2 The Special Cosserat Theory of Idealized Rods

2.1 Introduction

The next several weeks of class will be concerned with developing the static continuum elastic rod model of the DNA mini-circle cyclization problem that was outlined in the first days lectures.

The mathematics that will be introduced involves the calculus of variations, Hamiltonian formulations of the Euler-Lagrange equations, bifurcation theory, numerical methods for ODE two-point boundary value problems (including numerical symmetry breaking), and the theory of the second variation in isoperimetric calculus of variations problems.

We start with a description of the mechanics model, known as the Cosserat theory of rods, and the quaternion parameterization of proper rotation matrices. We then describe the mathematical methods, and finally we discuss how to fit the parameters of our model to describe DNA itself, and make some computations.

The sum total of all this material will occupy several weeks.

A *configuration* of a Cosserat rod is a parameterized space curve $\boldsymbol{r}(s)$ along with a parameterized family of right handed, orthonormal triads $\boldsymbol{d}_i(s)$, i.e. three unit vectors which for each s satisfy the constraints

$$\boldsymbol{d}_i \cdot \boldsymbol{d}_j = \delta_{ij}, \quad \boldsymbol{d}_3 = \boldsymbol{d}_1 \times \boldsymbol{d}_2$$

with δ_{ij} being the Kronecker delta function.

Throughout the course Roman indices will be assumed to run through the values 1, 2 and 3 with repeated indices summed, unless the context or text indicates otherwise.

These orthonormality relations mean exactly that the 3×3 cosine matrix $\boldsymbol{R}$ of components of the three vectors $\boldsymbol{d}_i$ with respect to any fixed (right handed) system of orthonormal basis vectors, say $\boldsymbol{e}_i$, is a proper rotation matrix (i.e. a rotation matrix with determinant $+1$) or (a representative of) an element of the group $SO(3)$.

Notice that the frame $\{\boldsymbol{d}_i\}$ is defined externally to the curve $\boldsymbol{r}$, i.e. the frame contains additional information to the centerline. Thus the $\{\boldsymbol{d}_i\}$ frame has a different status from say the Frenet-Serret frame which is made up from the tangent, normal and binormal to the curve, and which is defined by the curve $\boldsymbol{r}(s)$ and its derivatives.

The curve $\boldsymbol{r}$ should be thought of as the centerline of the rod (say the average of the two backbones in the case of DNA) while the frame $\{\boldsymbol{d}_i\}$ should be regarded as specifying the orientation of each cross-section of the rod (or say some smooth interpolation of the orientation of each set of base pairs in the case of DNA).

The parameter $\boldsymbol{s}$ will be given a precise physical interpretation in a little while, but is, roughly speaking, distance along the centerline of the rod.

2.2 Kinematics

With no loss of generality we may define the vector $\boldsymbol{v}(\boldsymbol{s})$ via $\boldsymbol{v}(\boldsymbol{s}) \equiv \boldsymbol{r}'(\boldsymbol{s})$ where $\boldsymbol{r}'$ denotes the derivative of $\boldsymbol{r}$ with respect to the parameter $\boldsymbol{s}$. And we will denote the components $\boldsymbol{v}_i$ of $\boldsymbol{v}$ with respect to the triad $\{\boldsymbol{d}_i\}$ by $\boldsymbol{v}_i \equiv \boldsymbol{v} \cdot \boldsymbol{d}_i$

As the $\boldsymbol{d}_i$ are an orthonormal basis they satisfy kinematic equations of the form

$$\boldsymbol{d}_i' = \boldsymbol{u} \times \boldsymbol{d}_i.$$

That is the vector $\boldsymbol{u}$ is an 'angular velocity' but representing 'evolution' with respect to the parameter $\boldsymbol{s}$ instead of with respect to time.

It is a simple exercise (do it!) to derive the existence of the vector $\boldsymbol{u}$, which is sometimes called the Darboux vector.

(Hint for one way: because the triad $\{\boldsymbol{d}_i\}$ is a basis at each $\boldsymbol{s}$ there exist functions $\boldsymbol{c}_{ij}(\boldsymbol{s})$ such that $\boldsymbol{d}_i' = \sum \boldsymbol{c}_{ij}\boldsymbol{d}_j$. Differentiating the orthonormality conditions on the $\boldsymbol{d}_i$ then give conditions on the $\boldsymbol{c}_{ij}$ which imply that there are only three independent non-zero $\boldsymbol{c}_{ij}$, and they have the property that they are equivalent to taking a vector cross-product, i.e. the Darboux vector exists.)

We also introduce the components of the Darboux vector $\boldsymbol{u}$ with respect to the basis $\{\boldsymbol{d}_i\}$ through

$$\boldsymbol{u}_i = \boldsymbol{u} \cdot \boldsymbol{d}_i$$

In fact one can see that

$$\boldsymbol{u}_i = \frac{1}{2}\boldsymbol{\epsilon}_{ijk}\boldsymbol{d}_j' \cdot \boldsymbol{d}_k$$

where $\boldsymbol{\epsilon}_{ijk}$ is the alternating tensor.

The components $\boldsymbol{u}_i(\boldsymbol{s})$ determine the frame $\{\boldsymbol{d}_i(\boldsymbol{s})\}$ (up to a single arbitrary rotation specifying initial conditions for the frame) through integration of the (nine, scalar) differential equations

$$\boldsymbol{d}_i' = \boldsymbol{u} \times \boldsymbol{d}_i = \sum_j \boldsymbol{u}_j \cdot (\boldsymbol{d}_j \times \boldsymbol{d}_i).$$

or

$$\boldsymbol{d}_i' = \epsilon_{jik} u_j \boldsymbol{d}_k$$

Once the frame $\{\boldsymbol{d}_i(s)\}$ is known, the three component $v_i(s)$ determine the centerline $\boldsymbol{r}(s)$ (up to a single arbitrary translation specifying initial conditions on the centerline) through integration of the equations

$$\boldsymbol{r}' = \boldsymbol{v} = \sum_i v_i \boldsymbol{d}_i$$

Thus the six scalar functions $u_i(s)$ and $v_i(s)$, determine the configuration of the rod up to a single arbitrary rigid body motion (i.e. a translation plus a rotation).

For this reason the six functions $u_i(s)$ and $v_i(s)$ can, and will, be called a set of *strains* for the rod.

There is certainly quite some arbitrariness in the definitions of the strains, which arises from an arbitrariness in the relation of the frame $\{\boldsymbol{d}_i(s)\}$, and the parameter s, to the material making up the rod.

2.2.1 The reference state and the adapted framing

This arbitrariness is usually removed by specifying particular *reference* strains $\hat{u}_i(s)$ and $\hat{v}_i(s)$ describing a *reference state.*

We will usually assume the reference state to be a *minimum energy* or *unstressed configuration* (the precise definitions of which will arise later).

Usually (but not necessarily) the parameter s is chosen to be arc-length along the centerline of the reference curve $\widehat{\boldsymbol{r}}(s)$, in which case the reference strains satisfy

$$|\widehat{\boldsymbol{r}}'| = |\widehat{\boldsymbol{v}}(s)| = 1,$$

and in any other configuration $|\boldsymbol{v}(s)| - 1$ is a measure of the local extension (or stretch) when positive, or of the local compression when negative.

Usually (but not necessarily) the reference unit vector field $\widehat{\boldsymbol{d}_3}(s)$ is chosen to be parallel to the tangent vector to the reference centerline $\widehat{\boldsymbol{r}}(s)$ so that provided s has been chosen to be arc-length in the reference state, we have for all s that

$$\widehat{\boldsymbol{r}}' = \widehat{\boldsymbol{d}_3}$$

The triad $\{\widehat{\boldsymbol{d}_i}(s)\}$ is then called an *adapted framing* of the curve $\widehat{\boldsymbol{r}}(s)$.

In this situation the $\mathbf{3}$-components of the strains are distinguished from the $\mathbf{1}$ and $\mathbf{2}$ components, and it makes sense to describe $\boldsymbol{v_1}$ and $\boldsymbol{v_2}$ as shear strains, and (somewhat imprecisely) $\boldsymbol{v_3}$ as a stretch. Similarly $\boldsymbol{u_1}$ and $\boldsymbol{u_2}$ are described as bending strains, while $\boldsymbol{u_3}$ will be described as the twist (strain).

As a matter of physical modelling, when interpreting mathematical rod models as an approximation of long slender deformable bodies, it is implicitly assumed that the whole three-dimensional shape of the body can be reconstructed to a good approximation from knowledge of the centerline and the frame, using a known, locally-defined reconstruction rule at each value of the parameter $\boldsymbol{s}$.

For example, in the context of DNA it is implicitly assumed that the location of each atom making up the $\boldsymbol{N}$th base pair can be well-approximated once the frame, and centerline at the arc-length corresponding to the $\boldsymbol{N}$th base are known.

The arbitrariness in the choice of reference strains $\hat{\boldsymbol{u}}_{\boldsymbol{i}}(\boldsymbol{s})$ and $\hat{\boldsymbol{v}}_{\boldsymbol{i}}(\boldsymbol{s})$ can be regarded as an arbitrariness in the definition of the local reconstruction rule, but once a reference configuration, or equivalently a local reconstruction rule at each value of $\boldsymbol{s}$, has been set, then the same local reconstruction rule is used in *any* configuration of the rod.

For this to be a reasonable approximation, it is implicit assumed that the differences between the strains in any physically relevant configuration and in the reference configuration, are quite small, and for this reason rod theory is sometimes described as a linearized or small strain theory.

However this does not mean that the differences in configurations are small, because small differences in strains over a relatively large distance $\boldsymbol{s}$ can give big differences in both centerline and frame. In particular rod theory keeps full geometrical nonlinearities.

2.2.2 The inextensible unshearable rod

Much of the classic work on rod theory, and in particular rod theories applied to modelling DNA, assume an inextensible, unshearable rod.

In an inextensible, unshearable rod, the strains $\boldsymbol{v_i}$ in any configuration equal the strains $\hat{\boldsymbol{v}}_{\boldsymbol{i}}$ in the reference configuration, i.e. they satisfy the constraints

$$v_i(s) \equiv \hat{v}_i(s).$$

For an inextensible, unshearable rod it almost always makes sense to choose the parameter s to be arc-length in the reference configuration, because it is then also arc-length in any configuration.

Similarly it almost always makes sense to choose the reference framing to be adapted so that $\widehat{\boldsymbol{r}}'(s) =_{3} (s)$, and

$$\hat{v}_1(s) = \hat{v}_2(s) = 0, \qquad \hat{v}_3(s) = 1$$

because the framing is then adapted in any configuration i.e. $\boldsymbol{r}'(s) = \boldsymbol{d}_3(s)$ in all allowed configurations.

In fact for an inextensible, unshearable rod the configuration space can be viewed as a centerline $\boldsymbol{r}(s)$ parametrized by arc-length s, along with a single unit vector field, $\boldsymbol{d}_1(s)$ say, which is everywhere orthogonal to the centerline $\boldsymbol{r}(s)$.

In other words a ribbon.

An orthonormal (adapted) right handed triad can then be constructed from any configuration of a ribbon from the definitions

$$\boldsymbol{d}_3 = \boldsymbol{r}', \qquad \boldsymbol{d}_2 \equiv \boldsymbol{d}_3 \times \boldsymbol{d}_1,$$

and the $\boldsymbol{u}_i$, as defined before, now form a full set of strains for the problem.

Even for an inextensible, unshearable rod, with an adapted framing, there remains some freedom as to which adapted frame is best chosen as the reference configuration. This is a topic that will be of some importance in our modelling of DNA, and we will return to it later.

2.3 Balance Laws

We now turn from the kinematics of rod models to a consideration of the balance laws, which will allow us to determine which configurations are possible equilibrium shapes.

The stresses exerted by the material on one side of the cross-section at s acting on the material on the other side of the cross-section can be averaged to yield a net force $\boldsymbol{n}(s)$ (of the material on the side s^+ acting on the material on the side s^-).

There is an arbitrary sign convention here, and as long as it is treated consistently it doesn't matter which convention you take.

Similarly the first moment of the stresses acting across the cross-section can averaged to yield a net moment $\boldsymbol{m}(s)$ (taken with the same sign convention as $\boldsymbol{n}(s)$).

As before $\boldsymbol{n}(s)$ and $\boldsymbol{m}(s)$ denote vectors, while $n_i(s)$ and $m_i(s)$ will denote components of these vectors with respect to the variable basis $\{\boldsymbol{d}_i(s)\}$.

It will be of importance to recall the elementary fact that because the frame $\{\boldsymbol{d}_i(s)\}$ is not constant, the component of the derivative $\boldsymbol{m}'(s)$ is not the derivative of the component $m_i'(s)$. Rather using the kinematics of the basis $\{\boldsymbol{d}_i(s)\}$, it can be seen that there is an additional cross-product term.

Explicitly

$$\boldsymbol{m}' = \left(\sum_i m_i \boldsymbol{d}_i\right)' = \sum_i \left(m_i' \boldsymbol{d}_i + m_i \boldsymbol{d}_i'\right)$$

$$= \sum_i \left(m_i' \boldsymbol{d}_i + m_i(\boldsymbol{u} \times \boldsymbol{d}_i)\right).$$

or

$$\boldsymbol{m}' \cdot \boldsymbol{d}_j = m_j' + (\boldsymbol{u} \times \boldsymbol{m}) \cdot \boldsymbol{d}_j$$

so that the jth component of $\boldsymbol{m}'$ is the derivative of the jth component plus the jth component of the cross-product $(\boldsymbol{u} \times \boldsymbol{m})$.

The coordinate free equilibrium equations are (in the absence of self-contact or external distributed forces):

$$\boldsymbol{n}'(s) = \boldsymbol{0},$$

and

$$\boldsymbol{m}'(s) + \boldsymbol{r}'(s) \times \boldsymbol{n}(s) = \boldsymbol{0}.$$

2.4 Constitutive Relations

Thus far our concept of a rod has implicitly used the assumption that the body is long and slender, in order that the kinematics yields a sensible description of the body. But we have not said anything about the properties of the material. How do we tell the difference between a long, slender body made of steel, of rubber, of chewing gum, of a viscous fluid such as treacle, or a segment of a DNA molecule?

The differences are quantified by specifying the *constitutive relations* of the rod, that is giving relations that determine the components of the stresses m_i and n_i in terms of the strains u_i and v_i (or vice versa).

We first give some simple examples of constitutive relations, and afterward discuss, how general the constitutive relations might reasonably be, and how reasonable the simple choices might be.

The choice of appropriate constitutive relations is one of the most crucial issues in using rod models to describe physical phenomena.

Perhaps the simplest choice of constitutive relation is a diagonal, linear (or at least affine) one. For each $i = 1, 2, 3$ (with no sum)

$$m_i = K_i(u_i - \hat{u}_i), \qquad n_i = A_i(v_i - \hat{v}_i).$$

Here, as before, the $\hat{u}_i(s)$ and $\hat{v}_i(s)$ are prescribed functions of s determining the reference shape. And the $K_i(s)$ and $A_i(s)$ are prescribed functions, which for reasons we shall see later, are usually assumed to be strictly positive.

Because of the form of the shift in the constitutive relations

$$m_i = K_i(u_i - \hat{u}_i), \qquad n_i = A_i(v_i - \hat{v}_i).$$

we see that the stresses in the reference state vanish, i.e. the reference state is unstressed.

More generally for positive coefficients $K_i(s)$ and $A_i(s)$ the stress-strain law can be inverted (here trivially) to yield

$$u_i = \hat{u}_i + m_i/K_i, \qquad v_i = \hat{v}_i + n_i/A_i.$$

In particular, invertibility implies (among other things) that when the reference state is stress-free it is the *only* unstressed state.

We will also use the fact that the constitutive laws are *hyper-elastic*, which means precisely that there is a scalar-valued function $W(u_i, v_i; s)$ of the six strains with the property that the constitutive relations can be written in the form of partial derivatives of W

$$m_i = W_{u_i}, \qquad n_i = W_{v_i}.$$

The linear diagonal constitutive laws described above are hyper-elastic with the quadratic strain-energy density function

$$\frac{1}{2}\sum_i \left\{K_i(u_i - \hat{u}_i)^2 + A_i(v_i - \hat{v}_i)^2\right\}$$

The elastic energy of the rod with strains $u_i(s)$ and $v_i(s)$ can be defined as

$$\int_0^L W(u_i, v_i; s)\, ds$$

and $\boldsymbol{W}$ is called the strain-energy density function.

In the diagonal, linearly elastic case we can see that the unstressed reference configuration is also the *unique minimum energy configuration.*

Much of the time we will be concerned with the case of inextensible, unshearable elastic rods. As discussed earlier the inextensibility and unshearability conditions are expressed by specifying the values of the strains $\boldsymbol{v_i}$ to always be their reference values $\hat{\boldsymbol{v}}_i$ (which are usually assumed to be the triple $(\boldsymbol{0},\boldsymbol{0},\boldsymbol{1})$).

For such rods the force $\boldsymbol{n}$ is a basic unknown, with no related constitutive relation, while the bending and twist strains $\boldsymbol{u_i}$ and components of the moments $\boldsymbol{m_i}$ are still related through constitutive relations, for example of the affine form:

$$\boldsymbol{m_i = K_i(u_i - \hat{u}_i)}, \qquad \boldsymbol{i = 1,\ 2,\ 3}\,,$$

where the $\boldsymbol{K_i(s) > 0}$ and the $\hat{\boldsymbol{u}}_i$ are as before.

We will also talk about *hyper-elastic*, inextensible, unshearable rods. For the affine constitutive relation on the previous page, the elastic strain energy of the rod is

$$\int_0^{\boldsymbol{L}} \frac{\boldsymbol{1}}{\boldsymbol{2}} \sum_{\boldsymbol{i}} \boldsymbol{K_i(u_i - \hat{u}_i)^2\, ds}$$

2.5 Equilibrium conditions

We can now describe the system of field equations (which for rods are ordinary differential equations) that must be solved to determine the equilibria of elastic rods. They are made up of the coupled system formed by the balance laws taken with the kinematic equations.

$$\boldsymbol{n'(s) = 0},$$

$$\boldsymbol{m'(s) + r'(s) \times n(s) = 0},$$

$$\boldsymbol{r' = v}, \qquad \boldsymbol{d_i' = u \times d_i}$$

plus some combination of constitutive relations, or unshearability or inextensibility constraints.

In fact when the balance laws are written in terms of components wrt the frame $\{\boldsymbol{d}_i\}$, they (apparently) decouple from the kinematic equations.

Balance laws become a closed system of 6 first-order ODE

$$n_i' + \epsilon_{ijk} u_j n_k = 0 \qquad m_i' + \epsilon_{ijk}(u_j m_k + v_j n_k) = 0,$$

after six of the twelve quantities m_i, n_i, u_i and v_i are eliminated using the constitutive relations

$$m_i = K_i(s)(u_i - \hat{u}_i(s)), \qquad n_i = A_i(s)(v_i - \hat{v}_i(s)).$$

(if the rod is inextensible and unshearable the unknowns n_i must be three of the variables retained).

Resulting equations are nonlinear with quadratic nonlinearities (because the constitutive relations were assumed to be linear).

When the constitutive laws are homogeneous (or autonomous) i.e. have no explicit s dependence, it is possible to find simple (and not so simple) closed form solutions of the balance laws, and then with known strains to integrate the kinematic equations to find explicit expressions for the associated centerline and director frame. See Exercise Session 1.

In general, both for non-homogeneous and homogeneous constitutive relations, it is necessary to solve the system numerically. Particularly because a typical rod problem involves two-point boundary conditions, i.e. some of the variables are given specified variables at $s = 0$, and some at $s = L$ say.

It should be recalled (or realized) that the numerical solution of *initial* value problems, i.e. problems where all the unknowns have specified values at one value of the independent variable s (usually $s = 0$), for smooth equations such as the above, and on a bounded interval, is essentially trivial given todays computers and knowledge of algorithms.

In contrast the numerical solution of two-point boundary value problems is still nontrivial (although of course the appropriate numerical methods are for the most part very well understood).

The other important observation is that even though the balance laws, and the kinematic equations appear to naturally decouple, the boundary conditions typically do not decouple, so that there is effectively little simplification.

For example perhaps the simplest boundary value problem for an elastic rod is the strut, where with a standard fixed basis e_i of R^3, the appropriate boundary conditions are

$$r(0) = 0 \quad d_i(0) = e_i, \quad m(L) = 0 \quad n(L) = -\lambda e_3$$

These do not reduce to an initial value problem for the stress components $\boldsymbol{m}_i$ and $\boldsymbol{n}_i$ because to know the components $\boldsymbol{n}_i(\boldsymbol{L})$ given the boundary condition $\boldsymbol{n}(\boldsymbol{L}) = -\lambda \boldsymbol{e}_3$ it would be necessary to know the relation between the directors $\boldsymbol{d}_i(\boldsymbol{L})$ and the basis vectors $\boldsymbol{e}_i$. Thus the problem is implicitly coupled.

(We remark parenthetically that the boundary conditions

$$\boldsymbol{r}(0) = 0 \quad \boldsymbol{d}_i(0) = \boldsymbol{e}_i, \quad \boldsymbol{m}(\boldsymbol{L}) = 0 \quad \boldsymbol{n}(\boldsymbol{L}) = -\lambda \boldsymbol{d}_3$$

do lead to a decoupled system for the stress components. This is a physically interesting, much studied problem with a so-called *follower* load. Its full solution behaviour is actually much more complicated than the strut.)

Another physically interesting set of boundary conditions that we will study in our preparation for modelling DNA are

$$\boldsymbol{r}(0) = 0 \quad \boldsymbol{r}(\boldsymbol{L}) \cdot \boldsymbol{e}_1 = \boldsymbol{r}(\boldsymbol{L}) \cdot \boldsymbol{e}_2 = 0, \quad \boldsymbol{r}(\boldsymbol{L}) \cdot \boldsymbol{e}_3 = \delta$$

$$\boldsymbol{d}_i(0) = \boldsymbol{d}_i(\boldsymbol{L}) = \boldsymbol{e}_i.$$

which have *no* boundary conditions whatsoever on the stress components.

Note that for an inextensible rod, there can be no solution for $\delta > \boldsymbol{L}$

Introduction for the Planar rod

We have introduced the kinematics, balance laws, and simple linear constitutive relations for a special Cosserat rod deforming in three dimensions, for both shearable, extensible rods, and for inextensible, unshearable rods.

In the exercises, you have considered some simple equilibria when the rod is uniform (i.e. the constitutive relations have no explicit dependence on the parameter $\boldsymbol{s}$) that can be calculated analytically in closed form. However for non-uniform rods, and for many two-point boundary value problems for rods it is necessary to use numerics to determine the equilibria.

We now turn to start a description of the combination of analysis and numerical methods that we will use to compute equilibria that model DNA.

For various reasons, for example handling boundary conditions efficiently, constructing explicit variational principles, it can sometimes be convenient to explicitly parametrize the directors $\{\boldsymbol{d}_i(\boldsymbol{s})\}$, which, as previously remarked, is equivalent to parametrizing the group of proper rotation matrices or $\boldsymbol{SO}(3)$.

The classic choice (as described in nauseating detail in infinitely many mechanics texts) for such a parametrization is some set of Euler angles. We will later describe a parametrization in terms of Euler parameters (or, more or less equivalently, quaternions or Cayley-Klein parameters). However the kinematics of $SO(3)$ is unavoidably complicated in some ways.

Consequently, so as to focus first on the mathematical and computational techniques that are pertinent to our study of rods we will start with a study of the much simpler case involving planar (untwisted) deformations of rods.

We will seek solutions of the full three dimensional equilibrium equations, which happen to be untwisted and to lie in a plane.

2.6 The planar rod example

That is (for example with deformations chosen to lie in the $(\boldsymbol{e}_1, \boldsymbol{e}_3)$ plane)

$$\boldsymbol{r}' \cdot \boldsymbol{e}_2 = 0 \quad \boldsymbol{d}_2(s) \equiv \boldsymbol{e}_2$$

or equivalently

$$v_2(s) \equiv 0$$

and the Darboux vector has the simple form

$$\boldsymbol{u} = u_2 \boldsymbol{d}_2 = u_2 \boldsymbol{e}_2$$

For such directors we can parametrize with a single angle $\phi(s)$

$$\boldsymbol{d}_1 = \cos\phi\, \boldsymbol{e}_1 - \sin\phi\, \boldsymbol{e}_3, \quad \boldsymbol{d}_3 = \sin\phi\, \boldsymbol{e}_1 + \cos\phi\, \boldsymbol{e}_3\,.$$

With this sign convention we find $\boldsymbol{d}_1 = \boldsymbol{e}_1$ and $\boldsymbol{d}_3 = \boldsymbol{e}_3$ when $\phi = 0$ and moreover

$$u_2 \equiv +\phi'$$

(In the $(\boldsymbol{e}_1, \boldsymbol{e}_3)$-plane positive ϕ then corresponds to a clockwise rotation of the $\boldsymbol{d}_i$, see figure.)

It is then easy to verify that provided

$$\hat{v}_2 \equiv 0 \qquad \hat{u}_1 \equiv \hat{u}_3 \equiv 0,$$

with the diagonal constitutive relation introduced earlier, such planar kinematics generates forces and moments of the special form

$$\boldsymbol{n} \cdot \boldsymbol{e}_2 \equiv 0 \qquad \boldsymbol{m}(s) \equiv m_2(s) \boldsymbol{e}_2$$

That is, the out of plane component of the force n vanishes, while the two in-plane components of the moment m vanish.

Notice that in general $m_2(s)$ denotes the component of m with respect to the director frame, but with the special kinematics introduced here $d_2(s) \equiv e_2$, so that $m_2(s)$ is also the component of m with respect to the fixed vector e_2.

Are there equilibria of this special form? To investigate this question we merely substitute our ansatz into the balance laws.

We find that the **2**-component of the force balance law is automatically satisfied, while both the d_1 and d_3 components (or equivalently here, both of the in-plane components e_1 and e_3) of the moment balance law are automatically satisfied.

The remaining system involves two scalar equations for the in-plane components (n_1, n_3) of the force n and a third scalar equation for the out-of-plane component m_2 of the moment m, Using the parametrization $\phi(s)$ of the directors, these equations can be written in the form

$$n_1' = -n_3\phi' \qquad n_3' = \phi' n_1,$$

$$m_2' = n_3 v_1 - n_1 v_3.$$

As we shall see in a moment, it can be convenient for the treatment of boundary conditions to retain the angle $\phi(s)$ as one of our basic unknowns. Thus we shall eliminate m_2 through the constitutive relation for bending which takes the form

$$m_2 = K_2(s)(\phi' - \hat{u}_2(s))$$

where as before $K_2(s) > 0$ and $\hat{u}_2(s)$ are given functions.

The above equations form a closed system when two of the four unknowns n_1, n_3, v_1, and v_3 are eliminated through the use of the shear and extension constitutive relations (or, equivalently, the shear and extension constitutive relations are appended as additional equations)

We shall be primarily concerned with the inextensible, unshearable case in which the basic unknowns are ϕ, n_1 and n_3 and

$$v_1 \equiv \hat{v}_1 \equiv 0, \qquad v_3 \equiv \hat{v}_3 \equiv 1.$$

and the governing equations reduce to

$$n_1' = -n_3\phi' \qquad n_3' = \phi' n_1,$$

$$[\boldsymbol{K}_2(s)(\phi' - \hat{u}_2(s))]' = -n_1.$$

$$\boldsymbol{r}' = \boldsymbol{d}_3 = \sin\phi\, \boldsymbol{e}_1 + \cos\phi\, \boldsymbol{e}_3\,.$$

For this system it is actually more traditional to introduce the components N_1 and N_3 of the force $\boldsymbol{n}$ with respect to the fixed basis vectors $\boldsymbol{e}_1$ and $\boldsymbol{e}_3$. Because the angle ϕ parametrizing the directors is one of our basic unknowns it is easy to achieve this explicitly through the invertible relationships (just familiar formulas from rotations in two-dimensions)

$$n_1 = N_1\, \cos\phi - N_3\, \sin\phi,$$

$$n_3 = N_1\, \sin\phi + N_3\, \cos\phi.$$

Then (verify this!) the governing equations reduce to

$$N_1' = 0, \qquad N_3' = 0,$$

$$[\boldsymbol{K}_2(s)(\phi' - \hat{u}_2(s))]' = N_3\, \sin\phi - N_1\, \cos\phi$$

$$\boldsymbol{r}' = \boldsymbol{d}_3 = \sin\phi\, \boldsymbol{e}_1 + \cos\phi\, \boldsymbol{e}_3.$$

which can be viewed as a single second-order ODE for the angle ϕ, containing two unknown constants N_1 and N_2 (but not the centerline $\boldsymbol{r}$), along with a quadrature to determine $\boldsymbol{r}$ once ϕ is known.

This decoupling into an equation for ϕ plus unknown constants, and a quadrature for $\boldsymbol{r}$ can be rather convenient, and is one of the reasons that we introduce a parametrization of the director frame $\{\boldsymbol{d}_i\}$. However it turns out that the decoupling only truly works for certain sets of boundary conditions. We will first consider a set of boundary conditions for which there is a genuine decoupling.

The *strut* boundary conditions for an inextensible, unshearable rod of arc-length $\mathbf{1}$ are

$$\phi(0) = 0, \quad m_2(1) \equiv \boldsymbol{K}_2(\phi' - \hat{u}_2)|_{s=1} = 0,$$

$$\boldsymbol{N_3(1) = -\lambda, \quad N_1(1) = 0}$$

$$\boldsymbol{r(0) = 0}$$

Here $\boldsymbol{\lambda}$ is a parameter (either positive or negative) which is specified. It represents the end-loading. The introduction of the minus sign is a convention meaning that the end-load is positive if the load is down the $\boldsymbol{e_3}$ axis.

More generally all the boundary conditions could have parameters on the right hand side instead of zero, although there is no interest in introducing additional parameters that merely generate planar rigid body motions of other solutions (e.g. parameters in the initial values for $\boldsymbol{r(0)}$ and $\boldsymbol{\phi(0)}$, in addition to parameters for *both* $\boldsymbol{N_3(1)}$ *and* $\boldsymbol{N_1(1)}$).

We will shortly introduce a discretization of the above two-point boundary value problem, which will lead to a discrete nonlinear system to be solved numerically. Actually we will introduce two different discretizations– a numerically naive one (a very rudimentary finite difference scheme) which we will use to explain the numerical solution procedures, and a numerically robust one (collocation) which we will compute with.

The computations will be carried out with a software package called VBM for solving the system obtained after discretization via collocation. VBM is a nice GUI and visualization package that in turn implements a continuation and bifurcation package called AUTO to generate numerical approximations to the solutions of the equilibrium boundary value problem.

We will use VBM as a black box (or at least a very dark grey box) code, first for the simple planar problem being considered now, and later for rod models more closely related to DNA. There will be no need for previous experience in programming (although if you have experience you will be able to do lots of extra nice things with VBM).

The first step is to discuss what bifurcation and continuation algorithms are, in order to have some idea about the output of the code.

We start with *continuation* algorithms for solving parameter dependent problems. And the first step in any continuation algorithm, is an explicitly known solution at some simple set of parameter values.

In what circumstances does the planar rod problem above have a simple explicit solution?

Recall that the balance laws are

$$\boldsymbol{N_1' = 0, \qquad N_3' = 0,}$$

$$[\boldsymbol{K_2}(s)(\phi' - \hat{\boldsymbol{u}}_2(s))]' = \boldsymbol{N_3} \sin\phi - \boldsymbol{N_1} \cos\phi$$

$$\boldsymbol{r}' = \boldsymbol{d_3} = \sin\phi\, \boldsymbol{e_1} + \cos\phi\, \boldsymbol{e_3}.$$

One case in which there is a simple solution is when $\boldsymbol{\lambda} = \mathbf{0}$. Then

$$\boldsymbol{N_1}(s) \equiv \boldsymbol{N_3}(s) \equiv \mathbf{0},$$

$$\phi(s) \equiv \int_0^s \hat{\boldsymbol{u}}_2(\sigma)\, d\sigma \quad \boldsymbol{r}(s) \equiv \int_0^s \cos\phi(\sigma)\, d\sigma.$$

Recall that the strut boundary value problem had the boundary conditions

$$\phi(0) = 0, \quad \boldsymbol{r}(0) = \mathbf{0}$$

and

$$\boldsymbol{N_3}(1) = -\boldsymbol{\lambda}, \quad \boldsymbol{N_1}(1) = \mathbf{0},$$

$$\boldsymbol{m_2}(1) \equiv [\boldsymbol{K_2}(\phi' - \hat{\boldsymbol{u}}_2)]|_{s=1} = 0,$$

where we have re-grouped the boundary conditions firstly into kinematic conditions, and secondly into external loading conditions.

Notice that we have used *all* of the boundary conditions in deriving the explicit representation of the solution to the boundary value problem.

Such a representation is called a *solution by quadrature*—all the variables are given in terms of indefinite integrals of known functions with explicit limits. In some sense the solution is not truly explicit unless the function $\hat{\boldsymbol{u}}_2(s)$ is simple enough that the quadratures can be carried out in closed form (e.g. $\hat{\boldsymbol{u}}_2$ a constant) but that is rarely important. For example continuation codes could easily use such a quadrature representation to generate a discretized solution of any required accuracy to use as a starting point.

The quadrature representation also contains interesting physical information. When all external loading vanishes, i.e. $\boldsymbol{\lambda} = \mathbf{0}$ in the above, then the rigid body transformation of the unstressed, minimum-energy shape that is uniquely defined by the kinematic boundary conditions is a solution of the equilibrium conditions.

Less obviously the constructive nature of the quadrature solution demonstrates that it represents the *unique* solution of the boundary value problem. Whenever $\boldsymbol{\lambda}$ is non-zero it is much harder to obtain such a uniqueness result. Indeed unless $|\boldsymbol{\lambda}|$ is sufficiently small, the boundary value problem has multiple solutions, and there is no uniqueness.

Exercise: What can be said about existence and uniqueness of solutions when the external loading conditions are assumed to be of the form

$$\boldsymbol{N_3(1) = 0, \quad N_1(1) = 0,}$$

$$\boldsymbol{m_2(1) \equiv [K_2(\phi' - \hat{u}_2)]|_{s=1} = \tau,}$$

i.e. no external force loading at the end, but instead an external torque loading $\boldsymbol{\tau d_2}$?

We will re-visit uniqueness (and non-uniqueness) results later after a variational formulation of the problem has been introduced and convexity arguments can be brought to bear.

We will also have need of another special case, namely when $\boldsymbol{\hat{u}_2(s) \equiv 0}$. Then the equilibrium equations reduce to

$$\boldsymbol{N_1' = 0\,, \qquad N_3' = 0,}$$

$$\boldsymbol{[K_2(s)\phi']' = N_3 \, \sin\phi - N_1 \, \cos\phi}$$

$$\boldsymbol{r' = d_3 = \sin\phi \, e_1 + \cos\phi \, e_3.}$$

and we see that there is a whole family of solutions satisfying the strut boundary conditions, namely

$$\boldsymbol{\phi(s) \equiv 0, \quad N_1 = 0, \quad N_3 = -\lambda}$$

$$\boldsymbol{r(s) = s \, e_3}$$

which is a solution for all values of the parameter λ.

Physically this means that when you lean straight down on a perfectly straight, inextensible (or in this case incompressible), upright rod, then the undeformed configuration is an equilibrium for any load.

What does your intuition tell you will happen for sufficiently large loads? Well the strut will start bending (or in another description will buckle). Or perhaps compression effects will become non-negligible (depending on the material or constitutive relation). Thus we see that in a parameter dependent problem we always have to be aware of limitations of the range and idealizations of the model.

In point of fact the inextensible model captures the phenomenon of buckling rather well through the phenomena of bifurcation and an associated loss of stability, and also through the inclusion of imperfections (for example the presence of a rather small $\hat{u}_2(s)$ representing a nearly straight unstressed shape). Both buckling and imperfections will be crucial in the DNA minicircle example.

To understand buckling and imperfections we will retreat to an extremely simple, finite dimensional model, which can be regarded as a finite difference approximation to the planar rod problem.

2.7 The Discrete Strut–problem formulation

Consider the energy

$$\sum_{i=0}^{n} k_{i+1}(\phi_{i+1} - \phi_i - \beta_{i+1})^2/2 + \lambda h \cos\phi_i$$

where λ, h, the β_{i+1}, and the k_{i+1}, $i = 0, \ldots, n$ are given, there are $n+2$ unknown variables ϕ_i, $i = 0, \ldots, n+1$, and we have the "boundary conditions"

$$\phi_0 = \beta_0, \quad k_{n+1}(\phi_{n+1} - \phi_n - \beta_{n+1}) = 0.$$

where β_0 is an additional parameter.

The quadratic terms are a model for the energy stored in n rotational springs, with spring constant k_i at the ith joint, $i = 1, \cdots, n$ and a minimum energy at the ith joint when the angles satisfy

$$\phi_i = \phi_{i-1} + \beta_i$$

The fact that $\phi_0 = \beta_0$ is prescribed is analogous to the boundary condition $\phi(0) = \beta_0$ in the continuous problem, while the "boundary condition" involving ϕ_{n+1} is a discretized version of the vanishing moment condition $K_2(1)(\phi'(1) - \hat{u}_2(1)) = 0$, which says that there is no moment applied at the top of the segment with nth joint.

We could equally well eliminate ϕ_0 and ϕ_{n+1} using the boundary conditions to leave a sum with index range $i = 1, \ldots, n$ with n unknowns (and redefining $\bar{\beta}_1 = \beta_0 + \beta_1$) and there would be only n quadratic terms to be summed.

The term in the energy involving λ and h models the potential energy of a vertical load (with the convention λ +ve acting downward) applied to the $(n+1)$st joint (or equivalently the end of the nth segment with each segment being of length h). Think of a bucket of water of weight λ hung from the $(n+1)$st joint, then the potential energy is just the height $\sum h \cos \phi_i$ times the weight λ.

In the formulation where ϕ_{n+1} is kept as a variable it seems a little strange that the load is not applied at the tip of the $(n+1)$st segment, but it is quite usual that in a finite difference approximation, it is necessary to introduce an additional mesh point 'beyond the boundary' in order to approximate derivative boundary conditions at the boundary. In any case, in this example eliminating ϕ_{n+1} eliminates any confusion.

The Minimum Energy State

We also see that in the case of a zero external load $\lambda = 0$ the minimum energy configuration is given by

$$\hat{\phi}_i \equiv \sum_{j=0}^{i} \beta_k, \qquad \forall i = 0, \ldots, n+1$$

where we have also included indices for the boundary variables in the summation range.

Thus we can see that the β_i/h can be identified as determining a piece-wise constant approximation of the unstressed strain $\hat{u}_2(s)$ in the continuous problem, or equivalently the β_i determine a piece-wise linear approximation with corner values $\hat{\phi}_i$, of the continuous unstressed shape $\hat{\phi}(s)$ determined by the property $\hat{\phi}'(s) = \hat{u}_2(s)$.

The Equilibrium conditions

Taking partial derivatives of the energy with respect to the variables ϕ_i, $i = 1, \ldots, n$ yields the stationarity, or first-order necessary conditions

$$
\begin{aligned}
-k_{i+1}(\phi_{i+1} - \phi_i - \beta_{i+1}) \quad &+ \\
k_i(\phi_i - \phi_{i-1} - \beta_i) \quad &- \quad \lambda h \sin \phi_i = 0
\end{aligned}
$$

which is a set of n nonlinear equations for the n unknowns ϕ_i, $i = 1, \ldots, n$, but the equations also contain many additional parameters, namely the single scalars λ and h and the sets of parameters k_i, $i = 1, \ldots, n$ and β_i, $i = 0, \ldots, n$.

Notice that the boundary condition relating ϕ_{n+1} to ϕ_n is exactly the additional first-order condition that is obtained by taking the partial derivative with respect to the variable ϕ_{n+1}. This is the finite dimensional manifestation of what is called *natural boundary conditions* in the calculus of variations (which we will see in a little while).

The boundary condition is actually implied by the variational principle. In contrast it makes no sense to take a partial derivative with respect to the variable ϕ_0 because the boundary condition gives ϕ_0 a definite value in terms of the given parameter β_0. Such boundary conditions are sometimes called *imposed*, and are not implied by the variational principle.

These equations could have been derived directly as the balance (about the ith joint) of moments acting on the ith segment. Thus we see that there is a *variational formulation* or *minimization principle* giving the moment balance equations. Notice that following common practice we talk about a minimization principle, but the balance law comes as the first-order stationary condition of all partial derivatives vanishing (i.e. the potential is only known to be flat, but it could be a local minimum, a local maximum, or some sort of saddle point).

It is certainly of considerable interest to know which of all the equilibria (solutions of the balance laws) are actually local minima (which we expect to be stable equilibria). The *type* of the critical point can be classified by studying the second derivative of the potential evaluated at a particular critical point.

The Second Variation of the Strut Problem

In the case of the discrete strut potential given above, the second derivative can be regarded as the $n \times n$ symmetric matrix whose (i, j)th entry is the second partial derivative of the potential with respect to the variables ϕ_i and ϕ_j (for $i, j = 1, \ldots, n$).

For the particular potential above, the matrix is tri-diagonal with a typical (3×3) block centred on the (i, i) entry being

$$\begin{pmatrix} (k_{i-1}+k_i)-\bar\lambda h\,\cos\bar\phi_{i-1} & -k_i & 0 \\ -k_i & (k_i+k_{i+1})-\bar\lambda h\,\cos\bar\phi_i & -k_{i+1} \\ 0 & -k_{i+1} & (k_{i+1}+k_{i+2}-\bar\lambda h\,\cos\bar\phi_{i+1}) \end{pmatrix}$$

and appropriately modified first and last blocks (after ϕ_0 and ϕ_{n+1} are eliminated using the boundary conditions).

Here $\bar\phi_i$ is a particular solution of the equilibrium conditions with $\lambda = \bar\lambda$ at which the second derivative (or second variation) is being computed. The dependence of the second derivative on the the solution $(\bar\phi_1, \ldots, \bar\phi_n, \bar\lambda)$ is particularly simple because the bending part of the energy is exactly quadratic.

Similarly because the terms involving the β_i in an expansion of the energy are all linear in the ϕ_i the parameters β_i do not appear explicitly in the second derivative (although the β_i certainly appear explicitly in the first-order conditions, and so implicitly in the second derivative through the exact form of the equilibrium $\bar\phi_i$). Because the second derivative matrix is symmetric, it has only real eigenvalues, and we will call the solution, $(\bar\phi_1, \ldots, \bar\phi_n, \bar\lambda)$ *non-degenerate* if its second derivative matrix is nonsingular, i.e. 0 is not an eigenvalue of the matrix.

For non-degenerate equilibria we define its *index* to be the number of negative eigenvalues of its second derivative matrix. In particular a Taylor series expansion of the (finite dimensional) potential shows that a non-degenerate equilibrium is a local minimum of the potential (and thus presumably stable) precisely if its index is zero.

We remark that classifying degenerate equilibria is more subtle, because the second derivative cannot be used to bound all the higher order terms. (Similarly in infinite dimensional problems the situation is not so straightforward even at non-degenerate equilibria.)

The exercise sessions will ask you to compute the second derivative at some particular solutions in the case $n = 2$.

In a little while we will turn to discussing how the above discrete system, can be interpreted as a discretization of the continuous system which will be recovered in a limit $n \to \infty$, $h \to 0$. But first we will discuss methods for solving nonlinear systems of parameter-dependent algebraic equations, and understanding typical behaviours.

Solution of nonlinear parameter-dependent equations

We would like to understand the solution set of the discretized strut as the parameters of the problem are varied. In particular we shall consider $\boldsymbol{\lambda}$ as the main parameter to be varied, with the stiffnesses k_i and the discretized unstressed shape β_i being for the most part regarded as given fixed parameters. In other words we want to know what are the possible equilibrium shapes as the end load is increased say from $\lambda = 0$.

Generally speaking it is hopeless to find anything other than very special solutions of nonlinear equations analytically, and we must instead turn to numerical methods.

Because we know that our problem has an associated variational formulation, we could implement a numerical method exploiting that structure, say steepest descent, i.e. take an iterative method that at each step marches down the gradient of the potential. Such methods have some very desirable features. For example they can reasonably be expected to converge to a local minimum of the potential from an essentially arbitrary starting point—which is an extremely powerful statement.

On the other hand minimization by steepest descent or other methods has its draw backs, for example often there is slow convergence in the neighbourhood of a local minimum. For our purposes the bigest disadvantage of minimization methods is that they can only be expected to find local minima, and as we shall see it can be rather instructive to find other types of (presumably unstable) equilibria.

The only practical alternative to a minimization method is a root-finding technique applied directly to the system of equations that arises as the first-order necessary conditions. If for the moment we freeze all parameters in the first-order conditions, the first-order conditions assume the general form

$$f(z) = 0$$

of n nonlinear equations in n unknowns z.

At this level of generality the only practical way to solve such systems numerically is some sort of iterative method starting from a sufficiently good guess for the solution. Amongst such iterative schemes, those that are some variant of Newton's method are extremely powerful.

Review of Newton's method

Let us therefore review Newton's method. Suppose that z^k is some good approximation to a root z^e of the governing equations. Then a Taylor series expansion gives us

$$f(z^e) = f(z^k) + J^k(z^e - z^k) + O(|z^k - z^e|^2)$$

where J^k denotes the $(n \times n)$ Jacobian matrix with (i, j)th entry $\partial f_i / \partial z_j$ evaluated at z^k.

Neglecting the higher order terms, and using the fact that by assumption $f(z^e) = 0$, we see that the next iterate $z^{(k+1)}$ defined via

$$z^{(k+1)} = z^k + w^k$$

with w^k defined to be the solution of the inhomogeneous linear system

$$J^k w^k = -f(z^k)$$

should be a very good estimate for z^e. In fact Newton's method converges quadratically once the approximation is sufficiently close (which is an extremely good feature).

Newton's method had two major potential difficulties. First it may be either computationally expensive or difficult to solve the linear system for the update w^k either because the matrix J^k is very large and dense, or is ill-conditioned or actually singular. Second, to start the iteration you need a first guess z^0, and often that guess has to be really quite good in order that the Newton iteration will converge.

Our particular problem has special features that can help with both of these potential difficulties. First because we have a variational principle we know that the first order conditions are of the form $f(z) = 0$, but with the vector field f equal to the gradient of the scalar potential, V say.

Thus $f_i(z) = \partial V / \partial z_i$. In particular the Jacobian matrix is just the symmetric matrix of second derivatives of the potential V, which for the planar rod problem is tri-diagonal (some such banded sparsity structure of the Jacobian is typical of problems that arise as discretizations of systems of ordinary differential equations). Linear systems involving banded matrices can be solved extremely efficiently (for example a banded Gaussian elimination), whenever they are not singular.

The second special feature that we shall exploit is the parameter dependence of the problem with an explicitly known solution at one parameter value (namely the unstressed shape when $\lambda = 0$). We shall apply *parameter continuation methods* in our numerical computations of rod equilibria.

Parameter Continuation Method

That is to say we shall use knowledge of previously computed solutions at nearby parameter values, to generate an accurate guess for the solution at the current parameter values, which is then corrected via a Newton-type iterative solve. The whole process can get started provided that a solution (or even just a good guess for the solution) is known at some specific set of parameter values.

Let us make this outline more concrete. Suppose therefore that we have a system of equations of the form

$$\boldsymbol{f}(\boldsymbol{z};\boldsymbol{\lambda}) = \mathbf{0}$$

and a known solution $(\boldsymbol{z}^0, \boldsymbol{\lambda}^0)$. Then provided that the Jacobian $\boldsymbol{J}(\boldsymbol{z}^0, \boldsymbol{\lambda}^0)$ is nonsingular the implicit function theorem guarantees that there is a neighbourhood $|\boldsymbol{\lambda} - \boldsymbol{\lambda}_0| < \boldsymbol{\epsilon}$, such that there is a unique *branch* of solutions passing through $(\boldsymbol{z}^0, \boldsymbol{\lambda}^0)$ which can be parametrized as $(\boldsymbol{z}(\boldsymbol{\lambda}), \boldsymbol{\lambda})$.

How might we compute such a branch of solutions? If we seek solutions for $\boldsymbol{\lambda} = \boldsymbol{\lambda}^0 + \boldsymbol{\lambda}^1$ with $\boldsymbol{\lambda}^1$ small, we can expand the solution as

$$\boldsymbol{z}(\boldsymbol{\lambda}) = \boldsymbol{z}^0 + \boldsymbol{\lambda}^1 \boldsymbol{z}^1 + \boldsymbol{O}(|\boldsymbol{\lambda}^1|^2)$$

and expand the equation

$$\boldsymbol{f}(\{\boldsymbol{z}^0 + \boldsymbol{\lambda}^1 \boldsymbol{z}^1 + \boldsymbol{O}(|\boldsymbol{\lambda}^1|^2)\}; \{\boldsymbol{\lambda}^0 + \boldsymbol{\lambda}^1\}) = \mathbf{0}$$

as

$$\boldsymbol{f}(\boldsymbol{z}^0;\boldsymbol{\lambda}^0) + \boldsymbol{\lambda}^1 \boldsymbol{f}_{\boldsymbol{z}}(\boldsymbol{z}^0;\boldsymbol{\lambda}^0)\boldsymbol{z}^1 + \boldsymbol{\lambda}^1 \boldsymbol{f}_{\boldsymbol{\lambda}}(\boldsymbol{z}^0;\boldsymbol{\lambda}^0) = \boldsymbol{O}(|\boldsymbol{\lambda}^1|^2)$$

But $\boldsymbol{f}(\boldsymbol{z}^0;\boldsymbol{\lambda}^0) = \mathbf{0}$ as $(\boldsymbol{z}^0;\boldsymbol{\lambda}^0)$ is known to be a solution, and (as a matter of notation) $\boldsymbol{f}_{\boldsymbol{z}}(\boldsymbol{z}^0;\boldsymbol{\lambda}^0)$ is just the Jacobian $\boldsymbol{J}(\boldsymbol{z}^0, \boldsymbol{\lambda}^0)$. Thus we see that by ignoring $\boldsymbol{O}(|\boldsymbol{\lambda}^1|^2)$ terms we obtain a nonhomogeneous linear system that can be solved for a unique solution $\boldsymbol{z}^1$ provided that the Jacobian $\boldsymbol{J}(\boldsymbol{z}^0, \boldsymbol{\lambda}^0)$ is nonsingular:

$$\boldsymbol{J}^0\, \boldsymbol{z}^1 = -\boldsymbol{f}^0_{\boldsymbol{\lambda}}$$

(where the superscript $\mathbf{0}$ on $\boldsymbol{J}$ and $-\boldsymbol{f}_{\boldsymbol{\lambda}}$ indicates that the functions are to be evaluated at $(\boldsymbol{z}^0, \boldsymbol{\lambda}^0)$).

Accordingly we can reasonably expect that for $\boldsymbol{\lambda}^1$ small a Newton iterative solve of the equations at the parameter value $\boldsymbol{\lambda} = \boldsymbol{\lambda}^0 + \boldsymbol{\lambda}^1$ will converge if we take as initial guess

$$\boldsymbol{z}^g = \boldsymbol{z}^0 + \boldsymbol{\lambda}^1\, \boldsymbol{z}^1$$

where z^1 is determined by a linear solve of the system with coefficients determined by the Jacobian J^0 and the function f_λ^0 evaluated at the previously computed solution (z^0, λ^0).

Once the Newton iterate has converged, and if the Jacobian is nonsingular at the new solution, we may repeat the process to generate a new initial guess at another nearby value of the parameter.

Thus we can generate a number of solutions at a discrete set of values of the parameter λ that are close to each other, and which approximate the branch of solutions $(z(\lambda), \lambda)$ through the initially known solution (z^0, λ^0).

This is an elementary (and rather naive) description of *numerical continuation.* As you can probably imagine there is a wealth of important detail in getting the method to work efficiently, but the above captures at least the flavour of the method that is implemented in the continuation code AUTO that we will use later on (via the graphics interface VBM).

We can see that the above method is problematic whenever we hit a solution with a singular or near singular Jacobian. The analysis of what happens at such *singular points* is the domain of *bifurcation theory.* The typical behaviour at simple singular points is that the branch of solutions persists, but there may be multiple branches of solutions, or there may be a fold point at which the parameter λ has a local extreme value along the branch of solutions.

We will illustrate these concepts by considering the case $n = 1$ of the discrete strut example. The $n = 1$ case can be solved in an essentially explicit fashion, but it is a little too simple to truly explain the perturbation expansion described above. You will consider the case $n = 2$ in the exercise session.

2.8 The continuous limit of the discrete strut

Introduction

Today we will briefly re-consider the n degree of freedom example, consider the limit $n \to \infty$ to recover the continuous case, do some bifurcation analysis on the continuous problem, and then describe a different discretization of the continuous problem (using collocation) that is used by the numerical continuation package AUTO.

Remember the expression for the discrete strut energy:

$$\sum_{i=0}^{n} k_{i+1}(\phi_{i+1} - \phi_i - \beta_{i+1})^2/2 + \lambda h \cos \phi_i$$

where λ, h, the β_{i+1}, and the k_{i+1}, $i = 0, \dots, n$ are given, there are $n+2$ unknown variables ϕ_i, $i = 0, \dots, n+1$, and we have the "boundary conditions"

$$\phi_0 = \beta_0, \quad k_{n+1}(\phi_{n+1} - \phi_n - \beta_{n+1}) = 0.$$

where β_0 is an additional parameter.

The Continuous Limit

1. Moment Balance Equation

The discrete moment balance equilibrium condition

$$-k_{i+1}(\phi_{i+1} - \phi_i - \beta_{i+1}) + k_i(\phi_i - \phi_{i-1} - \beta_i) - \lambda\, h\, \sin\phi_i = 0$$

can be recognized as a simple finite difference discretization of the continuum moment balance equation in the form

$$-\left(K_2(s)(\phi'(s) - \hat{u}_2(s))\right)' - \lambda \sin\phi(s) = 0$$

with discretization size $h = 1/n$ and

$$k_i \equiv K_2(i\, h)/h, \quad \beta_i \equiv \hat{u}_2(i\, h)\, h$$

.

For the extreme values of the index i we also recover a finite difference discretization of the boundary conditions

$$\phi(0) = 0 \quad K_2(1)(\phi'(1) - \hat{u}_2(1)) = 0$$

We can also see that the discrete energy

$$\sum_{i=0}^{n} k_{i+1}(\phi_{i+1} - \phi_i - \beta_{i+1})^2/2 + \lambda h\, \cos\phi_i$$

with discretization size $h = 1/n$ and

$$k_i \equiv K_2(i\, h)/h, \quad \beta_i \equiv \hat{u}_2(i\, h)\, h$$

is a finite difference approximation to the the continuous energy

$$\int_0^1 K_2(\phi' - \hat{u}_2)^2/2 + \lambda \cos\phi\; ds$$

(we will return to this observation in a little bit).

2. The Discrete Force Balance Equations

A question that we have so far ignored, is what happened to the force balance laws in our discrete variational formulation? It is reasonable that we can obtain a variational principle just in terms of the discretized angle, because for the strut problem, in the continuous version the force balance laws and the moment balance decouple.

Explicitly

$$N_1' = 0, \quad N_3' = 0$$

taken with the boundary conditions

$$N_1(1) = 0, \quad N_3(1) = -\lambda$$

imply $N_1(s) \equiv 0$ and $N_3(s) \equiv -\lambda$.

And the kinematic incompressibility and unshearability constraints and boundary conditions written in the form

$$x' = \sin\phi,\ x(0) = 0,\ z' = \cos\phi,\ z(0) = 0$$

(with the notation $\boldsymbol{x}(s) = x(s)\boldsymbol{e}_1 + z(s)\boldsymbol{e}_3$) allowed us to reduce the whole equilibrium problem to the solution for ϕ and m_2 of the (continuous) moment balance equation with boundary conditions and parameter λ, after which the centreline was reconstructed.

It will be instructive to describe the analogous procedure for the discretized system. The finite difference approximation to the kinematic constraint equations are for $i = 1, \ldots, n$

$$x_i - x_{i-1} = h\ \sin\phi_{i-1},\ x_0 = 0$$

$$z_i - z_{i-1} = h\ \cos\phi_{i-1},\ z_0 = 0$$

And the appropriate discretized variational principle is to 'minimize' the elastic potential energy in the form

$$\sum_{i=0}^{n} k_{i+1}(\phi_{i+1} - \phi_i - \beta_{i+1})^2/2 + \lambda z_n$$

subject to the same boundary conditions as before

$$\phi_0 = \beta_0,\ k_{n+1}(\phi_{n+1} - \phi_n - \beta_{n+1}) = 0.$$

but now also subject to the discretized kinematic constraints (and boundary conditions).

For such a finite dimensional constrained variational principle we have (at least formally) a Lagrange multiplier rule: to each constraint we can associate a Lagrange multiplier say ν_i and μ_i, $i = 1, \ldots, n$ and constrained critical points will satisfy the first-order conditions for the function

$$\sum_{i=0}^{n} k_{i+1}(\phi_{i+1} - \phi_i - \beta_{i+1})^2/2 + \nu_{i+1}(x_{i+1} - x_i - h\ \sin\phi_i) + \mu_{i+1}(z_{i+1} - z_i - h\ \cos\phi_i) + \lambda$$

with respect to the unknowns ϕ_i, x_i and z_i treated as independent variables (i.e. ignoring the constraint equations) for some values of the multipliers ν_i and μ_i. For consistency here we set $\nu_{n+1} = \mu_{n+1} = 0$ because there are only $2n$ constraints, and $2n$ unknowns x_i and z_i.

The first-order conditions with respect to x_i and z_i yield the system of equations

$$-\nu_{i+1} + \nu_i = 0, \quad i = 1, \ldots, (n-1), \quad \nu_n = 0,$$

and

$$-\mu_{i+1} + \mu_i = 0, \quad i = 1, \ldots, n-1, \quad \mu_n = -\lambda$$

which can be recognized as finite difference approximations to the two force balance equations, and their associated boundary conditions.

Just as in the continuous case it is trivial to solve the force balance conditions, and we obtain

$$\nu_i = 0, \quad \mu_i = -\lambda, \quad i = 1, \ldots, n$$

For the full variational principle, the first-order conditions for the ϕ_i regarded as being completely independent of the x_i and z_i, are

$$-k_{i+1}(\phi_{i+1} - \phi_i - \beta_{i+1}) + k_i(\phi_i - \phi_{i-1} - \beta_i) + \mu_{i+1}\, h\ \sin\phi_i - \nu_{i+1}\, h\ \cos\phi_i = 0.$$

And when we use the closed form expression to eliminate the multipliers, we re-obtain the first-order necessary conditions for the variational principle involving the angles only.

It should be noted that the strut boundary value problem is very special. Just about every piece of the problem enters in just the right way for this decoupling to arise. For different boundary conditions, or a z dependent potential term the problem remains coupled.

Nevertheless in general a variational principle is possible whenever the constitutive relations are hyper-elastic. The variational principle involves minimizing the internal elastic energy (or the integral of the strain energy density function) plus any additional potential terms which arise from end-loadings say, and subject to some set of imposed boundary conditions and constraints.
Exercise: Construct the variational principle for the strut boundary value problem (i.e. the same boundary conditions and end-loading), but now for a rod that has a diagonal shearable, extensible constitutive relation.

2.8.1 A Little Bifurcation Theory For The Continuous Problem

Recall that the equilibrium conditions are

$$N_1' = 0, \qquad N_3' = 0, \quad N_1(1) = 0, \quad N_3(1) = -\lambda$$

$$[K_2(s)(\phi' - \hat{u}_2(s))]' = N_3 \sin\phi - N_1 \cos\phi$$

$$\phi(0) = 0 \quad K_2(1)(\phi'(1) - \hat{u}_2(1)) = 0$$

$$x' = \sin\phi, \; x(0) = 0, \; z' = \cos\phi, \; z(0) = 0$$

When does this system have a singular Jacobian? Or in other words if we linearize about a known solution, when can the reuslting linear system have a non-trivial solution?

In general this is a different calculation for each known solution. We have two possibilities. First when $\lambda = 0$ we have the unstressed shape.

Exercise: Verify that the linearization at the unstressed shape is always nonsingular so that (assuming the appropriate infinite dimensional version of the implicit function theorem), there is locally a unique branch of solutions through the zero load configuration.

The second possibility is that when $\hat{u}_2(s) \equiv 0$, there is the upright straight solution for all λ, and we can linearize about each of those solutions. Motivated by the discrete problem we would anticipate finding λ vlues when there are non-trivial solutions (i.e. the Jacobian is singular). This can be shown to be true explicitly in the further special case $K_2(s) \equiv 0$ by calculating the sequence of λ_j, $j = 1, 2, \ldots$, called bifurcation points or buckling loads, at which the Jacobian is singular.

It can be shown that a countable sequence of λ_j, $j = 1, 2, \ldots$ of bifurcation points actually exists for any $K_2(s) > 0$, and you can estimate their location analytically, but you cannot (in general) find them explicitly. (This is the theory of ODE called spectral theory for Sturm-Liouville operators.) We will also return to this problem in a little while.

Of course we can numerically compute the bifurcation points to any reasonable accuracy with a good code.

2.8.2 The Collocation Discretization

We will discretize using (Gaussian) collocation. Roughly speaking this means approximating each unknown function by a continuous piece-wise polynomial of a specified degree (say 3 or 4 or 5) on a given number of sub-intervals (say 20 or 30).

When a function is approximated by a piece-wise polynomial, its derivative in the interior of the sub-intervals is also approximated, merely by differentiating the piece-wise polynomial. We get a finite dimensional system of equations to solve for the coefficients of the piece-wise polynomial merely by enforcing the differential equation at an appropriate number of interior points in each sub-interval.

Of course there is a big numerical analysis theory of collocation, giving error bounds and identifying the 'best' points at which to collocate (the Gaussian quadrature points of course!)

Collocation can be viewed as beoing intermediate between finite difference and finite element discretizations. For example piecewise linear collocation is simply related to first order differencing of derivatives. While collocation can also be viewed as a finite element method where the trial functions are piece-wise polynomial, and the test functions are Dirac deltas (centred on the collocation points).

For our later purposes Gaussian collocation is very nice because it can be shown to be equivalent to a symplectic Runge-Kutta method (say marching in from the left) for the values of the functions at the boundaries between sub-intervals. And when the underlying equations are Hamiltonian, symplectic R-K methods exactly conserve the value of any integrals of the continuous Hamiltonian system that are quadratic functions of the unknowns (ie the phase variables).

2.8.3 VBM and AUTO

The package AUTO implements collocation (even with an adapting mesh), exploits special linear algebra for the sparsity structure of the Jacobian associated with the collocation discretization, and has routines both for continuing around fold points in a branch of solutions, and for switching branches at bifuraction points.

As with any large code you must use AUTO with caution, until you get to know its strengths and limitations.

For example bifurcation points arise when the Jacobian is singular, but it is computationally expensive to monitor singularity. So AUTO tracks the sign of the determinant. This is much faster, and most of the time catches points where the Jacobian is singular. However in the atypical case that the Jacobian has a double zero eigenvalue, the Jacobian can go singular without the determinant vanishing and AUTO typically misses any such bifurcation points.

There are several other trade-offs between speed in the typical case, and robustness to even the most degenerate case, which can give misleading answers when the atypical cases happen to arise.

The resolution is to go ahead and use the code, but always to be questioning and testing it to the full extent possible.

Generally our experience is that AUTO is pretty robust.

VBM is an "easy to use" interface to AUTO that has two modes. (In fact VBM is designed to be coupled to many different continuation packages, but we will only use AUTO.)

The post-processing mode allows one to use interactive visualization to understand bifurcation diagrams, and particular solutions on the bifurcation diagram that have already been computed.

The computation mode allows you to interactively extend bifurcation diagrams through additional computations using AUTO.

In point of fact the planar strut examples you will see today are so simple that they don't really justify all the overhead you have to learn to use VBM. However the objective is to learn and build confidence in VBM/AUTO before using it in the three dimensional rod examples that we will use to model DNA.