

Elasticity & Discretization

Jan 14, 2014

Schedule

- First round schedule is almost complete
 - if you haven't picked a paper yet, email me ASAP!
- We'll have two round of presentations.
- For final third of the course, longer group discussions instead.

Elasticity

Elasticity

An elastic object is one that, when deformed, seeks to return to some reference or rest configuration.

Previously: discrete mass/spring models.

Today: more principled, “continuum mechanics” approach.


Generalize 1D elasticity (springs) to 3D objects.

Will roughly follow Sifakis’ SIGGRAPH course:


<http://run.usc.edu/femdefo/sifakis-courseNotes-TheoryAndDiscretization.pdf>

Continuum Mechanics

View the material under consideration as a continuous mass, rather than a collection of particles/atoms.



v.s.



Useful model for both solids and fluids.

Not always applicable: e.g., at tiny scales, during some kinds of fracture, etc.

Elasticity - Springs


Recall: (Linear) spring force is dictated by displacement, $\Delta x = L - L_0$, away from rest length (Hooke's law).

$$F = -k\Delta x$$

This follows from its potential energy:

$$U = \frac{1}{2}k(\Delta x)^2$$

i.e., force acts to drive potential energy towards zero.



Conservative Forces


A spring force is an example of a conservative force – it depends only on the current state (path-independent).

For such forces, given potential energy U , the force is $F = -\nabla U$.

Thus, we seek a potential energy that is zero when our 3D object is undeformed.

Elasticity – 3D

How can we generalize these notions to three-dimensional *volumes* of material?



First, need a way to describe 3D deformations.

Deformation Map


A function $\vec{\phi}$ that maps points from the reference configuration (\vec{X}) to current position in world space (\vec{x}).

$$\vec{\phi}: \mathbb{R}^3 \rightarrow \mathbb{R}^3$$

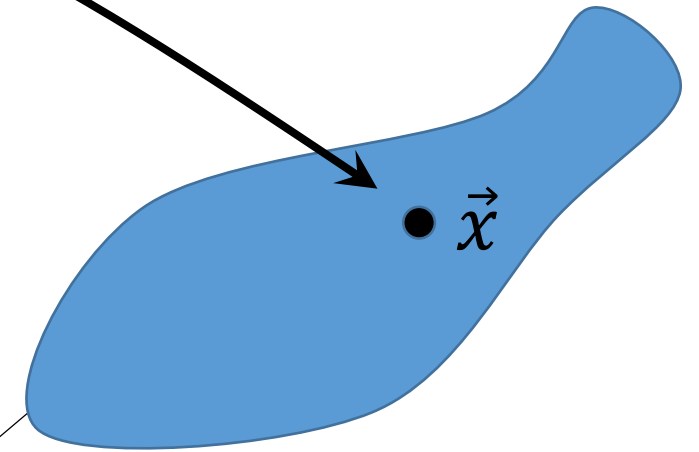
Similar to the state/transform of a rigid body, except points in the body may all have *different* transformations.

Deformation Map, $\vec{\phi}$

$$\vec{\phi}: \mathbb{R}^3 \rightarrow \mathbb{R}^3$$




Reference/rest/undeformed
configuration:




World/deformed configuration:

Deformation Gradient, $F = \frac{\partial \vec{\phi}}{\partial \vec{X}}$

$$\vec{\phi}: \mathbb{R}^3 \rightarrow \mathbb{R}^3$$



Reference/rest/undeformed configuration:



World/deformed configuration:

Deformation Gradient

Shape changes are indicated by local differences in the deformation map, ϕ .

This is reflected in the *deformation gradient* (wrt. reference configuration X .)

$$\mathbf{F} = \frac{\partial \vec{\phi}}{\partial \vec{X}} = \begin{pmatrix} \frac{\partial \phi_1}{\partial X_1} & \frac{\partial \phi_1}{\partial X_2} & \frac{\partial \phi_1}{\partial X_3} \\ \frac{\partial \phi_2}{\partial X_1} & \frac{\partial \phi_2}{\partial X_2} & \frac{\partial \phi_2}{\partial X_3} \\ \frac{\partial \phi_3}{\partial X_1} & \frac{\partial \phi_3}{\partial X_2} & \frac{\partial \phi_3}{\partial X_3} \end{pmatrix}$$

Deformation Gradient

For some offset position from \vec{X} , say $\vec{X} + \overrightarrow{dX}$, what is the corresponding world position?

$$\vec{x} + \overrightarrow{dx} = \vec{\phi}(\vec{X} + \overrightarrow{dX}) \approx \vec{\phi}(\vec{X}) + \frac{\partial \vec{\phi}}{\partial \vec{X}} \overrightarrow{dX} = \vec{x} + \mathbf{F} \overrightarrow{dX}$$

Taylor expand... Translation Relative deformation

Deformation gradient describes how particle positions have changed relative to one another.

Deformation Gradient

Examples of deformations:

Translation: $\vec{x} = \vec{\phi}(\vec{X}) = \vec{t} + \vec{X}$ implies $F = I$.

Uniform Scaling: $x = \phi(X) = sX$ implies $F = sI$.

Rotation: $x = \phi(X) = RX$ implies $F = R$.

One Possible Potential Energy

What if we use F directly to construct a potential energy?

$$U(F) = \frac{k}{2} \|F - I\|_F^2$$

Resulting forces will drive F towards I , i.e., a deformation that is (just) a translation.

What's wrong with this?

Strain Measures

Want a deformation measure that *ignores* rotation (and translation), but captures other deformations.

Can we extract this from F ?

Recall: Rotation matrices are *orthogonal*, $R^T R = I$.

So a useful measure is Green/Lagrange strain tensor, $E = \frac{1}{2} (F^T F - I)$.

Ignores translation *and* rotation, retains shear/stretch/compression.

Strain Measures

But, Green strain is nonlinear (quadratic), so more costly.

For *small* deformations, use small/infinitesimal/Cauchy strain:

$$\epsilon = \frac{1}{2} (F^T + F) - I$$

(A linearization of Green strain.)

Many other strain tensors exist (and these two have many names)...

Equations of Motion

Consider $F=ma$ for a small, continuous blob of material.

$$\int_{\Omega} F_{body} dX + \int_{\partial\Omega} T dS = \int_{\Omega} \rho \ddot{x} dX$$

F : body forces that act throughout the material (e.g. gravity, magnetism, etc.) force per unit *volume* (i.e., force density).

T : tractions, i.e., force per unit *area* acting on a surface.

Ω is the region of material being considered.

Traction


Traction T is the force (vector) per unit area on a small piece of surface.

$$\vec{T}(\vec{X}, \vec{n}) = \lim_{A \rightarrow 0} \frac{\vec{F}}{A}$$

Cauchy's postulate:

Traction is a function of position \vec{X} and normal \vec{n} .

i.e., doesn't depend on curvature or other properties.




Consists of normal/pressure component along \vec{n} , and tangential/shear components perpendicular to it.

Traction

Consider the *internal* traction on any slice through a volume of material.

This describes the forces acting on this plane between the two “sides”.

Note: $T(x,n) = -T(x,-n)$



Traction


We can characterize internal forces by considering tractions on 3 perpendicular slices (i.e., normals along x , y , z directions).

3 components per traction along 3 axes gives us 9 components.

This gives us the Cauchy stress tensor, σ .

Traction on any plane can be recovered with:

$$T = \sigma n$$



Stress

The 3x3 stress tensor describes the forces acting within a material.


$$\sigma = \begin{bmatrix} \sigma_{xx} & \sigma_{xy} & \sigma_{xz} \\ \sigma_{yx} & \sigma_{yy} & \sigma_{yz} \\ \sigma_{zx} & \sigma_{zy} & \sigma_{zz} \end{bmatrix}$$

σ can be shown to be symmetric, i.e., $\sigma_{yx} = \sigma_{xy}$, etc. (from conservation of angular momentum.)

Stress – Meaning?

Diagonal components correspond to compression/extension.

Off-diagonal components correspond to shears.



Equations of motion

$$\int_{\Omega} F_{body} dX + \int_{\partial\Omega} T dS = \int_{\Omega} \rho \ddot{x} dX$$

- Plug in $T = \sigma n \dots$


$$\int_{\Omega} F_{body} dX + \int_{\partial\Omega} \sigma n dS = \int_{\Omega} \rho \ddot{x} dX$$

- Integrate by parts (divergence theorem) to eliminate surface integral:

$$\int_{\Omega} F_{body} dX + \int_{\Omega} \nabla \cdot \sigma dX = \int_{\Omega} \rho \ddot{x} dX$$

In the limit, $F_{body} + \nabla \cdot \sigma = \rho \ddot{x}$, for every infinitesimal point.

Big Picture – So Far

- Deformation map ϕ describes map from rest to world state
 - Deformation gradient F describes deformations (minus translation)
 - Strains ϵ or E describe deformation (minus rotation)
 - Stress σ describes forces in material
 - PDE $F_{body} + \nabla \cdot \sigma = \rho \ddot{x}$ describes how to apply stress to get motion
 - (Later, will discretize the PDE to get discrete equations to solve.)
- 
- Missing step!

Constitutive models

Strain E/ ϵ describes deformations of a body.

Stress σ describes (resulting) forces within a body.

Constitutive models dictate the stress-strain *relationship* in a material. (Why rubber responds differently than concrete.)

i.e., Given some deformation, what stresses (forces) does it induce?

(e.g., linear spring force, $F=-kx$.)



Linear elasticity - simplest isotropic model

Hooke's law in 3D, for small strain, ϵ .

Potential Energy:

$$U(F) = \mu \epsilon : \epsilon + \frac{\lambda}{2} \text{tr}^2(\epsilon)$$

Stress:

$$\sigma = 2\mu\epsilon + \lambda \text{tr}(\epsilon)I$$

μ, λ are the *Lamé parameters*, one choice of “elastic moduli”.

“tr” is the trace operator (sum of diagonals)

“:” is a tensor double dot product, where $A:B = \text{tr}(A^T B)$

Linear Elasticity

- Can see it is linear by expressing in matrix/vector form
- Flatten 3x3 tensors ϵ and σ into vectors.
- Isotropy and symmetry of ϵ/σ reduce 81 coeffs down to 2 independent parameters.

$$\begin{bmatrix} \sigma_{xx} \\ \sigma_{xy} \\ \sigma_{xz} \\ \sigma_{yx} \\ \sigma_{yy} \\ \sigma_{yz} \\ \sigma_{zx} \\ \sigma_{zy} \\ \sigma_{zz} \end{bmatrix} = [9 \times 9 \text{ coefficient matrix}] \begin{bmatrix} \epsilon_{xx} \\ \epsilon_{xy} \\ \epsilon_{xz} \\ \epsilon_{yx} \\ \epsilon_{yy} \\ \epsilon_{yz} \\ \epsilon_{zx} \\ \epsilon_{zy} \\ \epsilon_{zz} \end{bmatrix}$$

Other Elastic moduli

More common/intuitive (but interconvertible) parameter pair is *Poisson's ratio*, ν , and *Young's modulus*, E , or Y . (Careful overloading E).

$$\mu = \frac{Y}{2(1 + \nu)}$$

and...

$$\lambda = \frac{Y\nu}{(1 + \nu)(1 - 2\nu)}$$


Elastic Moduli – Young's modulus

Young's modulus:

- Ratio of stress-to-strain along an axis.
- Should be consistent with linear spring.

Elastic moduli – Poisson's ratio

- Poisson's ratio is negative ratio of transverse to axial strain.
 - If stretched in one direction, how much does it compress in the others?
 - Expresses tendency to preserve volume.
 - 0.5 = incompressible (e.g., rubber)
 - 0 = no compression (e.g., cork)
- Negative is possible, though weird...



The “linear” in linear elasticity

- Describes the stress-strain relationship.
- But, strain itself could still be either linear (small strain, ϵ) or nonlinear (Green strain, E) in the deformation.

Use E instead of ϵ with the same equations gives:

$$U(F) = \mu E : E + \frac{\lambda}{2} \text{tr}^2(E)$$

Better for larger deformations/rotations. (AKA St. Venant Kirchhoff model)

Other models

- Corotational linear elasticity:
 - Try to factor out the rotational part of strain, treat the remainder with linear elasticity.
 - We'll see this idea in the "Interactive Virtual Materials" paper
- Neo-Hookean elasticity:
 - St.V-K breaks down under large compression (stops resisting)
 - Neo-Hookean is a nonlinear model that corrects this

Common Discretization Methods

Discretization

Need to turn our continuous model...

$$F_{body} + \nabla \cdot \sigma = \rho \ddot{x}$$

...into a discrete model that approximates it (and thus can be computed).

Common choices: Finite difference (FDM), finite volume (FVM), and finite element methods (FEM).

I'll give a brief flavour, but...

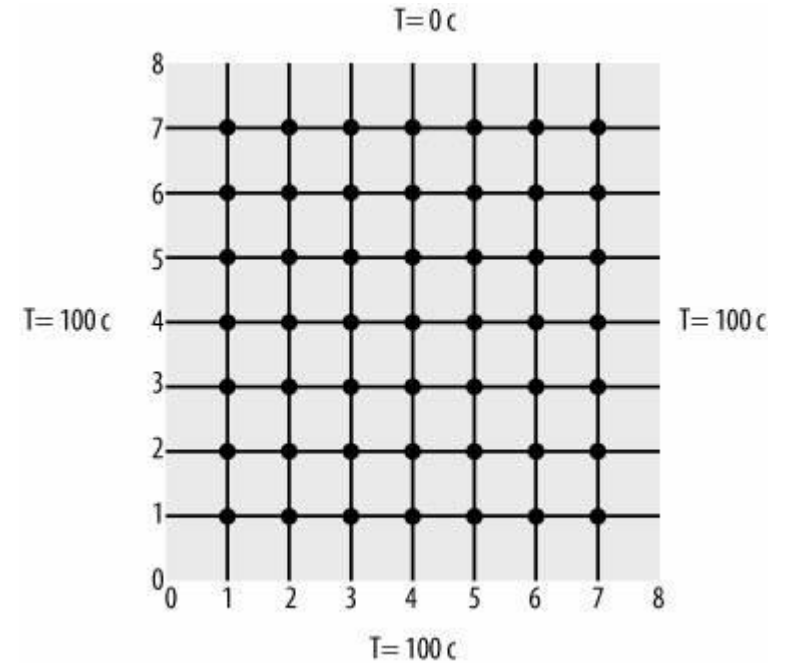
- There's a truly vast literature. (See e.g. Numerical PDE course, CS778.)

Discretization

- Notice: Time integration schemes (FE, RK2, BE, etc.) are just discretizations of time derivatives, along the 1D time axis.
- We will distinguish time discretization from spatial discretization.

Finite differences

Dice the domain into a grid of sample points holding the relevant data.



Replace all (continuous) derivatives with finite approximations.

e.g.,

$$\frac{dy}{dx} \approx \frac{y(x + \Delta x) - y(x)}{\Delta x}$$



Finite differences

Quite common for fluids... less so for elastic solids.

Some graphics papers use FDM for solids: e.g.,
“An efficient multigrid method for the simulation of high-resolution elastic solids”

Advantages: maybe simpler, grid structure offers various optimizations, cache coherent memory accesses...

Disadvantages: trickier for irregular shapes, boundary conditions



Finite volume

- Divide the domain up into a set of non-overlapping “control volumes.”
- Could be irregular, tetrahedra, hexahedra, general polyhedral, etc.
- Apply the relevant equations to each discrete control volume.




Figure from the *DistMesh* gallery:
http://persson.berkeley.edu/distmesh/gallery_images.html

Finite volume


Instead of differential/strong form, return to the integral form of equations...

$$\int_{\Omega} F_{body} dX + \int_{\Omega} \nabla \cdot \sigma dX = \int_{\Omega} \rho \ddot{x} dX$$

Convert divergence terms into *surface* integrals by divergence th'm.

e.g.

$$\int_{\Omega} \nabla \cdot \sigma dX = \int_{\partial\Omega} \sigma \cdot n dS$$
$$\approx \sum_{faces f} (\sigma_f \cdot n_f) L_f$$



Integrate remaining terms to get volume-averaged quantities per cell.

Finite volume

Convenient for conserved quantities:


- Exact “flow” leaving one cell enters the next.
 - E.g. liquid volume,

Easily applied to irregular shapes.

- ...**IF** you have a mesh of the domain.

Particularly common in fluids/CFD.

See e.g., “Finite Volume Methods. for the Simulation of Skeletal Muscle” for a nice step by step description of FVM applied to elasticity in graphics.



Finite element methods

Core idea: Can't solve the infinite dimensional, continuous problem – instead find a solution that we *can* represent, in some finite dimensional subspace.




Concretely, choose a representation of functions on a discrete mesh, and we'll try to find the “best” solution that it can describe.

Finite elements – basis functions

In 1D, consider the space of functions representable by (piecewise) linear interpolation on a set of grid points.

Just a linear combination of scaled and translated “hat” functions at each gridpoint, called a basis function.

Many others bases possible (e.g. higher order polynomials).



Finite elements

Then, any function u in this space can be described by:

$$u(x) = \sum_{k=1}^n u_k v_k(x)$$

where u_i are the coefficients, and $v_k(x)$ are the basis functions, (“hats” in our case.)


To find a solution to a problem, want to find the discrete set of coefficients, u_k . (Recover the actual shape by interpolation.)

Higher dimensional spaces


This generalizes to higher dimensions.

e.g., two dimensions:

2D mesh with numbered nodes.



2D linear basis function.



Finite elements

1D model problem: $\frac{d^2u}{dx^2} = f$ on $[0,1]$, with $u(0) = 0, u(1) = 0$.

For given f , find u .

For a proper solution, it will also be true that

$$\int \frac{d^2u}{dx^2} v dx = \int f v dx$$

for “test functions” v (that are smooth and satisfy the BC).

Finite elements

Integrate LHS by parts to get:

$$\int \frac{du}{dx} \frac{dv}{dx} dx = \int f v dx$$

This is called the *weak form* of the PDE.

Now, we will replace u , f , and v with our space of discrete, piece-wise linear functions.

Finite elements

Specifically:

- $u(x) = \sum_{k=1}^n u_k v_k(x)$
- $f(x) = \sum_{k=1}^n f_k v_k(x)$
- $v(x) = v_j(x)$ for $j = 1$ to n

From our simple (hat) basis functions, we can exactly find the inner products:

$$\langle v_j, v_k \rangle = \int v_j v_k dx$$
$$\phi(v_j v_k) = \int \frac{dv_j}{dx} \frac{dv_k}{dx} dx$$

Finite elements

Plug in, do some manipulation, and $\int \frac{du}{dx} \frac{dv}{dx} = \int f v$ becomes a set of n *discrete equations* of the form:

$$\sum_{k=1}^n u_k \underbrace{\phi(v_j, v_k)}_{\text{Known inner products}} = \sum_{k=1}^n f_k \underbrace{\langle v_j, v_k \rangle}_{\text{Known inner products}}$$

Unknown coefficients

Known input data

Final system


Letting \mathbf{u} be the vector of unknown coefficients, and \mathbf{b} the RHS vector, this becomes a matrix equation:

$$\mathbf{L}\mathbf{u} = \mathbf{b}$$

where the entries of L are just the $\phi(v_j, v_k)$'s we defined.

See paper “Graphical Modeling and Animation of Brittle Fracture” for details of an early application of FEM to elasticity in graphics.

Example



Discretization

- For FDM/FVM/FEM, much like mass springs, we get a (possibly nonlinear) system of equations to solve for data stored on a discrete mesh/grid.
- However:
 - we can use physically meaningful parameters.
 - as the mesh resolution increases, we can approach a true/real solution.
 - behaviour is independent of the mesh structure (under refinement!)