Discrete Differential Geometry and Physics of Elastic Curves

Discrete Differential Geometry and Physics of Elastic Curves

A DISSERTATION PRESENTED
BY
ANDREW MCCORMICK
TO
THE DEPARTMENT OF PHYSICS

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY
IN THE SUBJECT OF
PHYSICS

HARVARD UNIVERSITY
CAMBRIDGE, MASSACHUSETTS

AUGUST 2013

Thesis advisor: L. Mahadevan Andrew McCormick

Discrete Differential Geometry and Physics of Elastic Curves

Abstract

We develop a general computational model for a elastic rod which allows for extension and shear. The model, similar in mathematical construction to Cosserat rod theory, allows a wider variety of problems to be studied than previous models. In the first section we develop the continuous mathematical model, discretize the system to allow implementation on a computer, and then verify the model's output against classical buckling tests. We then develop a novel analytic solution for the critical buckling length of a vertically oriented, shearable elastic beam subject to gravity and show that the model's treatment of shear is correct. In the experimental section we analyze a number of different phenomena with the rod model. To begin, we explain the mechanical response of helically coiling tendrils. After self-collision is introduced, we explore the formation of plectonemes and solenoids in a highly extensible elastic string. We discuss a sheet adhering to a surface in several different regimes and use the rod model to discover a self-similarity solution in the low-damping limit. Physical entanglement is investigated in an experiment where randomly tumbled strings are used to derive scaling laws for the dynamics governing entanglement. Models for active

internal forces and anisotropic surface friction are introduced to explain the mechanics of a newly observed mode of snake locomotion. Finally, we extend the model from a single filament to an arbitrary number of strings and begin exploration into behavior of cloth, ponytails, and combing hair.

Contents

Listing of figures

To my parents, Frank and Christine.

Acknowledgments

I'd like to thank my advisor, the members of the Mahadevan group, and everyone else who has helped me persevere and prosper in graduate school.

1 Introduction

From bridges to dna to shoelaces, the ubiquitous elastic rod (equivalently filament or string) plays a large role in everyday life. Hence, it is unsurprising that elastic strings have been the subject of serious study since at least the mid 18^{th} century, when mathematicians like Euler and Lagrange began deriving results like the equations of motion for a simple elastic string [2, 26]. In this initial section, I will first give a brief history of the analytic models used to study elastic rods, then discuss the development of computational models, and finally discuss the innovations my model makes and how they connect to past works.

The history of analytic models for elastic rods is rich and long. While research continues in the area [37, 39, 49, 81, 89, 90], the fundamental work that is appropriate to an introduction is many

decades old, so there is little novel that I could add to previous previous historical summaries, i.e. [13, 26]. As such, I will here only give a brief qualitative summary, and direct the reader to the above sources for more details.

The history of analytic models for elastic rods is well described by a highly constrained initial model continuously generalized and relaxed over time. Kirchoff [53] and Clebsch [18] were early pioneers, viewing a rod as an assembly of short segments loaded by contact forces from adjacent segments. The displacements within segments were assumed to be small and continuity was expressed via a redudant system of four coordinates. Kirchoff's end result was a system of partial differential equations for the local deformations, whereas Clebsch's later work was a complete reduction from a three dimensional body to a one dimensional rod. Love [62] took a more modern point of view, dropping the redundant coordinate system in favor of one where the deformed state of the rod is formed by small local deformations wherein the cross sections remain undistorted and normal to the centerline. The equations of motion are then derived by essentially understanding how these cross sections rotate and twist as we move along the rod and through time. As with Kirchoff and Clebsch before him, Love began with a straight rod and generalized his equations to a arbirary curved rod. Cosserat [19] relaxed Love's assumption of the normality of the cross section, generalizing to a local-frame system that can handle shear, a feature previous models lacked.

On the computational side of things, the initial models [54, 91] were similar to general finite difference models, and in some sense are similar the the analytic models of Kirchoff and Clebsch. Somewhat analagous to Love's new coordinate system, a focus on the Bishop frame then produced models with more simple analytic and numerical equations [57, 71]. Further refinements and simplifications [7, 8] continued increasing the speed and simplicity of the models, though

the models still did not capture some phenomena, such as shear. My model generalizes the past ones, implementing a full simple Cosserat model where we keep track of both the centerline and a local frame along the rod. This formulation naturally allows shear and extension, while still having a simple Hamiltonian and fast, simple equations of motion.

1.1 Research Experiments

After defining a continuous mathematical model, discretizing it, and testing the computational version in 2, we move on to using the model to peform various experiments. We must often introduce new capabilities to the model, such as collision or friction, to handle the specific nature of each experiment. These additions are described at the beginnings of the relevant chapters. Some of the later experiments rely on a number of additions, leading to a natural ordering of the experiments where the simpler ones come first, introducing features that are used throughout the remainder.

For our first experiment, we consider the helical coiling of cucumber tendrils 3. The main story here is that the cucumber plants send up tendrils which wave around until they find a support, at which point they wrap themselves around it and latch on. They convert the previously straight tendril to a helical shape, which serves to pull the plant up as the support is fixed in space and the arclength of the tendril must remain unchanged. In our research we describe the biological structures within the tendril that allow this helical structure to form and use the rod model to help understand the mechanical properties of the resulting structure. For this problem, neither extension nor shear was relevant. However, it was necessary to be able to easily modify the internal material properties of the tendril and quickly find the resulting behavior, a task for which our rod

model was ideal.

Most, if not all, analytic methods in elasticity ignore the problem of self-intersection as the local nature of the phenomena makes analytic solutions nearly impossible. We develop a novel high-speed algorithm for determining whether two cylinders are intersecting and use the algorithm to develop robust collision models that prevent-self intersection 4.1. We use this newfound ability to study the formation of plectonemes, solenoids, and other complicated structures which can arise when an elastic string has a weight hung from it and then rotated. We find that our model gives results quite similar to those from an earlier physical experiment. We also use a slightly modified version of the experiment to explain the formation of solenoid structures in terms of twist, link, and writhe.

In another experiment, we examined the knotting dynamics of a randomly tumbling string using a proxy of entanglement called the average crossing number 5. In this experiment we again make extensive use of the collision model, but also make use of the model's ability to easily include other physical objects, such as a box through which the rope cannot pass (which we shake about). We find that dynamics of the problem are well explained by a scaling law which uses an energy balance argument to introduce a new energetic length scale and approximately reduce the problem to one of pure geometry.

A model for short-range adhesion is then introduced 6. Again, this effect is inherently local and hence is analytically intractable in the general case except for certain simple cases. We make use of this model to study the behavior of a sheet adhering to a surface in a number of limits. In various regimes, we analytically find scaling laws and self-similarity solutions and confirm computationally that they are correct.

Next we introduce a model for active forces and torques within the rod, as well as a friction model that allows for kinetic and static

friction behavior, spatial anisotropy, and rolling behavior 7.1. These two additions are then used to describe the mechanics behind a newly observed form of snake locomotion which we term trilateral locomotion. Finally, we generalize the model to allow multiple filaments at once 7.2. We give examples of practical problems which could be studied using the generalized model, such as cloth, ponytails, and combing hair. Some interesting phenomena are observed, such as jamming even in the absence of friction.

2

The Model

2.1 Definitions

We will denote vectors as bold-faced lower case letters ($\boldsymbol{x}$) and matrices as bold-faced upper case letters ($\boldsymbol{X}$). Operators will be denoted with upper case script letters ($\mathcal{X}$). Constants are given as plain face lower or upper case letters (a, or A). Derivatives are expressed as subscripts ($\boldsymbol{x}_t$) if there is no chance of confusion, and in tensor notation ($\frac{df^i}{dx^j}$) if necessary.

2.2 Geometry

We begin our description of the rod model by discussing how we quantify the rod's geometry, or shape. At each point along the rod we

consider both the point's position in space and the rod's local orientation, or local frame, both of which are also functions of time. This description of the rod is mathematically similar to that found in Cosserat rod theory [19, 43–45, 83] Mathematically, we define the centerline $\boldsymbol{x}(s,t)$ and the frame $\boldsymbol{Q}(s,t)$ as two mappings:

$$\begin{aligned} \boldsymbol{x} &: [0,L] \otimes \mathbb{R} \to \mathbb{R}^3 \\ \boldsymbol{Q} &: [0,L] \otimes \mathbb{R} \to SO(3) \end{aligned} \tag{2.1}$$

where the first index of each mapping gives the material coordinate (arclength for an unstretched rod) and the second index is the time. R^3 is just regular 3-dimensional space and $SO(3)$ is the standard rotation group.

Any element $\boldsymbol{Q} \in SO(3)$ can be written as an orthonormal 3 by 3 matrix of real numbers. If we think about each row vector as describing a particular axis, then the three rows together form a basis for R^3. Thus, we have a well defined local body frame at each point along the rod. A vector $\boldsymbol{a}^{lab}$ whose coordinates are expressed in the lab frame can be transformed into the body frame at material coordinate s and time t via the following relation:

$$\boldsymbol{a}^{body} = \boldsymbol{Q}(s,t)(\boldsymbol{a}^{lab} - \boldsymbol{x}(s,t)) \tag{2.2}$$

This transformation will prove very useful as we develop the model because some quanties, such as forces, are best considered in the lab frame, while others, such as torques, are naturally described in the local body frames.

We adopt row vector notation for $\boldsymbol{Q}$, writing

$$\boldsymbol{Q}(s,t) = \begin{pmatrix} \boldsymbol{d}_1(s,t) \\ \boldsymbol{d}_2(s,t) \\ \boldsymbol{d}_3(s,t) \end{pmatrix} \tag{2.3}$$

where, if the rod is unsheared, $\boldsymbol{d}_3$ points along the tangent vector of the centerline and $\boldsymbol{d}_1$ and $\boldsymbol{d}_2$ span the normal-binormal centerline plane. If the rod is sheared, then $\boldsymbol{d}_3$ may no longer coincide precisely with then tangent to the centerline x_s . We quantify this notion by defining the shear vector $\sigma : \ [0, L] \otimes \mathbb{R} \to \mathbb{R}^3$ as

$$\boldsymbol{\sigma} = \boldsymbol{Q}\boldsymbol{x}_s - \boldsymbol{d}_3 \tag{2.4}$$

in the local frame. Note that if $\boldsymbol{d}_3 = \boldsymbol{x}_s$ then $\boldsymbol{Q}\boldsymbol{x}_s = \boldsymbol{d}_3$ and hence $\boldsymbol{\sigma} = \boldsymbol{0}$ as expected.

2.3 Topology and the CFW Theorem

After describing the shape of our rod, a natural precursor to dicussing the rod's energetics is to discuss its topology. Although our model does not directly rely upon it, an extremely important theorem, the Călugăreanu-Fuller-White (CFW) theorem, explains the relationship between the fundamental topological quantities of our system. These quantities, Link, Writhe, and Twist, are largely analagous to their intuitive meaning, but there are some surprises so it is worthwhile to define them carefully.

2.3.1 Link

Imagine two curves in space. They might be completely separate, or they might be intertwined in some complicated way such that they

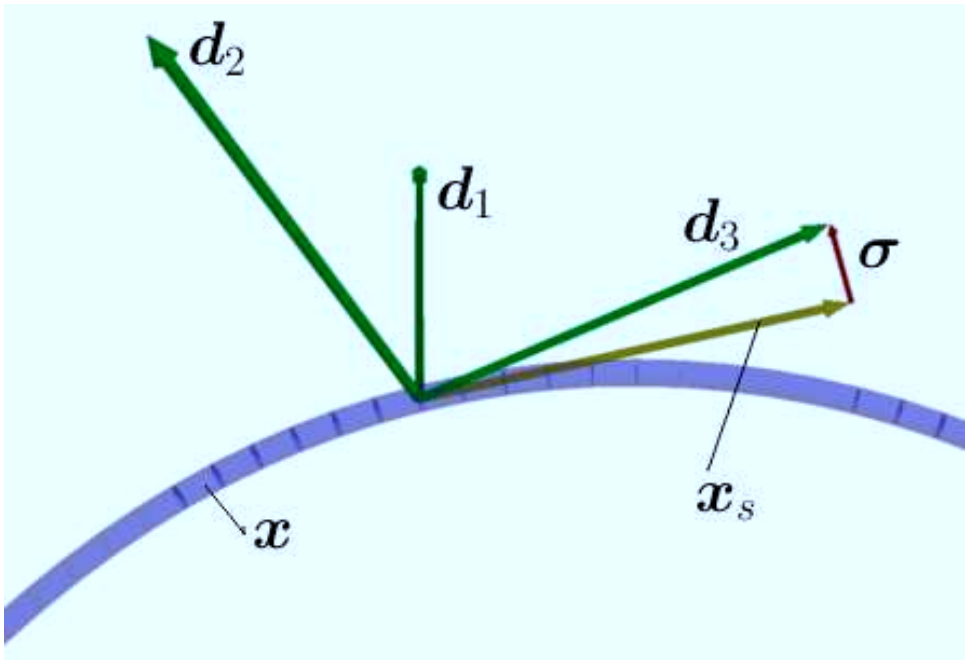

Figure 2.2.1: The rod is described by a centerline $\bar{\boldsymbol{x}}(s,t) \in \mathbb{R}^3$ and a local frame described by $\boldsymbol{Q}(s,t) = \{\boldsymbol{d}_i\} \in SO(3)$. If the rod is sheared then the 3rd director $\boldsymbol{d}_3$ of the local frame will no longer coincide with the tangent to the centerline, a notion we quantify by defining the shear vector $\boldsymbol{\sigma} = \boldsymbol{Q}\bar{\boldsymbol{x}}_s - \boldsymbol{d}_3$ in the local frame.

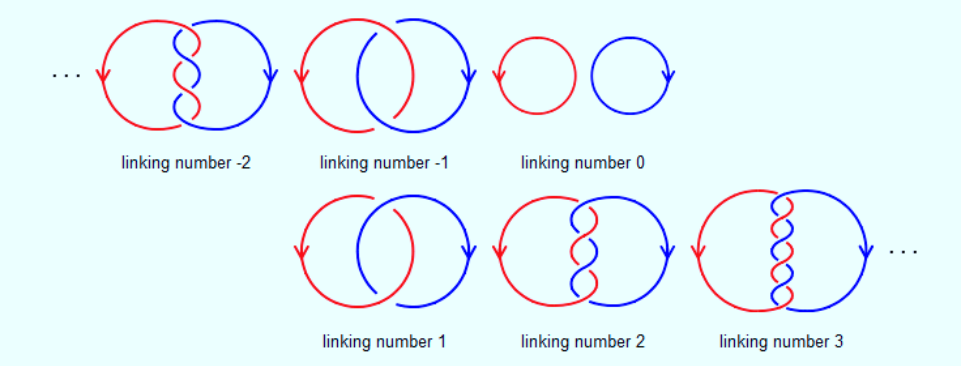

Figure 2.3.1: Two closed loops in space have a well-defined linking number that indicates how topologically connected they are. Note from the above examples [88] that the sign of the linking number depends upon the orientation of the curves.

are inseparable. To formalize this notion, we introduce the concept of a linking number of two curves. The linking number is a topological invariant – any pair of curves with the same linking number as another pair of curves can smoothly be transformed into the other pair without any crossings [46]. Some examples of pairs of curves with different linking numbers are given in figure 2.3.1.

If we are given a link diagram [75] (which is a 2-d projection of the curves onto a given plane) of two curves, such as the examples in figure 2.3.1, we can compute the linking number by first finding all of the intersections and labeling them positive or negative according to the following rule:

Subtracting the number of negative crossings from the number of positive crossings and then dividing the result by two gives the linking number.

If instead we are given two arbitrary non-intersecting differentiable curves $\mathbf{x}_1, \mathbf{x}_2 : S^1 \to \mathbb{R}^3$, we can calculate the linking number (Lk) via

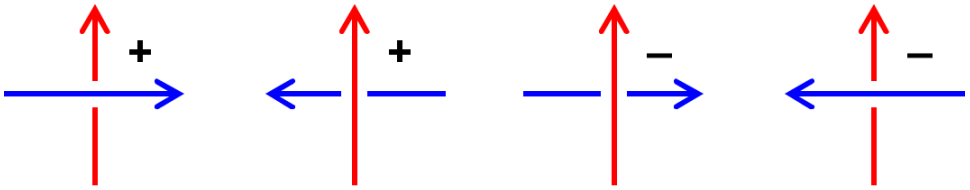

Figure 2.3.2: The linking number can be computed from a link diagram by assigning values to each crossings as shown above [88]. Alternatively it can be computed from the gauss linking integral 2.5.

the Gauss linking integral:

$$Lk = \frac{1}{4\pi} \oint_{\mathbf{x}_1} \oint_{\mathbf{x}_2} \frac{\mathbf{r}_1 - \mathbf{r}_2}{|\mathbf{r}_1 - \mathbf{r}_2|^3} \cdot (d\mathbf{r}_1 \times d\mathbf{r}_2) \tag{2.5}$$

This formula can be derived by an argument based on the signed counting of numbers as described above [72, 77].

2.3.2 Twist

Twist, denoted Tw, is simply the integrated local torsion (e.g. from the Frenet-Serret description):

$$Tw = \int_0^L \tau(s) ds \tag{2.6}$$

2.3.3 Writhe

Writhe is somewhat more difficult to understand then twist. Informally, writhe number measures how coiled or folded a curve is – hence the name writhe! One common way people intuitively think about writhe is as a measure of how much a curve deviates from being a plane curve, i.e. how "3-d" it is in a loose sense. For example, a helix with a large radius compared to its pitch cannot be projected into any 2-d plane without quashing important features; hence it is

highly "3-d" and has a large writhe. Conversely, a helix with a very small radius compared to its pitch has a small writhe, which we can intuitively understand because it nearly lies in a single dimension, therefore not being very "3-d". Coming up with an explicit formula for the writhe of a curve is difficult; indeed, we so far have not even formally defined it. The next section will clarify the issue.

2.3.4 The CFW Theorem

In a series of papers Calugreanu [15], White [87] and Fuller [30] developed the notions of link (Lk), writhe (Wr) and twist (Tw), and proved an important relationship between them:

$$Lk = Tw + Wr \tag{2.7}$$

which is now known as the CFW theorem. We can approximately derive this theorem by looking back to Gauss's definition of the linking number 2.5. Instead of two completely arbitrary curves in space, suppose our two curves are given by $\mathbf{r}_1(t) = \mathbf{x}(t)$ and $\mathbf{r}_2(t) = \mathbf{x}(t) + \epsilon\mathbf{u}(t)$, where $\epsilon \in \mathbb{R}$ and $\mathbf{x}, \mathbf{u} : \mathbb{R} \to \mathbb{R}^3$. If we take $\mathbf{u}$ to be, for example, the normal vector to $\mathbf{x}(t)$ then $\mathbf{r}_1$ and $\mathbf{r}_2$ are closed and everywhere disjoint, so we can apply the Gauss linking integral. If we then let ϵ go to zero we find that a (removable) sigularity arises, whereupon the integral can be split into two pieces. One piece, which we call the writhe, still involves a double integral, but this time over the *same* spacecurve:

$$\begin{aligned} Wr &= \frac{1}{4\pi} \oint_{\mathbf{x}} \oint_{\mathbf{x}} \frac{\mathbf{r}_1 - \mathbf{r}_2}{|\mathbf{r}_1 - \mathbf{r}_2|^3} \cdot (d\mathbf{r}_1 \times d\mathbf{r}_2) \\ &= \frac{1}{4\pi} \int_0^L \int_0^L (\mathbf{T}(s) \times \mathbf{T}(t)) \cdot \frac{\mathbf{x}(s) - \mathbf{x}(t)}{|\mathbf{x}(s) - \mathbf{x}(t)|^3} ds dt \end{aligned} \tag{2.8}$$

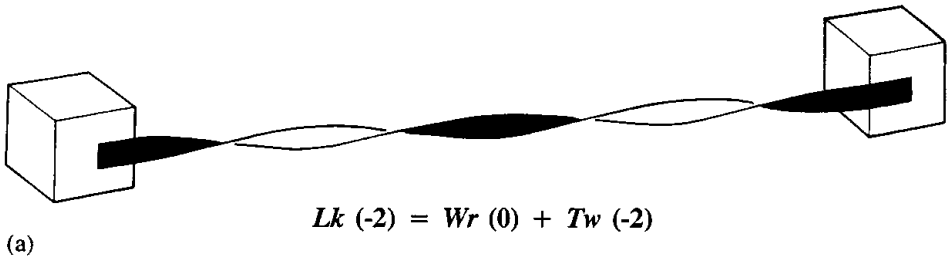

(a) $Lk\ (-2) = Wr\ (0) + Tw\ (-2)$

(b) $Lk\ (-2) = Wr\ (-2) + Tw\ (0)$

Figure 2.3.3: In the figure above we hold the ends of a ribbon twisted twice (a) fixed and move them ends together. We see that the two twists has been converted to two loops (maintaining the total link of 2) which is the essence of the CFW theorem 2.7 [84].

The other piece, which we call the twist, is precisely the same Tw as defined in 2.3.2.

When the centerline moves, the writhe will generally change; since the link is an invariant, the conservation law implied by the CFW theorem says that hence the twist must also change. This establishes that the centerline and the frame are not independent! A nice example of how twist and writhe can be converted is shown in figure 2.3.3, where an initially straight, twisted rubber band is converted into a helical, untwisted rubber shape.

In the case where τ is a constant, as Bergou assumed in [7], the change in τ as we vary the centerline is simply calculated as the total change in writhe (times negative one) divided by the length of the

rod. This calculation in the change of writhe is made fairly simple due to an equation by Fuller [31] which states that, excepting some pathological cases, the change in writhe between a curve with tangent $\mathbf{T}_1$ and a nearby curve with tangent $\mathbf{T}_2$ is given by

$$Wr(\mathbf{T}_2) - Wr(\mathbf{T}_1) = \frac{1}{2\pi} \int_0^L \frac{\mathbf{T}_1 \times \mathbf{T_2} \cdot \frac{d}{dt}(\mathbf{T}_1 + \mathbf{T}_2)}{1 + \mathbf{T}_1 \cdot \mathbf{T}_2} ds \tag{2.9}$$

Note that the above equation involves only a single integral, rather than the double integral of the direct Wr calculation. In actual practice, Bergou used the discrete version of this integral given by de Vries [24]. Other methods for calculating the writhe and link exist, see [36, 55]. I include these formulae because we can use them to help explain some of the phenomena we will see later on, e.g. in 4.

One potentially worrysome fact is that the CFW theorem technically only applies to closed rods. While there have been a number of attempts to extend the theorem to open rods [6], they all seem to be lacking a certain generality and the matter is far from clear. As such, it is a pleasant feature of our model that we do not need to directly rely on the CFW theorem even if make use of it to explain some of our experimental results.

2.4 Math Background

Before we move on to describing the rod's energies and the resulting equations of motions, we need to review some mathematical concepts which are not always studied in a typical undergrad and/or graduate program. The concepts herein are described with minimal formalism; we're interested in how to use them, not derive them!

2.4.1 Derivatives and the Lie Algebra

To derive the equations of motion for the system it will be necessary to take derivatives with respect to the two independent quantities, namely $\boldsymbol{x}$ and $\boldsymbol{Q}$. Derivatives w.r.t. $\boldsymbol{x}$ are the standard fare one sees in introductory calculus. Derivatives w.r.t. $\boldsymbol{Q}$ are considerably more complicated because $SO(3)$ has a more complex group structure than R^3. Understanding the meaning of a derivative of $\boldsymbol{Q} \in SO(3)$ begins by recalling that orthonormal matrices satisfy

$$\boldsymbol{Q}\boldsymbol{Q}^T = \boldsymbol{I} \tag{2.10}$$

Now take the derivative of each side of the above equation. The right hand side is then zero, and after applying the chain rule we have that

$$\begin{aligned} \boldsymbol{0} &= (\boldsymbol{Q}\boldsymbol{Q}^T)_s \\ &= \boldsymbol{Q}_s\boldsymbol{Q}^T + \boldsymbol{Q}\boldsymbol{Q}_s^T \\ &\stackrel{\text{def}}{=} \boldsymbol{K} + \boldsymbol{K}^T \\ \Rightarrow \boldsymbol{K} &= -\boldsymbol{K}^T \end{aligned} \tag{2.11}$$

where we have defined $\boldsymbol{K} = \boldsymbol{Q}_s\boldsymbol{Q}^T$ and noted that $\boldsymbol{K}$ is skew-symmetric. Using $\boldsymbol{K}$ we can write the derivative of $\boldsymbol{Q}$ as a skew-symmetric matrix times $\boldsymbol{Q}$ itself:

$$\boldsymbol{Q}_s = \boldsymbol{K}\boldsymbol{Q} \tag{2.12}$$

Formally we can say that each skew symmetric $\boldsymbol{K}$ belongs to the Lie Algebra $so(3)$ associated with $SO(3)$ [47, 78]. A Lie Algebra describes the group structure of the derivatives of another group. The fact that $so(3) \neq SO(3)$ is a consequence of the fact that $SO(3)$ is not a flat manifold, which follows because $SO(3)$ is diffeomorphic to S_3,

the 4-dimensional sphere. Conversely, R^3 is flat, and hence the Lie Algebra associated with R^3 is R^3 itself; this is why derivatives in first-year calculus are easy to understand.

2.4.2 The Skew Map

A useful manipulation is to note that a skew-symmetric 3 by 3 matrix has only 3 independent variables, and hence can be associated with an element of R^3. This association is defined by the skew-symmetric mapping $\mathcal{S} : \mathbb{R}^3 \to \mathbb{R}^{3\times 3}$

$$\mathcal{S}(x, y, z) = \begin{pmatrix} 0, & z, & -y \\ -z, & 0, & x \\ y, & -x, & 0 \end{pmatrix} \tag{2.13}$$

This mapping allows us to think about the derivatives of $\boldsymbol{Q}$ as vectors in the local body frame, and it is usually easier to think about changes in local frames as rotations about given axes rather than as matrices. An additional point to mention is that the analagous operation to two skew-symmetric matrices being multiplied is that of taking the cross product between corresponding vectors. More precisely, for any two skew symmetric matrices $\boldsymbol{A}$ and $\boldsymbol{B}$

$$\mathcal{S}^{-1}(\boldsymbol{AB}) = \mathcal{S}^{-1}(\boldsymbol{A}) \times \mathcal{S}^{-1}(\boldsymbol{B}) \tag{2.14}$$

2.4.3 Logarithms and Exponentials

The relation between a Lie Algebra $\mathfrak{g}$ and its associated group G is given by the exponential map, which says that

$$\forall g \in \mathfrak{g},\ e^g \in G \tag{2.15}$$

In the context of $G = SO(3)$, we have that the exponential of any skew-symmetric matrix is an orthogonal matrix. Using the definition

of a matrix exponential, which says that for every square matrix $\boldsymbol{M}$

$$e^{\boldsymbol{M}} = \boldsymbol{I} + \boldsymbol{M} + \frac{\boldsymbol{M} * \boldsymbol{M}}{2!} + \frac{\boldsymbol{M} * \boldsymbol{M} * \boldsymbol{M}}{3!} + \cdots \tag{2.16}$$

we can check that this is indeed true. The inverse mapping, i.e. the matrix logarithm $\ln \boldsymbol{A}$, can be computed by finding the eigenvalues matrix $\boldsymbol{V}$ of $\boldsymbol{A}$, computing the diagonal matrix $\boldsymbol{A}^{'} = \boldsymbol{V}^{-1}\boldsymbol{A}\boldsymbol{V}$, and finally

$$\ln \boldsymbol{A} = \boldsymbol{V}(\ln \boldsymbol{A}^{'})\boldsymbol{V}^{-1} \tag{2.17}$$

where the logarithm in the above equation in applied element-wise to the diagonal only.

Although the above expressions for the matrix exponential and log are always correct, they can be cumbersome. For $SO(3)$ there are faster ways to compute them: the so-called Rodriguez rotation formulas. Suppose we want to compute $e^{\boldsymbol{k}}$ for $\boldsymbol{k} \in so(3)$. We can view this as a rotation of angle $\theta = ||\boldsymbol{k}||$ about axis $\hat{k}$. Then we have that

$$e^{\boldsymbol{k}} = \mathbf{I} + \sin(\theta)\mathcal{S}(\hat{\boldsymbol{k}}) + (1 - \cos(\theta))(\mathcal{S}(\hat{\boldsymbol{k}})\mathcal{S}(\hat{\boldsymbol{k}})) \tag{2.18}$$

where $\boldsymbol{I}$ is the identity matrix. For the matrix logarithm of $\boldsymbol{R} \in SO(3)$, we compute the rotation angle as $\theta = \arccos(\frac{tr\boldsymbol{R}-1}{2})$ and then

$$\ln \boldsymbol{R} = \begin{cases} 0, & \theta = 0 \\ \frac{1}{2\sin(\theta)}(\boldsymbol{R} - \boldsymbol{R}^T), & \theta \neq 0 \end{cases} \tag{2.19}$$

There are many good references on geometry, frames, and rotational motion; see [4, 40, 59].

2.5 Energies

With the geometry and topology of our rod well defined and the proper mathematical tools in hand, we can move on to defining the energies of our rod. We assume that all energies are quadratic with respect to a particular strain. By strain we mean a derivative of $\boldsymbol{x}$ or $\boldsymbol{Q}$ in time or space (material coordinate). The strains in time give rise to rotational and translational energy, while the strains in the material coordinate rise to bending and shearing energy. These energies are all expressible as an integral of a quadratic function of the corresponding pointwise strain along the extent of the rod. We can express a general strain as $\boldsymbol{\epsilon}$ and the quadratic function as $\boldsymbol{\epsilon}^T \boldsymbol{G} \boldsymbol{\epsilon}$ where $\boldsymbol{G}$ is any symmetric, positive definite 3 by 3 matrix. Hence, we have

$$E = \int_0^L \boldsymbol{\epsilon}^T \boldsymbol{G} \boldsymbol{\epsilon} \, ds \tag{2.20}$$

as the general template for our energy functions.

Physics enters the problem when we decide which strains have energies associated with them. There may be more than one sensible way to assign an energy to a physical property, such as bending. For example, whereas a popular recent model [7] associated bending energy with the arclength strain of the centerline curve, in this model bending energy will be associated only with strains in the local body frames. (As an aside, Bergou's model in fact did not contain a local body frame as a free variable, so capturing bending energy with centerline strains was required.)

We choose to define four intrinsic energies:

$$\begin{aligned} V^{rot} &= \frac{1}{2}\int_0^L \boldsymbol{k}^T \boldsymbol{B}\boldsymbol{k}\ ds \\ V^{trans} &= \frac{1}{2}\int_0^L \boldsymbol{\sigma}^T \boldsymbol{S}\boldsymbol{\sigma}\ ds \\ T^{rot} &= \frac{1}{2}\int_0^L \boldsymbol{w}^T \boldsymbol{I}\boldsymbol{w}\ ds \\ T^{trans} &= \frac{1}{2}\int_0^L \rho \boldsymbol{v}^T \boldsymbol{v}\ ds \end{aligned} \tag{2.21}$$

where

- $\boldsymbol{k}$ is the curvature vector given by $\mathcal{S}^{-1}(\boldsymbol{K})$, where $\boldsymbol{K} = \boldsymbol{Q}_s\boldsymbol{Q}^T$
- $\boldsymbol{B}$ is the bending energy matrix
- $\boldsymbol{\sigma}$ is the shear vector given by $\boldsymbol{\sigma} = \boldsymbol{Q}\boldsymbol{x}_s - \boldsymbol{d}_3$
- $\boldsymbol{S}$ is the shearing energy matrix
- $\boldsymbol{w}$ is the rotational velocity vector given by $\mathcal{S}^{-1}(\boldsymbol{W})$, where $\boldsymbol{W} = \boldsymbol{Q}_t\boldsymbol{Q}^T$
- ρ is the mass per unit length of the rod
- $\boldsymbol{v}$ is the velocity of the centerline

The meaning of each component of $\boldsymbol{B}$ relies on the representation for $\boldsymbol{Q}$ we defined in (2.3). Namely, the $B_{1,1}$ component gives the bending stiffness in one principle direction, $B_{2,2}$ gives the bending stiffness in the other principal direction, and $B_{3,3}$ gives the twisting stiffness. Off-diagonal elements are possible, but are not found in uniform elastica. The components of $\boldsymbol{S}$ and $\boldsymbol{I}$ have analgously direction-dependent meanings. Note that in our formulation we have

a natural definition of curvature that is unambiguous, whereas when the curvature is defined based off the centerline there are a number of different possible definitions [20].

The energies in (2.21) hold for a naturally uncurved, unsheared rod whose material properties remain constant along the rod. We can generalize to the case where the rod has natural curvature $\boldsymbol{k}^0$, natural shear $\boldsymbol{\sigma}^0$, and material-coordindate dependent local properties by allowing $\boldsymbol{B}$, $\boldsymbol{S}$, and $\boldsymbol{I}$ to become functions.

$$\begin{aligned} V^{rot} &= \frac{1}{2}\int_0^L (\boldsymbol{k}-\boldsymbol{k}^0)^T \boldsymbol{B}(s)(\boldsymbol{k}-\boldsymbol{k}^0)\ ds \\ V^{trans} &= \frac{1}{2}\int_0^L (\boldsymbol{\sigma}-\boldsymbol{\sigma}^0)^T \boldsymbol{S}(s)(\boldsymbol{\sigma}-\boldsymbol{\sigma}^0)\ ds \\ T^{rot} &= \frac{1}{2}\int_0^L \boldsymbol{w}^T \boldsymbol{I}(s)\boldsymbol{w}\ ds \\ T^{trans} &= \frac{1}{2}\int_0^L \rho(s)\boldsymbol{v}^T\boldsymbol{v}\ ds \end{aligned} \tag{2.22}$$

Note that the rotational and translatational velocities have no possible "intrinsic" values, as that would not have any physical meaning. For the Hamilitionian evolution we will perform later, it is useful to re-write the kinetic energy terms using local translational momentum $\boldsymbol{p} = \rho\boldsymbol{v}$ and rotational momentum $\boldsymbol{l} = \boldsymbol{I}\boldsymbol{w}$

$$\begin{aligned} T^{rot} &= \frac{1}{2}\int_0^L \boldsymbol{l}^T\boldsymbol{I}^{-1}(s)\boldsymbol{l}\ ds \\ T^{trans} &= \frac{1}{2}\int_0^L \frac{\boldsymbol{p}^T\boldsymbol{p}}{\rho(s)}\ ds \end{aligned} \tag{2.23}$$

2.6 Forces and Torques

To make use of the energies defined above in our equations of motion, it is necessary to take the proper functional derivatives [21, 33]. The kinetic energy derivatives are straightforward: consider a small variation $\boldsymbol{\eta}$ to the rotational momenta $\boldsymbol{l}$ and compute

$$\begin{aligned} V^{rot}[\boldsymbol{l}+\boldsymbol{\eta}] &= \frac{1}{2}\int_0^L (\boldsymbol{l}^T\boldsymbol{I}^{-1}\boldsymbol{l} + \boldsymbol{l}^T\boldsymbol{I}^{-1}\boldsymbol{\eta} + \boldsymbol{\eta}^T\boldsymbol{I}^{-1}\boldsymbol{l} + +\boldsymbol{\eta}^T\boldsymbol{I}^{-1}\boldsymbol{\eta})\ ds \\ &\downarrow \\ V^{rot}[\boldsymbol{l}+\boldsymbol{\eta}] - V^{rot}[\boldsymbol{l}] &= \int_0^L (\boldsymbol{l}^T\boldsymbol{I}^{-1}\boldsymbol{\eta} + \mathcal{O}(\eta^2))\ ds \end{aligned} \tag{2.24}$$

where the material implication follows from the fact that $\boldsymbol{I}$ is symmetric, giving us that $\boldsymbol{l}^T\boldsymbol{I}^{-1}\boldsymbol{\eta} = \boldsymbol{\eta}^T\boldsymbol{I}^{-1}\boldsymbol{l}$. Similarly, if $\boldsymbol{\eta}$ now represents a small variation in the translational momenta $\boldsymbol{p}$

$$T^{trans}[\boldsymbol{p}+\boldsymbol{\eta}] - T^{trans}[\boldsymbol{p}] = \int_0^L (\frac{\boldsymbol{p}}{\rho}) + \mathcal{O}(\eta^2))\ ds \tag{2.25}$$

The above two equations don't immediately look that simple, but as we'll see later when we discuss Hamiltonian evolution that they essentially say that $\boldsymbol{x}_t = \boldsymbol{v}$ and $\boldsymbol{Q}_t = e^{\boldsymbol{w}}\boldsymbol{Q}$, precisely as we would expect.

The functional derivatives of the space-like strain energies are somewhat more complicated. In particular, when we consider a variation with respect to $\boldsymbol{Q}$ we cannot simply add a quantity $\boldsymbol{\eta}$ to $\boldsymbol{Q}$ analguously to before, because the new quantity $\boldsymbol{\eta} + \boldsymbol{Q}$ would not belong to $SO(3)$. Instead we need to use the exponential map and consider the variation in the form $e^{\boldsymbol{\eta}}\boldsymbol{Q}$, which we know belongs to $SO(3)$ because groups are closed under their operation. We can now

compute what the new curvature vector $\boldsymbol{k}$ will be after varying $\boldsymbol{Q}$ by $\boldsymbol{\eta}$

$$\begin{aligned}(e^{\boldsymbol{\eta}}\boldsymbol{Q})_s &= e^{\boldsymbol{\eta}}\boldsymbol{Q}_s + \boldsymbol{\eta}_s e^{\boldsymbol{\eta}}\boldsymbol{Q} \\ &= (1+\boldsymbol{\eta}+\cdots)\boldsymbol{k}\boldsymbol{Q} + \boldsymbol{\eta}_s((1+\boldsymbol{\eta}+\cdots)\boldsymbol{Q} \\ &= (\boldsymbol{k}+\boldsymbol{\eta}\times\boldsymbol{k}+\boldsymbol{\eta}_s+\mathcal{O}(\boldsymbol{\eta}^2))\boldsymbol{Q}\end{aligned} \tag{2.26}$$

With this result in hand we can now take the functional derivative of the bending energy

$$\begin{aligned}V^{rot}[e^{\boldsymbol{\eta}}\boldsymbol{Q}] - V^{rot}[\boldsymbol{Q}] &= \int_0^L ((\boldsymbol{k}-\boldsymbol{k}^0)^T\boldsymbol{B}\boldsymbol{\eta}_s + (\boldsymbol{k}-\boldsymbol{k}^0)^T\boldsymbol{B}(\boldsymbol{\eta}\times(\boldsymbol{k}-\boldsymbol{k}^0)) + \mathcal{O}(\boldsymbol{\eta}^2))\, ds \\ &= \int_0^L ((-((\boldsymbol{k}-\boldsymbol{k}^0)^T\boldsymbol{B})_s + (\boldsymbol{k}-\boldsymbol{k}^0)\times((\boldsymbol{k}-\boldsymbol{k}^0)^T\boldsymbol{B}))\cdot\boldsymbol{\eta} + \mathcal{O}(\boldsymbol{\eta}^2))\, ds\end{aligned} \tag{2.27}$$

Note that if we had naively taken the variation in $\boldsymbol{k}$ directly, that is simply taken $\boldsymbol{k}\to\boldsymbol{k}+\boldsymbol{\eta}$ we would not have picked up the second term in the final line above. We note this derivative is linear in η, so we can sensibly define the *pointwise* bending torque $\boldsymbol{\tau}$ as

$$\boldsymbol{\tau}^{rot}(s) = (-((\boldsymbol{k}-\boldsymbol{k}^0)^T\boldsymbol{B})_s + (\boldsymbol{k}-\boldsymbol{k}^0)\times((\boldsymbol{k}-\boldsymbol{k}^0)^T\boldsymbol{B}))(s) \tag{2.28}$$

Whereas the previous energies have been functionals of exactly one of $\boldsymbol{x}$ or $\boldsymbol{Q}$, the shear energy is a functional of both. Hence we'll need to take two functional derivatives. For the variation with respect to $\boldsymbol{Q}$

we have $\boldsymbol{\sigma}[e^{\boldsymbol{\eta}}\boldsymbol{Q}, \boldsymbol{x}] - \boldsymbol{\sigma}[\boldsymbol{Q}, \boldsymbol{x}] = \boldsymbol{\eta} \times \boldsymbol{Q}\boldsymbol{x}_s$ and thus

$$\begin{aligned} V^{trans}[e^{\boldsymbol{\eta}}\boldsymbol{Q}, \boldsymbol{x}] - V^{trans}[\boldsymbol{Q}, \boldsymbol{x}] &= \int_0^L \boldsymbol{\sigma}^T \boldsymbol{S}(\boldsymbol{\eta} \times (\boldsymbol{Q}\boldsymbol{x}_s))\ ds \\ &= \int_0^L (\boldsymbol{\sigma}^T \boldsymbol{S} \times \boldsymbol{Q}\boldsymbol{x}_s) \cdot \boldsymbol{\eta}\ ds \end{aligned} \tag{2.29}$$

For the variation with respect to $\boldsymbol{x}$ we have
$\boldsymbol{\sigma}[\boldsymbol{Q}, \boldsymbol{x} + \boldsymbol{\eta}] - \boldsymbol{\sigma}[\boldsymbol{Q}, \boldsymbol{x}] = \boldsymbol{Q}\boldsymbol{\eta}_s$, so we can compute

$$\begin{aligned} V^{trans}[\boldsymbol{Q}, \boldsymbol{x} + \boldsymbol{\eta}] - V^{trans}[\boldsymbol{Q}, \boldsymbol{x}] &= \int_0^L \boldsymbol{\sigma}^T \boldsymbol{S}\boldsymbol{Q}\boldsymbol{\eta}_s\ ds \\ &= \int_0^L -(\boldsymbol{\sigma}^T \boldsymbol{S}\boldsymbol{Q})_s) \cdot \boldsymbol{\eta}\ ds \end{aligned} \tag{2.30}$$

Again these two derivatives are linear in η (as they must be!) so we define pointwise shear torques and forces as

$$\begin{aligned} \boldsymbol{\tau}^{trans}(s) &= (\boldsymbol{\sigma}^T \boldsymbol{S} \times \boldsymbol{Q}\boldsymbol{x}_s)(s) \\ \boldsymbol{f}^{trans}(s) &= -((\boldsymbol{\sigma}^T \boldsymbol{S}\boldsymbol{Q})_s)(s) \end{aligned} \tag{2.31}$$

2.7 Time Evolution

We evolve the sytem through time using the Hamiltonian formulation [1], where we consider $\mathbf{x}$ and $\mathbf{Q}$ as the position coordinates and $\mathbf{p}$ and $\mathbf{w}$ as the momentum coordinates. We could integrate the system with any number of numerical schemes, from the simple explicit Euler to a more sophisticated Runge-Kutta model [14]. However, most of these schemes do not preserve the total energy of the system, and as energy preservation is a property we want for many reasons (an important one of which is to confirm the correctness of our derivatives and

implementation), we choose a specific class of integrators that *do* preserve the Hamiltonian energy: *symplectic integrators*. We will only give a brief explanation of these integrators in this section; for more details, see [10, 58, 63].

2.7.1 Sympletic Integrators

Let q denote the position coordinates of a system with Hamiltonian H and p denote the momentum coordinates. Symplectic integrators are designed for the numerical solution of Hamilton's equations

$$\begin{aligned} \dot{p} &= -\frac{\partial H}{\partial q} \\ \dot{q} &= \frac{\partial H}{\partial p} \end{aligned} \tag{2.32}$$

Formally, the time evolution of these equations is a symplectomorphism, which means that the two-form $dp \wedge dq$ is conserved. We won't make any further use of this formalism other than to note that it implies the Hamilitonian is conserved. Actually, to be precise, any finite-order symplectomorphism conserves a slightly modified Hamiltonian.

Let's now assume the Hamilonian is both separable, i.e. we have no terms that mix powers of p and q, and time-independent. Then, if we introduce the symbol $z = (p, q)$ we can write (2.32) quite simply using the Poisson bracket $\{\cdot, \cdot\}$ as

$$\dot{z} = \{z, H\} \tag{2.33}$$

(Recall that the Poisson bracket is defined as $\{f, g\} = \left[\frac{\partial f}{\partial q_i}\frac{\partial g}{\partial p_i} - \frac{\partial f}{\partial p_i}\frac{\partial g}{\partial q_i}\right]$). If we introduce an operator $D_H = \{\cdot, H\}$

then we can re-write the above as

$$\dot{z} = D_H z \tag{2.34}$$

Splitting the operator D_H into two parts $D_H = D_T + D_V$, where T is the kinetic energy and V is the potential, we can formally integrate the above equation, obtaining

$$z(t) = e^{tD_H} z(0) = e^{t(D_T + D_V)} z(0) \tag{2.35}$$

A k^{th} order symplectic integrator approximates the evolution operator $e^{t(D_T+D_V)}$ as

$$e^{t(D_T+D_V)} = \prod_{i=1}^{k} e^{tc_i D_T} e^{td_i D_V} + \mathcal{O}(t^{k+1}) \tag{2.36}$$

where $c_i, d_i \in \mathbb{R}$. To help understand what is going on, let's write out the effects of each operator

$$\begin{aligned} e^{tc_i D_T} \begin{pmatrix} q \\ p \end{pmatrix} &= \begin{pmatrix} q + tc_i \frac{\partial T}{\partial p} \\ p \end{pmatrix} \\ e^{td_i D_V} \begin{pmatrix} q \\ p \end{pmatrix} &= \begin{pmatrix} q \\ p - td_i \frac{\partial V}{\partial q} \end{pmatrix} \end{aligned} \tag{2.37}$$

Hence we need to choose the c_i, d_i such that (2.36) holds. The $k = 1$ case can be done by inspection, giving $c_1 = d_1 = 1$. This is the only first order SI integrator, and it is known as the symplectic Euler method. A useful second order case, known as the Verlet method, with $c_1 = c_2 = \frac{1}{2}, d_1 = 1, d_2 = 0$, can perhaps also be done by inspection, but for higher k more advanced techniques need to be used [92].

The code is written so as to allow an arbitrary order symplectic integrator to be used as long as the coefficients are given. In practice

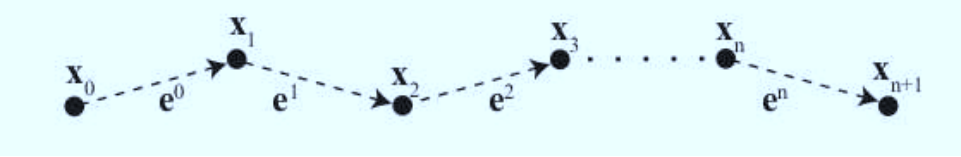

Figure 2.8.1: Quantities associated with the $n+1$ vertices are denoted with subscripts, while quantities associated with the edges n edges are denoted with superscripts.

the symplectic Euler method is most often used; increased accuracy is often less important than increased speed because we usually include damping terms anyway, and the computation time scales linearly with the order k.

2.8 Discretization

In order to actually implement the numerical time evolution scheme described in section 2.7 on a computer we need to discretize the forces and torques of section 2.6. Although not strictly necessary, it is also quite useful to discretize the energies of section 2.5, both as a check on our work and to further our understanding of how the energy flows over time between the different modes.

The centerline of the rod is described by $n+1$ vertices (in $\mathbb{R}^3$) $\boldsymbol{x}_0, \boldsymbol{x}_1, \ldots, \boldsymbol{x}_n$ connected by n edges $\boldsymbol{e}_0, \boldsymbol{e}_1, \ldots, \boldsymbol{e}_{n-1}$ where $\boldsymbol{e}^i = \boldsymbol{x}_{i+1} - \boldsymbol{x}_i$. See 2.8.1, taken from [7].

We adopt [7]'s notation, indexing quantities naturally ascribed to vertices with subscripts and quantities naturally ascribed to edges with superscripts. For each edge we have an attached local frame $\boldsymbol{Q}^0, \boldsymbol{Q}^1, \ldots, \boldsymbol{Q}^{n-1}$ where $\boldsymbol{Q}^i \in SO(3)$, i.e. $\boldsymbol{Q}^i$ is an orthonormal 3 by 3 real matrix.

In addition to the fundamental quantities above, we have many more variables that (potentially) can vary along the rod. Those

associated with the vertexes are

- $\boldsymbol{v}_i$: the lab-frame velocities
- $\boldsymbol{k}_i$: the curvatures
- $\boldsymbol{k}_i^0$: the intrinsic curvatures
- $\boldsymbol{B}_i$: the bending/twisting matrices
- d_i : the vornoi domains
- ρ_i : the mass densities

and those associated with the edges are

- $\boldsymbol{w}^i$: the angular velocities
- $\boldsymbol{\sigma}^i$: the shears
- $\boldsymbol{\sigma}_0^i$: the intrinsic shears
- $\boldsymbol{I}^i$: the inertia matrices
- $\boldsymbol{S}^i$: the shear matrices
- l^i : the rest lengths of the edges

2.8.1 Resolution Invariance

A central requirement for a good discretization is that the energies be (approximately) independent of the resolution/number of edges n: we want the difference between discrete energy and the continuous energy to scale as $\mathcal{O}(\frac{1}{n})$ so that the limits coincide as $n \to \infty$.

We begin by defining discrete curvature as the difference between adjacent frames

$$\begin{aligned} & e^{\boldsymbol{k}_i d_i} \boldsymbol{Q}_{i-1} = \boldsymbol{Q}_i \\ & \downarrow \\ & \boldsymbol{k}_i = \frac{\log(\boldsymbol{Q}_i \boldsymbol{Q}_{i-1}^T)}{d_i} \end{aligned} \tag{2.38}$$

There are two important things to note here. The first is that the curvature is only well defined on interior vertices if the rod is open, namely $\boldsymbol{x}_1$ through $\boldsymbol{x}_{n-1}$, because it measures the change between adjacent edges (if the rod is closed, then there is a curvature between e_0 and e_{n-1} we can sensibly define). The second is that the logarithm is dimensionless, so curvature has dimensions of $\frac{1}{L}$ as we would expect. Paying careful attention to the dimensions of all of the variables is critical to maintaining resolution invariance.

Discrete shear is defined on each edge, and is simply

$$\boldsymbol{\sigma}^i = \boldsymbol{Q}^i \frac{\boldsymbol{e}^i}{l^i} - \boldsymbol{d}_3^i \tag{2.39}$$

2.8.2 Discrete Energies

The discrete versions of the energies in (2.22) are

$$\begin{aligned} V^{rot} &= \frac{1}{2}\sum_{i=1}^{n-1}(\boldsymbol{k}_i - \boldsymbol{k}_i^0)^T \boldsymbol{B}_i(\boldsymbol{k}_i - \boldsymbol{k}_i^0)\ d_i \\ V^{trans} &= \frac{1}{2}\sum_{i=0}^{n-1}(\boldsymbol{\sigma}^i - \boldsymbol{\sigma}_0^i)^T \boldsymbol{S}^i(\boldsymbol{\sigma}^i - \boldsymbol{\sigma}_0^i)\ l^i \\ T^{rot} &= \frac{1}{2}\sum_{i=0}^{n-1}(\boldsymbol{w}^i)^T \boldsymbol{I}^i \boldsymbol{w}^i\ l_i \\ T^{trans} &= \frac{1}{2}\sum_{i=0}^{n}\rho_i \boldsymbol{v}_i^T \boldsymbol{v}_i\ d_i \end{aligned} \tag{2.40}$$

2.8.3 Discrete Forces and Torques

The discrete version of the forces and torques in section 2.6 are

$$\begin{aligned} \boldsymbol{\tau}_{rot}^i &= \Delta_i(\boldsymbol{k}^T\boldsymbol{B}) + \mathcal{A}_i(\boldsymbol{k} \times (\boldsymbol{k}^T\boldsymbol{B})\ d) \\ \boldsymbol{\tau}_{trans}^i &= (\boldsymbol{Q}^i\boldsymbol{e}^i) \times ((\boldsymbol{\sigma}^i)^T\boldsymbol{S}^i)) \\ \boldsymbol{f}_i^{trans} &= \Delta_i(\boldsymbol{\sigma}^T\boldsymbol{S}\boldsymbol{Q}) \end{aligned} \tag{2.41}$$

where the discrete operators Δ_i and $\mathcal{A}_i$ acting on a quantity X defined on the edges are given by

$$\begin{aligned} \Delta_i(X) &= \begin{cases} X_0, & i = 0 \\ X_i - X_{i-1}, & 0 < i < n \\ -X_n, & i = n \end{cases} \\ \mathcal{A}_i(X) &= \begin{cases} \frac{X_0}{2}, & i = 0 \\ \frac{X_i + X_{i-1}}{2}, & 0 < i < n \\ \frac{X_n}{2}, & i = n \end{cases} \end{aligned} \tag{2.42}$$

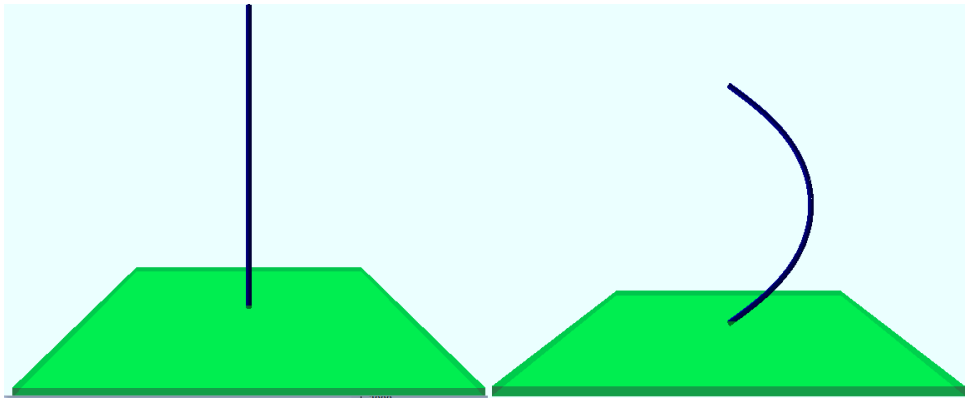

Figure 2.9.1: In this numeric experiment illustrating Euler buckling we have a homogeneous, inextensible vertically oriented rod whose ends are constrained to one the original vertical line but are free to rotate. Below the critical load threshold the rod remains straight as is shown above on the left. Above the threshold the first buckling mode, illustrated above on the right, become energetically favored so the rod deforms.

2.9 Validation

In developing and implementing a model, errors can occur at many different stages, from mistakes in derivation of the equations to bugs in the written code. Hence, in developing a new model it is imperative to find a set of tests where results from the model can be compared to known analytic results. Some classical, easily testable results for rods include Euler buckling, the Mitchell instability, and localized helical buckling. In performing these tests and analyzing the results, we will confirm that the core model is, to a high degree of confidence, correct and error free. Validation tests of additional features, such as friction, will be performed when those features are introduced.

2.9.1 Euler Buckling

Euler buckling involves a single isotropic, inextensible rod subjected to an axial load: see 2.9.1.

We are interested in how much axial load the rod can withstand before buckling under different boundary conditions. The analytical buckling threshold is [62]

$$F = \frac{\pi^2 EI}{(KL)^2} \tag{2.43}$$

where F is the threshold force, E is the modulus of elasticity, I is the moment of intertia, L is the length of the rod, and K is a constant which depends on the boundary conditions. If both ends are fixed spatially but free to rotate $K = 1$. There are numerous other possible boundary conditions, each with their own K factors, but we will only test the the $K = 1$ case.

In the context of our model, the bending constant α is equal to the modulus of elasticity times the moment of inertia. Thus, for the case with ends pinned in space but free to rotate we expect buckling at the critical force

$$F_c = \frac{\pi^2 \alpha}{L} \tag{2.44}$$

To test this phenomena numerically, I swept the phase space of F, α, and L and recorded when the beam remained straight and when it buckled. Figure 2.9.2 shows the results of varying the applied force versus alpha when $L = 1$ was held fixed. The red dots indicate that the rod buckled, the blue dots indicate that the rod remained straight, and the black line is the critical force given by (2.44). The agreement between the model and the theory is quite good.

Another slice of phase space is shown in figure 2.9.3, where force versus L at fixed $\alpha = 1$ is plotted. Again we see excellent agreement between theory and the model. We can conclude that the numerical model matches (up to some expected tolerances due to finite n and running times) the analytic Euler buckling behavior.

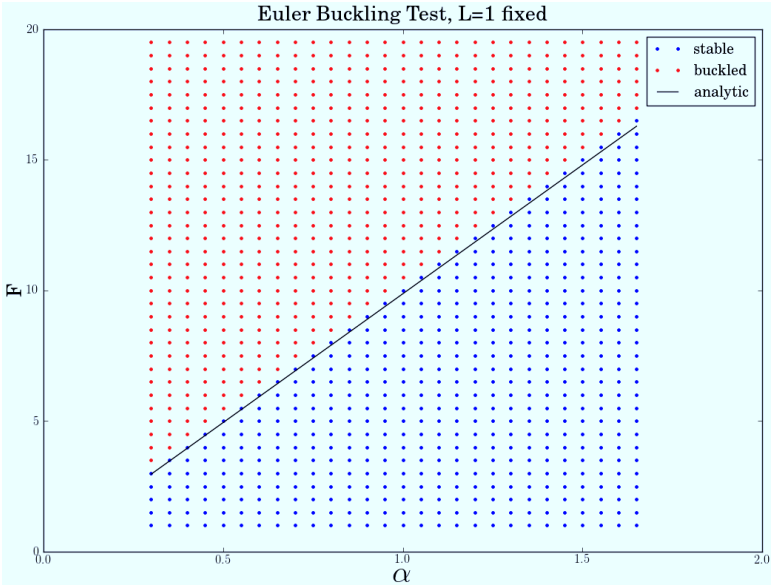

Figure 2.9.2: Results of the Euler buckling test, varying force and α with $L = 1$ fixed. Red means buckled, blue means straight, and the black line is the analytic result.

2.9.2 Mitchell Instability

To examine Mitchell's instability we begin with a single isotropic, inextensible rod in the shape of circle. We then cut the rod, twist one end a given number of times, and then glue the cut ends back together, leaving us with a circularly shaped rod with uniform twist rate. We then examine whether the rod is stable under small perturbations in the sense of whether it remains nearly circular or unstably buckles out of the plane: see 2.9.4.

One can analytically determine [38] that the instability occurs at a critical total twist:

$$\Phi_c = 2\pi\sqrt{3}/(\beta/\alpha) \tag{2.45}$$

Note that the critical total twist is independent of the total length of

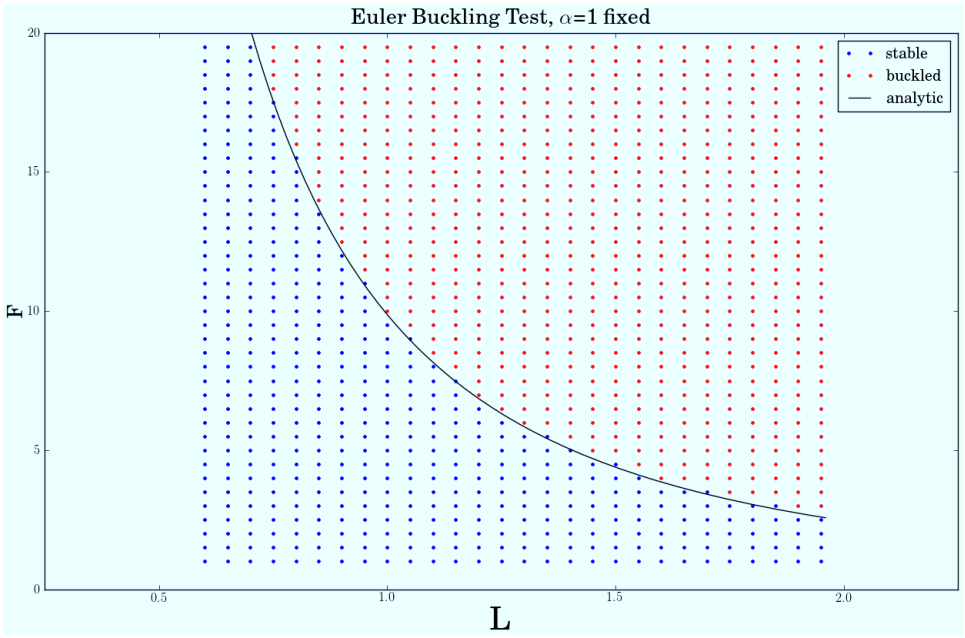

Figure 2.9.3: Results of the Euler buckling test, varying force and L with $\alpha = 1$ fixed. Red means buckled, blue means straight, and the black line is the analytic critical force.

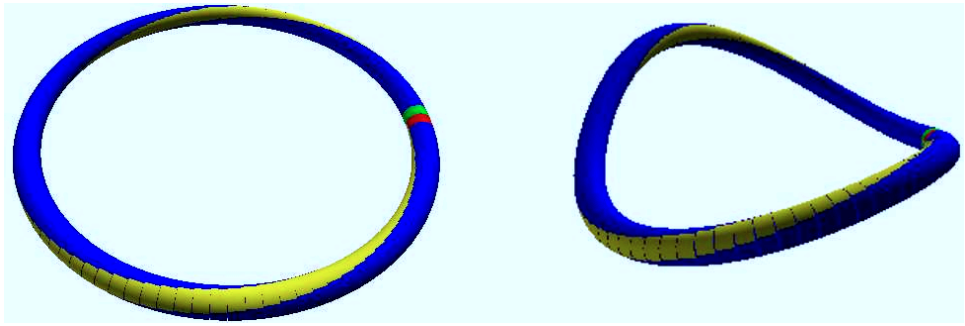

Figure 2.9.4: In Mitchell's buckling test, a homogeneous, inextensible straight rod is given a fixed amount of twist and then glued end-to-end. Below the critical twist the rod remains circular, as in the left case above. Above the threshold it deforms out of the plane, as seen in the right case above. The endpoints of the rod (where they were glued) are shown in red and green, and the twist is illustrated by the yellow stripe that runs straight along the undeformed rod.

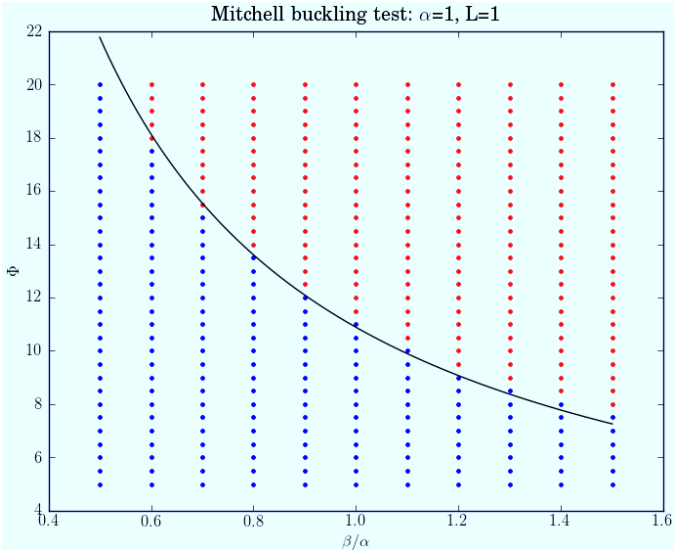

Figure 2.9.5: Varying Φ and β with $\alpha = 1$ and $L = 1$ fixed. Red means buckled, blue means stable, and the black line is the analytic critical total twist.

the rod. Equivalently the buckling depends on the global/integrated total twist rather than the local twist rate.

The first test run fixed $\alpha = 1$ and $L = 1$ and varied Φ and β: see figure 2.9.5. Another run still swept over Φ and β but changed the fixed constants to $\alpha = 5$ and $L = 2$: see figure 2.9.6. Up to some noise due to finite running times and finite n, the figures shows that the boundary between the stable region and the unstable region is precisely as predicted by 2.45. Note that the two figures look nearly identical, giving us confidence that the full system is, as is the analyic prediction, dependent only on the ratio of α to β, and independent of L.

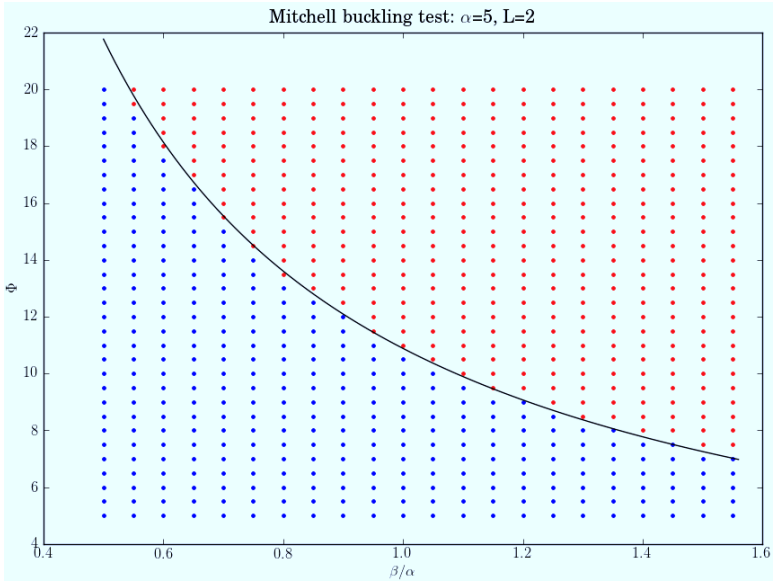

Figure 2.9.6: Varying Φ and β with $\alpha = 5$ and $L = 2$ fixed. Red means buckled, blue means stable, and the black line is the analytic critical total twist.

2.9.3 Localized Helical Buckling

In this test we again consider an single isotropic, inextensible rod. Unlike the Euler buckling case where we prescribe the load on the beam, here we prescribe how much the beam shrinks in the axial direction while we additionally apply a certain number of twists to the ends of the beam. We are interested in the resulting modulated helical shape the beam forms.

Defining $\phi(s) = \cos^{-1}(\mathbf{T}(s) \cdot \hat{x})$ as the measure of the angular deviation from the axial direction and $\phi_0 = \max_s \phi(s)$, the envelope of the resulting helix is given by [7] [84]

$$A(s) = \tanh^2\left(\frac{s}{s^*}\right) = \frac{\cos\phi(s) - \cos\phi_0}{1 - \cos\phi_0} \tag{2.46}$$

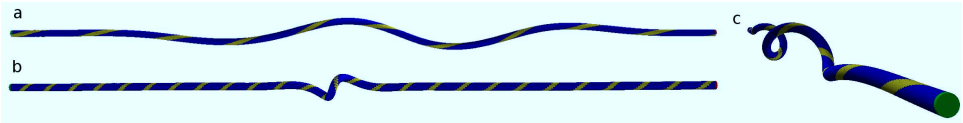

Figure 2.9.7: In the localized helical buckling test a homogeneous, inextensible rod has its ends slowly brought together while applying twist. At relatively small twists (a) a loose helical shape forms, whose shape can perhaps better be understood from the end-on view in (c). At higher twists, the deviation concentrates in a tighter swirl near the center (b).

where s^* is a characteristic length scale defined as

$$s^* = \frac{\beta m}{2\alpha}\sqrt{\frac{1-\cos\phi_0}{1+\cos\phi_0}} \tag{2.47}$$

where m is the magnitude of the torque applied to the ends, α is the bending stiffness, and β the twisting stiffness.

Following the test [7] performed to check their model, we consider a rod with rest length $L = 9.29$, $\alpha = 1.345$, $\beta = .789$ where we impose an an axial shortening of .3 units while applying 27 total twists (13.5 to each end). These numbers give a theoretical maximal angular deviation of $\phi_0 = .919$. As figure 2.9.8 shows, the results of the model are quite good, approaching the analytic solution as we increase n.

It is worth mentioning that in this test we are comparing the model's results with an entire analytically calculated shape, whereas the Euler and Mitchell tests only examined the onset of an instability. This provides support that the derivatives, etc. we have defined in the model match their analytic versions at high stresses in a manner that the Mitchell and Euler tests do not.

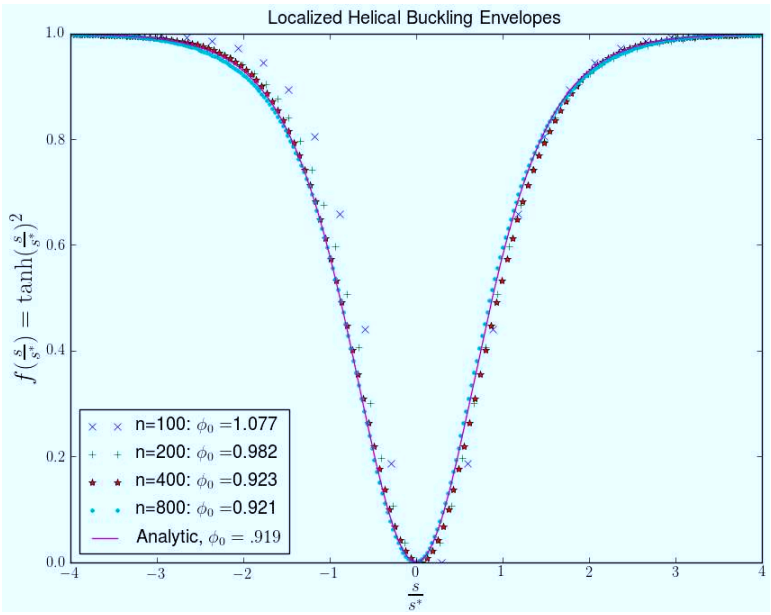

Figure 2.9.8: The envelope of the modulated helix arising from applying 27 turns while shortening the axial distance by .3 for a rod with $\alpha = 1.345$, $\beta = .789$, and $L = 9.29$ for various numbers of segments n. The maximal angular deviations approach the analytic value of $\phi_0 = 0.919$ as n increases.

3

Tendrils

With our model thoroughly tested (sucessfully!) we can move on to answering new questions. In a 2012 Sharon Gerbode, Josh Puzey, Maha and I published a paper [34] on the biology, mechanics, and behavior of the helical coiling of cucumber plants. It is natural to include this work as the first example of an application because, unlike the other applications, it does not require any extensions to the base model we've so far described. The paper is included in its entirety as the remainder of this chapter, with minor adaptions from its original format to better fit into the format of this dissertation. Note that for this paper, and for the others in this dissertation, I've amalgamated the references into a single bibliography so the references will appear somewhat different from the published versions.

3.1 Main Paper

How the cucumber tendril coils and overwinds

Sharon J. Gerbode, Joshua R. Puzey, Andrew G. McCormick, and L. Mahadevan

The helical coiling of plant tendrils has fascinated scientists for centuries, yet the mechanism of coiling remains elusive. Moreover, despite Darwin's widely accepted interpretation of coiled tendrils as soft springs, their mechanical behavior is unknown. Our experiments on cucumber tendrils demonstrate that tendril coiling occurs via asymmetric contraction of an internal 'fiber ribbon' of specialized cells. Furthermore, when pulled, both extracted fiber ribbons and old tendrils exhibit twistless overwinding rather than unwinding, with strong strain-stiffening at large extensions. We explain this behavior using simple geometric arguments, mathematical models of helical elastic filaments, and physical models composed of pre-strained rubber strips. Collectively, our study elucidates the origin of tendril coiling, quantifies Darwin's original proposal, and provides a basis for biomimetic twistless springs with tunable mechanical responses.

The transformation of an initially straight plant tendril into a helically coiled shape has captured the attention of numerous studies since the 1800's [22, 23, 40, 42, 51, 69], both from a mechanistic and a functional perspective. Tendrils serve as footholds for climbing plants, providing an alternative parasitic approach that circumvents the investment of valuable resources on independent structural support, allowing the plant to wend its way to sunlight and numerous ecological niches [50]. During climbing, an initially straight tendril searches for and attaches to a support (Fig. 3.1.1, A and B). Once tethered, the suspended portion of the tendril coils into a helical shape with at least two oppositely handed helices connected by a helical perversion (3.1.1, C and D, and Fig. 3.1.2, A and B), which Darwin first recognized as a topological necessity given the clamped, twistless boundary conditions at each end of the tendril [22]. This helical coiling axially shortens the tendril, hoisting the plant toward the attachment point (Fig. 3.1.1).

While many previous studies have investigated tendrils, the basic mechanism of tendril coiling has remained elusive. Historically, experimental studies have implicated asymmetric tendril growth as the mechanism for coiling, either via cell divisions and elongation [23, 61] or through the control of turgor pressure in the epidermal cells [64]. Theoretical treatments have incorporated intrinsic curvature or differential growth without addressing its origin in the plant or its macroscopic consequences on mechanical properties [40, 52, 66]. However, recent studies of tendril anatomy [11, 67] have provided a new twist by revealing an interior layer of specialized cells similar to the stiff, lignified gelatinous fiber (g-fiber) cells found in reaction wood [41]. In reaction wood, stiff g-fiber cells extend along tree branches, and the influx or efflux of water from these cells in conjunction with lignification in cell wall structures of oriented stiff cellulose microfibrils causes tissue elongation, contraction and

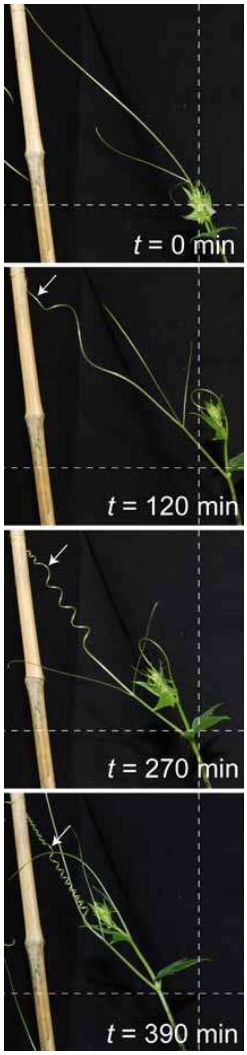

Figure 3.1.1: Cross-hairs mark the initial position of the shoot apex and highlight how tendril coiling winches the plant upward toward the support. (A) The tip of the tendril initially wraps around the supporting rod, providing secure attachment. (B) After 120 min, a helical perversion (arrow) begins to form in the suspended tendril. (C-D) The tendril shortens axially by coiling into a pair of helices with opposite handedness connected by a helical perversion (arrow).

morphosis to provide support for the branch[41]. The presence of a similar ribbon-like strip of g-fiber cells in tendrils suggests that the coiling of the soft tendril tissue may be driven by the morphosis of this internal, stiff 'fiber ribbon' [11].

We investigated the role of the fiber ribbon during tendril coiling in both *Cucumis sativus* (Cucumber) and *Echinocystis lobata* (Wild Cucumber) grown in long day conditions (details in 3.2). In cross sections of straight tendrils that have not yet attached to a support (Fig. 3.1.2A), a faint band of immature g-fiber cells is barely visible using darkfield microscopy (Fig. 3.1.2A, left inset). A second image of the same cross section taken under ultraviolet (UV) illumination shows no signal, indicating the absence of any lignification (Fig. 3.1.2A, right inset). In coiled tendrils (Fig. 3.1.2B), g-fiber cells are clearly visible under both darkfield illumination (Fig. 3.1.2B, upper left inset) and UV illumination (Fig. 3.1.2A, upper right inset), indicating prominent lignification. The fiber ribbon consists of approximately two cell layers, with the ventral layer on the inside of the helix showing increased lignification relative to the dorsal outer layer (Fig. 3.1.2B, lower insets), in agreement with previous observations of increased lignification on the stimulated side of the tendril [11, 67]. We extracted full fiber ribbons from coiled tendrils using an enzymatic treatment of fungal carbohydrolases (Driselase, Sigma) to break down the nonlignified epidermal tendril tissue (details in 3.2). Intriguingly, extracted fiber ribbons retain the helical morphology of a coiled tendril, and lengthwise cuts along the extracted fiber ribbon also reproduce the original shape (Fig. 3.1.2C, 3.2.1).

These observations suggest that tendril coiling occurs via asymmetric contraction of the fiber ribbon; the ventral side shrinks longitudinally relative to the dorsal side. The resulting relative

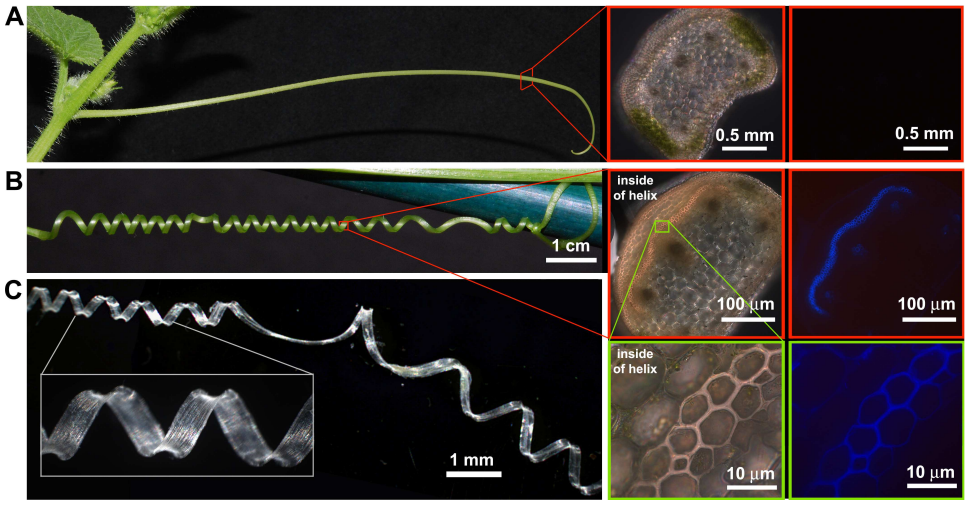

Figure 3.1.2: During coiling, a strip of specialized structural gelatinous fiber cells (the fiber ribbon) becomes lignified and asymmetrically contracts longitudinally. (A) A straight tendril that has not encountered a support lacks lignified gelatinous fiber (g-fiber) cells. Inset: Tendril cross section, darkfield (left) and UV autofluorescence (right) showing no lignin signal. (B) In coiled tendrils, the fully developed fiber ribbon consists of about two layers of highly lignified g-fiber cells extending along the entire length of the tendril. Inset: Tendril cross section, darkfield (left) and UV autofluorescence (right) showing a strong lignin signal in the fiber ribbon cells. Increased magnification (lower insets) reveals that the ventral cell layer is more lignified than the dorsal cell layer. (C) The extracted fiber ribbon retains the helical morphology of the coiled tendril. Inset: Higher magnification reveals long g-fiber cells extending along the fiber ribbon.

dorsiventral length difference gives the coiled fiber ribbon intrinsic curvature, which is retained even after lengthwise cuts because the entire ventral layer shrinks uniformly relative to the dorsal layer. The driving asymmetric contraction may be generated by a variety of dorsiventral asymmetries, including the observed differential lignification (Fig. 3.1.2B, lower right inset), possible variations in cellulose microfibril orientation as seen in reaction wood, or differential water affinities. For example, since lignin is hydrophobic, the ventral cells may expel more water during lignification, thereby driving increased cell contraction. This is supported by the observation that an extracted fiber ribbon from any coiled tendril passively shrinks and coils even further when dried, and regains its original shape when rehydrated. Dehydrated tendri ls also exhibit the same behavior since they are dominated by the stiff fiber ribbon. Although there may be multiple underlying sub-cellular regulatory processes for this asymmetry, our experiments show that the dorsiventrally asymmetric contraction of an initially straight g-fiber ribbon provides a biophysical mechanism for tendril coiling by imposing its curvature on the surrounding soft tendril tissue. The perversions in a doubly-supported tendril then follow naturally from the topological constraint imposed by the prevention of twist at its ends.

To gain intuition for the asymmetric contraction observed in tendril fiber ribbons, we reconstituted the mechanism using a physical model composed of two bonded, prestrained silicone rubber sheets, similar to rubber models for shaping sheets [3, 16, 76]. The first silicone sheet was uniaxially stretched along one direction, and the second silicone sheet of approximately equal thickness was then formed by spreading a layer of silicone adhesive onto the stretched sheet. After curing, thin strips were cut along the stretched direction, yielding ribbon-like bilayer silicone strips (Fig. 3.1.3A) with intrinsic curvature set by the relative strain, thickness, and stiffness of the two layers 3.2.3. Just

like tendril fiber ribbons, the initially straight physical models spontaneously form coiled configurations when released, with two opposite-handed helices connected by a helical perversion (Fig. 3.1.3A, left).

However, careful observations reveal an unexpected difference between the physical models and the tendril fiber ribbons. When clamped at both ends and pulled axially, a coiled physical model simply unwinds to its original flat, uncoiled state (Fig. 3.1.3A, SOM movie S4). However, in tendril fiber ribbons we observed a counterintuitive 'overwinding' behavior where the ribbon coils even further when pulled, adding turns on both sides of the helical perversion (Fig. 3.1.3A, right and SOM). Under high enough tension, the fiber ribbon finally unwinds, returning to a flat, uncoiled state as expected.

Inspired by the asymmetric lignification pattern in fiber ribbons, which suggests that the inner layer is stiffer and less extensible under tension, we added a relatively inextensible fabric ribbon to the inside of a coiled physical model to try to reproduce the overwinding behavior. Furthermore, to mimic lignified plant cells, which maintain their shape under compression, we added an incompressible copper wire to the exterior of the helix. These modifications increased the silicone model's effective bending stiffness relative to its twisting stiffness, making the intrinsic helix curvature harder to change while still allowing the model to twist around its centerline (3.2). The modified physical model exhibited significant overwinding (Fig. 3.1.3C; SOM, movie S6). Indeed, in the limit of infinite bending stiffness, the helix curvature can not change regardless of the tension applied, and a single helix with clamped ends can not axially extend when pulled. If, however, two such helices are connected by a helical perversion, then the perversion itself allows for rotation and thus the

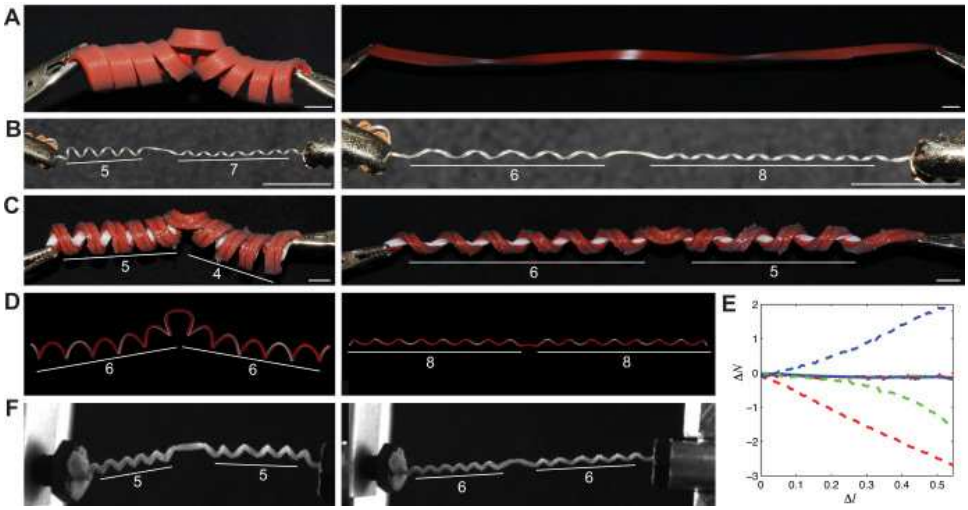

Figure 3.1.3: (A) A silicone twistless spring with lower bending stiffness than twisting stiffness exhibits unwinding when pulled, returning to its original flat shape. (B) A tendril that coils after contacting a support has at least one perversion, as shown here. When such a tendril is extended, it initially overwinds rather than unwinding; the number of turns in the left half increases from 5 to 6, while that in the right half change from 7 to 8. (C) Overwinding in the silicone model is induced by adding a relatively inextensible (under tension) fabric ribbon to the inside of a coiled silicone ribbon and an inextensible (under compression) copper wire to the exterior. These increase the bending stiffness while maintaining relatively low twisting stiffness. As a result, the twistless spring winds further when pulled, gaining one new turn on each side of the helical perversion (number of turns indicated in white). (D) When the ratio of the bending to twisting stiffness $B/C > 1$, numerical simulations of elastic helical filaments recapitulate this behavior consistent with physical and biological experiments. (E) Change in the number of turns in each helix ΔN is plotted versus scaled displacement Δl for B/C values 1/5 (red), 1 (green), 5 (blue). This reveals that overwinding occurs when $B > C$, and becomes more pronounced as bending stiffness increases relative to twisting stiffness. (F) Considerable overwinding is also observed in old tendrils, which are dried and flattened, with the ratio $B/C > 1$. All scale bars equal 1 cm.

ability to add helical turns. By overwinding, each helix can thus geometrically accommodate axial extension when pulled without varying its curvature (3.2, 3.2.3).

Of course, real tendril fiber ribbons do not have infinite bending stiffness, and eventually the overwinding behavior saturates and the helices unwind. In order to understand overwinding for the more general case of a filament or strip with different ratios of bending and twisting stiffness, we constructed a minimal mathematical model of the coiled fiber ribbon with a single helical perversion, using two equal length, elastic helical filaments of opposite handedness but identical intrinsic curvature k_0 and torsion w_0, and uniform bending stiffness B and twisting stiffness C, with the system initially at equilibrium (Fig. 3.1.3D, left). When the filament is pulled by its clamped ends, deviations in curvature and twist from their natural values lead to variations in the total energy of the filament due to bending and twisting energies (SOM text). Minimizing the energy of the extended filament numerically, we determine the filament shape and position as a function of the applied tension (Fig. 3.1.3D, right), while simultaneously measuring the number of turns in each helix to detect overwinding. We find that when $B/C < 1$, the composite tendril unwinds on extension, but when $B/C > 1$ the filament overwinds (Fig. 3.1.3D), and the maximum number of turns on each side of the perversion increases with increasing B/C (Fig. 3.1.3E) [62]. We note that for a circular cross section helical spring made of an isotropic material, $B/C = 1 + \nu$, with Poisson ratio ν normally in the range $0 < \nu \lesssim 0.5$ so that typical springs exhibit minimal overwinding when pulled and allowed to rotate at their ends.

The observation of overwinding in fiber ribbons naturally leads to the question of whether entire tendrils also overwind. Interestingly, we find that while fiber ribbons from any coiled tendril always overwind, recently coiled, fully hydrated tendrils (young tendrils) do not

overwind, but mature, dry tendrils (old tendrils) exhibit significant overwinding (Fig. 3.1.3F) and intermediate tendrils were variable in their overwinding behavior. The overwinding observed in old tendrils is likely due to the fact that as the tendril dries, the epidermal cells lose volume and the tendril flattens down to a ribbon-like shape much more similar to the internal fiber ribbon, so that $B/C > 1$. To investigate the mechanical and functional consequences of overwinding, we measured the force required to axially stretch tendrils using a custom force measurement setup (details in 3.2), increasing the axial displacement until mechanical failure occurred either within the tendril tissue or at the clamped ends. Force-extension curves were measured for a total of 20 tendrils. In each tendril, we first tested a segment containing the helical perversion and several turns on either side (perverted segment), and then tested a second segment with no perversion (clamped segment). A wide variety of mechanical responses were observed; the two most extreme force-extension curves were measured for the young tendril (red) and old tendril (blue) shown in Fig. 3.1.4A. To compare tendrils of different lengths and stiffnesses, and allow for comparison with our numerical model, we plotted the scaled force $\tilde{F}$, defined as the measured force F divided by the initial slope $\mathrm{d}F/\mathrm{d}l$ of the lower red curve as a function of the scaled displacement $\Delta l = (L - L_0)/S$, which is the instantaneous axial length L of the tendril relative to its relaxed axial length L_0, divided by its total arc length S. In the young tendril, the perverted segment (dotted red) is softer than the clamped segment of the same tendril (solid red), and becomes increasingly softer as Δl is increased (Fig. 3.1.4A). In contrast, the perverted segment of the old tendril (dotted blue) is initially softer than the clamped segment of the same tendril (solid blue), but becomes stiffer at large extensions. To account for differences in initial stiffness and compare the nonlinear portions of the force curves, we define the reduced force f by dividing

each force curve by its *own* initial slope. In Fig. 3.1.4B we plot the difference $\Delta f = f(\text{with perversion}) - f(\text{without perversion})$ to highlight the effect of the helical perversion on the mechanical response of the tendril. In the young tendril (red) where no overwinding occurs, Δf is always negative, indicating that the perversion consistently decreases the reduced force f necessary to stretch the tendril relative to the clamped case. However, in the old, overwinding tendril (blue), the perversion actually increases the reduced force needed to stretch the tendril as Δl increases.

To quantitatively elaborate the range of mechanical behaviors bounded by these two tendril measurements, we also calculated the force-extension curves for helical filaments numerically. In Fig. 3.1.4C, we plot the scaled force $\tilde{F}$ against the scaled displacement Δl for 5 numerical filaments with different ratios B/C of bending to twisting stiffness: 1/5 (red), 1 (green), and 5 (blue). For the soft filament with $B < C$, where no overwinding is observed, the presence of a helical perversion decreases the force needed to axially extend the filament, an effect that becomes more pronounced as the filament is extended. This is similar to what is seen in the young tendril (Fig. 3.1.4A and 4B). However, the force response qualitatively changes when B/C increases above about 3, and the filament exhibits overwinding with more than about 1 extra turn. Initially the presence of a perversion substantially decreases the force needed to stretch the filament, but at large extensions, the perversion actually increases the force needed for extension. This effect is best highlighted by plotting the difference $\Delta\tilde{F} = \tilde{F}(\text{with perversion}) - \tilde{F}(\text{without perversion})$ as shown in Fig. 3.1.4D. Whereas this difference is always negative for filaments with $B < C$, in overwinding filaments the difference $\Delta\tilde{F}$ is initially significantly negative but transitions to positive values as the filament is further stretched. This indicates that in overwinding

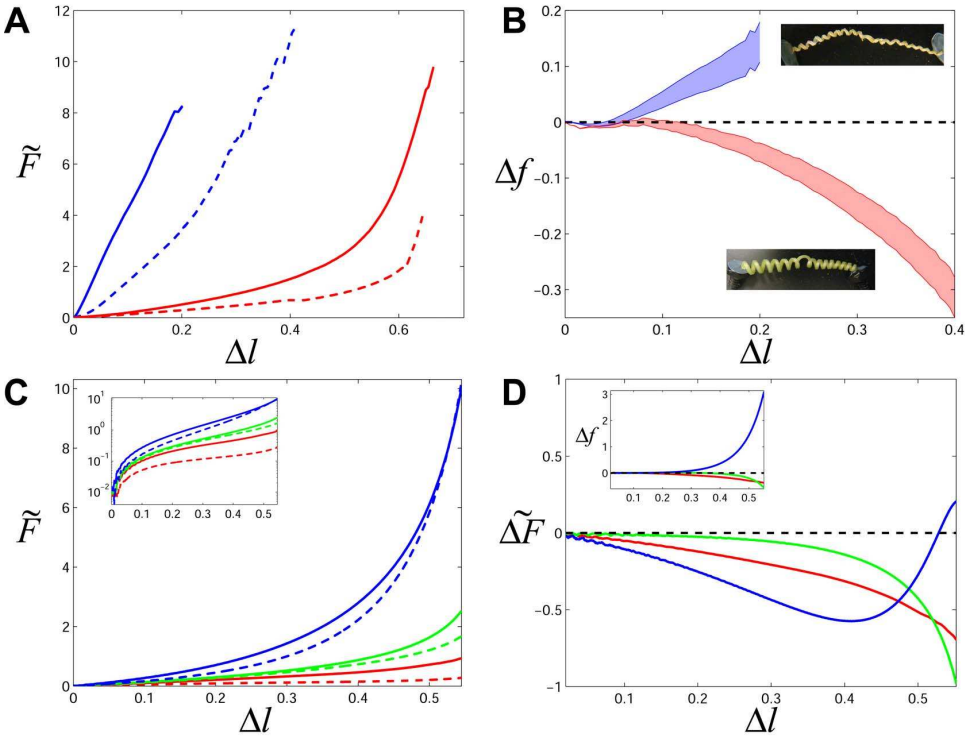

Figure 3.1.4: (A) Force extension curves for one young tendril that does not overwind (red curves) and one old tendril that exhibits significant overwinding (blue curves). Each tendril was separated into a segment containing the helical perversion (dotted curves) and a segment containing a single helix with no perversion (solid curves). The scaled force $\tilde{F}$, defined as the measured force divided by the initial slope of the dotted red curve, is plotted against the scaled displacement Δl for each case. (B) The difference in reduced force $\Delta f = f$(with perversion) $- f$(without perversion) is plotted against the scaled displacement Δl (young tendril in red; old tendril in blue). The shaded range in each curve indicates variations in the fitted initial slope value. Images of the two tendrils are shown next to the corresponding curves. (C) Scaled force $\tilde{F}$ vs. scaled displacement Δl is plotted for numerical filaments with B/C values 1/5 (red), 1 (green), 5 (blue). Dotted curves correspond to segments containing a perversion and the solid curves correspond to segments without a perversion. *Inset*: Log-linear plot of the same data. (D) The difference in scaled force $\Delta\tilde{F} = \tilde{F}$(with perversion) $-\tilde{F}$(without perversion) highlights the effect of the helical perversion as a function of B/C. For $B < C$, the presence of a helical perversion decreases the force needed to axially extend the filament, an effect that becomes more pronounced the further the filament is extended. For $B > C$, the perversion initially decreases the force needed to stretch the filament, but eventually increases the force needed as the filament is extended further. *Inset*: The difference in reduced force Δf is plotted against Δl for direct comparison with the experimental data.

filaments, the presence of a perversion initially softens the force response but eventually stiffens the filament relative to the clamped case. We note that this behavior is qualitatively different than that reported in earlier theoretical studies [40, 66], where overwinding was not observed in the range of B/C values studied. The inset in Fig. 3.1.4D shows the reduced difference in force Δf, allowing for direct comparison with the experimental tendril data. The trends shown in the numerical force extension curves are consistent with the observed experimental data, indicating that the unique force extension behavior shown in Fig. 3.1.4D explains the extremes observed in the two tendrils.

We have found that asymmetric contraction of a specialized 'fiber ribbon' of g-fiber cells drives tendril coiling in cucumbers. Furthermore, differential lignification in the two layers of the fiber ribbon increases the bending stiffness while maintaining relatively low twisting stiffness, leading to novel overwinding behavior in fiber ribbons and old tendrils. Our physical models allow us to recreate this behavior in simple prestrained composite rubber strips, while our geometric and mechanical theory and detailed computations explain them.

Taken together, our observations raise two classes of questions: at the evolutionary level, how widespread is this mechanism in other tendril-bearing species, and at the mechanical level, what might we learn about the functional principles of these soft twistless springs? Preliminary studies of *Passiflora* tendrils reveal an off-axis band of g-fibers, suggesting a qualitatively similar mechanism for coiling (SOM Fig. S4), although interestingly, both young and old coiled *Passiflora* tendrils exhibit pronounced overwinding (SOM Fig. S4, and movie S11). While Cucurbitaceae and Passifloraceae are distantly related members of the same phylogenetic clade, their tendrils have clearly evolved independently [82]. This opens the door for future

comparative studies between species, as well as mechanistic investigations of subcellular processes regulating asymmetric contraction. Functionally, the combination of mechanical asymmetry, helical perversion and the unusually large ratio of bending and twisting stiffness of the tendril creates an auto-adaptive spring, one that is initially soft because it can overwind and then stiffens strongly when deformed further. Darwin himself wrote that 'the tendril strikes some object, and quickly curls round ...contracts into a spire, dragging up the stem, and forming an excellent spring' [22]. Our quantitative understanding complements this proposal by showing how it happens and suggests methods to tune the linear and nonlinear response of the ubiquitous helical spring using a simple combination of geometry, pre-strain and material behavior.

3.2 Supplementary Material

3.2.1 Description of growth conditions

All plants were grown in the Department of Organismic and Evolutionary Biology greenhouses at Harvard University in long day conditions of 16 hour light/8 hour dark at a constant temperature of approximately 25°C.

3.2.2 Microscopy

A Zeiss AxioImager Z2 (Harvard Imagining Center) fitted with a Zeiss AxioCam Mrc digital camera was used for capturing darkfield and UV images. An Insight Spot camera mounted on a Leica Wild M10 dissecting scope was used to image fiber ribbons in figure 2C. Color images of tendrils, models, and fiber ribbons were taken with a Nikon D40x.

3.2.3 Fiber ribbon extraction

To extract fiber ribbons from whole tendrils, a solution of 2% Driselase (by weight) in phosphate buffered saline (PBS) was prepared. Since Driselase does not go completely into solution in PBS, the mixture was vortexed vigorously and then allowed to sit for 30 minutes. A fresh tendril was submerged in the Driselase supernatant and stored at 37°C overnight. After soaking overnight, the tendril was removed and briefly rinsed in fresh PBS. Forceps were used to gently slide the broken-down epidermal cells off of the fiber ribbon.

3.2.4 Force measurement apparatus

Tendrils were clamped at one end to a translation stage (Newport 426 crossed-roller bearing linear stage, with a Newport universal controller, model ESP 300) and clamped at the other end to a full bridge thin beam load cell (Omega LCL-113G) connected via a digital transmitter (Omegabus D1521) to RS-232. Custom LabVIEW software was used to move the translation stage in steps of 0.1 mm, measure force with the load cell, and simultaneously image the tendril with a digital ccd camera (Allied Vision Technologies Marlin) to monitor overwinding or unwinding behavior. Since both ends of the tendril were clamped, the ends were prevented from rotating during extension.

3.2.5 Intrinsic curvature of a prestrained bilayer ribbon

In cucumber tendrils, we observe that coiling occurs as a result of the morphosis of the fiber ribbon, a bilayer strip of long cells that extends along the entire length of the tendril. As one layer of the ribbon shrinks relative to the other, the composite ribbon develops intrinsic curvature. We have experimentally mimicked this process using prestrained bilayer rubber models as described in the manuscript.

Here we derive the intrinsic curvature k_0 of a bilayer rectangular cross-section ribbon with prestrain.

Consider two elastic ribbons of equal width t and heights h_1 and h_2 (see 3.2.2) that are combined to make a bilayer ribbon. The lower layer has Youngs modulus E_1 while the upper layer has Youngs modulus E_2. The lower layer is first stretched to a strain value of ϵ_1^* and is held in tension. The upper layer is not stretched, so that $\epsilon_2^* = 0$, and is permanently bonded to the lower layer. Then, when tension is released on the bonded bilayer ribbon, it relaxes to its equilibrium configuration with strains ϵ_1 and ϵ_2 and curvature k_0. The value of the intrinsic curvature k_0 is determined by the relative strain, Youngs moduli, and heights of each layer. Both force and torque are balanced in the bilayer ribbon:

$$\begin{aligned} \epsilon_1 E_1 h_1 &= -\epsilon_2 E_2 h_2 \qquad \text{(Force balance)} \\ F_1 h_1/2 + F_2 h_2/2 &= -k_0(E_1 I_1 + -E_2 I_2) \qquad \text{(Torque balance)} \end{aligned}$$

where $I_1 = (\frac{th_1^3}{12})$ and $I_2 = (\frac{th_2^3}{12})$ are, respectively, the moment of area of the lower and upper layers, and $F_1 = \epsilon_1 E_1 h_1 t$ and $F_2 = -\epsilon_2 E_2 h_2 t$ are the forces in each layer. Combining these two relations yields an expression for the curvature k_0:

$$k_0 = \frac{6(\epsilon_1 E_1 h_1^2 + \epsilon_2 E_2 h_2^2)}{E_1 h_1^3 + E_2 h_2^3}$$

3.2.6 Overwinding in the limit of infinite bending stiffness

Overwinding behavior can be intuitively understood in the limit of infinite bending stiffness relative to twisting stiffness by considering the simple case of a single helix with N turns, helical pitch p and

radius r (3.2.3A). This helix has uniform curvature k_0 and twist w_0 given by

$$\begin{aligned} k_0 &= \frac{r}{(r^2 + (p/2\pi)^2)} \\ w_0 &= \frac{p/2\pi}{(r^2 + (p/2\pi)^2)} \end{aligned}$$

The arclength S of this helix is related to the pitch, radius, and number of turns by $S = N\sqrt{(2\pi r)^2 + p^2}$ and its axial length L is given by $L = Np$. If this initial helix is slightly deformed to increase its axial length from L to $L + \delta L$ without increasing its total arclength S (ie. without stretching), we assume that it will take on a new helical shape with new pitch $p + \delta p$, radius $r + \delta r$, and number of turns $N + \delta N$. If the number of helical turns is held fixed by extending the helix with both ends clamped, then $\delta N = 0$. In this case, the axial deformation δL can be accommodated by changing the pitch and radius to reduce the curvature to $k_0 - \delta k$, ie. by flattening out the helix (3.2.3B).

However, in the limit of infinite bending stiffness relative to twisting stiffness, the helix curvature cannot change, so $\delta k = 0$. In this case, the axial deformation δL can still be achieved by allowing one end of the helix to rotate, increasing the number of turns by δN. In this case, the pitch and radius both change so as to maintain constant curvature k_0, and the number of turns increases correspondingly in order to accommodate the total helix arclength S given the new pitch and radius (Fig. S2C). This increase in the number of turns is precisely what occurs in the experimentally observed overwinding.

3.2.7 Prediction of overwinding with finite bending stiffness

For helical filaments with finite bending and twisting stiffness, the effect of overwinding can again be predicted for the simple case of a single helix. This prediction follows the calculation outlined in Love [cite]. Again we consider an initial helix with pitch p, radius r, and arclength S, but now the helix has finite bending stiffness B and twisting stiffness C. A helix with pitch p and radius r is equivalently defined by its pitch angle $\alpha = \arctan(p/2\pi r)$. In terms of this pitch angle and the arclength, the axial length of the helix is $L = S\sin(\alpha)$. The helix can be described in cylindrical coordinates as

$$\begin{aligned} r(s) &= r \\ \theta(s) &= s\cos(\alpha)/r \\ z(s) &= s\sin(\alpha) \end{aligned}$$

Here the radial coordinate is a constant r, the angular coordinate θ varies from 0 to $\vartheta = S\cos\alpha/r$ at the other end, the vertical coordinate varies from 0 at one end to L at the other, and $N = \vartheta/2\pi$ is the total number of turns in the helix.

We now consider small deformations of the helix due to a force F and torque τ applied at both ends. In response to these, the helix will take on a new equilibrium shape, which we assume is a helix with the same arclength S, but new pitch angle $\alpha + \delta\alpha$ and radius $r + \delta r$. The new shape is determined by the equations of mechanical equilibrium for force and torque:

$$
\begin{aligned}
F &= \frac{1}{Sr^2}\left[\delta L(C\cos^2\alpha + B\sin^2\alpha) + 2\pi r\ \delta N(C-B)\sin\alpha\cos\alpha\right] \\
\tau &= \frac{1}{Sr}\left[\delta L(C-B)\sin\alpha\cos\alpha + 2\pi r\ \delta N(C\sin^2\alpha + B\cos^2\alpha)\right]
\end{aligned}
$$

where δN is the change in the total number of turns. Negative δN indicates unwinding, while positive δN corresponds to overwinding. In order to determine whether overwinding occurs when the helix is extended to length $L + \delta L$ under only an applied axial force, we set the applied torque $\tau = 0$ and solve for δN and δL in terms of F, l, r, B, and C. With two equations and two unknowns, we attain the desired relations:

$$
\begin{aligned}
\delta L &= Sr^2(\sin^2\alpha/B + \cos^2\alpha/C)F \\
\delta N &= \frac{Sr}{2\pi}\sin\alpha\cos\alpha(\frac{1}{C} - \frac{1}{B})F
\end{aligned}
$$

Since the helix pitch angle α is between 0 and $\pi/2$, the factor of $\sin\alpha\cos\alpha$ is always positive. Thus, this prediction for small deformations shows that the change in number of turns is positive if the bending stiffness B is greater than the twisting stiffness C, and negative if $B < C$, consistent with our observations of overwinding in cucumber tendrils and numerical simulations. We note that since this linearized prediction assumes small changes in shape, it cannot recreate the eventual unwinding that is observed in the numerical simulations and the experiments.

Figure 3.2.1: The same fiber ribbon is progressively cut lengthwise to yield narrower and narrower strips. Arrows indicate helical perversion in all panels. Left: The intact fiber ribbon. Middle: Resulting fiber ribbon from one lengthwise cut. Right: The fiber ribbon from the middle panel is cut lengthwise again, yielding an even narrower ribbon with the same shape.

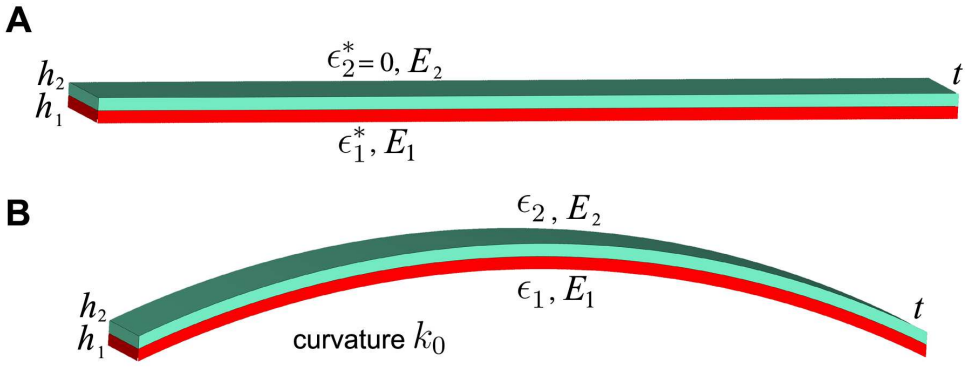

Figure 3.2.2: (A) Two elastic ribbons of equal width t and heights h_1 (lower red) and h_2 (upper blue) are combined to make a bilayer ribbon. The lower layer is stretched to a strain value of ϵ_1^* and held in tension, while the upper layer is not stretched, so that $\epsilon_2^* = 0$. (B) When tension is released on the bilayer ribbon, it relaxes to its equilibrium configuration with strains ϵ_1 and ϵ_2 and curvature k_0.

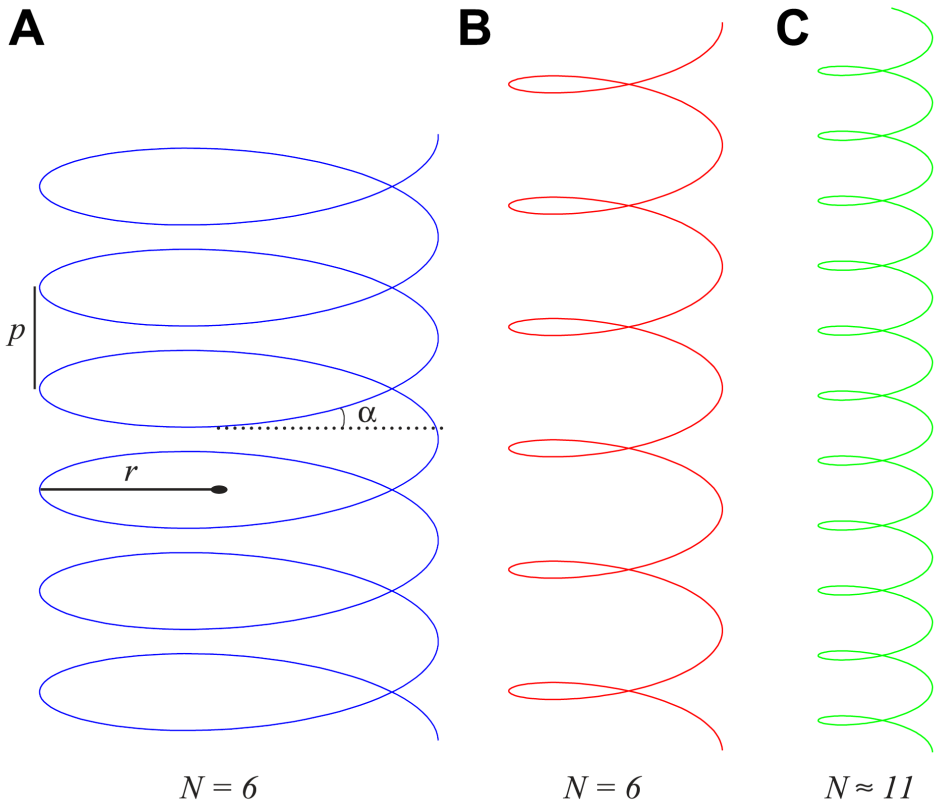

Figure 3.2.3: (A) A helix with pitch p, radius r, pitch angle $\alpha = \arctan(p/2\pi r)$ and axial length L. (B) A deformed helix with increased axial length $L + \delta L$, achieved by maintaining a constant number of helical turns (ie. enforcing no unwinding or overwinding). (C) A deformed helix with increased axial length achieved in the limit of infinite bending stiffness by maintaining constant curvature k.

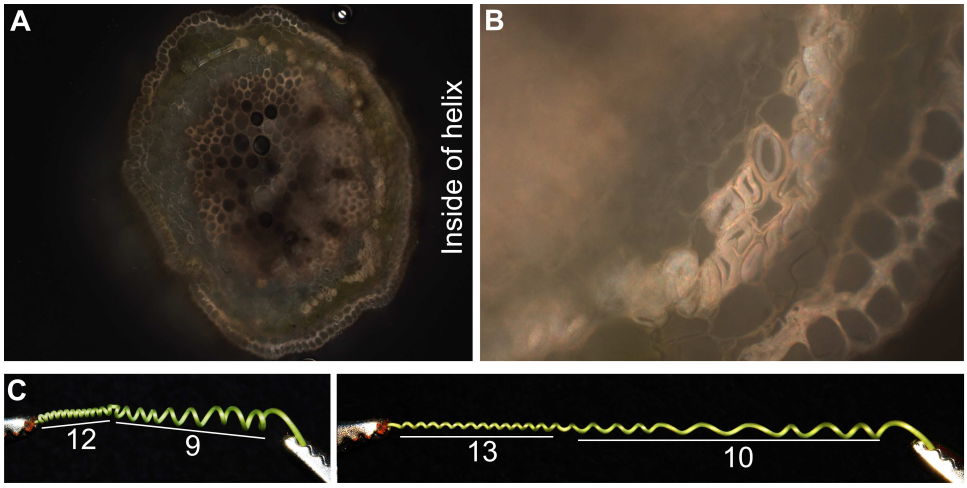

Figure 3.2.4: (A) Cross section of coiled *Passiflora sp.* tendril. Gelatinous fiber (g-fiber) cells are present on the ventral side of the tendril, ie. on the inside of the helix. (B) Magnified view of g-fiber cells. Artifactual detachment of the gelatinous fiber from the secondary wall, as seen in panel B has been shown to be a distinctive characteristic of g-fiber cells [17, 67]. (C) A coiled *Passiflora sp.* is shown in its relaxed configuration. When tension is applied, the tendril overwinds and adds additional turns on both sides of the helical perversion.

4

Solenoids and Plectonemes

In this chapter we're interested in understanding the configurations that arise when a extensible string is twisted. Whereas the tendril question could be fully addressed witht he basic model, for this problem we're going to need our first extension to the model: self-collision.

4.1 Self-Collision

Solenoids and plectonemes arise because the rod cannot pass through itself (and hence reduce its link!). To study them, therefore, we'll need to introduce a self-collision model which enforces impenetrability. There are two primary ways in which we could implement the self-collision model. The first is a projection-based method where two

overlapping cylinders (e.g. [7]) are essentially moved apart until they are no longer overlapping. This option is similar to a hard constraint or boundary condition, and makes particularly good sense when modeling things like metal equipment that cannot interpenetrate any measurable amount. The trouble with this method is that enforcing the hard constraint makes the system very stiff, requiring the use of very small timesteps or a highly sophisticated projection algorithm (or most likely, both). Moreover, many of the systems we are interested in happen to actually be squishy, where we would like some amount of deformation/penetration to be possible.

That brings us to the second option, which is to assign an energy to the interpenetration. We could do something complicated like calculate the volume of the overlap between the two cylinders, but the simplest model is to calculate the minimum distance between the two cylinder centerlines and base the repulsive force on that. Considerable has been done in this area [32, 80, 86], but the algorithms used are complex and hence computationally expensive for dynamic simulation. In searching for a simpler model we surprisingly find that something as simple as an algorithm for finding that minimum distance between two cylinders efficiently doesn't seem to be available in the literature, so I created my own algorithm. Essentially what we do is linearly parameterize each of the cylinder's centerlines, as $c_i(h) = x_i + (x_{i+1} - x_i)h$, and then find the two points on those lines that are minimally apart:

$$d_{min}^{ij} = \min_{t,s \in [0,1]} ||c_i(s) - c_j(t)|| \tag{4.1}$$

In the above equation we're finding the minimum distance between the i^{th} and j^{th} segment. Self-collision is inherently an $O(n^2)$ algorithm (unless you do something really fancy with binary bounding boxes, etc.), so in larger systems self-collision will come to dominate the

running time. As such, there are a number of important tricks we can use to speed up the calculation of the optimization above. These tricks are best shown in the actual code, which is given in the coding section 8.4.

Once we have the minimum distance vector $\boldsymbol{d}_{min}^{ij}$ (which points from the closest point on cylinder i to the closest point on cylinder j) between the two segments we can apply the repulsive forces. If $\boldsymbol{d}_{min}^{ij}$ is larger than the sum of the two radii, i.e. $\boldsymbol{d}_{min}^{ij} > r_i + r_j$ then the segments are not in collision and no forces are applied. Otherwise we calculate the scalar overlap as $\gamma_{ij} = r_i + r_j - \boldsymbol{d}_{min}^{ij}$ and then apply forces

$$\begin{aligned} \boldsymbol{f_i} &= -\hat{\boldsymbol{d}}_{min}^{ij}(\zeta_{soft}\gamma_{ij} + \zeta_{hard}\gamma_{ij}^3) \\ \boldsymbol{f_j} &= \hat{\boldsymbol{d}}_{min}^{ij}(\zeta_{soft}\gamma_{ij} + \zeta_{hard}\gamma_{ij}^3) \end{aligned} \tag{4.2}$$

The functional form of the repuslive forms above was somewhat arbitrarily chosen, but in practice the combination of the linear repulsive force ($\propto \zeta_{soft}$) emulates gives a nice "squishyness" while the "hard core" term ($\propto \zeta_{hard}$) prevents the filament from ever crossing through itself.

4.2 Solenoid Experiments and Shearable Buckling

While our model is able to pass the classical buckling tests which assume inextensible, unshearable rods, it can also generalize to extensible, shearable rods and explain behaviors like those shown in 4.2.1. Unfortunately, there aren't any classical buckling tests for shearable rods. To verify the model's behavior, then, we'll have to develop our own analytic buckling test and then compare the model

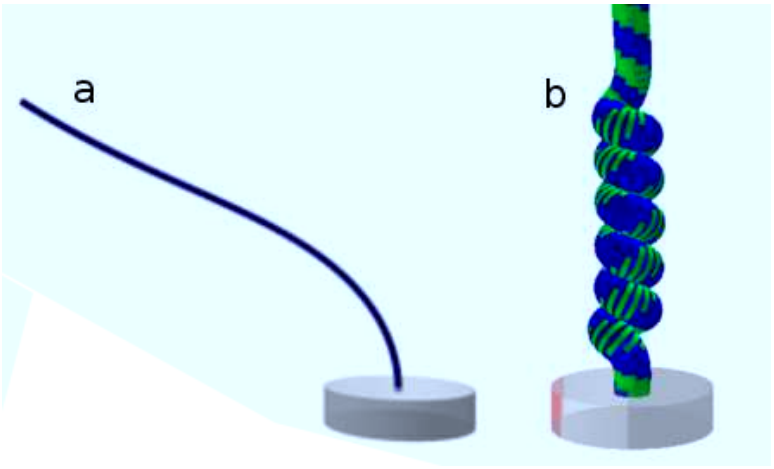

Figure 4.2.1: In (a) a shearable rod gravitationally self-buckles, forming a shape qualitatively different than the unshearable version's (see also 4.2.2). (b) shows a extensible rod deformed into a solenoid, a conformation not seen in inextensible rods.

with those results.

Consider the static gravitational self buckling of a vertically standing rod that is *shearable* (but not extensible). We consider an isotropic rod, so we can think of the stiffness matrices as diagonal constants and construct the dimensionless parameter $\epsilon = \frac{B}{Sr^2}$ as the scaled ratio of the bending and shearing stiffnesses. Our goal is to derive, as a function of ϵ, the critical length at which the rod first buckles. We can accomplish this by performing a linear stability calculation as follows.

Suppose that the rod's deviation from its vertical rest state is attributable to two effects: $y_b(x)$ the deviation due to bending, and $y_s(x)$, the deviation due to shear, where $x \in [0, L]$ (see fig. 2a). We linearly approximate the curvature as $y_b''(x)$ and the shear as $y_s'(x)$, leading to the following potential energy:
$E = \int_0^L \frac{B}{2}(y_b''(x))^2 + \frac{S}{2}(y_s'(x))^2 + \rho(L-x)\sqrt{1-(y_b'(x)+y_s'(x))}\,dx$. By

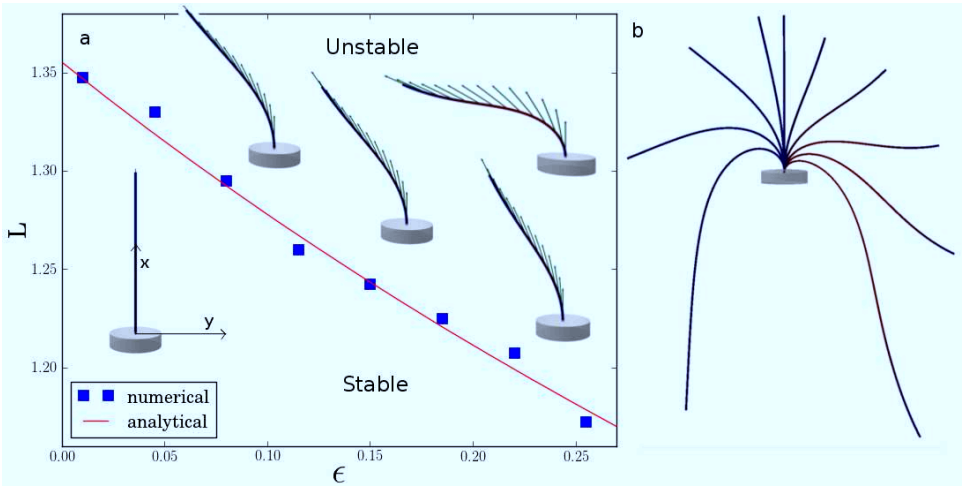

Figure 4.2.2: A filament of variable length is fixed in the ground and subjected to gravity. In (a), the impact of the shear mode is quantified by dimensionless parameter by $\epsilon = \frac{B}{Sr^2}$, the scaled ratio between the bending stiffness and the shear stiffness. In the $\epsilon \to 0$ limit we recover the unshearable result, finding the (appropriately scaled) critical length $L_c \approx 1.36$. The figure shows both the results from the full simulation and the results from a linear stability calculation. Included in the figure are a number of pictures showing the rest configuration of a filament with length and ϵ corresponding to the location of the filament's base on the graph. The shear is illustrated by the green arrows, which give the normal to the cross-section along the filament. In (b), we compare the well-known case of unshearable rods of different lengths bending under gravity ($\epsilon = .001$, vertical: L = 1.35, left side: 1.36, 1.4, 1.6, 2.8) with the new case of highly shearable rods ($\epsilon = .25$, right side: L = 1.21, 1.25, 1.36, 1.6, 2.5). The redness of the rod represents the magnitude of the shear vector.

taking variations of this energy with respect to y_s and y_b, integrating by parts and then applying the fundamental theorem of variational calculus, we obtain the following coupled equations

$$By_b''''(x) + (\rho(L-x)(y_b' + y_s'))' = 0$$
$$-(Sy_s')' + (\rho(L-x)(y_b' + y_s'))' = 0$$

with boundary conditions $y_x(0) = \mathfrak{d}_x y_b(0) = \mathfrak{d}_{xx} y_b(L) = \mathfrak{d}_{xxx} y_b(L) = 0$ and $y_s(0) = \mathfrak{d}_x y_s(L) = 0$.

After subtracting these equations and integrating, we find that $y_s' = -\frac{B}{S} y_b''' = -\epsilon y_b'''$. This lets us reduce the system to one equation $y_b'''(x) + \frac{(L-x)y_b'(x)}{\frac{B}{\rho} - \epsilon(L-x)} = 0$. As $\epsilon \to 0$ we recover well-known case of the unshearable beam [56] with critical buckling length $L_c = \left(\frac{9BJ_0^2}{4\rho g \pi r^2}\right)^{\frac{1}{3}}$ where J_0 is the first zero of the Bessel function of order $-\frac{1}{3}$. For non-zero ϵ we must use a numeric method to solve the problem. Taking $z = y_b'$ we are left with a second order eigenvalue problem, which we then solve using standard finite difference methods. The results from this calculation and the full numerical model are displayed in 4.2.2. The match is quite good, lending support to the both our stability calculation and the model's treatment of shear.

Having gained confidence in our treatment of shear, we can move on to focusing on the model's treatment of extensibility. To accomplish this we numerically replicate the experiment in [35], where the phenomena arising are critically dependent on the extensibility of the rod. An extensible elastic rod hung is from a platform with a weight rigidly attached to the bottom, causing it to stretch a certain amount. We then rotate the weight a chosen amount and observe the response of the rod. There are four qualitiatively different responses: remaining straight, forming a plectoneme, forming a solenoid, and forming a combination plectoneme-solenoid.

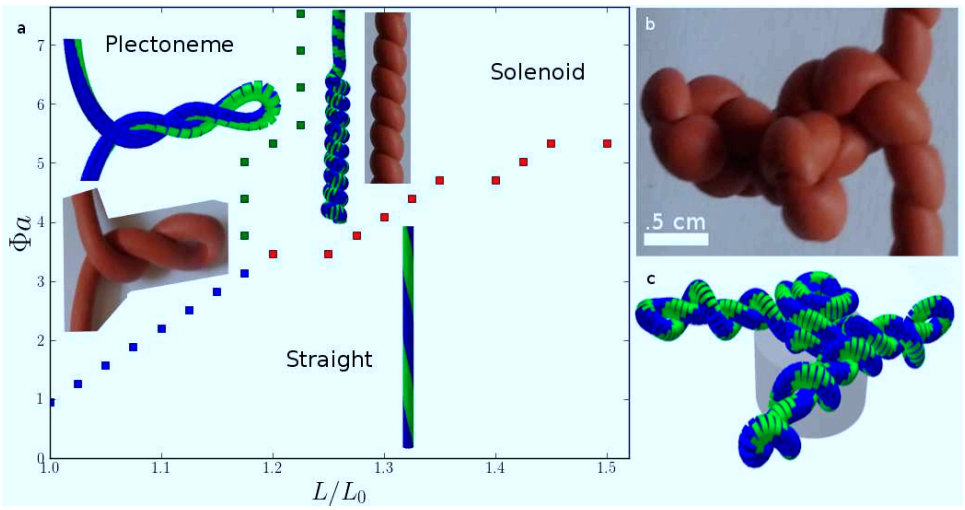

Figure 4.2.3: A simple twisting experiment (a) was performed, using a weight to prestretch a radius a =1/16", 2 MPa Young's modulus silicon rubber filament. When unstretched or lightly stretched the filament forms a plectoneme shape above a critical twist (shown: $\lambda \approx 0$ and the twist density $\Phi a \approx 0.6$). Above a critical initial stretch the solenoid deformation mode (shown: $\lambda \approx 0.1$, $\Phi a \approx 1.5$) appears instead of the plectoneme, where tightly packed coils tend to form at one end of the string and steadily proliferate to cover the entire string as the twisting continues. (b) and (c) give examples of a fourth phase not shown in (a): a combination plectoneme-solenoid which arises if a pully packed solenoid continues to be twisted (shown: $\lambda \approx 0.1$, $\Phi a \approx 3.0$).

We see that there are three distinct regions in the phase graph, each corresponding to a particular configuration; moreover, the graph is qualitatively identical to the corresponding graph in [35]. The region corresponding to the plectoneme-solenoid phase is off the scale in this graph, occuring at twist densities much higher ($\Phi a \approx 20$) than the region shown. Real-world and numerical examples of plectoneme-solenoid phase configurations are shown in figs. 4b and 4c.

We can also consider the solenoid/plectoneme transitions in the context of the CFW theorem. Although our model does not explicitly use the theorem, it can still help us understand the behavior from another perspective. Consider a setup similar to that in [35], except that instead of fixing the load, we instead fix the axial extension. Hence, if as we apply twist the rod deforms from a striaght configuration, it must be increasing its length. By rotating the ends of the rod, we are externally setting the link of the rod. The rod will then naturally choose a twist and writhe that minimizes its energy. The key idea here is that as the twist gets sufficently high the rod is willing to stretch and bend, and pay the resulting energy cost, to increase its writhe, which in turn decreases the twist as the CFW theorem requires.

Looking at 4.2.4, we see that the rod remains straight up until a critical twist, at which point it deforms into a single loop shape. Note how the twist decreases, resulting in the corresponding jump in the energy savings even though the rod had to stretch to allow the loop. As we continue to add link, the deformation migrates to the bottom, forming a more complicated single loop structure until near the end it transitions into a solenoid-type deformation, again reducing the twist and increasing the energy savings. If we continue twisting, this solenoid phase will add more and more loops until it entirely consumes the straight region.

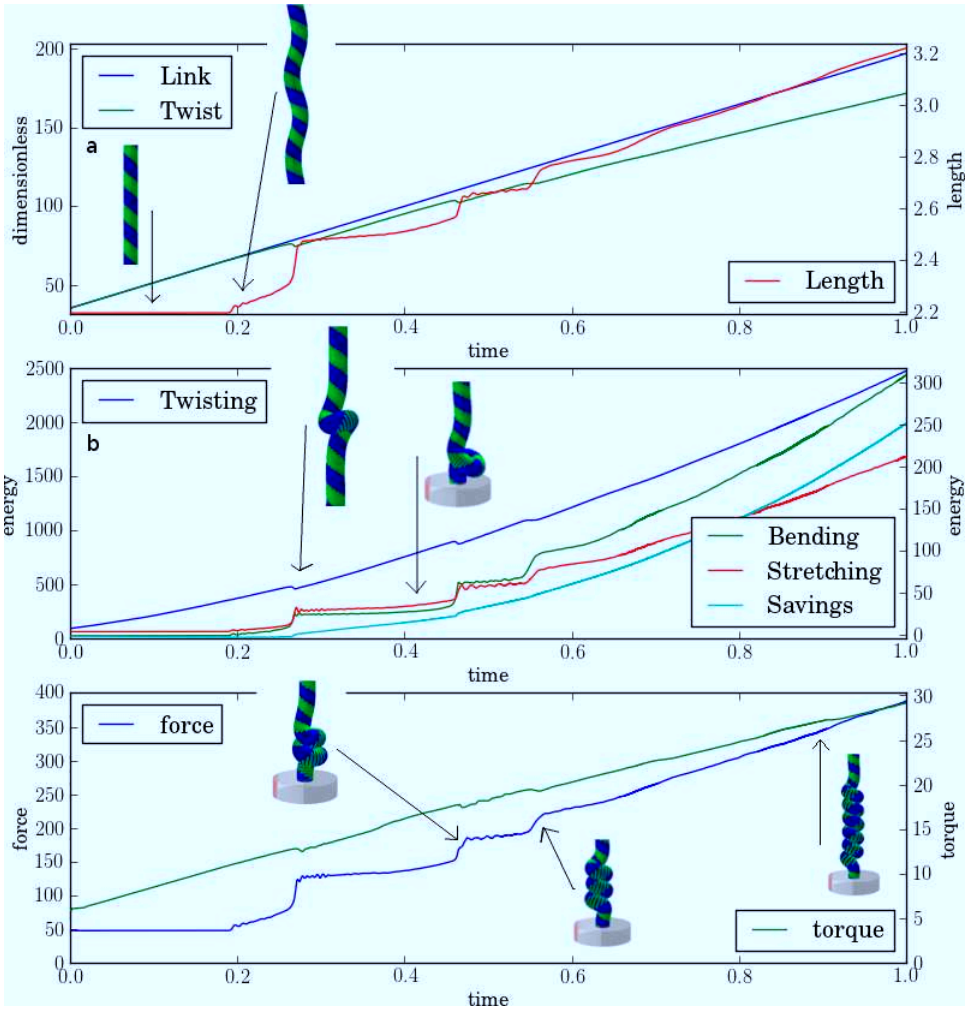

Figure 4.2.4: In this version of the twisting experiment instead of hanging a weight from the rod we fix the axial extension to $\lambda = 1.1$ and rotate both ends at a constant rate. In (a) we show the link, twist, and length of the rod over time; note that with fixed axial extensions, any deformations necessarily cause stretching. Fig. (b) shows the bending, twisting, and stretching energies of the rod, as well as the energy savings, defined as the what a the energy a straight, underformed rod would have minus rod's actual total energy. The basic idea is that when the twist becomes sufficiently high, the rod can lower its overall energy by converting some twist into writhe: the reduction in twisting energy outweighs the increase in bending energy and stretching energy. In (c) we show the external force and torque required at the endpoints. In each of the figures we illustrate the various phases the rod travels through as the link increases.

5

Entanglement

We've probably all experienced the following problem: whenever you put your earbuds into your pocket and walk around, they inevitably seem to come out all tangled together. Less familiar, but rather more important, corrolaries of this problem abound in biology and medicine [25, 27, 28, 65, 68]. Understanding the dynamics of the knotting process, and how they vary as a function of the relevant parameters, could have substantial practical utility. Yet, despite over a century of work in knot theory [9, 60], little is known about physical knotting processes, and purely analytic approaches appear untractable.

5.1 Simulating Entanglement

Motivated by previous work in the area [73], for which Raymer and Smith won an Ig Nobel prize, here we computationally simulate the following experiment: a homogeneous, inextensible filament of length L, radius r, bending stiffness α and mass per unit length ρ is droped into a box with dimensions 1 x 1 x L and then randomly shaken with amplitude A and frequency ω. After a set amount of time we pull the ends of the rope in opposite directions and determine whether a knot has formed. provides snapshots of the process and gives examples of the types of knots that can arise.

Note that, unlike previous experiments/simulations with vibrating chains [5, 48] or self-avoiding random walks [29, 70, 79] we are simulating a full physical system. This necessarily means that there will be many more parameters here than in the other experiments. One of our goals in this paper will be to discover the key effective parameters that will allow us to make general conclusions about knotting in this same fashion as the previous experiments do.

Determining the existance of and/or classifying a knot is non-trivial. While methods exist to exactly classify knots [73], they tend to be computationally expensive and difficult to implement. A simpler method is based on projecting a knot onto a plane and counting the number of crossings. Averaging over all possible planes gives the average crossing number (ACN) [12]. While not itself a topological invariant, the ACN can be directly computed as

$$ACN = \frac{1}{4\pi}\int_0^L \int_0^L \frac{|(x'(s) \times x'(t)) \cdot (x(s) - x(t))|}{|(x(s) - x(t))|^3}\, ds\, dt \qquad (5.1)$$

and gives a upper bound on the minimum crossing number (MCN), which *is* invariant but is not directly computable. Note that the ACN

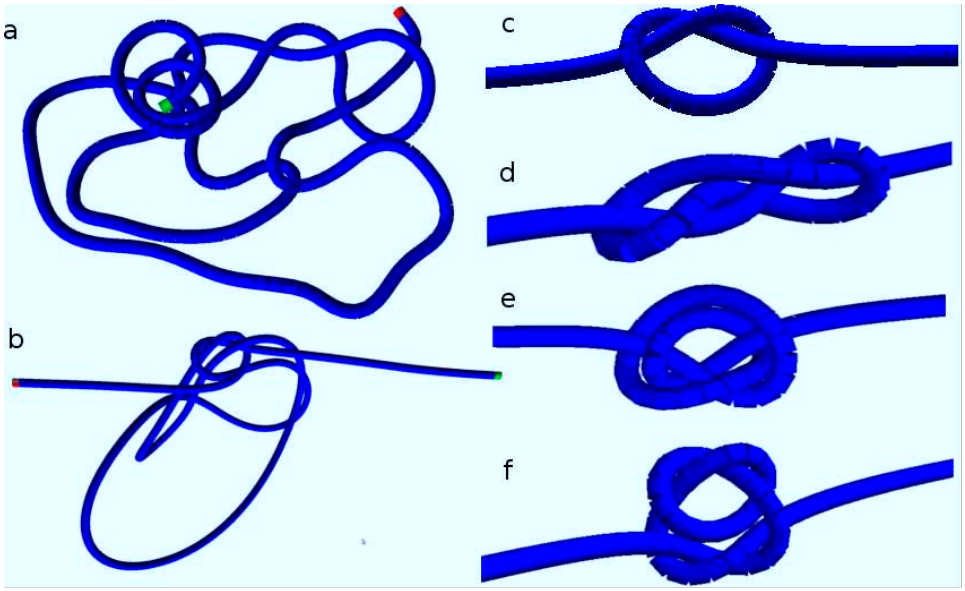

Figure 5.1.1: A homogeneous, inextensible elastic filament is dropped into a square bin and then shaken randomly for a set period of time. In (a) we see a typical configuration. Gravity is present and tends to keep the rope near the bottem of the bin so the configuration is largely planar. After the shaking period ends the container dissapears and the ends are pulled in opposite directions (b) to determine if there are any knots. ((b), surprisingly, is unknotted.) Many, many kinds of knots arise from this procedure; (c)-(f) show examples.

of an unknotted rope that is bunched together tightly in space will be non-zero; we pull the rope taught so as to eliminate as many spurious crossings as possible and bring the ACN as close as possible to the lower MCN bound.

We will return to the ACN shortly, but to begin we instead consider the rope to be in one of only two states: unknotted and knotted (defined as ACN >2). To keep things as simple as possible we consider the steady-state limit where the rope has been tumbling for an amount of time much larger than any relevant timescale. (This was confirmed by doubling the tumbling time and getting identical results.) Hence the relevant variables in the problem reduce to A, α, ρ, and $l = \frac{L}{r}$, the dimensionless ratio between the filament length and radius. We're interested in the probability that the rope ends in a knotted state as a function of those four variables. To approach this rather difficult problem, we can argue that the knotting behavior of the rope is governed largely by two energies: the bending energy of the rope and the kinetic energy enduced by the shaking. Balancing these two energies in scaling form gives $\alpha\kappa_e^2 \approx \rho(A\omega)^2$ where $(A\omega)$ is taken as the velocity scale. Rearranging this equation gives the energetic length scale as $r_e = \frac{1}{\kappa_e} = \sqrt{\frac{\alpha}{\rho}}\frac{1}{A\omega}$. To the extent that the energy balance determines the behavior of the system, we expect that the dimensionless geometric parameter $\frac{l}{r_e}$ should fully determine the knotting probability.

To test this expectation we varied A, α, ρ, and l, performing 100 trials for each tuple. We summarize the results in 5.1.2, which shows the probability of a knot forming for each type of rope at the given effective length. The effective length is defined as the ratio $\frac{l}{r_e}$ minus a constant γ, where γ represents the fact that there is a minimum length of rope below which it is not possible to form a knot. (γ was determined by reducing rope length until no knots ever formed.) The collapse is relatively good considering that we reduced four parameters

into a single one using a simple scaling argument. Empirically we see that the knot probability initially increasing linearly with effective length before gradually asymptoting to 1 as $l_e \to \infty$. The simplest form which accomodates this behavior is an exponential such that $p_k(l_e) = 1 - e^{-\zeta l_e}$, where $\zeta \in \mathbb{R}$ is found from the slope.

We can also look at the knotting process as system in two states, unknotted (state u) and knotted (state k), with a transition rate τ_1 from unknotted to knotted and τ_2 from knotted to unknotted. In steady-state detailed balance implies that the observed probability of being knotted is related to the transition rates via $p_u(l_e) = \frac{1}{\frac{\tau_2}{\tau_1}+1}$. Therefore, knowing $p_u(l_e)$ as we do gives us $\frac{\tau_2}{\tau_1}$ as a function of the effective length l_e, but doesn't tell us anything about τ_1 or τ_2 individually.

To shed some light on τ_1, we move from the steady-state regime to the short-time regime where we look for the first time that a knot has formed during each trial. To do this, every 5000 timesteps we pull the string straight to check if its knotted. If so, we stop and record the time; if not, we return the rope to the configuration it was in before pulling and continue tumbling. Note that this is impossible, or at least highly impractical, to do in a physical experiment setting. Here we chose a single A, α, ρ tuple and varied l, performing 100 runs for each length. In 5.1.3(a) we show the mean first transition time versus effective length. The mean is fit quite well by $\mu(l_e) \propto \frac{1}{l_e}$. Since the transition out of the unknotted state takes the form of an exponential random variable with parameter $\tau_1(l_e)$ and the mean of an exponential r.v. with parameter θ is $\frac{1}{\theta}$, this implies that $\tau_1(l_e)$ is linear in l_e: $\tau_1(l_e) = \lambda l_e$. The linearity is further confirmed by 5.1.3(b), where we plot the cumulative distribution functions for the first transition time, where time has been scaled by l_e.

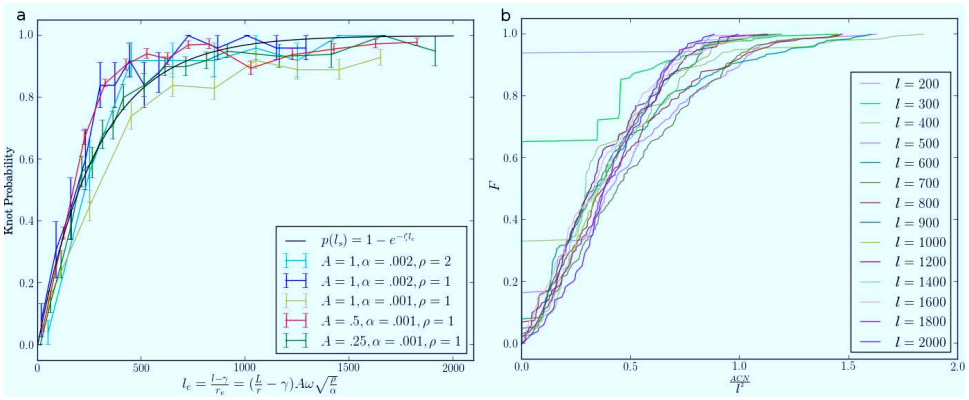

Figure 5.1.2: We claim the fundamental scaling of the system arises from the balance between the bending energy in the rod and the kinetic energy imparted by the shaking motion, leading to the energtic length scale $r_e = \sqrt{\frac{\alpha}{\rho}}\frac{1}{A\omega}$ where A is the shaking amplitude, ω the shaking velocity, ρ the mass per unit length, and α the bending stiffness. To test this scaling we varied A, α, and ρ. For each set of parameters we varied the filament length and performed 100 trials at each length, shaking until a steady state has been reached. At the end of each trial we record whether the filament is knotted, which we define as a ACN over 2. In (a) we show the results where the x-axis has been scaled by the energetic length scale. The collapse is relatively good, and is fit well by a exponential $p(l) = 1 - e^{\zeta l}$, a consequence we explore in the text. In (b) we've recorded the actual ACN from each trial at a single tuple of A, α, and ρ, rather than only whether the filament is knotted. The results are displayed as cumulative distribution functions where the x-axis has been scaled by l^2. The resulting collapse implies that indeed the average ACN scales as l^2, which we attribute to multiple prime knots forming over time and moving together.

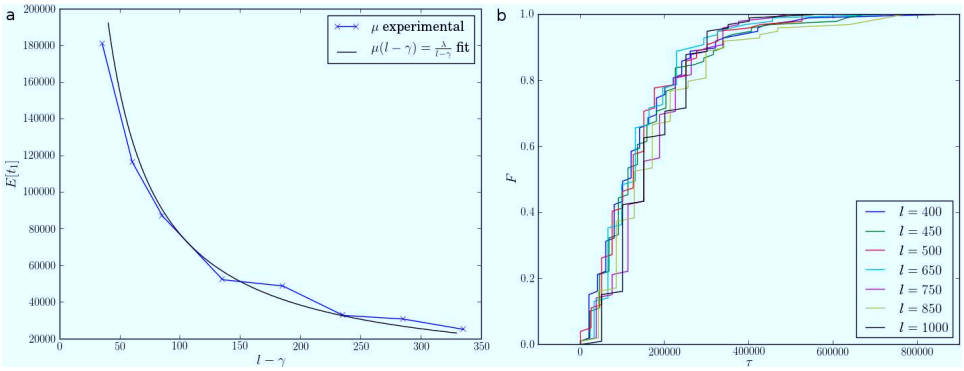

Figure 5.1.3: Whereas before we considered the steady-state behavior of the tumbled filaments, here we are interested in determining when the first transition occurs to a knotted state. Here we fix A, α, ρ and vary l, performing 100 trials per length. Within each trial every 5000 timesteps we pull the string straight to check if its knotted. If so, we stop and record the time; if not, we return the rope to the configuration it was in before pulling and continue tumbling. In (a) we show the mean first transition time as a function of effective rope length. If we suppose the first transition time is distributed as an exponential random variable with parameter λl, where $\lambda \in \mathbb{R}$, then the expectation time for the first transition is $\frac{1}{\lambda l}$, which we see fits the data well. In (b) we consider the cumulative distribution function for the first transition time. Rescaling time by $\tau = \frac{t}{l}$ we get a good collapse of the cdfs, as we would expect if indeed the parameter of the exponential variable scales with l.

Combining our empirical fit for the steady-state knotting probability with our formula from the state-transition picture, we find $1 - e^{\zeta l_e} = \frac{1}{\frac{\tau_2}{\tau_1}+1} \rightarrow \tau_2 = \frac{\lambda l_e}{e^{\zeta l_e}-1}$. This is actually a rather interesting result, as it implies that for small l_e the unknotting rate is constant ($= \frac{\lambda}{\zeta}$) and exponentially shrinks as l_e increases. A possible physical interpretation of this is that there is a constant entropic process that tries to unknot the system, but as length increases the unknotting becomes blocked by the physical constrant that the rope cannot self-intersect.

So far we've focused on the binary question of whether or not a knot exists; we now want to extend our thinking to include the complexity of the observed knots, which we quantify by their ACN. Consider 5.1.4, which shows a particular rope's ACN means versus length in the short time scale (a) and the steady-state (b). For the shorter lengths in the short time-scale we see that the average ACN is well fit by a line. This makes sense considering that we have established that the initial knotting transition has rate linearly proportional to the effective rate, so in a small period of time it is likely only one knot has formed. However, for longer strings we see that the mean ACN deviates considerably from the low-length trend line and begins to have an increasingly quadratic shape. The quadratic behavior arises from two effects. The first the probability of having more than one knot at a short time is much higher when the transition rate is higher (i.e. l_e is larger). The second is that seperate knots have a tendency to cluster together, as you can see from the following experiment: take a string and tie two trefoil knots in it at different places on the string. Then take the ends and begin to pull on and shake them around randomly. You'll likely find that the the knots come together and form a single cluster. This confluence of the knots explains the quadratic scaling since the number of knots we expect scales like l_e from τ_1 but, assuming they come together, the ACN

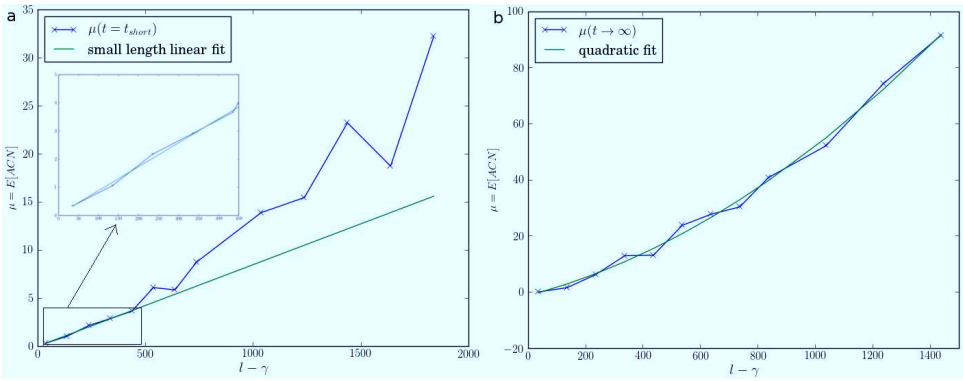

Figure 5.1.4: We've looked at steady-state behavior and short-time behavior; we can also consider the transition between the two. Ideally we would look at the temporal evolution of the ACN distribution functions, but overcoming sampling error would require prohibitively long computation times. Instead, we consider the distribution of ACN versus length at an intermediate time (a), and steady state (b). For small lengths in (a) we see that the ACN is linear, while at longer lengths increasingly runs above the short-length trend line. This contrasts with (b), where the quadratic dependence of the ACN on l appears evident at all lengths. A simple explanation for these results is as follows: at short times, the low knot formation rate of the small ropes means it is highly unlikely that more than one knot will have formed (and hence the ACN $\approx l$), whereas for the longer ropes it grows increasingly likely. Under the assumption that the multiple knots tend to clump up together, multiple knots will lead to a ACN $\approx l^2$ scaling. In (b) even the shorter filaments have had time to form multiple knots, hence their ACN too scales like l^2.

scales like the number of knots squared. At long times, however, even the smaller ropes with their lower knotting rates have had sufficient time to form multiple knots, and hence their ACN also scales like l_e^2. This conclusion is further corroborated by fig. 2b, which shows that indeed the cumulative distribution functions of the ACN collapse when scaled by l_e^2.

6

Adhesion

A natural next addition to the model which has many applications is adhesion, a phenomena relevant to many biological and mechanical systems (TODO). We model adhesion as some sort of potential energy that is a function of just the distance between the interacting bodies. In Rod-Rod adhesion we can use the collision model's minimum distance algorithm to find how close the two cylinders are, so applying the forces is a simply matter of choosing the energetic form. Alternatively, we can have the rod adhere to some external surface, such as a plane, as we do in the remainder of this chapter.

6.1 Hydrodynamic Adhesion

Consider a one–dimensional elastic plate immersed in an ambient fluid and is attracted to a substrate through a short–range potential, as shown in the schematic in figure 6.1.1. Given the properties of the plate and of the intervening fluid, as well as the interaction potential, we have to determine the distance of the plate from the substrate $h(x,t)$, where x is the coordinate along the length of the plate and t is time. We assume the interaction potential to be described by a generalized Lennard–Jones function

$$\Phi(s) = 4\left(\frac{1}{s^{2m}} - \frac{1}{s^m}\right), \tag{6.1}$$

with a parameter m. This form for the potential enters the dynamics of the bending of the plate as

$$p(x,t) = Bh_{xxxx} + \frac{A}{\epsilon}\Phi'\left(\frac{h}{\epsilon}\right), \tag{6.2}$$

where B is the bending stiffness of the plate, A is the adhesion energy per unit length between the substrate and the plate, ϵ is the interaction distance and p is the hydrodynamic pressure in the thin gap between the plate and the substrate. The potential has a minimum as $2^{1/m}$, where Φ takes the value -1 and rapidly decays to zero for $h \gg \epsilon$. Typically, the adhesive interaction between the plate and the substrate is short–ranged, implying that ϵ is much smaller than typical length scale in the problem. Thus it is desirable to label the region over which Φ is appreciably non–zero as the contact region and formulate effective conditions to be applied at the edge of this region called the contact point. This is the central goal of this letter. We see that the condition depends not only on whether the situation

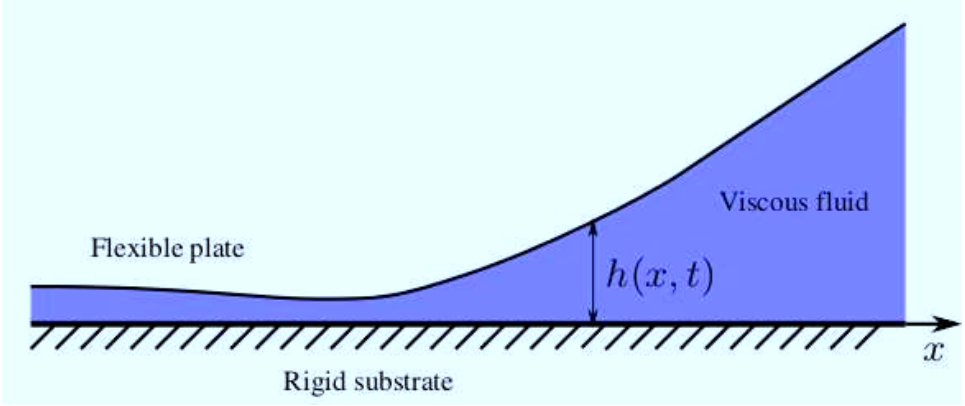

Figure 6.1.1: Schematic setup for a flexible plate adhering to a substrate in the presence of an intervening fluid layer.

is static or dynamic, but also on the nature of the dynamics.

6.2 Static Case

We start with the static case, described by Landau and Lifshitz, because of its simplicity. In the static case the fluid pressure $p \equiv 0$ and thus (6.2) simplifies to the ODE

$$Bh_{xxxx} + \frac{A}{\epsilon}\Phi'\left(\frac{h}{\epsilon}\right) = 0. \tag{6.3}$$

This can be cast into a variational form as equivalent to minimizing the total energy

$$E[h] = \int_0^L \frac{Bh_{xx}^2}{2} + A\Phi\left(\frac{h}{\epsilon}\right) \mathrm{d}x, \tag{6.4}$$

where the plate extends from $x = 0$ to L. Clearly, if h is dynamically free at the ends of the plate (i.e. $h_{xx} = h_{xxx} = 0$ at $x = 0, L$), then the minimum in $E[h]$ occurs for $h \equiv 2^{1/m}\epsilon$, meaning that the whole plate is in contact with the substrate. To eliminate that possibility, we hold

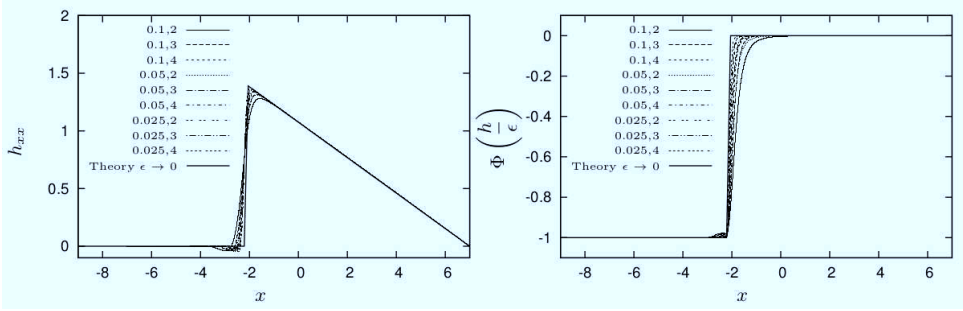

Figure 6.2.1: Curvature of the plate (left) and the adhesive potential (right). Legend shows various parameter pairs (ϵ, m).

the plate a distance $h_{\max}$ away from the substrate at the right end $x = L$, with the left end free. The details of the right boundary condition are not important for our discussion as long as it leads to a contact region. It is easy to imagine holding the plate so far away from the substrate, or applying such a large force on the plate at that end, so that the plate completely loses contact with the substrate; we assume in our discussion that such is not the case. Our strategy in this letter is to consider smaller and smaller values of ϵ, solve (6.3) numerically and analyze the ensuing limit $\epsilon \to 0$. In particular, we consider the right end to be hinged (i.e. $h_{xx} = 0$) at $h_{\max} = 40$ for a plate of length $L = 14$.

The curvature of the plate h_{xx} and the adhesive potential Φ from the numerical solutions of (6.3) for various ϵ and m are shown in figure 6.2.1. A small region near $x \approx -2$ develops to the left of which the curvature is zero and $h \approx 2^{1/m}\epsilon$ and to the right of it $\Phi(h/\epsilon) \approx 0$. Using dominant balances, the length of this region can be estimated to be $O((B\epsilon^2/A)^{1/4})$, while $h = O(\epsilon)$. As $\epsilon \to 0$, this region get narrower and narrower, eventually reducing to a point $x = x_c$, which defines the point of contact.

The effective conditions at the point of contact may be derived

from the variational principle (6.4) as follows (Landau & Lifschitz). In the energy integral, the bending term derives its value from the region $x > x_c$, while the adhesion term is non–zero only for $x < x_c$. Thus the integral can be split into

$$E[h] = \int_0^{x_c} A\Phi\left(\frac{h\ell}{\epsilon}\right) \mathrm{d}x + \int_{x_c}^{L} B\frac{(\partial_{xx}h)^2}{2}\mathrm{d}x. \tag{6.5}$$

Variations with respect to h of this energy lead to $h = 2^{1/m}\epsilon \to 0$ for $x < x_c$ and $Bh_{xxxx} = 0$ for $x > x_c$. At $x = x_c$, using the scalings with ϵ in the inner region (i.e. $h = O(\epsilon)$ and $\partial_x = O(\epsilon^{-1/2})$), $h = h_x = 0$ but h_{xx} approaches a finite value. This value can be determined by applying variations with respect to x_c in (6.5), and the extra condition provides the value of x_c. Perturbing $x_c = x_{c*} + \delta x_c$, where x_{c*} corresponds to the minimum and δx_c is a test perturbation, we can write the resulting perturbation in E as

$$\delta E = -\delta x_c \left(A + B\frac{h_{xx}^2}{2}\right) + \int_{x_c}^{L} B\delta h_{xx}h_{xx}\ \mathrm{d}x, \tag{6.6}$$

δh being the induced perturbation in h owing to the perturbation in x_c. δh satisfies $\delta h_{xxxx} = 0$ with $\delta h + \delta x_c h_x = 0$ and $\delta h_x + \delta x_c h_{xx} = 0$ at $x = x_c$, while $\delta h = \delta h_{xx} = 0$ at $x = L$. Simplifying the bending integral in (6.6) by parts and using the boundary conditions on δh leads to $\delta E = \delta x_c(Bh_{xx}^2/2 - A)$. Setting this first variation in δx_c to zero leads to the bending moment condition

$$Bh_{xx} = \sqrt{2AB}. \tag{6.7}$$

An analytical solution can now be obtained in the limit $\epsilon \to 0$:

$$h = \sqrt{\frac{A}{2B}}(x - x_c)^2\left(1 - \frac{x - x_c}{3(L - x_c)}\right), \tag{6.8}$$

where

$$(L - x_c)^2 = \frac{3h_{\max}}{2}\sqrt{\frac{2B}{A}}. \tag{6.9}$$

The numerical limiting procedure is observed to approach this solution, as shown in figure 6.2.1.

6.3 Dynamic Case

Does this condition change in the dynamic case and how? To answer this question, we modify the system slightly; we consider a plate initially inclined to the substrate with slope α with its left end adhering to the substrate. Namely, $h(x,0) = 2^{1/m}\epsilon + \alpha x$. The plate is attracted towards the substrate, but it is resisted by the interveneing fluid that has to drain. We model this drainage by a lubrication approximation, exploiting the thinness of the film compared to the x-length scale. Summarily, this approximation implies that the hydrodynamic pressure $p(x,t)$ satisfies the approximate x-momentum balance $\mu u_{yy} = p_x$, where μ is the fluid viscosity and $u(x,y,t)$ is the x-component of the fluid velocity field, y being the coordinate normal to the substrate. The fluid in the gap is incompressible $u_x + v_y = 0$, where $v(x,y,t)$ is the y-component of the velocity field. These two equations, along with the kinematic boundary condition $h_t + uh_x = v$, stating that the fluid velocity at the plate matches with the velocity of the plate, leads to

$$12\mu h_t = \left(h^3 p_x\right)_x. \tag{6.10}$$

We use $p = 0$ at $x = 0, L$ as boundary conditions applying to (6.10). The system under consideration is now equations (6.2), (6.10) with the potential Φ in (6.1) subject to boundary conditions

$h_{xx} = h_{xxx} = p = 0$ at $x = 0, L$. The parameters in the system are A, B, μ, ϵ, α, m and L. The number of parameters can be reduced by non-dimensionalizing the system using the length scale $\ell = \sqrt{B/A}$ for x and h, the time scale $\mu\ell^3/B$ for t, and the scale B/ℓ^3 for pressure. This simplifies the system to

$$p = h_{xxxx} + \frac{1}{\sigma}\Phi'\left(\frac{h}{\sigma}\right) \quad \text{and} \tag{6.11}$$

$$12h_t = \left(h^3 p_x\right)_x, \tag{6.12}$$

where $\sigma = \epsilon/\ell$ is the non–dimensional adhesion length scale. We choose representative values of the other parameters α, m, σ and the dimensionless L and time–march (6.11-6.12) numerically starting from the initial condition $h(x,0) = 2^{1/m}\epsilon + \alpha x$ and $p = 0$ to get a preliminary idea of the ensuing dynamics. Figure 6.3.1 show the results; a dynamic contact zone forms with the plate making contact to the left of the point, i.e. $\Phi \neq 0$ only to the left of the zone. This zone moves with a constant speed to the right. It is also numerically observed that all the fluid displaced from underneath the plate in the process is accumulated in a bulge immediately to the right of the contact zone. Moreover, the shape of the deformed plate to the right of this zone at various times appear self–similar. This prompt us to look for a solution of the form.

$$h(x,t) = t^\beta f(\eta), \quad p(x,t) = t^\kappa g(\eta), \quad \eta = \frac{x - ct}{t^\gamma}. \tag{6.13}$$

The exponents β, κ, γ are determined using the governing equations and volume conservation. We assume $\gamma < 1$, subject to subsequent verification, so that the time derivative in (6.12) is approximated as $h_t = -t^{\beta-\gamma}cf' + t^{\beta-1}(\beta f - \gamma\eta f')$ only by the first term for large t. The governing equations (6.11-6.12) to the right of the contact zone

where $\Phi \approx 0$ then give

$$g = f'''', \quad \kappa = \beta - 4\gamma, \tag{6.14}$$

$$-12cf' = (f^3 g')', \quad \beta - \gamma = 3\beta + \kappa - 2\gamma. \tag{6.15}$$

Also the accumulated volume in the bulge is $\alpha(ct)^2/2$, giving

$$\int_0^\infty f(\eta)\mathrm{d}\eta = \frac{\alpha c^2}{2}, \quad \beta + \gamma = 2. \tag{6.16}$$

This set of three equations in three unknowns gives $\beta = 5/4$, $\kappa = -7/4$ and $\gamma = 3/4$. Moreover, (6.14–6.15) can be simplified further to yield

$$f^2 f'''' = -12c. \tag{6.17}$$

The function f is plotted in figure 6.3.1 to verify our similarity hypothesis and agreement can be observed. One useful feature of this similarity solution is that the small η asymptotics can be analytically derived to be

$$f \sim kc^{1/3}\eta^{5/3}, \quad g \sim \frac{40}{81}kc^{1/3}\eta^{-7/3} \quad \text{for} \quad \eta \ll 1, \tag{6.18}$$

where $k = 9(70^{-1/3})$. Due to this scaling, close to the contact zone $h \propto t^{5/4}\left(\frac{x-ct}{t^{3/4}}\right)^{5/3} = (x-ct)^{5/3}$, which is purely steadily propagating. This allowed Rieutord, et al to derive the propagating speed without recourse to the similarity solution. This power law scaling is cut off by an inner scale $\delta = \sqrt{\epsilon}$ determined by taking $h = O(\epsilon)$, so that the adhesion potential is non–zero. Thus, it can be seen that as $\epsilon \to 0$, the outer solution satisfies $h = h_x = 0$ at the contact point, but h_{xx} grows like $\epsilon^{-1/6}$. So far the analysis is silent about the speed c, which we will derive using energy conservation and see to be also dependent

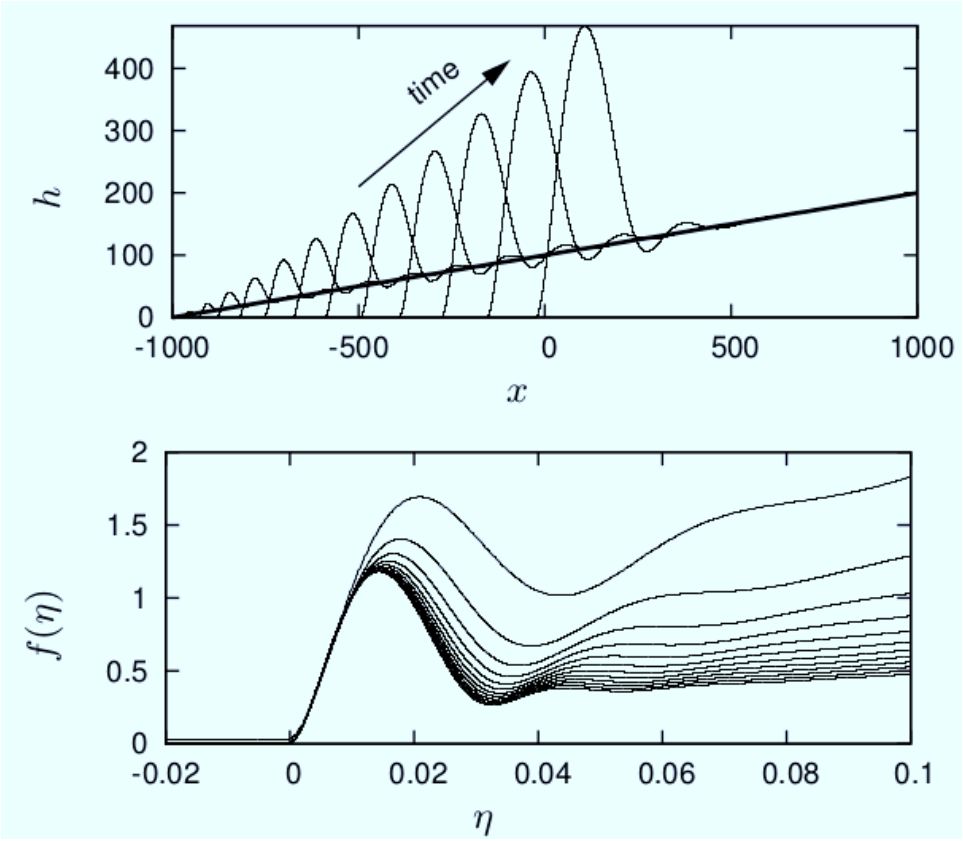

Figure 6.3.1: Dynamics of the bonding process starting from an inclined plate. Top panel shows snapshots in time of the shape of the plate for $\sigma = 0.1$, $L = 2000$, $m = 2$, $\alpha = 0.1$. Bottom panel shows the collapse of these shapes onto a universal self–similar curve described in (6.13).

on inner scale.

An analogue of the energy equation can be derived to find the speed c by multiplying (6.11) by h_t, multiplying (6.12) by p and subtracting the two results to get

$$\frac{\mathrm{d}}{\mathrm{d}t}\int_0^L \frac{h_{xx}^2}{2} + \Phi\left(\frac{h}{\sigma}\right)\mathrm{d}x = -\int_0^L \frac{h^3}{12}p_x^2\ \mathrm{d}x. \tag{6.19}$$

Not only does this equation show that the dynamics evolve towards decreasing the total energy $E[h]$, but also gives a handle on the rate at which they happen. Substituting the similarity scalings in (6.19) yields the various terms to be

$$\frac{\mathrm{d}}{\mathrm{d}t}\int_0^L \frac{h_{xx}^2}{2}\mathrm{d}x = \frac{\mathrm{d}}{\mathrm{d}t}t^{1/4}\int_{\eta_{\min}}^{\infty}\frac{f''^2}{2}\ \mathrm{d}\eta, \tag{6.20}$$

$$\frac{\mathrm{d}}{\mathrm{d}t}\int_0^L \Phi\left(\frac{h}{\sigma}\right)\ \mathrm{d}x = -c, \tag{6.21}$$

$$-\int_0^L \frac{h^3}{12}p_x^2\ \mathrm{d}x = -t^{-1/2}\int_{\eta_{\min}}^{\infty}\frac{f^3g'^2}{12}\ \mathrm{d}\eta, \tag{6.22}$$

where $\eta_{\min} = O(\delta/t^{3/4})$ signifies the inner scale cut–off. The integrand in (6.20) diverges for small η, but the integrand is not only bounded but also approaches zero like $\epsilon^{1/6}$ as $\epsilon \to 0$. Thus, the gain in adhesion energy denoted in (6.21) is not spent on increasing the bending energy, but is instead dissipated through viscous effects in (6.22). The integrand in the last term is also divergent, and so is the integral. Hence for small ϵ, the integral derives its value from the lower limit. Substituting (6.18) in (6.22) yields the term to independent of time and proportional to $-18c^{5/3}/k\delta^{2/3}$ in magnitude. The constant of proportionality can only be found numerically by matching with the inner solution (and is independent of the particular short–range potential?); we find it to be $D = ...$ Substituting this back in (6.19)

gives the propagation speed to be

$$c = \left(\frac{k}{18D}\right)^{3/2} \delta = \left(\frac{k}{18D}\right)^{3/2} \epsilon^{1/2}. \tag{6.23}$$

Dimensionally,

$$c_{\text{dimensional}} = \left(\frac{k}{18D}\right)^{3/2} \frac{A^{5/4}\epsilon^{1/2}}{B^{1/4}\mu}. \tag{6.24}$$

In the context of synaptic contact, A and ϵ model the adhesive interaction between the neural synapses, B is the bending stiffness of the cell membrane and μ is the viscosity of intercellular fluid. Using representative values of the parameters ($A = 10^{-2}$ J/m^2, $\epsilon = 10^{-10}$ m, $B = 10^{-19}$ N m and $\mu = 10^{-3}$ Pa s), $c_{\text{dimensional}} =$

6.4 No Damping Case

We can also consider adhesion in the limit where there is no damping due to hydrodynamic terms or otherwise. In this case we will have that the power flowing in from the adhesive potential will continuously increase the kinetic and bending energy of the sheet:

$$\begin{aligned} E^{adh} &= -\int_0^L A\Phi(\frac{h}{\epsilon})\ ds \\ E^{bend} &= \int_0^L B\kappa^2\ ds \\ E^{kin} &= \int_0^L \rho \boldsymbol{v}^2\ ds \end{aligned} \tag{6.25}$$

where B is the bending stiffness, κ is the curvature (h_{xx} in the small angle limit), ρ is the mass per unit length, and $\boldsymbol{v}$ is the velocity of the centerline.

Our goal is to derive the scaling behavior of the system, in

particular the velocity of the contact point. Assume that the rod begins at an angle ϕ to the surface, such that $h(x,0) = \phi x$. Suppose the sheet has already fully adhered a distance d, has non-zero curvature in a boundary layer region of width ξ, and remains unperturbed from its original configuration for $x > d + \xi$, and we are in the limit $\xi \ll d \ll L$; see Fig. 4a. Then the kinetic and bending energy terms are non-zero only in the boundary layer region. We now make the assumption that the curvature scaling in the boundary layer will be that of the static limit, namely $\kappa_0 = \sqrt{A/B}$ and that the velocity in the x-direction is neglible. Then in the $\epsilon \to 0$ limit we have

$$\begin{aligned} E^{adh} &\approxeq -Ad \\ E^{bend} &\approxeq B\kappa_0^2 \xi \approxeq A\xi \\ E^{kin} &\approxeq \rho h_t^2 \xi \end{aligned} \tag{6.26}$$

To extract the steady state power balance we need the derivatives of these energies in time:

$$\begin{aligned} P^{adh} &\approxeq -Ad_t \\ P^{bend} &\approxeq A\xi_t \\ P^{kin} &\approxeq \rho(h_t^2 \xi_t + h_t h_{tt} \xi) \approxeq \rho h_t^2 \xi \end{aligned} \tag{6.27}$$

where the last equality is a dominance assumption that can be checked later.

Now, in the time it takes the contact point to travel a distance ξ the sheet has vertically traveled a distance $\phi(d + \xi) \approxeq \phi d$, so $\frac{\xi}{d_t} \approxeq \frac{\phi d}{h_t}$. Moreover, we have that the rod rises ϕd over a distance ξ, so we have $\kappa_0 \xi^2 \approxeq \phi d \to \xi \approxeq \sqrt{\phi d/\kappa_0}$. Hence we find $\xi_t \approxeq \sqrt{\phi/d\kappa_0} d_t$.

Using the above relations, and noting that the only consistent balance is actually between the adhesive and kinetic powers, we find

that

$$\begin{aligned} -Ad_t + \rho(\frac{dd_t}{\xi})^2\sqrt{\frac{\phi}{d\kappa_0}}d_t &\approxeq 0 \\ \Rightarrow d_t &\approxeq \frac{A^{\frac{3}{8}}B^{\frac{1}{8}}}{\rho^{\frac{1}{2}}\phi^{\frac{1}{8}}d^{\frac{1}{4}}} \end{aligned} \tag{6.28}$$

Intruigingly, the velocity scales inversely with d, so the contact velocity is actually slowing down: $d(t) \propto t^{\frac{4}{5}}$. We can also directly obtain the scaling laws for the x and y axes: $\bar{x} = xt^{\frac{-2}{5}} - d$ and $\bar{y} = yt^{\frac{-4}{5}}$. Under these transformations we expect the curve to be self-similar with respect to time.

To test these scaling relations we performed simulations using a rod model which is equivalent to our reduced-dimension sheet. 6.4.1(b) shows the height of a rod above the surface for several fixed times. 6.4.1(c) displays the collapsed curves after the scaling trasformations are applied. Another way to test the scaling law is to vary each parameter A, B, ρ independently and record the resulting contact point velocities. Fitting these velocities versus the parameters on a log-log scale gives then gives us the scaling (see the SI). Similarly, for a fixed A, B, ρ we can find d's power law dependence on t. The results of these experiments are:

quantity	analytic	numeric
d	.800	.805
A	.375	.303
B	.125	.139
ρ	.500	.500

We see that the results are reasonably good given our simple approximation to the full non-linear problem. That the scaling law does not match perfectly for A is unsurprising, since numerically we must use a finite length scale for the adhesion potential.

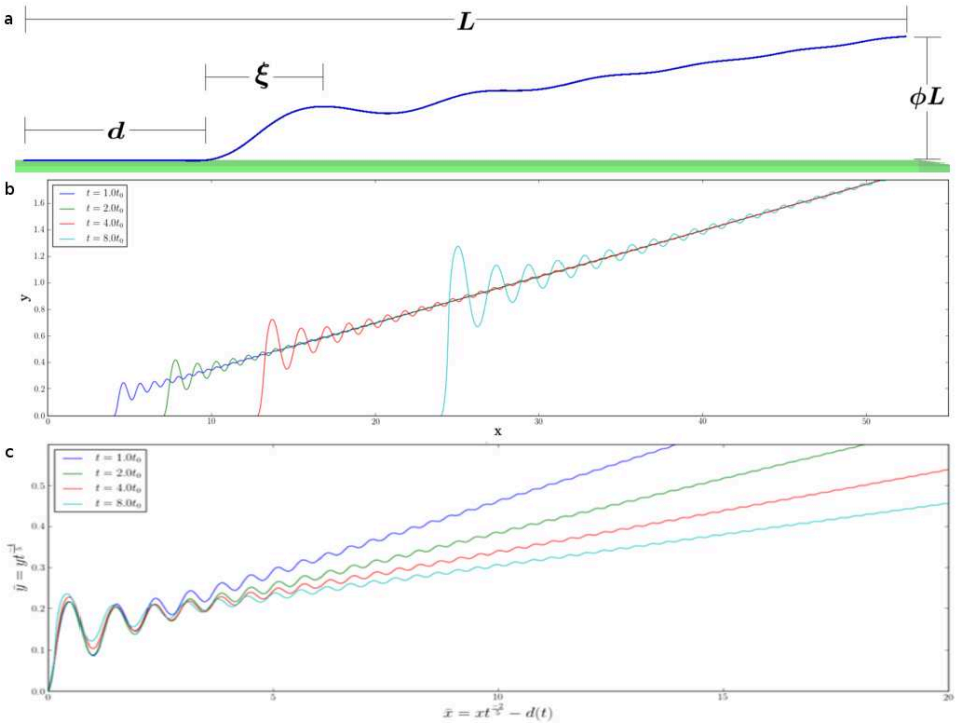

Figure 6.4.1: In (a) we show a frame from the rod model simulation of the no-damping adhesion case. We assume the rod has adhered a distance d, has a transition region of length ξ, and the remainder of the length L rod remains approximately in its original straight configuration oriented at angle ϕ to the horizontal. In (b), we show the rod height versus x for several fixed times. In (c) we translate the curves in (b) so that the origin is at the contact point and then scale the x and y axes as predicted by the scaling analysis. The collapse is quite good for the first few wavelengths beyond that contact point; as the scaled distace grows large the curves separate, as expected, because the scaling law only applies to the transition region.

7

Future Directions

In addition to the base model and the extensions previously discussed, the model has been built to handle a number of additional phenomena, including friction, active forces, and multiple filament systems. These phenomena are grouped here because although they have been implmented and tested, they have not been studied to the extent that the material in previous chapters has been. As such, each of these areas would make an excellent direction for future research.

7.1 Active Forces

The muscular contractions of snakes can be represented as an internal torque acting on the rod. These torques can be implemented in the model by using an appropriate custom ExternalForces 8.9 object.

While any arbitrary function could be used, a form that has biological relevance is to consider the contraction localized at some point on the snake and then falling off over some length scale σ. In particular, a natural form for a contractive force is

$$\boldsymbol{A} \exp \frac{s - s_0 - vt}{2\sigma^2} \tag{7.1}$$

which defines a contractive wave centered about the point with material coordinate s_0 and moving with a velocity v. We can have multiple forces acting at once, and in such a way build up any complicated motion we want.

7.1.1 Friction

While active forces could be handled with the tools already made, including friction requires some additions to our general model. Friction models can be very complicated and can require a lot of computational power [85]. Moreover, they generally are difficult to formulate as a term in the Lagrangian [74], so we instead describe the forces directly.. As with our collision model, we want to find the simplest possible model that has the features we need, namely the ability to account for anisotropy (e.g. snakes slide forward more easily than backwards), a distinction between kinetic and static friction, and a distinction between sliding and rolling friction.

First let's ignore rolling friction, which we can do we assuming that the rod is sliding on a flat surface in the direction of its tangent. The first things we calculate are the relative velocity between the rod and surface, $\boldsymbol{v_r}$, and the magnitude of the normal force preventing interpenetration, F_n. If the magnitude of the relative velocity is above a threshold v_{static} then we consider the friction to be kinetic, and otherwise static. In the kinetic case, we simply use the Coulumb

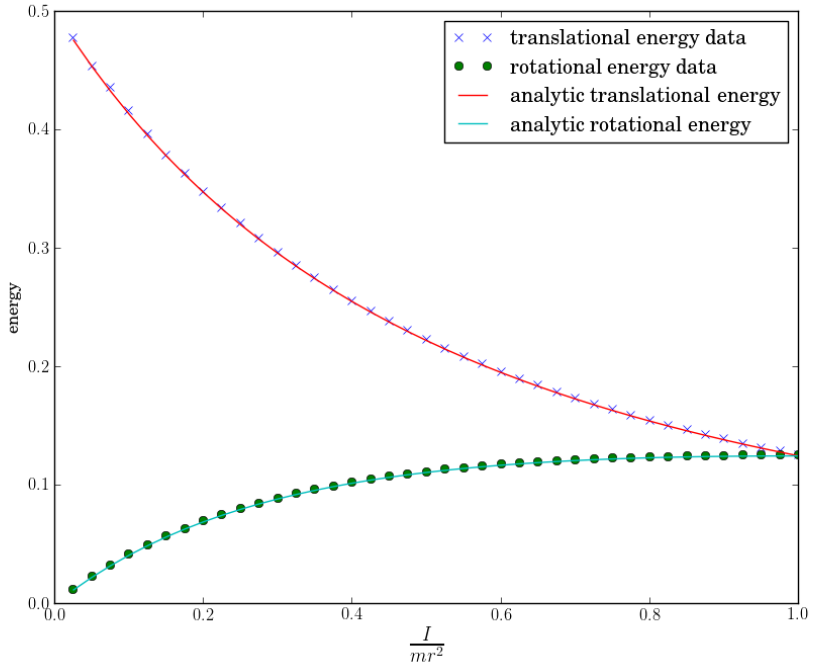

Figure 7.1.1: To test our model for rolling friction, we set a mass $m = 1$ cylinder translating at velocity $v(0) = 1$ across a frictional surface while varying the effective moment of inertia $\gamma = \frac{I}{mr^2}$. Using rolling steady-state condition that $wr = v$, we can derive that the final kinetic energies should be $U_{trans} = \frac{mv(0)^2}{2}\frac{1}{(1+\gamma)^2}$ and $U_{rot} = \frac{mv(0)^2}{2}\frac{\gamma}{(1+\gamma)^2}$. The model matches these precisely, so we can conclude our rolling friction model is working correctly.

model for friction, which gives the friction force as

$$\boldsymbol{F_f} = -\mu \hat{\boldsymbol{v}} F_n \tag{7.2}$$

If we wish to include anistropy, then we can general the friction coefficient from a scalar to a function of the angle between the material direction (say, the tangent) and the relative velocity. In the static case, instead of applying a force which acts to reduce the relative velocity, we instead record the maximal force the friction could produce, which would be $\mu_{static} F_n$, and then at the end of the timestep once we've added up all the forces acting on the cylinder if the magnitude of the other forces is less than the maximum possible static force, we set the forces to zero. In other words, static friction acts so as to reduce the other forces acting on the cylinder to zero (up to its maximum possible strength). If the static friction is not enough to reduce the forces to zero, then the cylinder will begin to move and generally speaking kinetic friction will take over at the next time step.

Adding in rolling friction now, we now consider the tangential velocity of the cylinder along the surface seperately from the radial direction. The tangential component is done just as we did before when we ignored rolling friction. For the radial component, we need to also take into account the rotation of the cylinder. The model we are using for the rolling friction is that friction acts so as to make the instanteous velocity at the point of contact be the same between the two surfaces. That is, a 1 meter cylinder translating to the right at 2π meters per second while rotating 1 rotating per second would not experience any frictional force. So, all we do is set the relative radial velocity to be the translational relative velocity minus the rotational velocity at the surface: $\boldsymbol{v}_{||} = (\boldsymbol{v_r} - \hat{\boldsymbol{t}} \boldsymbol{v_r} \cdot \hat{\boldsymbol{t}}) - r\boldsymbol{\omega} \cdot \hat{\boldsymbol{t}}$. If the magnitude of this velocity is above a threshold, we apply the frictional force as above for the static sliding block. If it is below, then we do nothing.

We can test our model for rolling friction via the following simple experiment: consider a rod with total mass m that begins with velocity $v(0)$ in the radial direction, and suppose that we vary a parameter γ which sets the relative moment of inertia: $\gamma = \frac{I}{mr^2}$. From our definition of rolling friction, we know that after friction has acted the steady state velocity and rotational speed will satisfy $wr = v$. We also have that if the frictional force f acts over a time period t the torque is $\tau = rf$ and we have that $w(t) = w(0) + \frac{rft}{I} = \frac{rft}{I}$ and $v(t) = v(0) - \frac{ft}{m}$. Combining these equations, we can arrive at the final energies we expect to have when we've reached rolling equilibrium as $U_{trans} = \frac{mv(0)^2}{2}\frac{1}{(1+\gamma)^2}$ and $U_{rot} = \frac{mv(0)^2}{2}\frac{\gamma}{(1+\gamma)^2}$. We performed a simple experiment where we varied γ and recorded the resulting energies; as 7.1.1 shows, the model matches the analytic calculations perfectly.

7.1.2 Trilinear Locomotion

With the models for active forces and anisotropic friction working, we were able to help explain the new form of snake locomotion discovered by Bruce Young, which we dubbed trilinear locomotion based off its unique shape. Our brief description of this motion finishes off this chapter.

Snakes exhibit a wide variety of forms of locomotion, from periodic modes such as undulation, sidewinding, and rectilinear, to transient modes such as concertina and striking. In this paper we describe a new transient form of motion observed in juvenile yellow anacondas (Eunectes notaeus) which we term trilinear locomotion. This mode gives the snake a very fast forward velocity, often used defensively to escape a threatening situation.

The starting position for the motion is a s-shape where the snake aligns itself into three roughly co-linear segments separated by curved portions which are lifted off the ground 7.1.2. As the motion begins,

the curved segments effectively act as counter-rotating wheels, propelling the head of the snake forward while maintaining the s-shape. Interestingly, the counter-rotation means that the middle part of the snake remains almost stationary, making it appear as though the snake is "flowing" through the shape.

To make this motion possible, it is necessary for the snake to be able to lift its curved portions off the ground. If it did not, then due to the anisotropy of the snakes scales (where moving foward is much easier than sideways or backwards), the friction in the curved regions would slow the snake down drastically and prevent trilinear locomotion from being an effective escape mechanism. To lift itself off the ground, the snake must be very strong compared to its weight. Anacondas are extremely strong snakes, and this strength may be why they are the only snakes in which this mode has been observed. Moreover, as a snake gets larger, its strength scales only with the 2/3rds power of its mass, meaning that the ratio of strength to mass decreases; this implies that older, bigger snakes would have difficulty performing trilinear locomotion because they will struggle to lift themselves off the ground. Indeed, this is what we observe: only young anacondas have been observed to perform trilinear locomotion.

In the computer simulation, the motion is driven by two traveling contraction waves proceeding from head to tail. They have the form $\boldsymbol{A} \exp \frac{(x-x_0-vt)}{2\sigma^2}$ where $\boldsymbol{A}$ is an amplitude vector, x is the arclength position along the snake, x_0 is a constant saying where the wave starts, v is the speed of the wave, t is time, and σ gives the width of the contraction. To effectively produce forward motion, it was indeed seen that it was necessary to provide enough muscular torque to lift the curved sections off the ground. If the contraction was purely in the plane, the snake barely moved, if at all. We can also concluude that trilinear locomotion is a fairly "advanced" form of motion: to effectively replicate trilinear locomotion it was necessary to carefully

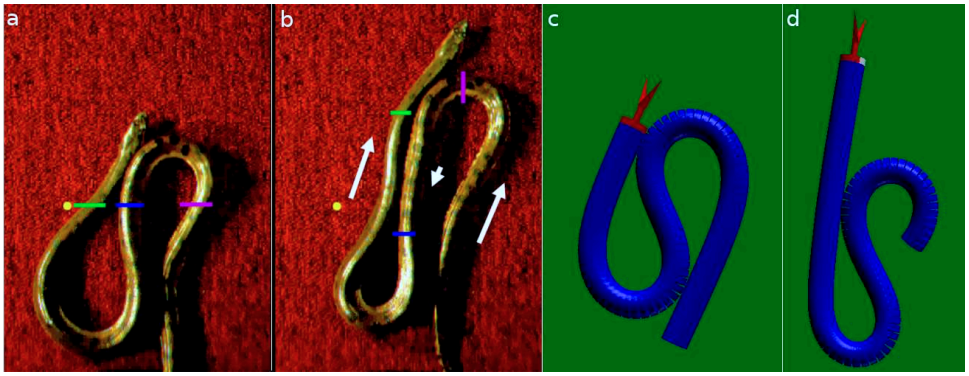

Figure 7.1.2: A juvenile yellow anaconda (Eunectes notaeus) performing trilateral locomotion (a, b). A reference point (yellow dot) was used to align three body marks at $t = 0$ (a). Image processing was used to calculate the position and orientation of these marks at subsequent times (e.g. b, $t = .24s$), demonstrating that: the tail segment moves rostally at high velocity, rotating to become the middle segment; the middle segment has nearly zero velocity, rotating in the opposite direction to become the head segment; the head segment is propelled rostally at high velocity. This kinematic pattern is summarized by the white arrows which indicate the direction, and relative magnitude, of movement of the three body segments. A computational rod model (c, d) was used to replicate the motion, showing that that it can be reproduced with simple muscular contractions.

chose each parameter, a task performed by iteratively changing parameters and observing the resulting motion of the snake.

We've seen that trilinear locomotion relies on the snake having sufficient strength to lift itself off the ground. The snake's anisotropic friction with the substrate clearly also plays a large role. An interesting question then arises: for a given snake and a given surface, what is the most effective way for the snake to achieve a certain velocity (say, fast enough to run from danger)? Exploring this question would be an interesting way to unify all the known forms of snake locomotion onto a single phase space.

7.2 Multiple Filaments

So far our experiments have focused on a single filament. However, nothing prevents us from simultaneously simulating multiple filaments at once. Indeed, the Polymer 8.6 class is designed to function with an arbitrary number of simultaneous rods. The most obvious addition we need with multiple filaments is a collision model so they can physically interact with each other. Adding this functionality turns out to be quite simple, as we can re-use our self-collision model with only minor changes. In fact, we can use the Interaction 8.6 class to model arbitrary interactions between two or more rods.

7.2.1 Cloth and Ponytails

There are many areas involving multiple filaments that are ripe for exploration. We could explore the properties of cloth, where many fibers are woven together to form an effective surface 7.2.1. We could numerically investigate an area there has been a surprisingly large amount of analytical work on: ponytails 7.2.2. The possibilities really are endless, but each requires considerable time.

7.2.2 Combing Hair

The one multiple-filament situation we have had time to quantitatively analyze is a problem we all know well: combing your hair. In our experiment we consider a minimalistic version of haircombing, where two entwined helical strands of opposite handedness are pierced by a tine from the comb and then the tine is pulled downward 7.2.3. For some combination of parameters the tine will remain stuck in the hair even at large times, which we call jamming, and for other combinations it fully pass through the hair.

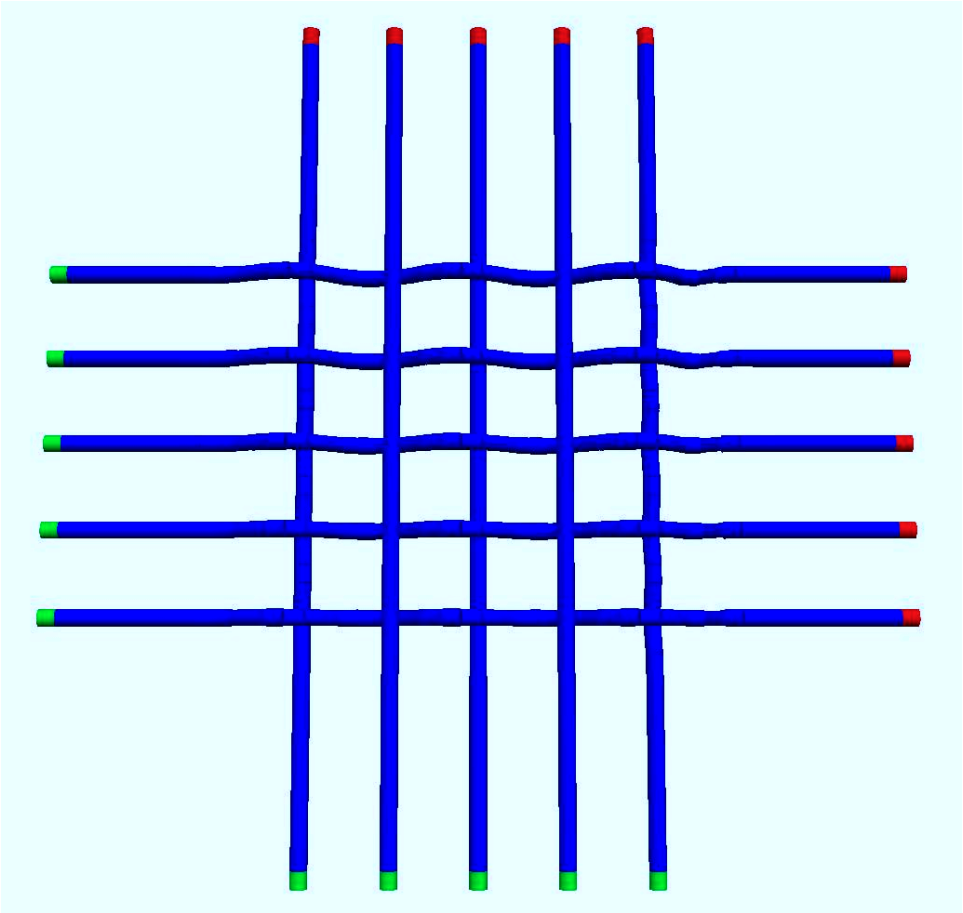

Figure 7.2.1: Cloth can be simulated precisely as a interwoven set of strings.

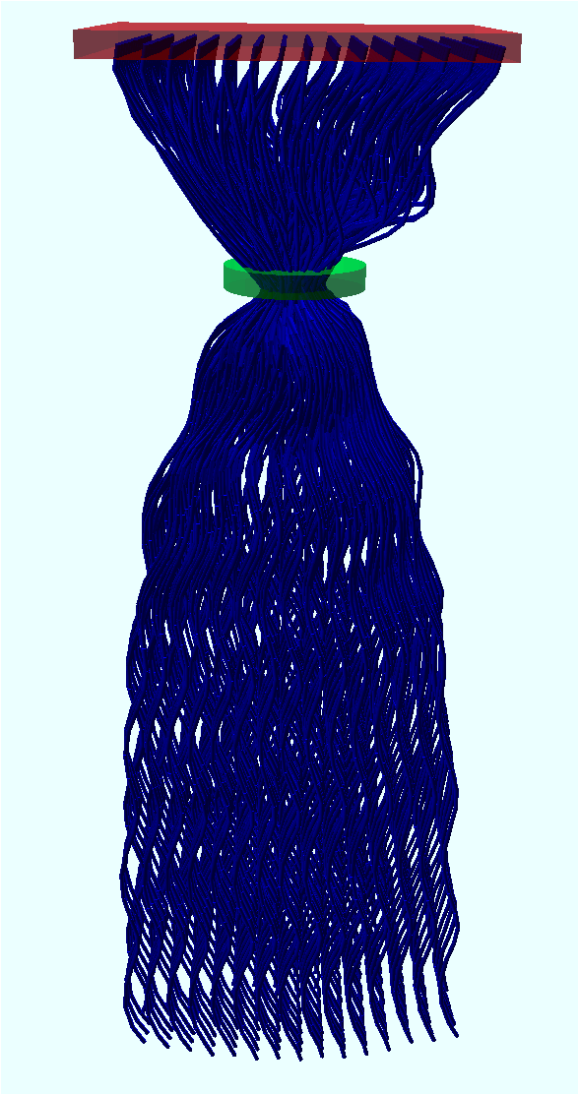

Figure 7.2.2: The model can handle an arbitrary number of strings. This ponytail is comprised of 225 hairs attached at the head (red) and constained by a hairband (green).

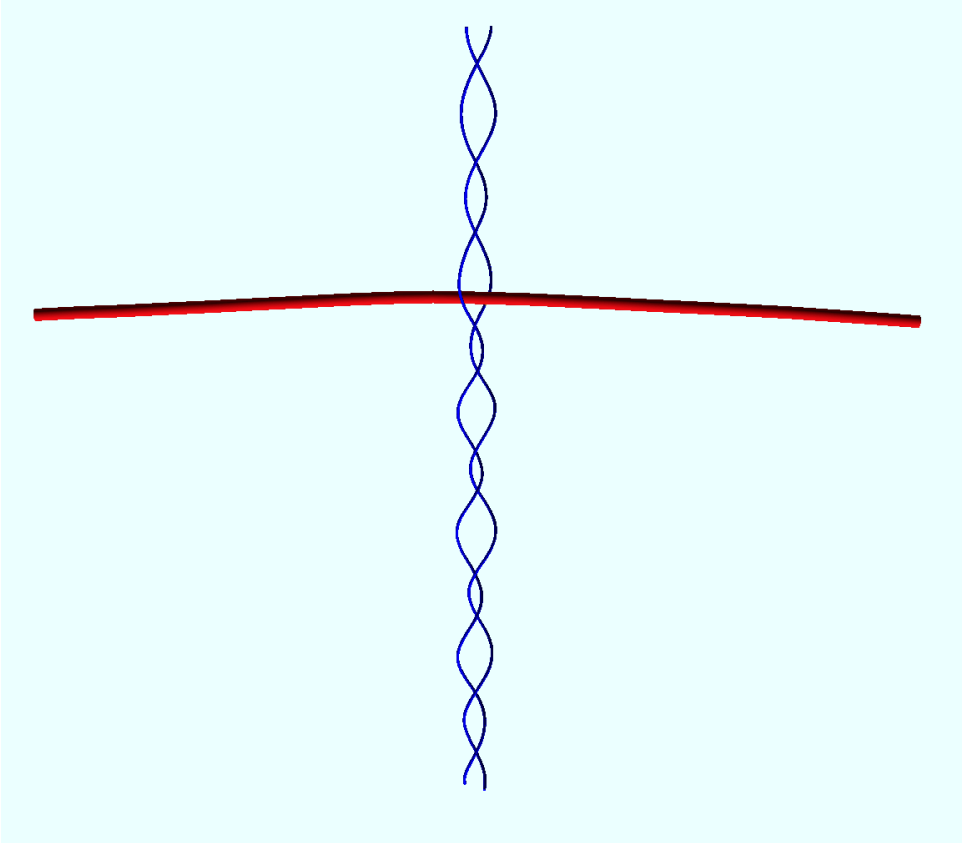

Figure 7.2.3: A tine from a comb intersections two helical strands of hair. In the case shown here, the comb is not pulled down sufficiently hard and remains stuck in the hair. The surprising fact is that this jamming occurs even without friction; it is purely energetic.

The relevant parameters in the system are the bending stiffness of the hair, α, the radius of the helix, r_{helix}, the pitch of the helix P, the radius of the hair r_{hair}, and the radius of the brush tine, r_{brush}.

To understand the jamming behavior we varied a number of parameters in the combing system, recording each time whether the comb made it through the hair or got jammed. We expect that the length of the hair might have some effect on the jamming, so in these runs we use relatively long hair (20 periods) in hopes that we are in

the long-scale regime where changing length of the hair won't matter. (Clearly if your hair was only half a loop long you could always untangle it.) Later we will investigate the role that hair length plays. In 7.2.4 a typical sweep is shown, where we varied the bending stiffness α of the hair and swept the pulling force from zero to 5 units, recording whether it broke free (blue dot) or remained jammed (red dot). The green line helps to illustrate the phase transition. In 7.2.5 we've combined six different sweeps into one graph, where in the blue dots we have digitized the green line transition and then applied a linear fit (red line).

From the right side of 7.2.5 it appears evident tha the radii of the hair and the brush have relatively little effect on the jamming behavior. This is somewhat surprising, as one might expect that, for example, the angle of contact between the brush and the hair it is pushing against might influence the ability to push through a knot. Contrastingly, the parameters on the left hand side clearly have strong impacts on the jamming transition. The commonality between these parameters is that each of them effects the bending energy of the helical hair. Our qualitative conclusion, hence, is that the system is largely explaining by an energetic barrier effect resulting in what we see as a jamming phenomenon.

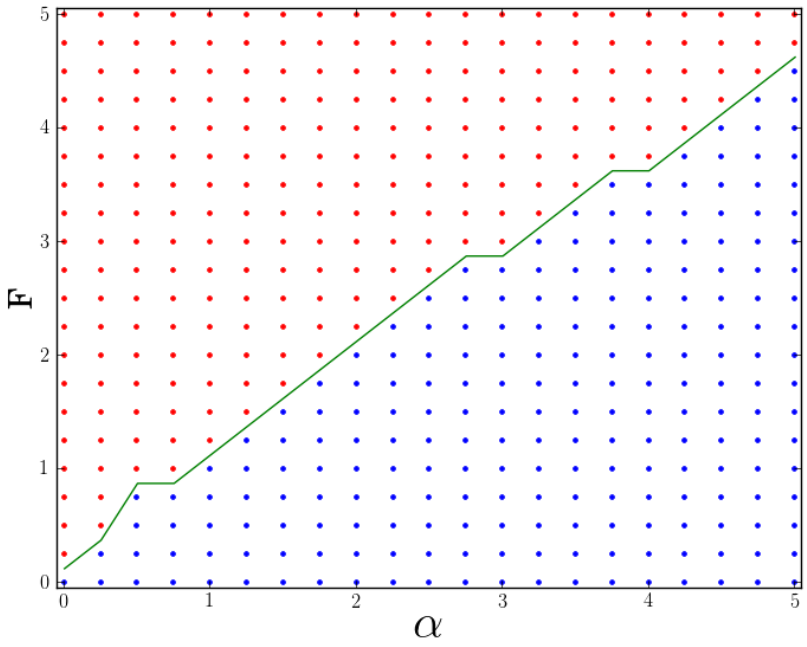

Figure 7.2.4: We sweep a plane in the combing phase space, varying the hair bending stiffness and the tine pulling force F, recording whether it broke free (blue dot) or remained jammed (red dot). The green line helps to illustrate the phase transition.

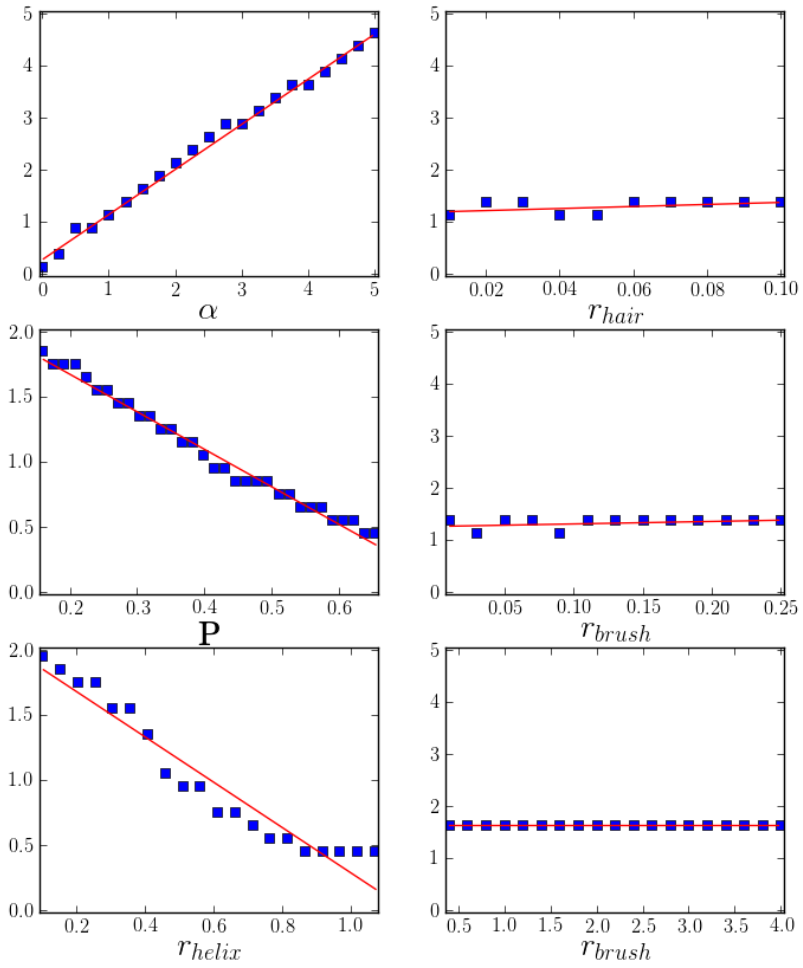

Figure 7.2.5: In each subgraph, we vary a single parameter on the x-axis and xothe tine pulling force on the y-axis. The blue dots represent the tranition between the free (below) and jammed (above) phases. On the left side we see that the variables which effect the bending energy of the hair in its rest state have a definite effect on the force required to free the comb, whereas the graphs on the right show that the the radius of the hair and brush have essentially no effect.

8 Coding

The core model is written in C++, while the visualization and data processing are in python. In this section I give a very brief overview of the most important pieces of the code. The actual source code is available in the supplementary information.

8.1 Preliminaries

Throughout the code I make extensive use of the STL vector construct rather than raw pointers to minimize memory management issues. That said, the reader should be reasonably familiar with standard c++ constructs, particularly pointers and function pointers. The basic orginization of the code is an object-oriented simulation system sitting on top of a fast low-level math package.

8.2 Basic Linear Algebra System (BLAS)

The key to the model's high performance was the implementation of a BLAS (Basic Linear Algebra System) specifically written for 3-vectors and 3x3 matrices. The specific implementation chosen has the additional benefit of being able to have the code closely resemble actual mathematical statements; for example, we use operator overloading to allow matrices A and B to be written as A * B rather than something like matrixMultiply(A, B).

8.2.1 Vector3 Class: Vector3.cpp

Our most basic class is the 3-vector. We equip it with all of the functions you would expect and overload operators to match common mathematical behavior.

8.2.2 Matrix3 Class: Matrix3.cpp

Building upon the Vector3 class we implement the Matrix3 class, which represents a general 3x3 real matrix. We provide fast implementations of common matrix routines like determinant, inverse, etc., as well as defining the operations between matrices and vectors.

8.3 Vector Functions: vectorFunctions.cpp

The code makes extensive use of the STL vector class, greatly simplying the memory management issues one can encounter in c++ when using raw pointers. In the vectorFunctions.cpp file I continue the principle of using operator overloading to continue having the coding statements be as close as possible to their mathematical analogs. For example, if we have two stl::vector's whose elements are Vector3's that we label A and B, then we define A * B to be a new

vector where elementwise we take the * operator (which is the cross-product). Using the new c++0x construct decltype allows us to make great use of general meta-templating to shorten the necessary code. As in the Vector3 and Matrix3 classes, the functions here are highly optimized to maximize performance.

8.4 COLLISION: COLLISION.CPP

One of the computationally most intensive parts of the simulation is self-collision. We use a fast algorithm that calculates the minimum vector between two centerlines representing the cylinders we are checking. Note that this algorithm also works for collision detection between two different filaments.

8.5 ROD CLASS: ROD.CPP

The Rod class contains all of the data needed to describe a filament and calculate all of the internal forces acting on it.

8.6 POLYMER CLASS: POLYMER.CPP

The Polymer class holds a set of rods, applies boundary conditions, external forces, and interactions between the rods. In essence the Polymer class is the main "engine" of the code. You pass it vectors of rods, boundary conditions, external forces, and interaction forces, and then it knows how to evolve the system forward in time. Terms refencing the data parameter of the Rod objects below are used to output data for specific experiments and are not part of the general functionality of the Polymer.

8.7 PolymerIntegrator Class:Polymer_integrator.cpp

While the Polymer class is the general engine of the system, the PolymerIntegrator performs the core mathematics. In particular, it performs the sympletic evolution of the rods and directly applies the BCs, etc.

8.8 RodBC Class: Rod_boundaryConditions.cpp

The RodBC class allows us to apply arbirary boundary conditions to a rod, including those that are a function of time or random. These conditions are enforced at each time step by the PolymerIntegrator class.

8.9 ExternalForces Class: Rod_externalForces.cpp

The ExternalForces class allows us to apply arbitrary external forces to rods. These forces are calculated right after the rod's internal forces are calculated and are added to them.

8.10 Interaction Class: Interaction.cpp

The Interaction Class is designed to allow interactions between two or more rods, for example allowing them to collide. It also is used to represent interactions between the rod and external objects such as planes or spheres.

8.11 An Experiment Example: SolenoidExpt.cpp

Here I give an example of how we use the classes we've built up to conduct a specific experiment. Our setup here is from the solenoid

and plectoneme experiment. We have a vertically hanging rod with a weight hanging from the bottom that acts to stretch it out. During the course of the experiment we will apply a certain number of twists to the bottom of the rod (the top is fixed) and observe the response of the rod. We then want to categorize the shape into one of three types: straight, plectonemal, or solenoidal.

We run one instance of the experiment by calling solenoidRun, passing in the how many cylinders to use, the length of the rod, the hanging force, and the twist to apply. After setting up some reasonable constants, we initialize a rod via compressedRod that represents our hanging rod. (In this case, it is anti-compressed, but thats ok.) We then define the boundary conditions for the frame which represents the twisting motion we are going to apply. We leave the boundary condition for the starting x position as default-initizialized, which the program intreprets as being a free boundary condition within the EndpointBC boundary condition. Now, we also want the weight to remain fixed directly below the support it is hanging from, so we also create a PinXY boundary condition and then use a MultipleBC boundary condition to apply both the EndpointBC and PinXY boundary conditions simultaneously.

Similarly for the external forces, we create and EndpointForces to represent the hanging weight, a RandomForces to break the symmetry, a ZPlanesForce to prevent the rod from passing through the plane from which it is hanging, and a MultipleForces object to apply all of the these forces simulataneously.

Since we only have one rod in this experiment, we push only one reference onto the rodptrs vector. Similarly, since we've defined no Interactions for this experiment, we create a interactions vector but don't push anything onto it. Next we create a PolymerIntegrator object. Note that we pass it a "1", which means we are going to a 1st

order integrator. Usually the first order integrator is sufficient; higher order integrators are more accurate but must calculate the forces multiple times per time step so they are slower.

After constructing all the elements of our experiment, we create a new Polymer object and run its simulate method to evolve the system forward in time. Once it is finished running we can use the visualization tools to watch what happened during the experiment (snapshots are periodically taken and saved to disk), or we could, for example, pass the evolved Rod object to the classifyShape function which would tell us what final form the rod was in.

References

[1] R. Abraham and J. E. Marsden. *Foundations of Mechanics*. London: Benjamin-Cummings, 1978.

[2] S. Antman. *Nonlinear problems of elasticity*. Springer, New York, 2005.

[3] Shahaf Armon, Efi Efrati, Raz Kupferman, and Eran Sharon. *Science*, 333(6050):1726, 2011.

[4] J. Baillieul and M. Levi. Rotational elastic dynamics. *Physica D: Nonlinear Phenomena*, 27(1):43, 1987.

[5] A. Belmonte, M. J. Shelley, S. T. Eldakar, and C. H. Wiggins. *Phys. Rev. Lett.*, 87:114301, 2001.

[6] M. A. Berger and C. Prior. *J. Phys. A: Math. Gen.*, 39:8321, 2006.

[7] M. Bergou, M. Wardetzky, S. Robinson, B. Audoly, and E. Grinspun. *ACM TOG*, 27:63, 2008.

[8] M. Bergou, B. Audoly, E. Vouga, M. Wardetzky, and E. Grinspun. *ACM Trans. Graph. 29*, 4:116, 2010.

[9] J. S. Birman. New points of view in knot theory. *Bull. Amer. Math. Soc.*, 28(2):253–287, 1993.

[10] Nawaf Bou-Rabee. *Hamilton-Pontryagin Integrators on Lie Groups*. PhD thesis, California Institute of Technology, 2007.

[11] Andrew J. Bowling and Kevin C. Vaughn. *American Journal of Botany*, 96(4):719–727, 2009.

[12] G. Buck and J. Simon. *PNAS A*, 469(2148):4024, 2012.

[13] D. A. Burton and R. W. Tucker. *Mechanics of Generalized Continua*. Springer New York, 2010.

[14] J. C. Butcher. *Numerical Methods for Ordinary Differential Equations*. New York: John Wiley and Sons, 2003.

[15] G. Calugreanu. *Rev. Math. Pures Appl.*, 4:5, 1959.

[16] Z. Chen, C. Majidi, D. J. Srolovitz, and M. Haataja. *Applied Physics Letters*, 98(1), 2011.

[17] B Clair, B Thibaut, and J Sugiyama. *Journal of Wood Science*, 51(3):218, 2005.

[18] A. Clebsch. *Theorie de l'elasticite des corps solides*. Dunod, 1883.

[19] E. Cosserat and F. Cosserat. *Theorie des Corps deformables*. Hermann et Fils, Paris, 1909.

[20] D. Couerjolly, S. Minuet, and L. Tougne. *Visual Form 2001*, page 303, 2001.

[21] R. Courant and D. Hilbert. *Methods of Mathematical Physics Vol. 1*. Interscience Publishers, 1953.

[22] C. Darwin. *On the movements and habits of climbing plants*. John Murray, 1865.

[23] RH Dastur and GA Kapadia. *Annals of Botany*, 45(178):279, 1931.

[24] R. de Vries. *J. Chem. Phys.*, 122:064905, 2005.

[25] M. Delbruck. Knotting problems in biology. *Proc. Symp. Appl. Math.*, 14, 1962.

[26] E. H. Dill. Kirchhoff's theory of rods. *Archive for the History of Exact Sciences*, 44(1):1, 1992.

[27] C. M. Dobson. Protein folding and misfolding. *Nature*, 426 (6968):884–890, 2003.

[28] T. Frede, C. Renner, C. Budair, Y. Adbel-Salam, and J. Rassweiler. Geometry of laparoscopic suturing and knotting techniques. *Journal of endourology*, 13(3):191–198, 1999.

[29] H. L. Frisch and E. Wasserman. *J. Am. Chem. Soc.*, 83: 3789–3795, 1961.

[30] F. B. Fuller. *Proc. Natl. Acad. Sci. USA*, 68:815, 1971.

[31] F. B. Fuller. *Proc. Natl. Acad. Sci. USA*, 78:3557, 1978.

[32] M. Gascuel. An implicit formulation for precise contact modeling between flexible solids. In *Proceedings of the 20th annual conference on Computer graphics and interactive techniques*, 1993.

[33] I. M. Gelfand and S. V. Fomin. *Calculus of variations*. Dover Publications, 2000.

[34] S.J. Gerbode, J. R. Puzey, A. G. McCormick, and L. Mahadevan. *Science*, 337:1087, 2012.

[35] A. Ghatak and L. Mahadevan. *Phys. Rev. Let.*, 65:057801, 2005.

[36] R. E. Goldstein, T. R. Powers, and C. H. Wiggins. Viscous nonlinear dynamics of twist and writhe. *Phys. Rev. Let.*, 80(23): 5232, 1998.

[37] R. E. Goldstein, A. Goriely, G. Huber, and C. W. Wolgemuth. Bistable helices. *Physical Review Letters*, 84(7):1631–4, 2000.

[38] A. Goriely. *Journal of Elasticity*, 84:281, 2006.

[39] A. Goriely and M. Tabor. New amplitude equations for thin elastic rods. *Physical Review Letters*, 77(17):3537–40, 1996.

[40] A Goriely and M Tabor. *Physical Review Letters*, 80(7):1564, 1998.

[41] Luna et al. Goswami. *Plant Journal*, 56(4):531, 2008.

[42] A. Gray. *Structural botany: or organography on the basis of morphology; to which is added the principles of taxonomy and phytography, and a glossary of botanical terms.* Ivison, Blakeman, Taylor, 1880.

[43] A. E. Green. A general theory of rods. *Proceedings of the Royal Society of London, Series A. Mathematical and Physical Sciences*, 293(1433):145–155, 1966.

[44] A. E. Green and R. S. Rivlin. *Arch. Rat. Mech. Anal.*, 16:65, 1964.

[45] A. E. Green and R. S. Rivlin. *Arch. Rat. Mech. Anal.*, 17:113, 1964.

[46] N. Habegger and X. S. Lin. The classification of links up to homotopy. *Journal of the American Mathematical Society*, 3(2), 1990.

[47] B. C. Hall. *Lie Groups, Lie Algebras, and Representations: An Elementary Introduction.* Springer, 2006.

[48] J. Hickford, S. duPont, and J. Eggers. *Phys. Rev. E*, 74:052101, 2006.

[49] K. A. Hoffman and T. I. Seidman. A variational rod model with a singular nonlocal potential. *Archive for rational mechanics and analysis*, 200(1):255–284, 2011.

[50] Sandrine Isnard and Wendy K. Silk. *American Journal of Botany*, 96(7):1205, 2009.

[51] MJ Jaffe and AW Galston. *Annual Review of Plant Physiology*, 19:417, 1968.

[52] J. B. Keller. *Lectures on mathematics in the life sciences*, 13:257, 1980.

[53] G. Kirchhoff. *J. firr Mathematik*, 56:285, 1859.

[54] I. Klapper. *J. Comp. Phys.*, 125:325, 1996.

[55] K. Klenin and J. Langowski. *Biopolymers*, 54, 2000.

[56] L. D. Landau and E.M. Lifschitz. *Elasticity Theory*. Pergamon, Oxford, 1986.

[57] J. Langer and D. A. Singer. *SIAM Review 38*, 4:605, 1996.

[58] M. Leok. An overview of lie group variational integrators and their applications to optimal control. In *International Conference on Scientific Computation and Differential Equations*, 2007.

[59] M. Levi. Composition of rotations and transport. *Nonlinearity*, 9:413, 1996.

[60] W. R. Lickorish. *An introduction to knot theory*. Springer, 1997.

[61] Henrietta Lisk. *Botanical Gazette*, 78(1):85, 1924.

[62] A. E. H. Love. *A treatise on the mathematical theory of elasticity, Second Edition*. Cambridge University Press, Cambridge, 1906.

[63] E. Hairer C. Lubich and G. Wanner. *Geometric Numerical Integration: Structure-Preserving Algorithms for Ordinary Differential Equations*. Springer, New York, 2006.

[64] D. T. MacDougal. *Annals of Botany*, os-10(3):373, 1896. URL http://aob.oxfordjournals.org/content/os-10/3/373.short.

[65] A. L. Mallam, E. R. Morris, and S. E. Jackson. Exploring knotting mechanisms in protein folding. *Proceedings of the National Academy of Sciences*, 105(48):18740–18745, 2008.

[66] T McMillen and A Goriely. *Journal of Nonlinear Science*, 12(3): 241, 2002.

[67] Christopher G. Meloche, J. Paul Knox, and Kevin C. Vaughn. *Planta*, 225(2):485, 2007.

[68] C. Micheletti, J. R. Banavar, A. Maritan, and F. Seno. Protein structures and optimal folding from a geometrical variational principle. *Physical Review Letters*, 82(16):3372, 1999.

[69] Hugo von Mohl. *Principles of the anatomy and physiology of the vegetable cell*. London John Van Voorst, 1852.

[70] J. P. J. Muchels and F. W. Wiegel. *Phys. Lett. A*, 90:381–384, 1982.

[71] D. K. Pai. *CGF*, 21:3, 2002.

[72] L. S. Pontryagin. Smooth manifolds and their applications in homotopy theory. *Transl. Amer. Math. Soc.*, 11:1–114, 1959.

[73] D. M. Raymer and D. E. Smith. *PNAS*, 104(42):16432, 2007.

[74] F. Riewe. Nonconservative lagrangian and hamiltonian mechanics. *Phys. Rev. Let. E*, 53(2), 1996.

[75] D. Rolfsen. *Knots and Links*. AMS Chelsea Pub., 2003.

[76] Thierry Savin, Natasza A. Kurpios, Amy E. Shyer, Patricia Florescu, Haiyi Liang, L. Mahadevan, and Clifford J. Tabin. *Nature*, 476:57, 2011.

[77] H. Seifert and W. Threlfall. *Lehrbuch der topologie*. Leipzig and Berlin: BG Teubner, 1934.

[78] J. Serre. *Lie Algebras and Lie Groups*. Springer, 2006.

[79] M. K. Shimamura and T. Deguchi. *Phys. Rev. E*, 66:040801, 2002.

[80] J. Spillman and M. Teschner. *Eurographics 2008*, 27(2), 2008.

[81] A. R. Srinivasa. *International Journal of Plasticity*, 17(9): 1189–1214, 2001.

[82] P. F. Stevens. http://www.mobot.org/MOBOT/Research/APweb/welcome.html, 2008.

[83] R. A. Toupin. *Arch. Rat. Mech. Anal.*, 17:85, 1964.

[84] G. H. M. van der Heijden and J. M. T. Thompson. *Nonlinear dynamics*, 21:71, 2000.

[85] M. Weiss and F. Elmer. Dry friction in the frenkel-kontorova-tomlinson model: Static properties. *Physical Review B*, 53(11):7539, 1996.

[86] T. P. Westcott, I. Tobias, and W. K. Olson. *J. Chem. Phys.*, 107 (10), 1997.

[87] J. H. White. *Am. J. Math.*, 91:693, 1969.

[88] Wikipedia. URL en.wikipedia.org/wiki/Linking_number.

[89] C. W. Wolgemuth, T. R. Powers, and R. E. Goldstein. Twirling and whirling: Viscous dynamics of rotating elastic filaments. *Physical Review Letters*, 84(7):1623–6, 2000.

[90] C. W. Wolgemuth, R. E. Goldstein, and T. R. Powers. Dynamic supercoiling bifurcations of growing elastic filaments. *Physica D: Nonlinear Phenomena*, 190(3):266–89, 2004.

[91] Y. Yang, I. Tobias, and W.K. Olson. *J. Chem. Phys. 98*, 2:1673, 1993.

[92] H. Yoshida. *Phys. Lett. A*, 150:150, 1990.