monomers are more likely to be in the same knotted arc. More precisely, we find that the StA writhe representation is superior to all other descriptors, with a localisation accuracy of 93%, confirming its potential usefulness as a tool to help in knot localisation tasks. For instance, in Fig. 4(D) we report the prediction and ground truth for the 4_1 knot shown in Fig. 4(A) and (B). In this case, the StA writhe perfectly agrees with the KymoKnot ground truth, whereas the *XYZ* and unsigned StA writhe yield less accurate localisation predictions.

In the ESI† (Fig. S9), we also used our StA-trained RNN model to track the unknotting of a 5_1 knot tied on an open curve. Despite the fact that the algorithm was not trained on open curves, the results were surprisingly accurate. The model can be seen to clearly detect the presence of short knotted arcs even at the final step before complete unknotting. Once again, we find that the StA-trained model is largely superior to the *XYZ*-trained model.

Overall, our results highlight the power of StA and StS writhe in not only classifying but also localising knots. We acknowledge that our results are non-exhaustive and more work will be needed in the future to find the best architectures and models to optimally solve these tasks.

3 Conclusions

In conclusion, we have discovered that local "segment-to-all" and "segment-to-segment" writhe (eqn (1)) are geometric descriptors

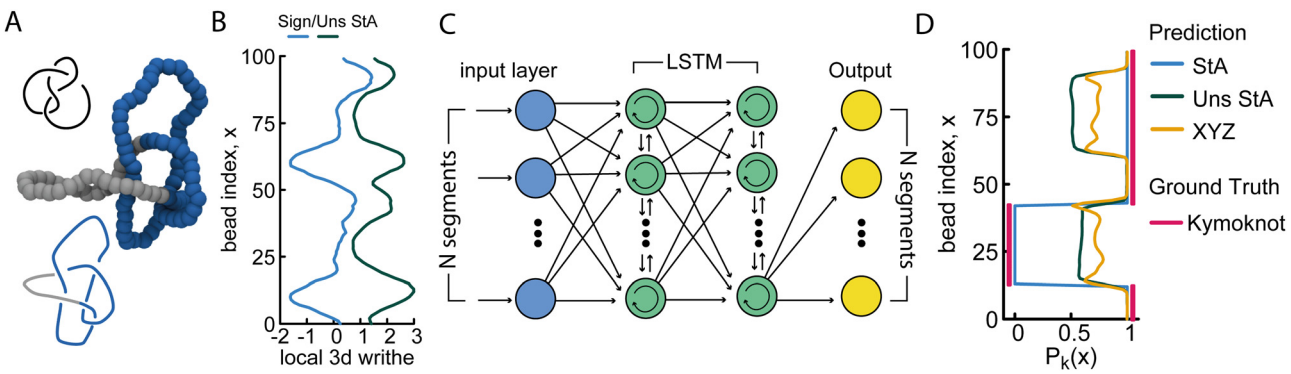

Fig. 4 (A) Example of a 4_1 knot where the knotted core is localised within ~80% of the contour. (B) Signed and unsigned StA writhe profiles for the conformation shown in (A). (C) Sketch of an LSTM (recurrent) NN, encoding the sequential information of the segments. (D) Profile of the knot probability $P_k(x)$ as a function of bead index x, as predicted by the RNN with different geometric features. The ground truth was generated using KymoKnot.

 This journal is © The Royal Society of Chemistry 2024

Open Access Article. Published on 20 October 2023. Downloaded on 5/28/2026 7:00:45 PM.
This article is licensed under a Creative Commons Attribution 3.0 Unported Licence.

of a curve that contain information about its underlying topology. Our AI-driven approach can classify, using a single quantity, complex knot topologies that would otherwise be impossible to disentangle using a single algebraic invariant. More specifically, we demonstrated, for the first time, that NNs can utilise the information encoded in StA and StS writhe to classify the curve topology significantly more accurately than what can be achieved using the Cartesian coordinates of the curve's segments or other local geometric quantities (Fig. 1). We hypothesise that our NNs trained on local 3D writhe representations may numerically encode a new type of geometric topological invariant. This conjecture is supported by the fact that even a simple FFNN architecture can distinguish the topology of knot mutants and composites that share several algebraic knot polynomials (Fig. 2). Finally, we showed that our new proposed geometric feature (eqn (1)) is robust to more complex knots than the ones tackled in the literature so far; indeed, we have managed to classify all 250 prime knots up to 10-crossings with 95% accuracy (Fig. 3). We argue that deeper NN or convolutional NN may be able to push this result further, to >10-crossings knots.

We stress that this method only requires a snapshot of a knot embedding with a list of 3D coordinates for each polymer segment and is trained on thermal conformations under a readily tunable temperature. For this reason, it will require longer training for longer polymers but should be essentially insensitive to the number of non-essential crossings, as shown by the excellent accuracy achieved in spherically confined polymers.[24] This feature is in marked contrast to standard knot topology algorithms, that take 2D projections and need to compute matrices as big as the number of crossings in a given projection, irrespective of whether they are essential or not.[1] Finally, we show that by deploying recurrent NNs, our geometric StA descriptor can also solve knot localisation problems (Fig. 4). More work will be needed in the future to determine optimal NN architectures.

We note that though we do not have a full understanding of how the NNs are using StA and StS writhe features to identify knots, we hypothesise that they are classifying the patterns of consecutive maxima and minima, thus capturing the entanglement of pairs of segments, accounting for their chirality and magnitude. This argument directly suggests that employing a distance map between segments or other geometric "unsigned" representations will yield lower accuracies, due to the fact that they do not capture the chiral nature of the entanglements between segments. For these reasons, we believe that StS (or StA) representations are possibly the best features to connect the geometry of a given curve embedding, to its underlying topology. A possible limitation of this method is that it is restricted to pairwise entanglement. Generalising the Gauss linking number to higher-order relations is itself an active field of research, and it is foreseeable that a local version of the Milnor triple linking number[40] may be used to generate 3D tensors of Brunnian links, for example.

In conclusion, we established that StS/StA-trained NNs are powerful tools to accurately classify and localise knots in thermally equilibrated curves. Importantly, knot classification and localisation are achieved without any explicit calculation of Alexander or other algebraic invariants. We propose that the local writhe – once fed through deep NNs – yields an accurate map from the configurational space of a curve to its underlying topology. The approach we reported in our paper naturally lends itself to be applied to protein folding,[34,41] DNA[42,43] and, in general, entanglements in open curves and complex systems.[20,21,36,44–47] We hope that our results will also inspire mathematicians and topologists to formulate new topological invariants based on the geometrical embeddings of knotted curves.

Conflicts of interest

There are no conflicts to declare.

Acknowledgements

DM and FC thanks the Royal Society for support through a University Research Fellowship. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement no. 947918, TAP). JLS thanks the Physics of Life network for a student summer bursary in 2021 providing initial funds for this research. The authors thank Enzo Orlandini, Marco Baiesi, Pawel Dabrowski-Tumanski, Ken Millett, and Luigi del Debbio for insightful discussions. The authors also acknowledge the contribution of the COST Action Eutopia, CA17139.

Notes and references

1 C. Adams, in *The Knot Book*, W. H. Freeman, 1994.
2 W. Thomson, *Proc. R. Soc. Edinburgh*, 1867, **6**, 95–105.
3 P. G. Tait, *Trans. R. Soc. Edinburgh*, 1885, **32**, 493–506.
4 J. I. M. Hoste, M. Thistlethwaite and J. Weeks, *Math. Intell.*, 1998, **20**, 33.
5 J. W. Alexander, *Trans. Am. Math. Soc.*, 1928, **30**, 275.
6 L. Kauffman, *Formal Knot Theory*, Princeton University Press, 1983.
7 J. Hoste, *Pac. J. Math.*, 1986, **124**, 295–320.
8 W. P. Thurston, *Three-Dimensional Geometry and Topology*, Princeton University Press, 1997, vol. 1.
9 M. Culler, N. M. Dunfield, M. Goerner and J. R. Weeks, SnapPy, a computer program for studying the geometry and topology of 3-manifolds, Available at **https://snappy.computop.org**.
10 C. M. Gordon and J. Luecke, *Bull. Amer. Math. Soc.*, 1989, **20**, 83–87.
11 A. Stasiak, V. Katritch and L. Kauffman, *Ideal Knots*, World Scientific, 1998.
12 V. Jejjala, A. Kar and O. Parrikar, *Phys. Lett. B: Nucl. Elem. Part. High-Energy Phys.*, 2019, **799**, 135033.
13 A. Davies, P. Veličković, L. Buesing, S. Blackwell, D. Zheng, N. Tomašev, R. Tanburn, P. Battaglia, C. Blundell, A. Juhász, M. Lackenby, G. Williamson, D. Hassabis and P. Kohli, *Nature*, 2021, **600**, 70–74.

Open Access Article. Published on 20 October 2023. Downloaded on 5/28/2026 7:00:45 PM.
This article is licensed under a Creative Commons Attribution 3.0 Unported Licence.

14 H. Murakami, J. Murakami, M. Okamoto, T. Takata and Y. Yokota, *Exper. Math.*, 2002, **11**, 427–435.
15 J. W. Milnor, *Ann. Math.*, 1950, **52**, 248–257.
16 A. Stasiak, V. Katritch, J. Bednar, D. Michoud and J. Dubochet, *Nature*, 1996, **384**, 122.
17 D. Michieletto, D. Marenduzzo and E. Orlandini, *Proc. Natl. Acad. Sci. U. S. A.*, 2015, **112**, E5471–E5477.
18 V. Katritch, W. Olson and P. Pieranski, *Nature*, 1997, **388**, 148–151.
19 D. Michieletto and M. S. Turner, *Proc. Natl. Acad. Sci. U. S. A.*, 2016, **113**, 5195–5200.
20 R. Stano, C. N. Likos and J. Smrek, *Soft Matter*, 2022, **19**, 17–30.
21 R. Everaers, S. K. Sukumaran, G. S. Grest, C. Svaneborg, A. Sivasubramanian and K. Kremer, *Science*, 2004, **303**, 823–826.
22 J. Smrek, J. Garamella, R. Robertson-Anderson and D. Michieletto, *Sci. Adv.*, 2021, **7**, eabf9260.
23 O. Vandans, K. Yang, Z. Wu and L. Dai, *Phys. Rev. E*, 2020, **101**, 1–10.
24 A. Braghetto, S. Kundu, M. Baiesi and E. Orlandini, *Macromolecules*, 2023, **56**, 2899–2909.
25 D. Michieletto, *Soft Matter*, 2016, **12**, 9485–9500.
26 J. L. Sleiman, R. H. Burton, M. Caraglio, Y. A. Gutierrez Fosado and D. Michieletto, *ACS Polym. Au*, 2022, **2**, 341–350.
27 M. R. Dennis and J. H. Hannay, *Proc. R. Soc. A*, 2005, **461**, 3245–3254.
28 A. Stasiak, V. Katritch, J. Bednar, D. Michoud and J. Dubochet, *Nature*, 1996, **384**, 122.
29 S. Plimpton, *J. Comput. Phys.*, 1995, **117**, 1–19.
30 K. Kremer and G. S. Grest, *J. Chem. Phys.*, 1990, **92**, 5057.
31 L. Tubiana, G. Polles, E. Orlandini and C. Micheletti, *Eur. Phys. J. E: Soft Matter Biol. Phys.*, 2018, **41**, 1–7.
32 T. O'Malley, E. Bursztein, J. Long, F. Chollet, H. Jin and L. Invernizzi *et al.* Keras Tuner, **https://github.com/keras-team/keras-tuner**, 2019.
33 P. Dabrowski-Tumanski, A. Stasiak and J. I. Sulkowska, *PLoS One*, 2016, **11**, 1–14.
34 M. Giulini and R. Potestio, *Interface Focus*, 2019, **9**, 20190003.
35 A. R. Klotz, B. W. Soh and P. S. Doyle, *Phys. Rev. Lett.*, 2018, **120**, 188003.
36 M. Caraglio, F. Baldovin, B. Marcone, E. Orlandini and A. L. Stella, *ACS Macro Lett.*, 2019, **8**, 576–581.
37 M. Caraglio, B. Marcone, F. Baldovin, E. Orlandini and A. L. Stella, *Polymers*, 2020, **12**, 1–19.
38 B. W. Soh, A. R. Klotz, R. M. Robertson-Anderson and P. S. Doyle, *Phys. Rev. Lett.*, 2019, **123**, 048002.
39 G. Tsoumakas and I. Katakis, *Int. J. Data Warehous. Min.*, 2007, **3**, 1–13.
40 M. Polyak, *C. R. Acad. Sci. Paris*, 1997, **325**, 77–82.
41 P. Dabrowski-Tumanski and J. I. Sulkowska, *Proc. Natl. Acad. Sci. U. S. A.*, 2017, **114**, 3415–3420.
42 J. T. Siebert, A. N. Kivel, L. P. Atkinson, T. J. Stevens, E. D. Laue and P. Virnau, *Polymers*, 2017, **9**, 1–10.
43 D. Goundaroulis, E. Lieberman Aiden and A. Stasiak, *Biophys. J.*, 2020, **118**, 2268–2279.
44 M. Dennis, R. King, B. Jack, K. O'Holleran and M. Padgett, *Nat. Phys.*, 2010, **6**, 118–121.
45 F. Landuzzi, T. Nakamura, D. Michieletto and T. Sakaue, *Phys. Rev. Res.*, 2020, **2**, 033529.
46 A. Rosa, J. Smrek, M. S. Turner and D. Michieletto, *ACS Macro Lett.*, 2020, **9**, 743–748.
47 T. Herschberg, K. Pifer and E. Panagiotou, *Comput. Phys. Commun.*, 2023, **286**, 108639.

 This journal is © The Royal Society of Chemistry 2024