

Multistable Elastic Knots Supplementary Material

Michele Vidulis, Yingying Ren, Julian Panetta, Eitan Grinspun, Mark Pauly

Periodic rod implementation

A discrete *open* rod can be represented by a polyline with $n + 2$ nodes and $n + 1$ edges. Such a rod has a total of $4n + 7$ degrees of freedom, 3 for each nodal position $\mathbf{x}_i \in \mathbb{R}^3$, $i \in \{0, \dots, n + 1\}$ and 1 for each material frame angle variable $\theta^i \in \mathbb{R}$, $i \in \{0, \dots, n\}$. A discrete *closed* rod can be thought of as an open rod in which the first and the last edges are identified. A discrete closed rod with n nodes and edges has in principle $4n$ degrees of freedom. To control the twisting moments acting between the ends of the open rod, we introduce an additional degree of freedom $\Theta = \theta^n - \theta^0$ that specifies the relative rotation of the first and last material frames (see Figure 1).

Our implementation joins the beginning and end of the rod via a linear change of variables—ensuring agreement of both the centerline and adapted reference frames—while halving the stretching stiffness to avoid double-counting the corresponding energy contribution. The following transformation maps the closed to the open rod variables:

$$\begin{bmatrix} x_0 \\ y_0 \\ z_0 \\ \dots \\ x_{n+1} \\ y_{n+1} \\ z_{n+1} \\ \theta^0 \\ \dots \\ \theta^n \end{bmatrix} = \underbrace{\begin{pmatrix} I_6 & 0 & 0 & 0 & 0 \\ 0 & I_{3(n-2)} & 0 & 0 & 0 \\ I_6 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & I_{n-1} & 0 \\ 0 & 0 & 1 & 0 & 1 \end{pmatrix}}_J \begin{bmatrix} x_0 \\ y_0 \\ z_0 \\ \dots \\ x_{n-1} \\ y_{n-1} \\ z_{n-1} \\ \theta^0 \\ \dots \\ \theta^{n-1} \\ \Theta \end{bmatrix},$$

where $J \in \mathbb{R}^{(4n+7) \times (4n+1)}$, $\mathbf{x}_i = [x_i, y_i, z_i]^T$, and I_k is the k -dimensional identity.

This parametrization of a closed discrete rod allows us to impose two kinds of boundary conditions that affect the twisting behavior of the rod: freely-rotating or clamped ends. In the first case, the Θ variable is a degree of freedom, and therefore the link of the framed curve can change during the simulation. In the second case, Θ is held constant, so that link is preserved, while twist and writhe can be traded according to Călugăreanu’s relation (equation (3) in Appendix C of the accompanying paper). Note that positional and tangent continuity of the centerline across the connection is guaranteed in both cases.

We collect the position variables and material frame angle variables, Θ included, in the two arrays $\mathbf{x} \in \mathbb{R}^{n \times 3}$ and $\boldsymbol{\theta} \in \mathbb{R}^{n+1}$, respectively. The total elastic energy can be expressed as a sum of bending, stretching and twisting contributions:

$$E_{\text{rod}}(\mathbf{x}, \boldsymbol{\theta}) = E_{\text{bend}}(\mathbf{x}) + E_{\text{stretch}}(\mathbf{x}) + E_{\text{twist}}(\mathbf{x}, \boldsymbol{\theta}).$$

We refer to [Bergou et al., 2008, Bergou et al., 2010] for the exact formulation of the energy terms.




Figure 1 The mapping between an open rod and a closed periodic rod. The first and the last edges are identified, and the relative rotation of their cross-sections around the common tangent is encoded by $\Theta = \theta^n - \theta^0$. The reference and material frames are shown in gray and blue, respectively.

Adaptive Sampling Strategy

To efficiently explore the set of equilibrium configurations of knotted rods, we adapt the number of samples drawn during the simulation. In the scope of our work, a fixed amount of computational resources is better invested in exploring the equilibrium states of knot types for which small or no information is available. For example, simulating the relaxation of an elastic unknot will (most likely¹) result in a circle. Repeated sampling will not provide additional information about different equilibria.

Given n equilibrium states, the amount of novel information a new state k brings to the set of known equilibria can be estimated as $I_k = 1 - \max_i s_{ki}$, where s_{ki} , $i \in \{1, \dots, n\}$ is the similarity between states k and i . A second criterion to measure the relevance of state k is the ratio between its elastic energy E_k and E_{\min} , the lowest energy currently attained by the knot type. Low-energy states can be preferable compared to tangled high-energy configurations because they are easier to be observed in physical prototypes.

The score S summarizing the information content of a batch of m equilibrium states is defined as

$$S = \frac{1}{m} \sum_{k=1}^m I_k \frac{E_{\min}}{E_k}.$$

Figure 2 shows the robustness of the adaptive sampling strategy when $m = 10$, and new batches are sampled until $S < \bar{S} = 0.1$ (at least 50 states are always sampled).

Note that other choices of hyperparameters are possible. This freedom allows the user to control the trade-off between computational cost and depth of exploration, depending on the application. The values we adopt are picked to ensure the simulation finishes within reasonable time on a single workstation.

References

- [Bergou et al., 2010] Bergou, M., Audoly, B., Vouga, E., Wardetzky, M., and Grinspun, E. (2010). Discrete viscous threads. *ACM Transactions on graphics (TOG)*, 29(4):1–10.
- [Bergou et al., 2008] Bergou, M., Wardetzky, M., Robinson, S., Audoly, B., and Grinspun, E. (2008). Discrete elastic rods. In *ACM SIGGRAPH 2008 papers*, pages 1–12.
- [Coleman and Swigon, 2004] Coleman, B. and Swigon, D. (2004). Theory of self contact in Kirchhoff rods with applications to supercoiling of knotted and unknotted DNA plasmids. *Philosophical transactions. Series A, Mathematical, physical, and engineering sciences*, 362:1281–99.

¹Other twist-free configurations that locally minimize energy exist, as reported in [Coleman and Swigon, 2004], but have much higher energy. A few instances can be found in our accompanying dataset.

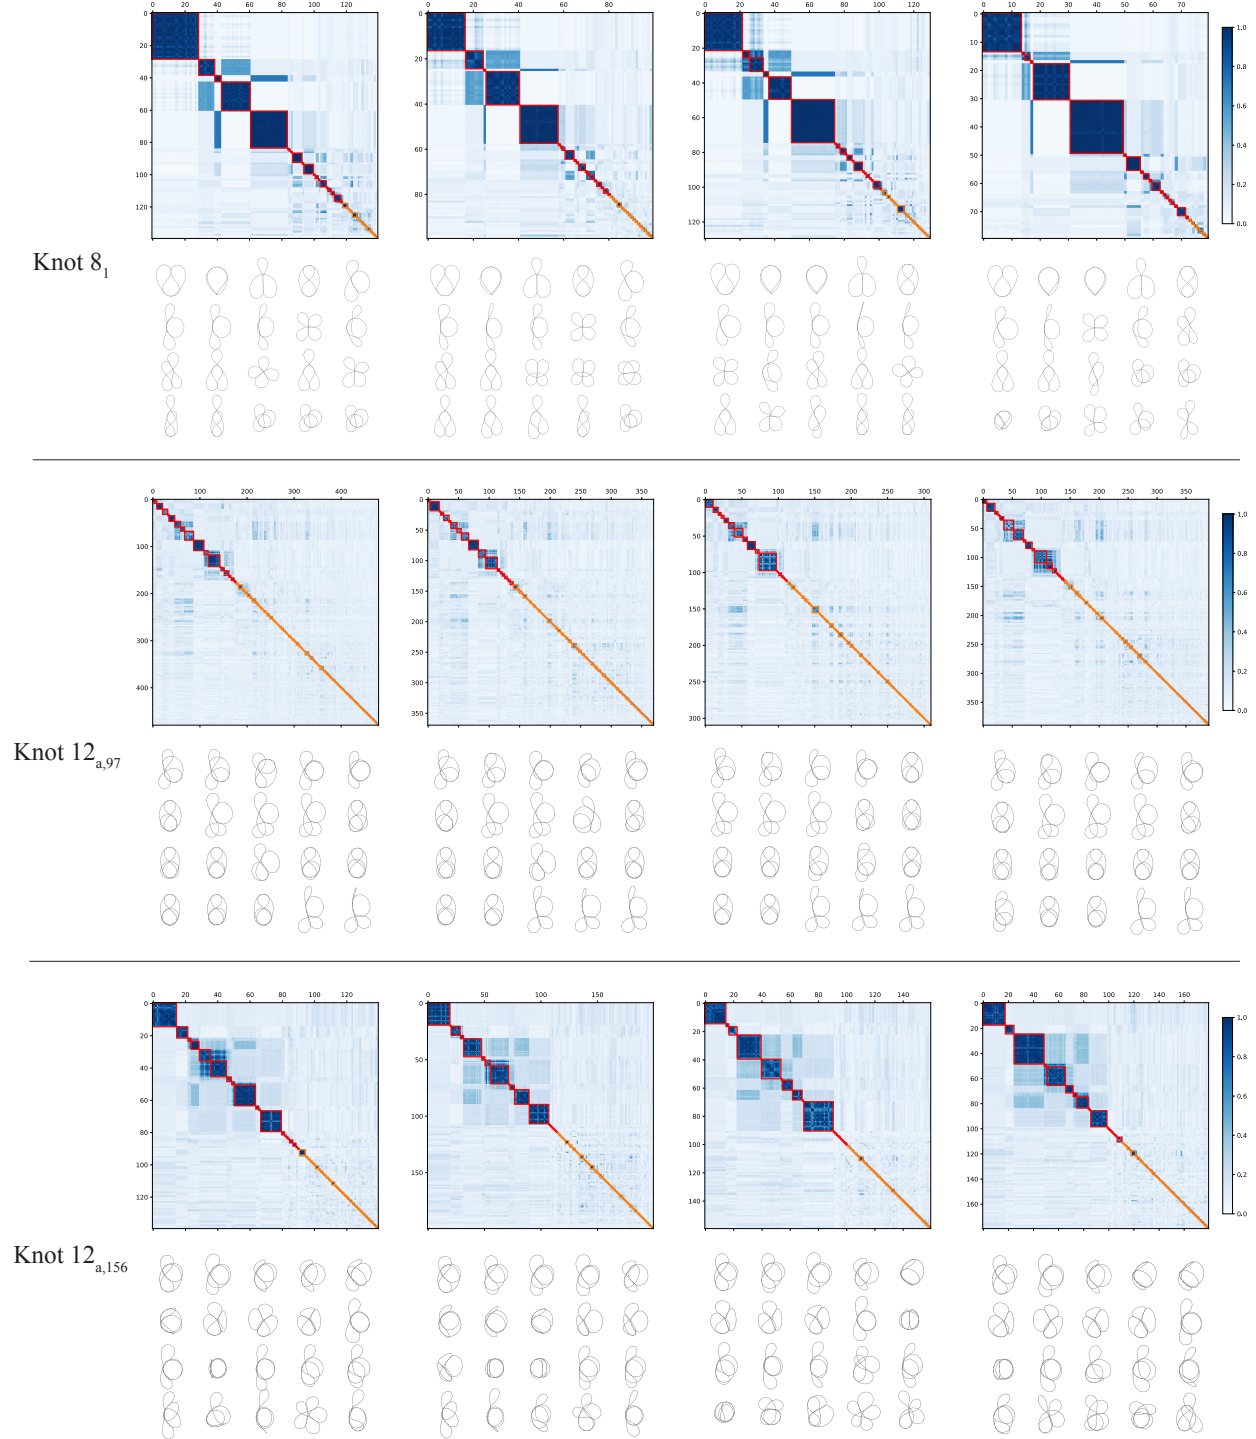


Figure 2 Four different runs of our simulation pipeline for three distinct knot types. The similarity matrices are displayed with the estimated clusters highlighted as diagonal blocks. The clusters are sorted by increasing average energy. A representative for each of the 20 lowest-energy clusters, corresponding to the red blocks, is shown below each matrix. The number of samples, which can change across different runs and is controlled by our heuristics, has an average value of 112.5, 387.5, and 177.5 for knots 8_1 , $12_{a,97}$, and $12_{a,156}$, respectively. To define clusters, dendrograms are cut at 0.25, 0.7, and 0.5 distance, depending on the knot type. The consistent presence of the same equilibrium types across different runs demonstrates the robustness of our sampling strategy.