Tencers: Tension-Constrained Elastic Rods

LILIANE-JOY DANDY*, EPFL, Switzerland
MICHELE VIDULIS*, EPFL, Switzerland
YINGYING REN, EPFL, Switzerland
MARK PAULY, EPFL, Switzerland

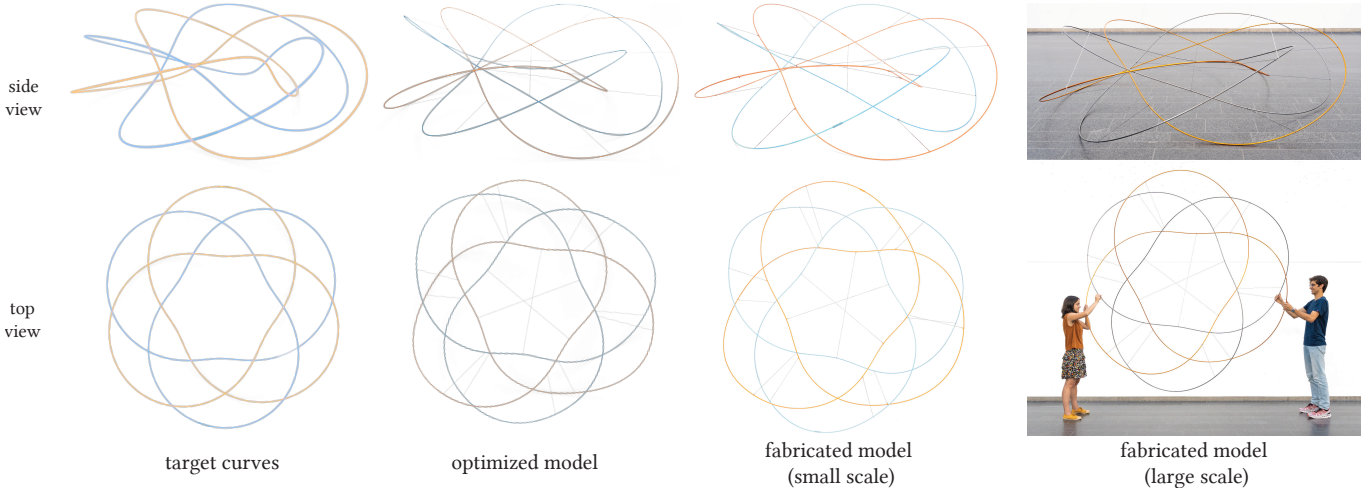

Fig. 1. *Tencers* on a torus. This model is composed of two interleaved trefoils with identical target curves of the (2,3)-torus knot with a relative rotation of $\pi/3$. We attach 15 nylon cables passing between the two knots with no cables needed within the same knot. The small and large scale models are fabricated from two straight glass-fiber rods and aluminum poles, respectively. Their bounding box diagonals measure 105 cm and 387 cm.

We study ensembles of elastic rods that are tensioned by a small set of inextensible cables. The cables induce forces that deform the initially straight, but flexible rods into 3D space curves at equilibrium. Rods can be open or closed, knotted, and arranged in arbitrary topologies. We specifically focus on equilibrium states with no contacts among rods. Our setup can thus be seen as a generalization of classical tensegrities that are composed of rigid rods and tensile cables, to also support rods that elastically deform. We show how this generalization leads to a rich design space, where complex target shapes can be achieved with a small set of elastic rods. To explore this space, we present an inverse design optimization algorithm that solves for the length and placement of cables such that the equilibrium state of the rod network best approximates a given set of input curves. We introduce appropriate sparsity terms to minimize the number of required cables, which significantly simplifies fabrication. Using our algorithm, we explore new classes of bending-active 3D structures, including elastic tensegrity knots that only require a few internal cables. We design and fabricate several physical models from basic materials that attain complex 3D shapes with unique structural properties.

CCS Concepts: • **Applied computing** → **Computer-aided design**; • **Computing methodologies** → *Modeling and simulation.*

Additional Key Words and Phrases: computational design, physics-based simulation, elastic structures, fabrication

ACM Reference Format:
Liliane-Joy Dandy, Michele Vidulis, Yingying Ren, and Mark Pauly. 2024. Tencers: Tension-Constrained Elastic Rods. 1, 1 (September 2024), 13 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

Digital fabrication has revolutionized how we make physical objects with numerous techniques to shape materials, such as additive or subtractive manufacturing, injection molding, or robotic forming. Most of these techniques aim to create a desired 3D shape at rest with minimal internal stress. A typical example is sheet metal forming, where the material is plastically deformed by pressing against an optimized rigid mold.

A complementary strategy is to shape structures based on elastic deformation. A prominent example is bending-active gridshells, primarily used in architectural designs [Becker et al. 2023; Lienhard 2014; Panetta et al. 2019; Pillwein et al. 2020; Soriano et al. 2019]. These structures can be assembled from simple straight or planar elastic components that are then deployed through specific actuation forces to achieve a desired doubly curved 3D target state.

*joint first authors.

Authors' addresses: Liliane-Joy Dandy, liliane-joy.dandy@epfl.ch, EPFL, Switzerland; Michele Vidulis, michele.vidulis@epfl.ch, EPFL, Switzerland; Yingying Ren, yingying.ren@epfl.ch, EPFL, Switzerland; Mark Pauly, mark.pauly@epfl.ch, EPFL, Switzerland.

XXXX-XXXX/2024/9-ART $15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

Here, we study a specific class of such elastic structures that we coin *tencers*, short for **ten**sion-**c**onstrained **e**lastic **r**ods. These lightweight structures consist of initially straight, flexible rods that are shaped through inextensible cables connected to the rods or external anchor points. A *tencer* is in static equilibrium when the tension forces in the cables exactly cancel the bending, stretching, and twisting forces in the rods. The resulting 3D shapes of the rods then emerge as a consequence of all the forces acting on them.

The global coupling of forces makes manual design of *tencers* challenging, as a small change to the location or length of cables can have a global impact on the equilibrium state. Furthermore, we aim for as few cables as possible in our designs to avoid visual clutter and simplify assembly. As a consequence, every cable needs to be carefully positioned to contribute as much as possible to the final equilibrium shape. The resulting topological structure of the network of cables can be non-intuitive and does not easily emerge in some incremental process where cables are added one at a time.

We therefore propose a simulation-supported inverse design algorithm for *tencers*. Our method takes as input a set of target 3D curves and produces as output the locations and lengths of cables such that the equilibrium state of the elastic rods best approximates the target curves. We first initialize our algorithm with a dense network of non-zero rest length springs that keep the rods close to the desired target. Our algorithm proceeds from this dense network and applies a greedy decimation heuristic until only tensile forces are left. A sparsifying optimization then reduces the spring network to a minimum size, while guaranteeing that all the springs stay in tension and a faithful approximation of the target curves is retained.

Our formulation supports arbitrary rod cross-sections and material properties. Rods can be open or closed and combined in knotted or linked arrangements. We are specifically interested in equilibrium states with no contacts among rods, i.e. each rod is only constrained by cables but no self-contacts or contacts with other rods. These configurations can be considered as a generalization of tensegrities, which are commonly composed of rigid rods in compression and flexible cables in tension. Allowing the rods to deform fundamentally changes the underlying kinematics and thus warrants a different computational approach than those commonly used for tensegrities, as we will discuss in more detail below.

Contributions. Our main contribution is a physics-based computational pipeline for inverse design of *tencers*. Our algorithm solves for the length and attachment points of cables such that the resulting equilibrium state approximates a set of target 3D space curves. We achieve this goal by implementing a nested optimization that tracks the equilibrium state using a physics simulation while updating the cable parameters. The number of required cables is minimized through a sparsity objective, which allows balancing the complexity of the assembly with the approximation error of the target curves. We highlight the performance of our approach with a series of design studies that illustrate the rich design space of *tencers*. We fabricate several physical prototypes that show how complex 3D shapes can be shaped from standard materials such as straight glass-fiber or aluminum poles and nylon cords. The source code of our algorithm and all the 3D models we show in the paper can be downloaded at https://go.epfl.ch/tencers.

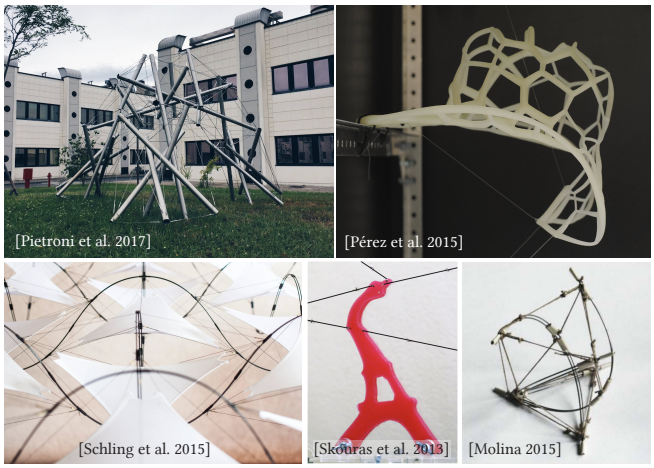

Fig. 2. Selected previous work: Position-based tensegrity design [Pietroni et al. 2017], design of flexible rod meshes [Pérez et al. 2015], bending-activated tensegrity [Schling et al. 2015], actuated deformable characters [Skouras et al. 2013] (image courtesy of Computer Graphics Laboratory, ETH Zurich), elastically bent tensegrity structures [Molina 2015].

2 RELATED WORK

Our method builds upon previous work on computational design of constrained elastic systems. Optimization and physics-based simulation are extensively used in such design processes to provide user control and enable predictive modeling. One class of design tools aims at shaping slender curvilinear elements to conform to a prescribed target (see also Figure 2). If the target is itself one-dimensional, methods were proposed to make elastic rods [Derouet-Jourdan et al. 2010; Hafner and Bickel 2021, 2023] and metallic wires [Hsiao et al. 2018] achieve the desired shape. Hair design [Hsu et al. 2023] represents a practical application of this line of research in computer graphics. Surfaces have been approximated by rigid wire assemblies [Garg et al. 2014; Miguel et al. 2016], gridshells [Becker et al. 2023; Lienhard 2014; Panetta et al. 2019; Pillwein et al. 2020; Soriano et al. 2019], elastic curve networks [Neveu et al. 2022], elastic ribbons [Liu et al. 2023; Ren et al. 2021], and rods coupled with elastic membranes [Pérez et al. 2017]. Our setting differs in that we solve for a set of cables that constrain initially straight elastic rods to deform into a set of prescribed target curves. Cables and springs have also been used as means to actuate and deploy folded surfaces [Kilian et al. 2017], release stress in architectural glass shells [Laccone et al. 2020], and levitate rigid objects while staying hidden to the viewer [Kushner et al. 2021].

In the context of character animation, tension-constrained elastic bodies have been studied by Skouras and colleagues [2013]. The aim of the authors is to find a sparse set of forces that, when exerted on an elastic solid by an external fixed frame, make the object conform to a desired set of target poses. Despite having a similar ultimate goal, in our setup we look for *pairs* of tensile forces exerted by cables anchored to the structure at both endpoints. This constraint restricts the design space, as most force distributions would not result in a self-standing structure. We then devise an initialization strategy based on a dense spring network which is reminiscent of

the Ground Structure Method [Dorn et al. 1964] commonly used in truss optimization. The dense initialization allows to compute a self-standing equilibrium state that we use to warm-start the optimizer. The resulting quadratic complexity of the network requires the introduction of a greedy pre-processing step to ensure the subsequent global optimization problem becomes tractable.

Tensegrities. Also related to our work are *tensegrities*, a class of structures that combine rigid straight elements in compression (struts) with tensile elements (cables). Despite their seeming simplicity and elegance, the form-finding problem for tensegrity structures poses significant challenges in terms of topology and shape optimization. The mechanical engineering community has addressed these challenges using a multiplicity of methods, ranging from deterministic force-based computations [Zhang and Ohsaki 2006] to more stochastic approaches, like genetic algorithms [Koohestani 2012; Lee et al. 2017] for shape space exploration, and Monte Carlo simulations to evaluate stability [Li et al. 2010].

In general, these form-finding methods make it difficult for the user to actively influence the output of the computation, and are thus not ideal for design. Tachi [2013] addressed this issue and proposed a design framework that allows for interactive editing. Gauge and colleagues [2014] show how known modules can be assembled into structures with a given shape. More recently, Pietroni and colleagues [2017] tackled the more general problem of designing stable tensegrities that minimize deviations from a user input.

The basic assumption in tensegrities is that the compressive elements are sufficiently stiff to not deform noticeably. This allows abstracting these elements as rigid bodies, which significantly simplifies algorithm design. In contrast, elastic deformation of rods is essential for *tencers*. We note that in principle, our approach can also be used for the design of classical tensegrities. This, however, would require drastically increasing the material stiffness of the rods, which would lead to stability issues in the optimization. Our method is thus not suitable for efficient design of classical tensegrities.

Tensegrities have been widely exploited in robotics and material science to improve impact resistance [Pajunen et al. 2019; Rieffel and Mouret 2018], as locomotion devices [Paul et al. 2006; Rhodes et al. 2019; Spiegel et al. 2023], and as actuators to control shape morphing behavior [Caluwaerts et al. 2014; Liu et al. 2017]. For an in-depth overview of recent research on tensegrity structures we refer to [Micheletti and Podio-Guidugli 2022]. We envisage similar application potential for *tencers*, with additional functionalities that become realizable due to the elastic deformation of the rods.

Elastic curves. The study of the shapes that physical curves can attain is often facilitated by formulating suitable energy-minimization problems [Strzelecki and von der Mosel 2017]. Different energy functional exists that represent, for example, electrostatic charge [Fukuhara 1988], self-contact interactions in a thick wire [Gonzalez and Maddocks 1999], or elastic energy [Langer and Singer 1984]. The canonical shape of a knot (i.e. its lowest-energy embedding with respect to a certain energy measure) is usually the object of interest of knot simulations [Buck and Orloff 1993]. However, as knot energies are in general non-convex, multi-stable behavior can emerge [Furrer et al. 2000; Gilsbach et al. 2021; Vidulis et al. 2023], making the computation of a global energy minimum a hard task.

Repulsive knot energies tend to remove self-contacts from the curve embedding [Yu et al. 2021]. On the other hand, it is known that a closed, knotted bending-resistant wire cannot rest in stable equilibrium without points of self-contact [Langer and Singer 1985]. Several analytical and numerical studies provide theoretical and empirical evidence of this fact [Bartels and Reiter 2021; Coleman and Swigon 2004; Diao et al. 2021; Gerlach et al. 2017]. One motivation of our work is to eliminate such self-contacts in physical knots to create *elastic tensegrity knots*.

3 OVERVIEW

Figure 3 provides an overview of our computational pipeline for inverse design of *tencers*. The algorithm takes as input a set of smooth 3D curves. These define the targets that the initially straight elastic rods should deform into when constrained by the set of cables that we aim to find. We simulate this deformation using the Discrete Elastic Rods model [Bergou et al. 2010, 2008; Kaldor et al. 2010] that represents rods as discrete polylines. Here, we focus on naturally straight rods with uniform, circular cross-section, but our method is agnostic to the rod model and can handle curved rest states and arbitrary cross-sections.

We model cables as linear springs with a specific rest length and stiffness that are connected to the rods or to external anchor points. Given the rest quantities of the rods and springs, we compute the equilibrium state of the *tencer* as detailed in Section 4.

Our inverse design optimization algorithm then solves for the number, lengths, and attachment points of the springs such that the deformed rods at equilibrium approximate the target curves within a user-specified tolerance.

We initialize the optimization by setting the deformed locations of the initially straight rods to match the target curves. This results in a non-zero elastic energy and corresponding forces acting on the nodes of the rod polyline. We compensate for these forces by adding a dense set of non-zero rest length springs connecting all pairs of nodes of the discrete curve (Section 5.1).

The algorithm now proceeds in two stages as illustrated in Figure 3. We first run a greedy algorithm that iterates between computing an equilibrium state and pruning all springs that are in compression. Once no compressed springs remain, springs with the lowest tensile force are removed, assuming that this spring contributes the least to the equilibrium state, until a user-defined distance from the target is reached (Section 5.2). While this greedy strategy can drastically reduce the number of springs, the local nature of the algorithm, in general, does not yield an optimal distribution of springs. We therefore perform a global, sparsity-inducing optimization in the second stage that tracks the equilibrium in its inner loop, while updating spring stiffness and attachment points (Section 5.3). This method only becomes tractable due to the significant reduction in springs achieved by the greedy decimation of the previous stage.

4 SIMULATION

An essential component of our inverse design optimization is accurate tracking of the equilibrium state of a *tencer*. By definition, the structure is in equilibrium when the configuration corresponds to a (local) minimum of its elastic energy.

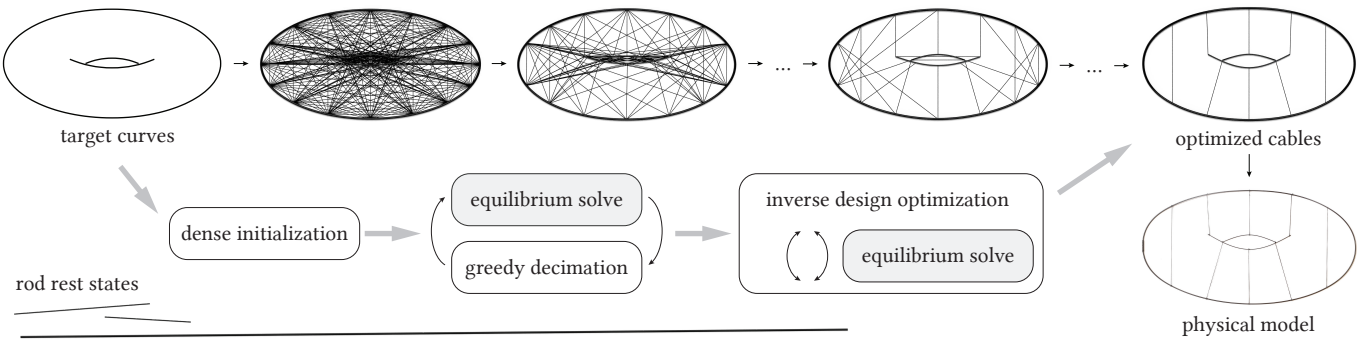

Fig. 3. Our computational pipeline takes as input a set of target curves and finds a network of cables that, when attached to straight elastic rods, force them to deform into the target curves. Our approach starts with a dense network of springs on the target curves. We then run an iterative decimation to obtain a first reduction in the number of springs, followed by a nested global optimization that further removes springs using a sparsity norm, while reducing the deviation from the targets. The physical model of this 2D line art sketch has been constructed from three straight glass-fiber rods connected with eight nylon cables.

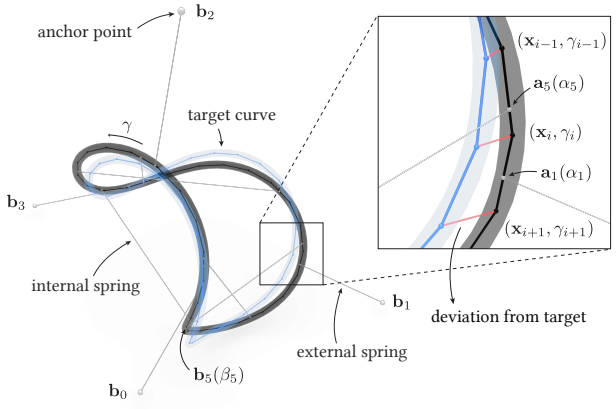

Fig. 4. Notation and terminology. A *tencer* consists of elastic rods constrained by tensile springs to approximate a given target curve. Attachment points on rods are parametrized with respect to the DER nodes.

Rods. Our simulation framework uses the Discrete Elastic Rod (DER) model [Bergou et al. 2010, 2008]. The centerline of an open DER is modeled by the 3D polyline connecting n nodes $\{\mathbf{x}_i\}_{i=0}^{n-1}$, which we collectively refer to as $\mathbf{x} \in \mathbb{R}^{n\times 3}$. The energy of a rod is the sum of a stretching, a bending, and a twisting term:

$$E_{\text{rod}}(\mathbf{x}, \boldsymbol{\theta}) = E_{\text{stretch}}(\mathbf{x}) + E_{\text{bend}}(\mathbf{x}, \boldsymbol{\theta}) + E_{\text{twist}}(\mathbf{x}, \boldsymbol{\theta}), \qquad (1)$$

where $\boldsymbol{\theta} := [\theta_0, \ldots, \theta_{n-2}] \in \mathbb{R}^{n-1}$ are the angles encoding the local orientation of the cross-section on each rod edge. The cross-section rotation is measured with respect to a reference frame updated via parallel transport. We refer to the work of Bergou and colleagues [2010; 2008] for more details on the energy terms. To handle closed rods we use the approach proposed by Vidulis et al. [2023], which only entails minor, straightforward adaptations in the notation.

We note that, while the nodal positions appear in our design objective that measures the distance from a target curve (see Section 5.3), the twist variables $\boldsymbol{\theta}$ only influence the design problem through the computation of the equilibrium state. In practice, our design algorithm is agnostic to the parametrization used in the rod model: Any model relying on an explicit representation of the rod's centerline can be used in our framework.

Cables. While rods bend and can resist compressive forces, we assume that cables only sustain tensile forces and thus are always straight. In our simulation, we model cables as linear springs, whose behavior is defined by their stiffness and rest length. While these springs can also sustain compressive forces, we ensure in our inverse design optimization that only springs with tensile forces remain in the final structure (see Section 5).

At equilibrium, the tension exerted by a linear spring can be equivalently parametrized by either rest length or stiffness, while the other quantity is held fixed. We can therefore replace all remaining springs in the final optimized network with inextensible cables (springs with infinite stiffness) by setting their rest length equal to the distance between the attachment points at equilibrium.

Tencer representation. We consider *tencers* composed of R rods and $S = S_E + S_I$ springs, where S_E *external* springs connect a rod to a fixed spatial location and S_I *internal* springs span between two rods (potentially the same). We denote the set of simulation degrees of freedom, i.e. the rod vertex coordinates, by $\mathbf{x} \in \mathbb{R}^{n\times 3}$, where $n = n_0 + \cdots + n_{R-1}$, n_r is the number of vertices of rod r, and we overloaded the symbol $\mathbf{x}$ previously used for a single rod.

We parametrize the position of internal spring endpoints as a function of the arc-length of the polygonal centerline curve. To introduce the notation, we consider an elastic rod defined by the set of n nodes $\{\mathbf{x}_i\}_{i=0}^{n-1}$. A generalization to multiple rods is straightforward. In the rest configuration, the arc-length position of node i is given by $\gamma_i = \sum_{j=0}^{i-1} \|\mathbf{x}_j - \mathbf{x}_{j+1}\|$, with $\gamma_0 = 0$. For an open rod, $\gamma_{n-1} = L$, with L denoting the total length of the rod at rest. If the rod is closed, we augment the set of nodes with a duplicate of the first node, i.e. $\mathbf{x}_n = \mathbf{x}_0$, such that $\gamma_n = L$. A spring endpoint attached to the rod is parametrized by the piecewise linear function $\alpha \mapsto \mathbf{a}(\alpha)$ s.t.

$$\mathbf{a}(\alpha) = \mathbf{x}_i + \frac{\alpha - \gamma_i}{\gamma_{i+1} - \gamma_i}(\mathbf{x}_{i+1} - \mathbf{x}_i) \quad \text{for } \alpha \in [\gamma_i, \gamma_{i+1}]. \qquad (2)$$

See Figure 4 for a schematic representation of this parametrization.

We define the energy of an internal spring as:

$$E_{\text{spring}}(\mathbf{x}; k, l, \alpha, \beta) = \frac{1}{2} k (\|\mathbf{a}(\alpha) - \mathbf{b}(\beta)\| - l)^2, \tag{3}$$

where k is the spring stiffness, l is its rest length, $\mathbf{a}$ and $\mathbf{b}$ are the spring endpoints, and α and β are the corresponding arc-length parameters. Note that the dependency on the degrees of freedom $\mathbf{x}$ is implicitly encoded in the definition of the endpoints $\mathbf{a}$ and $\mathbf{b}$. In case the spring is external, we simply drop the dependency on the arc-length parameter β associated to the second endpoint.

The explicit parametrization of the spring attachment points effectively permits sliding motion while guaranteeing that each spring anchor lies on the rod centerline. The representation is compact, with two variables for each spring, and results in a sparse energy Hessian thanks to local support stencils, thus preserving the computational efficiency typical of DER. An alternative approach to enforcing sliding motion with local stencils relies on distance-field-based regularization, and has been proposed by Skouras and colleagues [2013]. In our setting, given the one-dimensional nature of curves embedded in three-dimensional space, we prefer an explicit parametrization that uses a single variable per-point instead of three. This choice also eliminates the need of a weight parameter that tunes the regularization strength and controls the accuracy of the curve membership constraint, which in our model is always satisfied exactly. We note that a similar choice was made to model sliding contacts in rod assemblies for yarn-level cloth simulation by Sánchez-Banderas and coauthors [2020].

The total energy of the system is given by:

$$E(\mathbf{x}; \mathbf{p}) = \sum_{r=0}^{R-1} E_{\text{rod}_r}(\mathbf{x}) + \sum_{s=0}^{S-1} E_{\text{spring}_s}(\mathbf{x}; \mathbf{p}), \tag{4}$$

where $\mathbf{p} = [\mathbf{k}, \mathbf{l}, \mathbf{u}]$ is the set of design parameters, i.e. the spring stiffnesses $\mathbf{k} = [k_1, \ldots, k_S] \in \mathbb{R}^S$, rest lengths $\mathbf{l} = [l_1, \ldots, l_S] \in \mathbb{R}^S$, and attachment points along the rods' centerlines denoted as $\mathbf{u} = [\alpha_0, \alpha_1, \ldots, \alpha_{S_E-1}, \alpha_{S_E}, \beta_{S_E}, \ldots, \alpha_{S-1}, \beta_{S-1}] \in \mathbb{R}^{S_E+2S_I}$.

We can then compute an equilibrium state corresponding to a specific choice of the design parameters by minimizing the energy with respect to $\mathbf{x}$:

$$\mathbf{x}^*(\mathbf{p}) = \arg\min_{\mathbf{x}} E(\mathbf{x}; \mathbf{p}). \tag{5}$$

We refer to Section 5.3 for details on how to solve this optimization problem.

5 INVERSE DESIGN OPTIMIZATION

With equilibrium simulation in place, we can now define our inverse design optimization that aims at finding a small set of springs such that the equilibrium position of a *tencer* approximates the target curves within a user-defined threshold.

5.1 Initialization

For each individual curve of the set of target curves, we initialize a separate rod of the same length with straight rest shape. We define the initial deformed positions $\mathbf{x}$ by sampling the input curves uniformly at a user-specific resolution. Springs are initialized by creating a dense network that connects all pairs of deformed rod vertices. The rest length l of each spring is set to the initial distance of the spring endpoints with a high initial stiffness parameter (see below). These springs effectively act like stiff beams that hold the elastic rod in position. The equilibrium state of the rods is then computed according to Equation (5). Due to the large number of stiff constraints enforced by the springs, the resulting deviation of the rods from their target position is negligible at this stage.

5.2 Iterative Decimation of Springs

Pruning compressive springs. Given the initial equilibrium state, we first prune all springs whose length at equilibrium is smaller than their rest length. This removes all springs that experience compressive forces, which do not correspond to the action of a physical cable. Since this modification changes the force balance, we re-run the equilibrium solve and iterate the pruning.

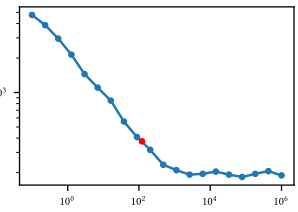

The inset figure shows a typical curve representing the number of remaining springs after pruning as a function of the initial spring stiffness. We found that a value of $k = 100\frac{Ec^2}{L}$ (red dot), where E is the rods' Young modulus, c their thickness and L their length, provides a good compromise that allows removing many compressive springs while pinning the rods very close to the target state. At the same time, we avoid extreme stiffnesses that would cause numerical difficulties. This curve was generated from the (5,3)-torus knot in Figure 10; the other design studies show analogous behavior.

External Frame. If the internal springs that remain after removing all compressed springs are not sufficient to retain the equilibrium shape within the approximation tolerance, the user is given the option to add external springs that connect each rod node to external anchor points. These anchors can be placed arbitrarily as long as their convex hull completely encloses the target curves. For a 3D shape, this can be achieved with a minimum of four points. External anchors might also be part of the design specification, for instance, if a *tencer* is suspended in a pre-defined space.

Greedy Decimation. After the elimination of all compressive springs, there can still be many springs with low tensile force contribution. We thus proceed by greedily removing the spring with the lowest force, which we assume has the smallest contribution to the global equilibrium state. After this removal, we recompute the equilibrium and prune any springs that have switched from being tensioned to being compressed. We terminate this greedy decimation once the approximation error has exceeded the user-specified tolerance. In practice, we found that setting this threshold to 5e-6·Ln, with n the number of vertices and L the rod length, strikes a good balance between target approximation and running time.

5.3 Global Optimization

Greedy decimation leads to a drastic reduction in the number of springs but does not necessarily find a satisfactory set of springs (see also Figure 5). We thus further optimize the spring network in a global optimization that proceeds in two steps. We first reduce

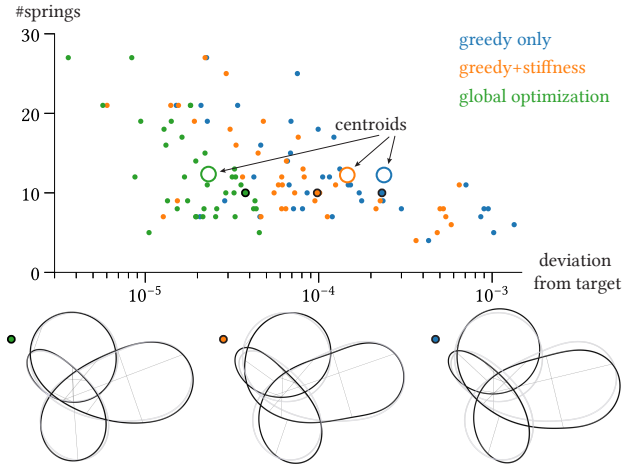

Fig. 5. The global optimization significantly improves the results compared to greedy decimation only. For models shown in Figure 14, we run both methods until they reach the same number of springs to compare the deviation from the target. The orange dots indicate the result after greedy optimization with an additional global stiffness optimization but without the sparsity objective. We highlight an example model in the bottom row illustrating the different quality of the results.

the number of springs by pushing certain spring stiffnesses to zero using a sparsity objective. Such springs can be removed from the *tencer* without incurring a change in the equilibrium shape. We then optimize for the attachment points of the springs to further reduce the deviation from the target curves.

We achieve this goal by formulating an objective function that is composed of two terms as

$$\mathcal{J}(\mathbf{p}) = \mathcal{T}(\mathbf{x}^*(\mathbf{p})) + \mathcal{S}(\mathbf{p}). \tag{6}$$

The term $\mathcal{T}$ measures the distance between the current equilibrium state $\mathbf{x}^*(\mathbf{p})$ and the target curves as

$$\mathcal{T}(\mathbf{x}) = \sum_{i=0}^{n-1} \|\mathbf{x}_i - \bar{\mathbf{x}}_i\|^2, \tag{7}$$

where $\bar{\mathbf{x}}_i$ indicates the vertex coordinates of the discretized target curves. The term $\mathcal{S}$ enforces spring sparsification:

$$\mathcal{S}(\mathbf{p}) = \omega_e \sum_{s=1}^{S_E} k_s^q + \omega_i \sum_{s=S_E+1}^{S_E+S_I} k_s^q \tag{8}$$

where ω_e and ω_i are weights that control the relative importance of sparsity for external and internal springs, respectively, in relation to the shape optimization term. In practice, these weights are first set to 0 to get as close to the target as possible, then set to 1 and iteratively increased by a factor of 10 to explore several designs until an appropriate trade-off between the number of springs and the distance to the target is achieved (see Figure 6). For structures with an external frame, external springs are sparsified first.

Ideally, we would want to use an L_0-norm so that $\mathcal{S}$ simply counts the number of springs with non-zero stiffness. Unfortunately, the L_0-norm is discontinuous. Following Skouras and coauthors [2013], we then use an L_q-norm, with $q = \frac{1}{2}$, to trade exact counting for differentiability. Since L_q-norms with $q < 1$ are still not differentiable at zero, we perform the change of variables $K = k^q$ proposed by Panetta et al. [2019] to eliminate the singularity. We also set non-negativity bounds on $\mathbf{k}$ to constrain the design space to physically meaningful solutions. Figure 6 illustrates the effect of the sparsification term $\mathcal{S}$.

The pairwise distance objective defined in (7) is sensitive to rigid motion. For self-standing shapes, whose global position and orientation are not fixed by external anchor points, we align the target to the shape at every iteration during optimization. To rigidly register the curves, we translate their centers of mass and perform closest-point matching with known pairing by computing a singular value decomposition [Sorkine-Hornung and Rabinovich 2017].

Spring rest length. Since the equilibrium solve is now in the inner loop of the optimization, we cannot easily prune compressive springs without creating discontinuities. Instead, we use the following strategy: After the greedy decimation stage, we set the rest lengths of all springs to zero and adapt the spring stiffnesses to obtain the same force profile of the current equilibrium state $\mathbf{x}^*$ (see also the paragraph on cables in Section 4). This ensures that $\mathbf{x}^*$ remains an equilibrium state also for the zero-rest-length springs. Since such springs will always be in tension, we do not need to prune them from the structure. The reduction in springs will be solely handled by the sparsity objective.

One additional issue arises. While the adaptation of spring stiffnesses ensures that the force profile remains identical when switching to springs with zero rest length, the equilibrium state is no longer guaranteed to be stable, i.e. can be at a saddle point of the elastic energy. This is problematic since the solver will then immediately escape to a different configuration with lower elastic energy, but potentially significant deviations from the target curves. We detect this issue by analyzing the eigenvalues of the energy Hessian after having set the rest lengths to zero. Negative eigenvalues indicate that we are at a saddle point. In this case, we search for the smallest rest lengths that result in a positive definite Hessian, since the smaller the rest length is, the less likely it is for springs to get compressed. We test 99 values, uniformly spaced from 1% to 99% of the current rest lengths. This check only takes a few of seconds on the largest designs. In our experiments, we generally observe that we only require minor increases in rest length to obtain a stable equilibrium. This ensures that we will be able to replace the springs in the final structure with inextensible cables. We comment on a few exceptions in Section 6.3.

Spring Attachment Sliding. So far, we keep the spring attachment points fixed, while optimizing the spring stiffnesses using the sparsity term to further reduce the number of springs. Attachment points have been constrained to the vertices of the DER polyline selected during initialization. We now apply an additional optimization step that helps reduce the approximation error by letting the spring attachment points slide along the rod centerlines. In this step, we remove the sparsity term $\mathcal{S}$ from the objective function and only optimize $\mathcal{T}$ over both $\mathbf{u}$ and $\mathbf{k}$.

For open rods, we constrain each arc-length position of each spring endpoint to lie in the interval $[0, L]$, where L is the length of the corresponding rod centerline. This prevents the springs from

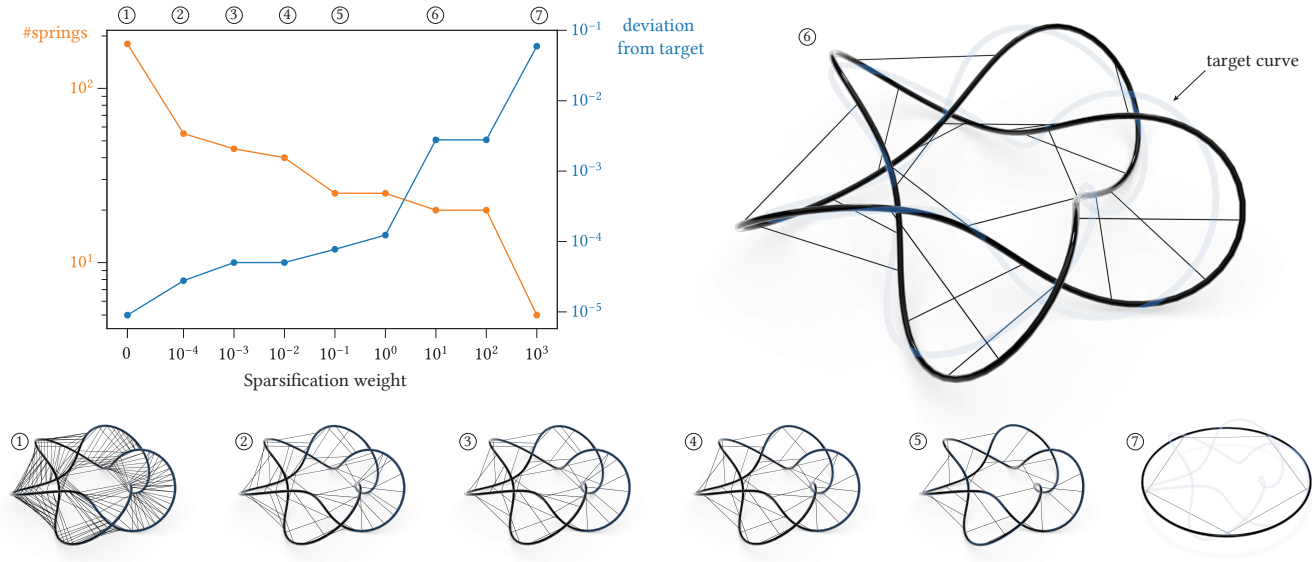

Fig. 6. Effect of the sparsification weight ω_i on the number of springs and the deviation from the target curve, shown with transparency. Our algorithm allows exploring different *tencer* designs to find the most appropriate trade-off. The example shows a (5,2)-torus knot.

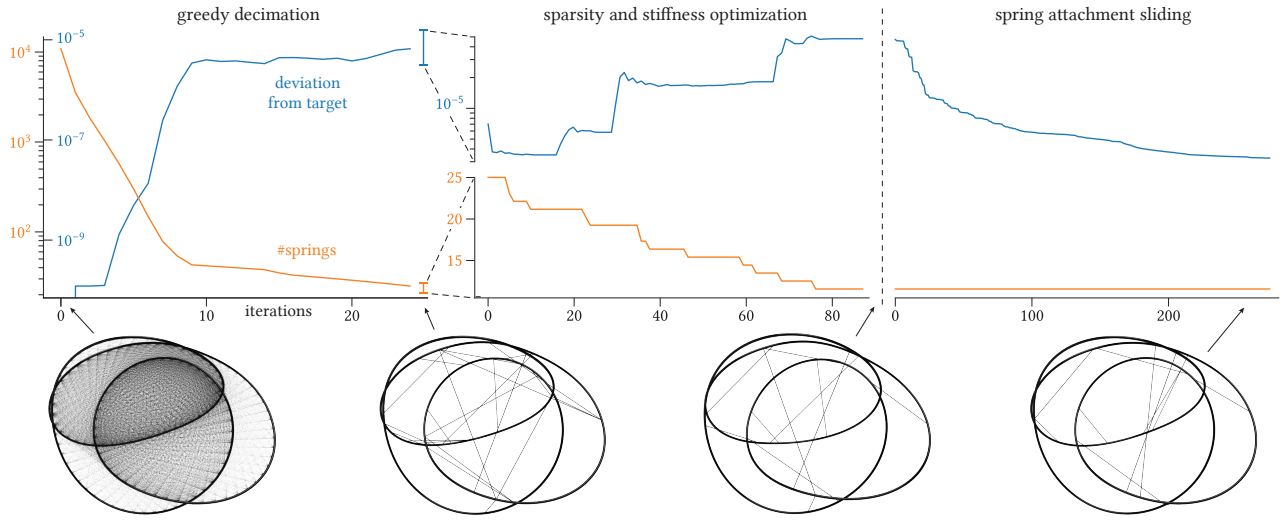

Fig. 7. A typical evolution of our optimization pipeline. We start from a dense network of springs that almost perfectly reproduces the target curves. The greedy method and global optimization reduce the number of springs at the cost of an increase in the deviation from the target curves. The final stage reduces this deviation by fine-tuning the spring attachments and letting them slide along the rods.

sliding off the rod. For closed rods, we measure the positions modulo L to allow cyclic sliding. As illustrated in Figure 7, this additional spring sliding step reduces the approximation error noticeably.

Numerical Solver. As the distance-to-target term $\mathcal{T}$ is evaluated at equilibrium, we execute two nested optimization loops. The inner loop solves for the equilibrium corresponding to a given set of springs, while the outer loop minimizes the objective by improving the design parameters $\mathbf{p}$. For the inner equilibrium solve we use the Newton-based optimizer originally proposed by Panetta and coauthors [2019]. For rods with a few hundred vertices at most, we empirically select a convergence tolerance of 1e-7, and observe that setting the maximum number of iterations to 1000 is sufficient in most cases. To eliminate the zero-energy rigid motion of the system at each Newton step, we shift the Hessian eigenvalues by a small amount (1e-6 in our experiments). We refer to [Panetta et al. 2019] for more details on the Newton-CG trust region method, implemented in the commercial solver Knitro [Byrd et al. 2006], that we use to solve the outer design loop. In our experiments, we set the trust region radius to 5e-3, the gradient convergence tolerance to

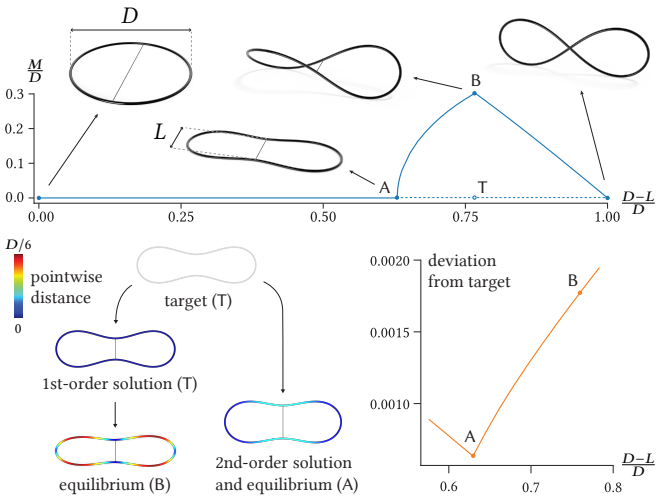

Fig. 8. *Top*: An elastic ring buckles into 3D as the length L of a cable spanning its diameter D is reduced. M indicates the smallest side of the tightest bounding box. As L approaches zero, the rod flattens into a self-intersecting eight-shaped configuration. The solid and dashed lines represent stable and unstable equilibria, respectively. *Bottom*: If one of the 2D unstable equilibria is chosen as target (T), the first-order optimizer returns an exact match, regardless of the stable equilibrium for that cable length being far (B); the second-order optimizer converges to the closest stable equilibrium (A). The measured deviation from the target (right) confirms that the second-order solution is the correct stable minimizer.

1e-5 and maximum number of iterations of the outer loop to 1000, as we observe that a few hundreds of iterations at most are generally sufficient for the algorithm to converge. To ensure reproducibility, the data and the source code implementing our method can be found at https://go.epfl.ch/tencers.

Equilibrium conditions. During optimization, we enforce second-order equilibrium conditions as hard constraints by solving Equation (5). We notice that first-order equilibrium conditions enforced as soft constraints are often sufficient to simulate material systems that are less prone to buckling [Skouras et al. 2013]. Computing stable equilibrium states, however, is crucial for *tencers*, as saddle points of the energy are frequently encountered, and the corresponding configurations would buckle away from the desired target if fabricated. Figure 8 shows a simple *tencer* in which this buckling behavior occurs. We consider an elastic ring with a cable spanning its diameter. As the cable is shortened, the 2D configuration becomes unstable, and the ring warps out of plane. If one of the unstable equilibria is used as target for a design problem, first-order equilibrium conditions would misleadingly result in a perfect match, even though the physical structure would settle into a significantly different 3D state. Second-order conditions allow computing the stable state closest to the target. Analogous buckling behavior was frequently observed in all our designs. We notice how exactly enforcing equilibrium conditions also provides an advantage in interactive design exploration applications, as every state the optimizer visits is a physically stable state ready to be fabricated.

6 RESULTS AND DISCUSSION

In this section, we show a number of design studies created with our optimization framework. Please refer to the supplemental video to better appreciate the 3D shape of these models. We also present experiments to evaluate our approach, and discuss limitations and directions for future work.

Fabrication. To validate the predictive accuracy of our simulation, we fabricated four models, see Figures 1, 3, and 11. Overall, we see great agreement between simulation and physical models, in line with several previous works that have validated the accuracy of the DER model. The pool of fabricated designs encompasses the whole set of features *tencers* support: 2D and 3D targets, made of a single or multiple, closed or open curves; the rods can have different material properties; the shape can be self-standing or suspended from an external frame. We refer to the supplemental video for more comparisons between physical models and simulation results.

Three prototypes are built from elastic glass fiber rods with a diameter of 2 mm (the 2D torus also contains two 1 mm rods), tensioned by nylon cables. The fourth prototype (Figure 1, right) is made from aluminum poles with a diameter of 8.5 mm and polyester cables. For the models that include glass fiber closed rods, the rod extremities are inserted into a spring pin and glued with thick cyanoacrylate glue. The cables are attached to the rods using clove hitch knots. After all cables have been positioned, their endpoints are glued to the rods to prevent sliding.

Symmetry constraints. We provide an additional feature for symmetric curves by allowing the user to specify symmetry constraints on the springs. These consist of identity constraints on the symmetric springs' stiffnesses and affine equality constraints for the endpoint positions. They are enforced through a change of variables in the design objective and can represent different types of symmetry, such as reflection (see Figure 15 and 13) or rotational (Figure 1, 6, 9, 10, and 12).

Performance. We sample target curves in a range of 100 to 400 samples depending on the curve complexity. Runtime varies significantly across models due to the intricacies of the Newton line-search method. We measured runtimes between a few seconds and 40 minutes on a 64-Core AMD Ryzen Threadripper 3990X Processor with 128 GB of RAM. See Table 1 for detailed timing data.

6.1 Design Studies

Figure 9 shows the simplest possible knot, the *trefoil*. Without cables, the minimum energy state for a rod with straight rest shape is a double-covered circle with continuous self-contacts [Gerlach et al. 2017]. We specify three different analytical curves with trefoil knot topology as targets and run our optimization to find the placement and lengths of cables that force the rod towards the desired curve.

Figure 10 shows a *tencer* that approximates a 5-3 torus knot, modeled by the analytical function shown in the figure. The parameter p counts how many times the curve winds around the axis of rotational symmetry of the torus, while q counts how many times the curve winds around the inner circle of the torus. The same function is used in Figure 1, where two interleaved trefoil torus knots mutually suspend each other.

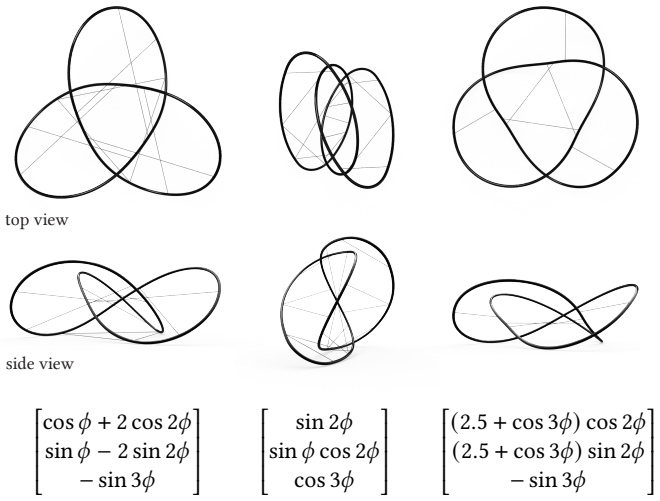

Fig. 9. Three different analytical functions define the target curves of the trefoil knot, formed from a single straight elastic rod.

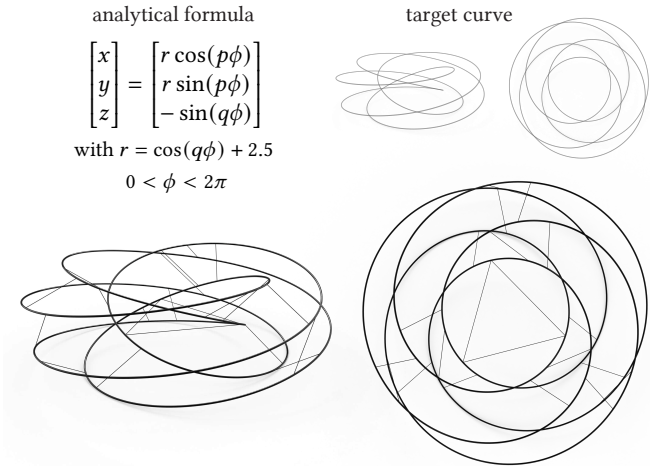

Fig. 10. An elastic rod is tensioned by 21 cables to form a (5,3)-torus knot. The target curve is given by the analytical formula of the (p, q)-torus knot.

Figure 11 shows a single open rod that is deformed into a helix. This structure cannot retain its shape with internal springs only, so the algorithm adds external springs to suspend the structure from a cubic frame specified by the user.

Figure 12 shows the Borromean rings, three closed rods that form a Brunnian link [Berrick et al. 2006]. Such a link cannot be separated, but breaks apart into separate unknotted loops when any one of its components is removed. While a single closed rod would settle into a circle at equilibrium, it is known that the Borromean ring topology cannot be realized with exact circles [Freedman and Skora 1987]. For the model at the top of Figure 12, we thus prescribe as target curves three ellipses that can be embedded without self-intersections. Cables are constrained to only pass between different rods to create mutual dependency between the rings in the equilibrium state. The free-form variation shown on the right of the same figure requires more cables, including cables within the same rods.

We created some designs based on 2D and 3D sketches that illustrate the potential of *tencers* for freeform design. We adapt our pipeline to 2D examples by pinning the z-coordinates of all variables. Figure 3 shows a *tencer* composed of three separate rods that reproduce a 2D line-sketch of a 3D torus. The resulting structure is planar. The two inner rods have a cross-section with 1 mm diameter, while the outer rod has a diameter of 2 mm. Given the mild curvature of the target curves, a coarse discretization with only 26 total nodes is sufficient to model the shape accurately during optimization. This choice results in an efficient, almost-interactive update rate, with the whole design pipeline taking only 10 seconds (see Table 1). The final result is then forward-simulated using 100 nodes per rod to confirm its stability. The maximum distance between the physical model and the optimized result is 0.88% of the bounding box diagonal.

In Figure 13 we approximate the curves of the SIGGRAPH Logo. For this model with open rods, external anchor points are needed to constrain the endpoints of the rods.

Figure 15 shows a design study with a *tencer* composed of three closed rods that approximate the head of an owl. The optimized cables respect the reflection symmetry of the target curves.

6.2 Evaluation

The above studies give an indication of the flexibility for design provided by our framework. To gain further insights into the design space of *tencers* and the performance of our algorithm, we conducted

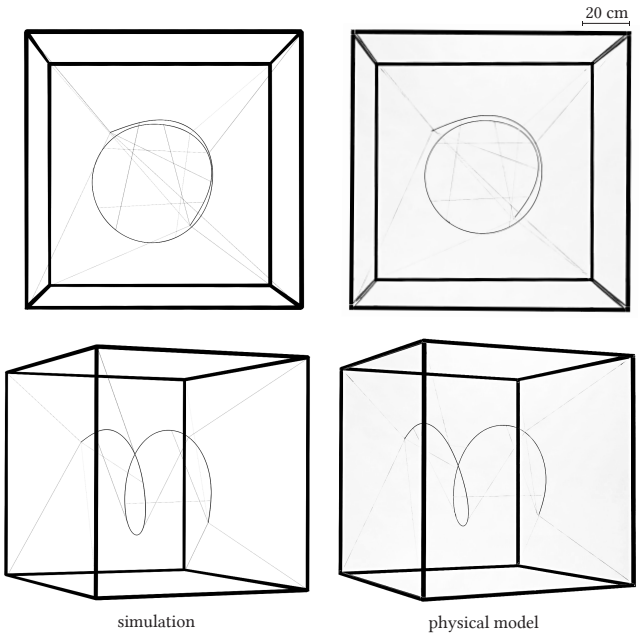

Fig. 11. A 3D helix formed by a straight elastic rod suspended from a frame, viewed from two different angles (top and bottom). The target curve is a cylindrical helix with a diameter of 50 cm and a pitch of 33 cm. The rod is constrained by 6 internal and 8 external cables. We refer to the accompanying video for a 3D view of the model.

Table 1. Statistics of examples shown in the paper: Number of rods, number of vertices per rod, rod lengths in cm, and timings for the greedy decimation, spring sparsification and spring attachment sliding phase, provided in seconds. Max. Error measures the maximum distance between the simulated curves and the target, in % of the model bounding box. The value of the rods' Young modulus is set to 40 GPa, although the final result is insensitive to this value as the simulation, in particular the initial spring stiffness, is scaled. The thickness of the rod is set to 2 mm, except for the 2D torus that also contains two 1 mm rods.

				Running times (s)			
Model (Figure)	#rods	#vertices	Rod lengths (cm)	Greedy	Sparsification	Sliding	Max. Error
Twin trefoils (1)	2	102,102	398,398	34	40	22	5.4
2D torus (3)	3	16,5,5	96,16,10	0.3	8	0.8	0.7
Trefoil (9, left)	1	102	198	1	9	1	0.8
Trefoil (9, center)	1	102	198	1	27	9	1.7
Trefoil (9, right)	1	102	198	0.6	11	2	5.7
Torus knot (10)	1	201	437	37	16	2	1.1
Helix (11)	1	100	96	0.7	69	7	5.3
Borromean rings (12, left)	3	100 (×3)	36 (×3)	3	6	40	0.3
Borromean rings (12, right)	3	100 (×3)	85 (×3)	22	928	51	0.5
SIGGRAPH logo (13)	6	37,43,27 (×2)	33,43,40 (×2)	2	31	18	1.6
Owl (15)	3	64,256,24	305,699,42	12	1513	715	7.6
Treble key	1	136	35	6	228	5	2.8

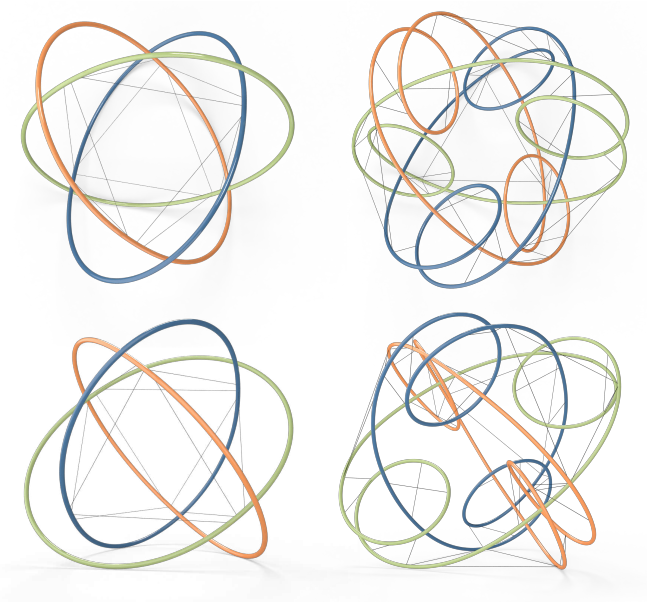

Fig. 12. Two topological links with interconnected components. Left: A *tencer* with three ellipses arranged as Borromean rings, realized with 12 cables. Right: A *tencer* with geometrically more complex rings requires 48 cables. The top and bottom images show different views of the same model.

two additional experiments. In a first study, we generated 50 target curves by computing the forward equilibrium of a closed rod constrained by different numbers of random springs with zero rest length. While only representing a small sample of the infinite space of realizable *tencer* knots, these curves exhibit geometric and topological diversity as illustrated in Figure 14. The known, ground-truth solution of these target curves facilitates a quantitative analysis of the performance of our optimization algorithm. We observe that our method was successful in recovering the target curves, albeit often with a different set of strings. This indicates that there can be multiple solutions that provide good approximations of the same target curve. This offers additional design freedom to satisfy other potential constraints, for example on the maximum length of cables or to exclude certain regions that cables can pass through.

For our second study, we generate visually pleasing 3D curves by applying the repulsive-curves method of [Yu et al. 2021] on a set of knot topologies. While the resulting target curves are smooth and have relatively low curvature, we cannot assume that an accurate *tencer* with few cables exists. Figure 16 shows the results produced by our method. Overall, these experiments indicate that *tencers* offer a rich design space with many opportunities for creative exploration.

6.3 Limitations and Future Work

Our evaluation indicates that our optimization provides an effective tool for the design of *tencer* structures. Since our objective is nonlinear and non-convex, our gradient-based optimization in general converges only to a local minimum. While performing well in practice, we cannot guarantee that our sparsity optimization finds the globally optimal solution with the smallest number of cables. More fundamentally, we do not yet have a comprehensive understanding of the design space of *tencers*. Specifically, we cannot predict if a given target curve is realizable with internal cables only or how many cables are needed to reach a certain approximation error.

Currently, the number of cables in the output is controlled indirectly by adapting the sparsity weights. Due to the above-mentioned local minima, it is not always guaranteed that a higher sparsity weight leads to fewer cables, which complicates the direct exploration by tuning this parameter. In the future, we aim to implement an automatic weight tuning scheme that allows directly specifying the desired number of cables. The algorithm should then automatically adapt the sparsity weights to reach this target.

The instability that can occur when we transition from springs with finite rest length to zero rest length in the global optimization

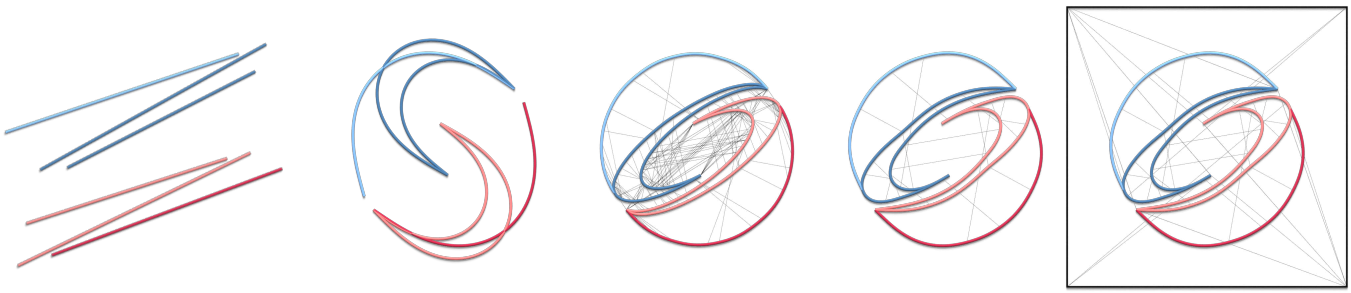

Fig. 13. Six initially straight rods are deformed to match the SIGGRAPH logo. The endpoints of the six rods are pinned in place. These constraints alone are not sufficient for faithfully reproducing the shape. Our algorithm is able to find a network of only 20 internal cables that faithfully approximates the input. The rod endpoints can then be fixed using additional cables attached to an outer frame.

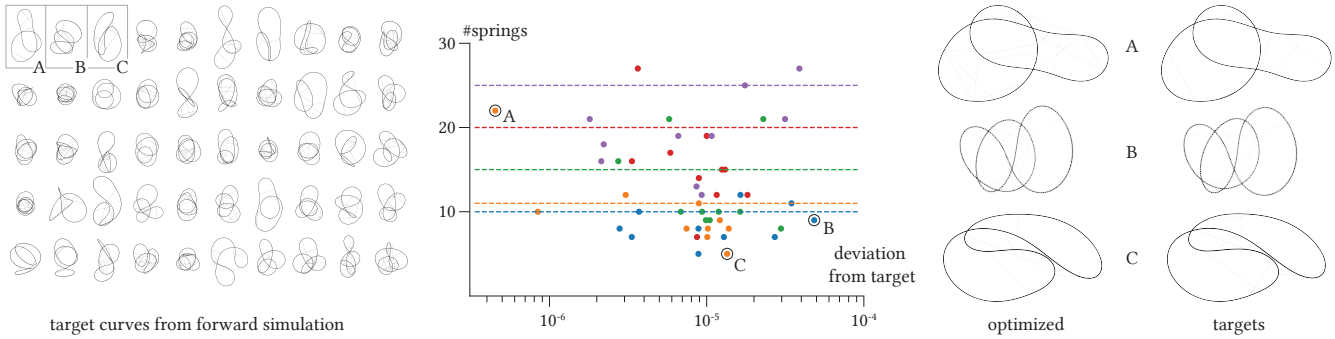

Fig. 14. A set of target curves (left) that have been generated with forward simulation of random cable attachments on a closed rod. In the plot (middle), colored horizontal lines indicate the number of cables used to create these targets. The dots show the number of cables in the corresponding optimized solution. On the x-axis we plot the approximation error. Two highlighted model show examples where our optimization produced significantly more (A) or fewer (C) cables than the targets. All models retained the target shape to within the tolerance threshold with model B having the highest deviation from the target.

requires gradually increasing rest lengths until the equilibrium state is a stable local minimum (see Section 5.3). At this point, we have no guarantee that this increase does not create springs in compression during optimization. In such a case, which occurred for 3 of the 50 curves used for the method evaluation shown in Figure 14, we terminate the global optimization early to ensure fairness and consistency across the random sample. Instabilities were resolved as discussed in Section 5.3 for all the other design studies.

When a *tencer* is made of multiple rods, there is no guarantee that the spring network resulting from the optimization procedure blocks all zero-energy relative rigid motion between rods. While this situation did not occur in any of our results, we can easily detect such a case by checking if the energy Hessian has any zero eigenvalues. More cables can then be added to block all inter-rod rigid motions.

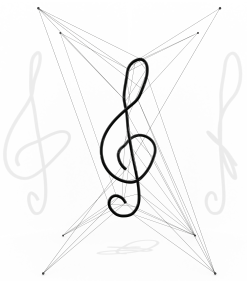

In the inset, we show an example where a 2D curve has been manually transformed into a 3D target curve. The orthogonal shadow projections give an indication of the 3D shape. While the optimization finds a solution, the number of cables is high, even when we allow significant deviations from the input. Here a more satisfactory solution could be found if our shape preservation term was not formulated directly on the 3D vertex positions, but only on their 2D projections,

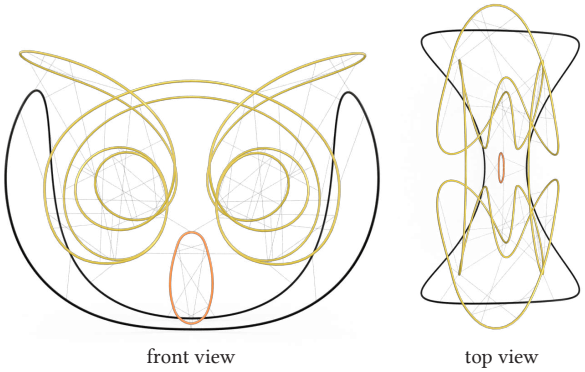

Fig. 15. A *tencer* composed of three closed rods constrained by 72 cables. The target curves are from a 3D sketch of an owl's head. The optimized *tencer* approximates the target shape closely.

since the design intent is mainly to reproduce the 2D curve from a single point of view. This would increase the search space and potentially allow for a solution with significantly fewer cables.

While our design optimization exclusively considers stable equilibrium states as candidate results (forward), we do not analyze the stability of the found local minimum with respect to small inaccuracies in the assembly of cable positions and lengths (inverse). In practice, the models we fabricated did not show such instability and all settled in the predicted equilibrium state.

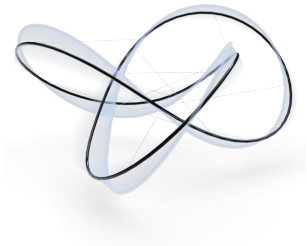

The inset shows a vibrational mode of a *tencer*, illustrated as an offset curve proportional to the mode magnitude. A more in-depth analysis of the deformation behavior could be a basis for future algorithms that, in addition to geometric fidelity, also optimize the *tencer* for structural behavior. This will be important for applications, for example in architecture, or if *tencers* are used as building blocks for mechanical metamaterials.

Other interesting areas of future work include actuated systems with potential applications in robotics. For example, *tencers* could be used to design string-actuated robots that dynamically deform to create locomotion, absorb shocks, and adapt the robot's shape to different environments. Tensegrity-based robots are classically built using elastic cables. Our system adds compliance in the rods, introducing new morphing behavior and material properties. Studying buckling instabilities to trigger fast and drastic shape transformations is another avenue of future research. We also envision fabrics replacing or complementing tensile cables, with interesting potential applications in adaptive facades.

Finally, *tencers* are interesting objects of theoretical study. We mention a few open questions: For an elastic knot of a specific knot type, what is the minimal number of cables needed to separate all self-contacts? Does a *tencer* have multiple stable equilibrium states? If so, is a transition between such states possible without inter-penetrations of rods or cables?

7 CONCLUSIONS

We introduced *tencers*, tensioned-constraint elastic rods, as a versatile new class of lightweight, bending-active structures. The complex balance of elastic bending and compression forces in the rods and tensile forces induced by the cables makes manual design of *tencers* highly challenging. Our optimization algorithm addresses this challenge and provides an effective framework for design exploration.

ACKNOWLEDGMENTS

We thank Filip Goč and Florin Isvoranu for their help in fabricating the physical models. This research was supported by the Swiss National Science Foundation (Grant 188582).

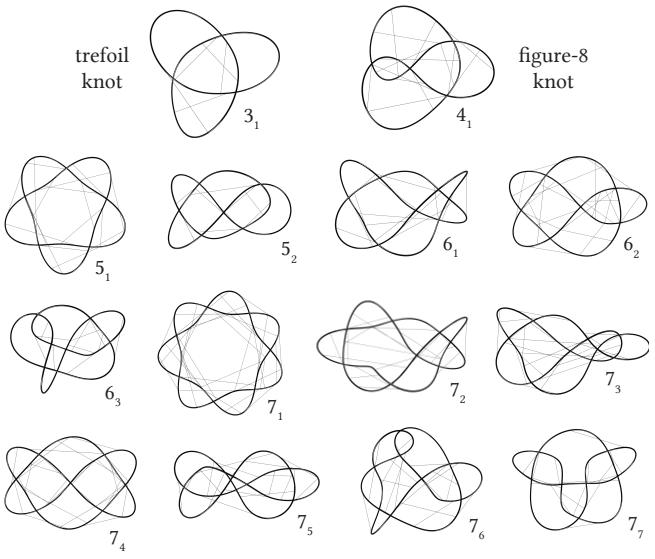

Knot	3_1	4_1	5_1	5_2	6_1	6_2	6_3	7_1	7_2	7_3	7_4	7_5	7_6	7_7
#springs	5	14	17	9	17	22	12	35	21	23	26	19	20	17
Max. Error	2.59	1.95	0.26	3.59	0.65	0.37	0.80	0.33	0.82	2.67	0.68	1.55	0.34	2.11

Fig. 16. *Tencers* optimized for target curves generated with the repulsive-curves algorithm of [Yu et al. 2021]. We show the first 14 knots of common knot tables labeled in Alexander-Briggs notation [Adams 2004] where the first number indicates the crossing number of the knot, while the subscript is an arbitrary index. The table shows the number of springs and the maximum deviation from the target in % of the bounding box diagonal.

REFERENCES

C. C. Adams. 2004. *The Knot Book: An Elementary Introduction to the Mathematical Theory of Knots.* American Mathematical Soc.

S. Bartels and P. Reiter. 2021. Stability of a Simple Scheme for the Approximation of Elastic Knots and Self-Avoiding Inextensible Curves. *Math. Comp.* 90, 330 (July 2021), 1499–1526. https://doi.org/10.1090/mcom/3633

Q. Becker, S. Suzuki, Y. Ren, D. Pellis, J. Panetta, and M. Pauly. 2023. C-Shells: Deployable Gridshells with Curved Beams. *ACM Transactions on Graphics* 42, 6 (Dec. 2023), 1–17. https://doi.org/10.1145/3618366

M. Bergou, B. Audoly, E. Vouga, M. Wardetzky, and E. Grinspun. 2010. Discrete viscous threads. *ACM Trans. Graph.* 29, 4, Article 116 (jul 2010), 10 pages. https://doi.org/10.1145/1778765.1778853

M. Bergou, M. Wardetzky, S. Robinson, B. Audoly, and E. Grinspun. 2008. Discrete elastic rods. *ACM Trans. Graph.* 27, 3 (aug 2008), 1–12. https://doi.org/10.1145/1360612.1360662

A. J. Berrick, F. Cohen, Y. L. Wong, and J. Wu. 2006. Configurations, braids, and homotopy groups. *Journal of the American Mathematical Society* 19 (04 2006). https://doi.org/10.1090/S0894-0347-05-00507-2

G. Buck and J. Orloff. 1993. Computing Canonical Conformations for Knots. *Topology and its Applications* 51, 3 (July 1993), 247–253. https://doi.org/10.1016/0166-8641(93)90079-S

R. H. Byrd, J. Nocedal, and R. A. Waltz. 2006. KNITRO: An integrated package for nonlinear optimization. In *Large-Scale Nonlinear Optimization*, G. di Pillo and M. Roma (Eds.). Springer, 35–59.

K. Caluwaerts, J. Despraz, A. Işçen, A. P. Sabelhaus, J. Bruce, B. Schrauwen, and V. SunSpiral. 2014. Design and Control of Compliant Tensegrity Robots through Simulation and Hardware Validation. *Journal of The Royal Society Interface* 11, 98 (Sept. 2014), 20140520. https://doi.org/10.1098/rsif.2014.0520

B. Coleman and D. Swigon. 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 (Aug. 2004), 1281–99. https://doi.org/10.1098/rsta.2004.1393

A. Derouet-Jourdan, F. Bertails-Descoubes, and J. Thollot. 2010. Stable Inverse Dynamic Curves. *ACM Transactions on Graphics* 29, 6 (Dec. 2010), 1–10. https://doi.org/10.1145/1882261.1866159

Y. Diao, C. Ernst, and P. Reiter. 2021. Knots with Equal Bridge Index and Braid Index. *Journal of Knot Theory and Its Ramifications* (2021). https://doi.org/10.1142/

s0218216521500759
W. S. Dorn, R. E. Gomory, and H. J. Greenberg. 1964. Automatic Design of Optimal Structures. *Journal de Mecanique* 3 (1964), 25–52.
M. H. Freedman and R. Skora. 1987. Strange actions of groups on spheres. *Journal of Differential Geometry* 25, 1 (1987), 75 – 98. https://doi.org/10.4310/jdg/1214440725
S. Fukuhara. 1988. Energy of a knot. In *A Fête of Topology*. Elsevier, 443–451.
P. B. Furrer, R. S. Manning, and J. H. Maddocks. 2000. DNA Rings with Multiple Energy Minima. *Biophysical Journal* 79, 1 (July 2000), 116–136. https://doi.org/10.1016/S0006-3495(00)76277-1
A. Garg, A. O. Sageman-Furnas, B. Deng, Y. Yue, E. Grinspun, M. Pauly, and M. Wardetzky. 2014. Wire Mesh Design. *ACM Transactions on Graphics* 33, 4 (July 2014), 1–12. https://doi.org/10.1145/2601097.2601106
D. Gauge, S. Coros, S. Mani, and B. Thomaszewski. 2014. Interactive Design of Modular Tensegrity Characters. *Eurographics/ ACM SIGGRAPH Symposium on Computer Animation* (2014), 8 pages. https://doi.org/10.2312/SCA.20141131
H. Gerlach, P. Reiter, and H. von der Mosel. 2017. The Elastic Trefoil Is the Twice Covered Circle. *Archive for Rational Mechanics and Analysis* 225, 1 (July 2017), 89–139. https://doi.org/10.1007/s00205-017-1100-9
A. Gilsbach, P. Reiter, and H. von der Mosel. 2021. Symmetric Elastic Knots. *Math. Ann.* 385 (May 2021), 811–844. https://doi.org/10.1007/s00208-021-02346-9
O. Gonzalez and J. H. Maddocks. 1999. Global Curvature, Thickness, and the Ideal Shapes of Knots. *Proceedings of the National Academy of Sciences* 96, 9 (April 1999), 4769–4773. https://doi.org/10.1073/pnas.96.9.4769
C. Hafner and B. Bickel. 2021. The Design Space of Plane Elastic Curves. *ACM Transactions on Graphics* 40, 4 (Aug. 2021), 1–20. https://doi.org/10.1145/3450626.3459800
C. Hafner and B. Bickel. 2023. The Design Space of Kirchhoff Rods. *ACM Transactions on Graphics* 42, 5 (Oct. 2023), 1–20. https://doi.org/10.1145/3606033
K.-W. Hsiao, J.-B. Huang, and H.-K. Chu. 2018. Multi-View Wire Art. *ACM Transactions on Graphics* 37, 6 (Dec. 2018), 1–11. https://doi.org/10.1145/3272127.3275070
J. Hsu, T. Wang, Z. Pan, X. Gao, C. Yuksel, and K. Wu. 2023. Sag-Free Initialization for Strand-Based Hybrid Hair Simulation. *ACM Trans. Graph.* 42, 4, Article 74 (jul 2023), 14 pages. https://doi.org/10.1145/3592143
J. M. Kaldor, D. L. James, and S. Marschner. 2010. Efficient Yarn-Based Cloth with Adaptive Contact Linearization. *ACM Transactions on Graphics* 29, 4 (July 2010), 1–10. https://doi.org/10.1145/1778765.1778842
M. Kilian, A. Monszpart, and N. J. Mitra. 2017. String Actuated Curved Folded Surfaces. *ACM Transactions on Graphics* 36, 4 (Aug. 2017), 1. https://doi.org/10.1145/3072959.3015460
K. Koohestani. 2012. Form-Finding of Tensegrity Structures via Genetic Algorithm. *International Journal of Solids and Structures* 49, 5 (March 2012), 739–747. https://doi.org/10.1016/j.ijsolstr.2011.11.015
S. Kushner, R. Ulinski, K. Singh, D. I. Levin, and A. Jacobson. 2021. Levitating Rigid Objects with Hidden Rods and Wires. *Computer Graphics Forum* 40, 2 (May 2021), 221–230. https://doi.org/10.1111/cgf.142627
F. Laccone, L. Malomo, M. Froli, P. Cignoni, and N. Pietroni. 2020. Automatic Design of Cable-Tensioned Glass Shells. *Computer Graphics Forum* 39, 1 (Feb. 2020), 260–273. https://doi.org/10.1111/cgf.13801
J. Langer and D. A. Singer. 1984. Knotted Elastic Curves in R3. *Journal of the London Mathematical Society* s2-30, 3 (Dec. 1984), 512–520. https://doi.org/10.1112/jlms/s2-30.3.512
J. Langer and D. A. Singer. 1985. Curve Straightening and a Minimax Argument for Closed Elastic Curves. *Topology* 24, 1 (Jan. 1985), 75–88. https://doi.org/10.1016/0040-9383(85)90046-1
S. Lee, J. Lee, and J. Kang. 2017. A Genetic Algorithm Based Form-finding of Tensegrity Structures with Multiple Self-stress States. *Journal of Asian Architecture and Building Engineering* 16, 1 (Jan. 2017), 155–162. https://doi.org/10.3130/jaabe.16.155
Y. Li, X.-Q. Feng, Y.-P. Cao, and H. Gao. 2010. A Monte Carlo Form-Finding Method for Large Scale Regular and Irregular Tensegrity Structures. *International Journal of Solids and Structures* 47, 14-15 (July 2010), 1888–1898. https://doi.org/10.1016/j.ijsolstr.2010.03.026
J. Lienhard. 2014. *Bending-Active Structures: Form-finding Strategies Using Elastic Deformation in Static and Kinetic Systems and the Structural Potentials Therein.* Universität Stuttgart Inst. f. Tragkonstr.
D. Liu, D. Pellis, Y.-C. Chiang, F. Rist, J. Wallner, and H. Pottmann. 2023. Deployable Strip Structures. *ACM Transactions on Graphics* 42, 4 (Aug. 2023), 1–16. https://doi.org/10.1145/3592393
K. Liu, J. Wu, G. H. Paulino, and H. J. Qi. 2017. Programmable Deployment of Tensegrity Structures by Stimulus-Responsive Polymers. *Scientific Reports* 7, 1 (June 2017), 3511. https://doi.org/10.1038/s41598-017-03412-6
A. Micheletti and P. Podio-Guidugli. 2022. Seventy Years of Tensegrities (and Counting). *Archive of Applied Mechanics* 92, 9 (Sept. 2022), 2525–2548. https://doi.org/10.1007/s00419-022-02192-4
E. Miguel, M. Lepoutre, and B. Bickel. 2016. Computational Design of Stable Planar-Rod Structures. *ACM Transactions on Graphics* 35, 4 (July 2016), 1–11. https://doi.org/10.1145/2897824.2925978
P. Molina. 2015. Bent-Tensegrity Structures. https://issuu.com/pabloantunamolina/docs/article. Accessed: 2024-05-19.
W. Neveu, I. Puhachov, B. Thomaszewski, and M. Bessmeltsev. 2022. Stability-Aware Simplification of Curve Networks. In *ACM SIGGRAPH 2022 Conference Proceedings (SIGGRAPH '22)*. Association for Computing Machinery, New York, NY, USA, Article 20, 9 pages. https://doi.org/10.1145/3528233.3530711
K. Pajunen, P. Johanns, R. K. Pal, J. J. Rimoli, and C. Daraio. 2019. Design and Impact Response of 3D-printable Tensegrity-Inspired Structures. *Materials & Design* 182 (Nov. 2019), 107966. https://doi.org/10.1016/j.matdes.2019.107966
J. Panetta, M. Konaković-Luković, F. Isvoranu, E. Bouleau, and M. Pauly. 2019. X-Shells: A New Class of Deployable Beam Structures. *ACM Transactions on Graphics* 38, 4 (July 2019), 1–15. https://doi.org/10.1145/3306346.3323040
C. Paul, F. J. Valero-Cuevas, and H. Lipson. 2006. Design and Control of Tensegrity Robots for Locomotion. *IEEE Transactions on Robotics* 22, 5 (Oct. 2006), 944–957. https://doi.org/10.1109/TRO.2006.878980
J. Pérez, M. A. Otaduy, and B. Thomaszewski. 2017. Computational Design and Automated Fabrication of Kirchhoff-Plateau Surfaces. *ACM Transactions on Graphics* 36, 4 (Aug. 2017), 1–12. https://doi.org/10.1145/3072959.3073695
J. Pérez, B. Thomaszewski, S. Coros, B. Bickel, J. A. Canabal, R. Sumner, and M. A. Otaduy. 2015. Design and Fabrication of Flexible Rod Meshes. *ACM Transactions on Graphics* 34, 4 (July 2015), 1–12. https://doi.org/10.1145/2766998
N. Pietroni, M. Tarini, A. Vaxman, D. Panozzo, and P. Cignoni. 2017. Position-Based Tensegrity Design. *ACM Transactions on Graphics* 36, 6 (Dec. 2017), 1–14. https://doi.org/10.1145/3130800.3130809
S. Pillwein, K. Leimer, M. Birsak, and P. Musialski. 2020. On Elastic Geodesic Grids and Their Planar to Spatial Deployment. *ACM Transactions on Graphics* 39, 4 (July 2020). https://doi.org/10.1145/3386569.3392490 arXiv:2007.00201
Y. Ren, J. Panetta, T. Chen, F. Isvoranu, S. Poincloux, C. Brandt, A. Martin, and M. Pauly. 2021. 3D Weaving with Curved Ribbons. *ACM Transactions on Graphics* 40, 4 (July 2021), 127:1–127:15. https://doi.org/10.1145/3450626.3459788
T. Rhodes, C. Gotberg, and V. Vikas. 2019. Compact Shape Morphing Tensegrity Robots Capable of Locomotion. *Frontiers in Robotics and AI* 6 (Nov. 2019), 111. https://doi.org/10.3389/frobt.2019.00111
J. Rieffel and J.-B. Mouret. 2018. Adaptive and Resilient Soft Tensegrity Robots. *Soft Robotics* 5, 3 (June 2018), 318–329. https://doi.org/10.1089/soro.2017.0066
R. M. Sánchez-Banderas, A. Rodríguez, H. Barreiro, and M. A. Otaduy. 2020. Robust Eulerian-on-Lagrangian Rods. *ACM Transactions on Graphics* 39, 4 (July 2020). https://doi.org/10.1145/3386569.3392489
E. Schling, R. Barthel, A. Ihde, J. Tutsch, and S. Huth. 2015. Bending-Activated Tensegrity. In *Proceedings of IASS Annual Symposia*. International Association for Shell and Spatial Structures (IASS), 1–15.
M. Skouras, B. Thomaszewski, S. Coros, B. Bickel, and M. Gross. 2013. Computational Design of Actuated Deformable Characters. *ACM Transactions on Graphics* 32, 4 (July 2013), 1–10. https://doi.org/10.1145/2461912.2461979
E. Soriano, R. Sastre, and D. Boixader. 2019. G-shells:flat collapsible geodesic mechanisms for gridshells. In *Proceedings of IASS Annual Symposia*. International Association for Shell and Spatial Structures (IASS). http://hdl.handle.net/2117/330614
O. Sorkine-Hornung and M. Rabinovich. 2017. Least-Squares Rigid Motion Using SVD. (Jan. 2017), 5. https://igl.ethz.ch/projects/ARAP/svd_rot.pdf Accessed: 2024-05-19.
S. Spiegel, J. Sun, and J. Zhao. 2023. A Shape-Changing Wheeling and Jumping Robot Using Tensegrity Wheels and Bistable Mechanism. *IEEE/ASME Transactions on Mechatronics* 28, 4 (Aug. 2023), 2073–2082. https://doi.org/10.1109/TMECH.2023.3276933
P. Strzelecki and H. von der Mosel. 2017. Geometric Curvature Energies: Facts, Trends, and Open Problems. In *New Directions in Geometric and Applied Knot Theory*, Philipp Reiter, Simon Blatt, and Armin Schikorra (Eds.). De Gruyter Open, 8–35. https://doi.org/10.1515/9783110571493-002
T. Tachi. 2013. Interactive Freeform Design of Tensegrity. In *Advances in Architectural Geometry 2012*. Springer Vienna, Vienna, 259–268. https://doi.org/10.1007/978-3-7091-1251-9_21
M. Vidulis, Y. Ren, J. Panetta, E. Grinspun, and M. Pauly. 2023. Computational Exploration of Multistable Elastic Knots. *ACM Transactions on Graphics* 42, 4 (Aug. 2023), 1–16. https://doi.org/10.1145/3592399
C. Yu, H. Schumacher, and K. Crane. 2021. Repulsive Curves. *ACM Transactions on Graphics* 40, 2 (April 2021), 1–21. https://doi.org/10.1145/3439429
J. Y. Zhang and M. Ohsaki. 2006. Adaptive Force Density Method for Form-Finding Problem of Tensegrity Structures. *International Journal of Solids and Structures* 43, 18-19 (Sept. 2006), 5658–5673. https://doi.org/10.1016/j.ijsolstr.2005.10.011