Variational Methods & Optimal Control

lecture 17

Matthew Roughan
<matthew.roughan@adelaide.edu.au>

Discipline of Applied Mathematics
School of Mathematical Sciences
University of Adelaide

April 14, 2016

Non-fixed end point problems

What happens when we don't fix the end-points of an extremal? In this case **natural boundary conditions** are automatically introduced, and these can allow us to solve the E-L equations.

Non-fixed end point problems

What happens when we don't fix the boundary points?

There are real problems like this, for instance

- a freely supported beam
  end points fixed, but not derivatives
- a beam supported at only one end
  one end point and derivative fixed, other free
- shortest path between two curves
  end points lie of curves, but not fixed
- rocket changing between two orbits
  end points lie on curves, and path is tangent
  to the two orbits.

We then get **natural boundary conditions**

Free end points: Fixed x, Free y and/or y'

First we'll consider what happens when we allow y or y' to vary at the end-points, but we still keep the x values of the end-points fixed at x_0 and x_1.

Example: freely supported beam

Freely supported beam

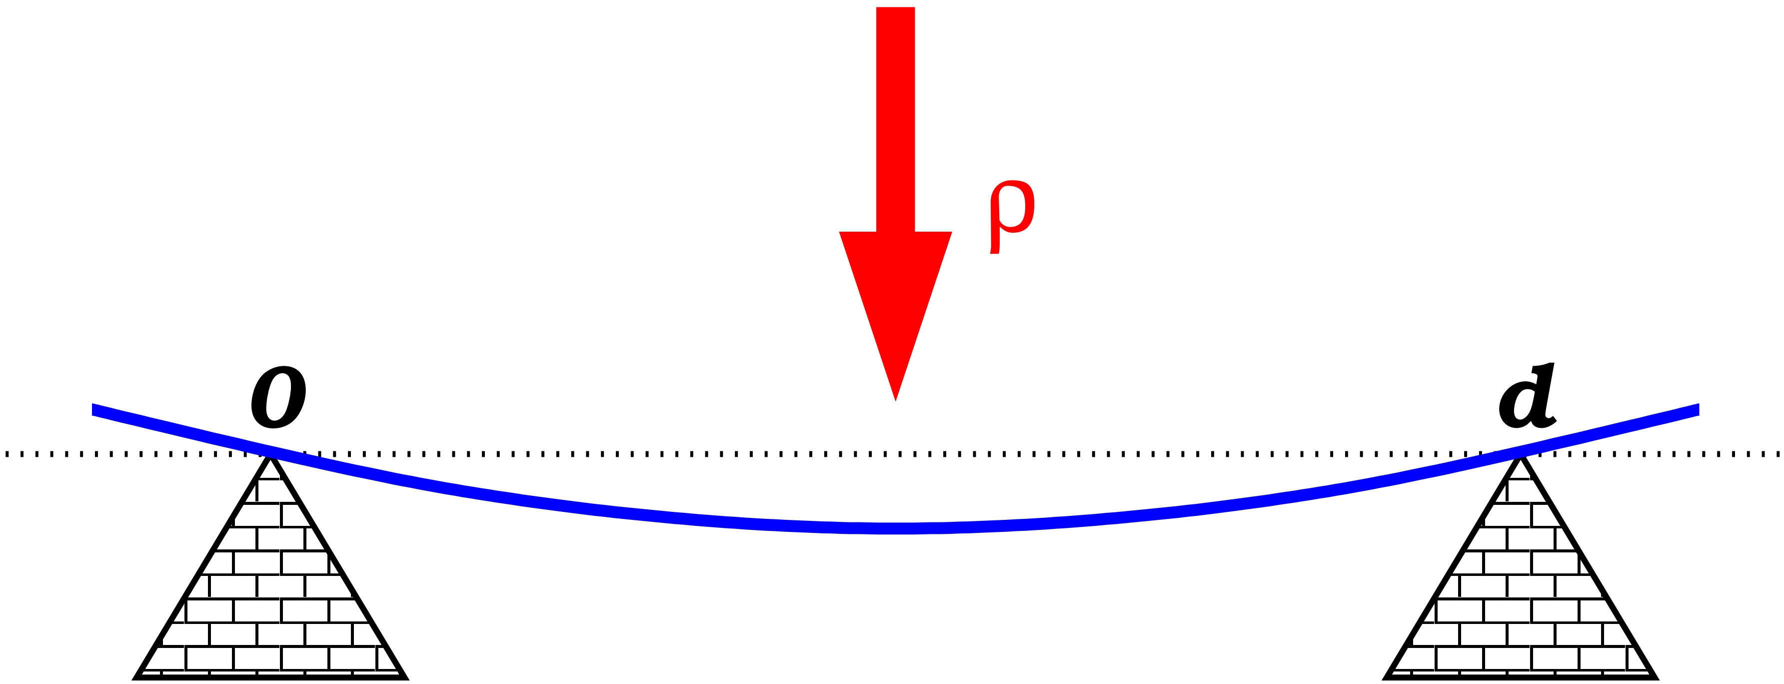

For the beam problems considered before, we had to specify the derivative at the boundary, but here it can vary.

Example: beam fixed at one end point

Beam fixed at one end point

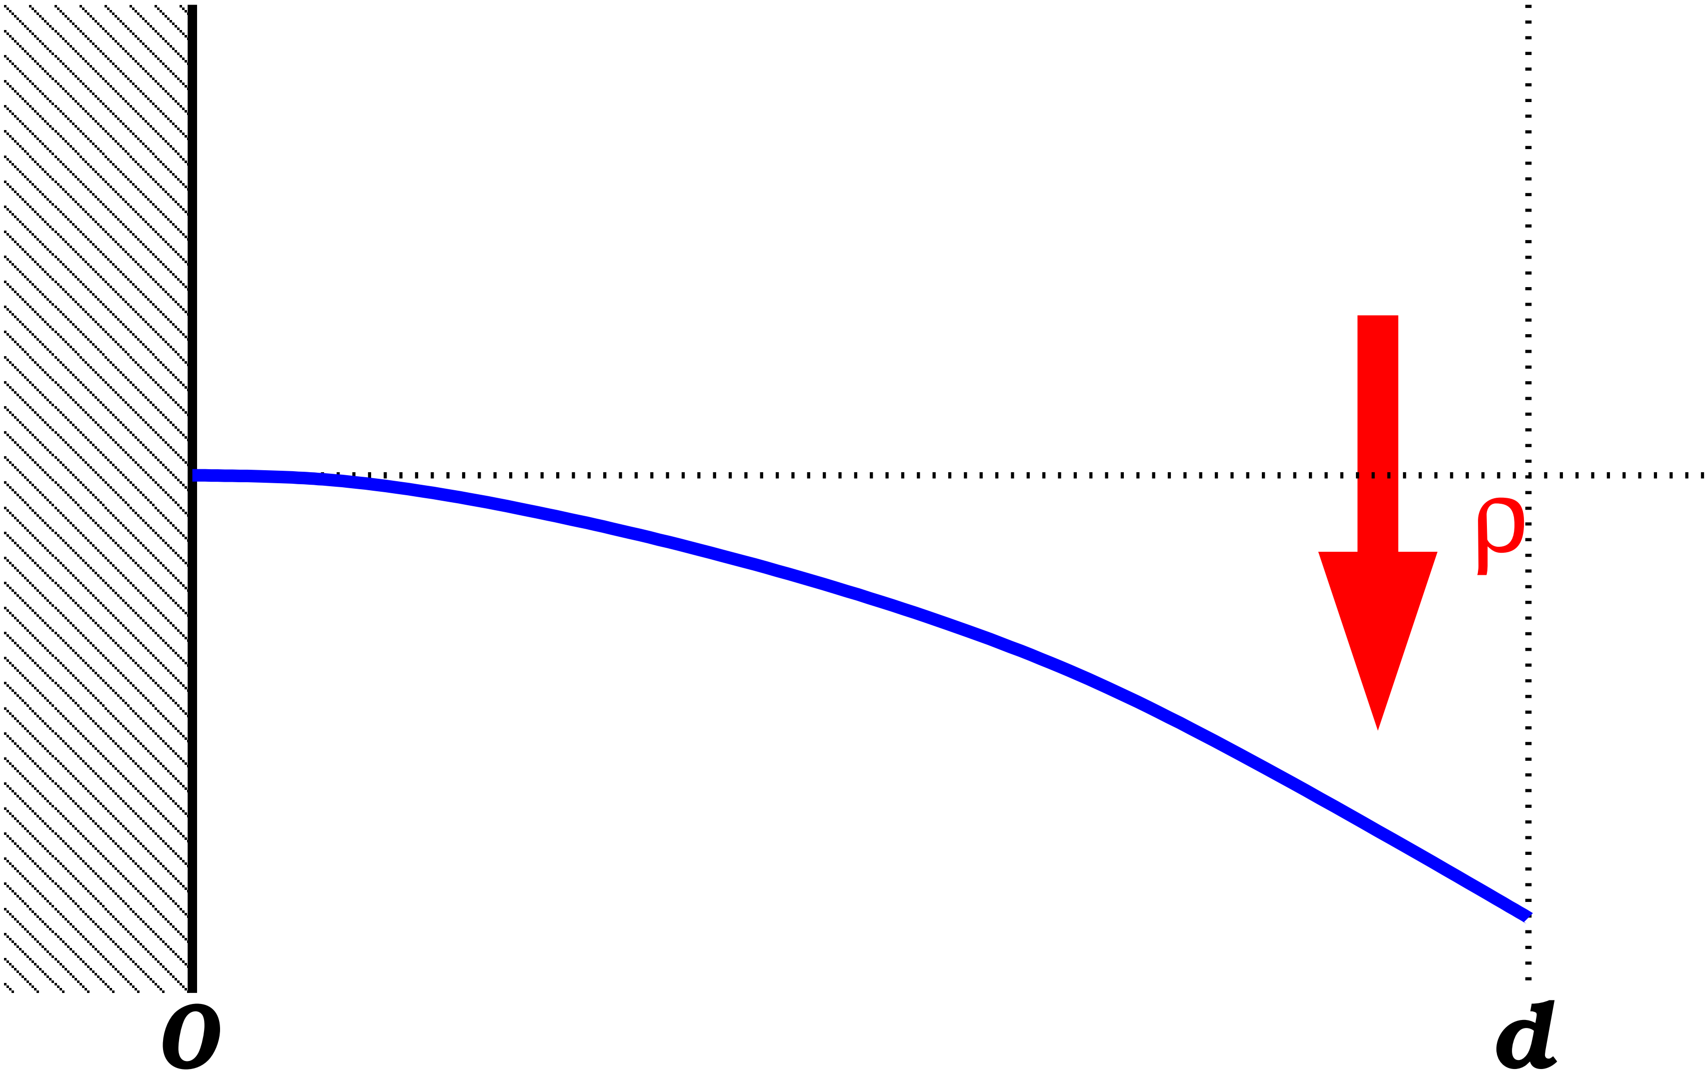

Perturbation again

We approach this the same way we did with all other variational problems, we perturb the curve and examine the First Variation, but this time, we allow $y(x_0)$ and $y(x_1)$ to vary as well.

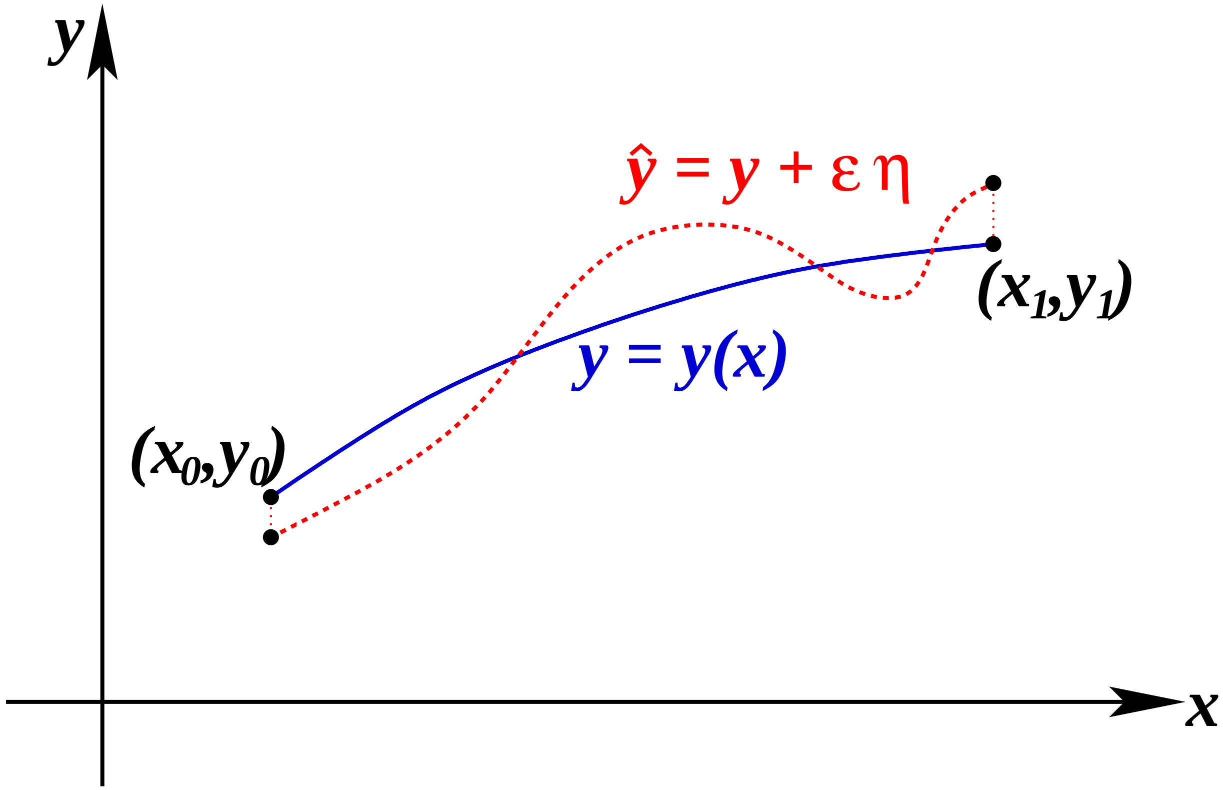

Space of Perturbations

Now the space $\mathcal{H}$ of perturbations η contains functions whose value at x_0 and x_1 is no longer zero.

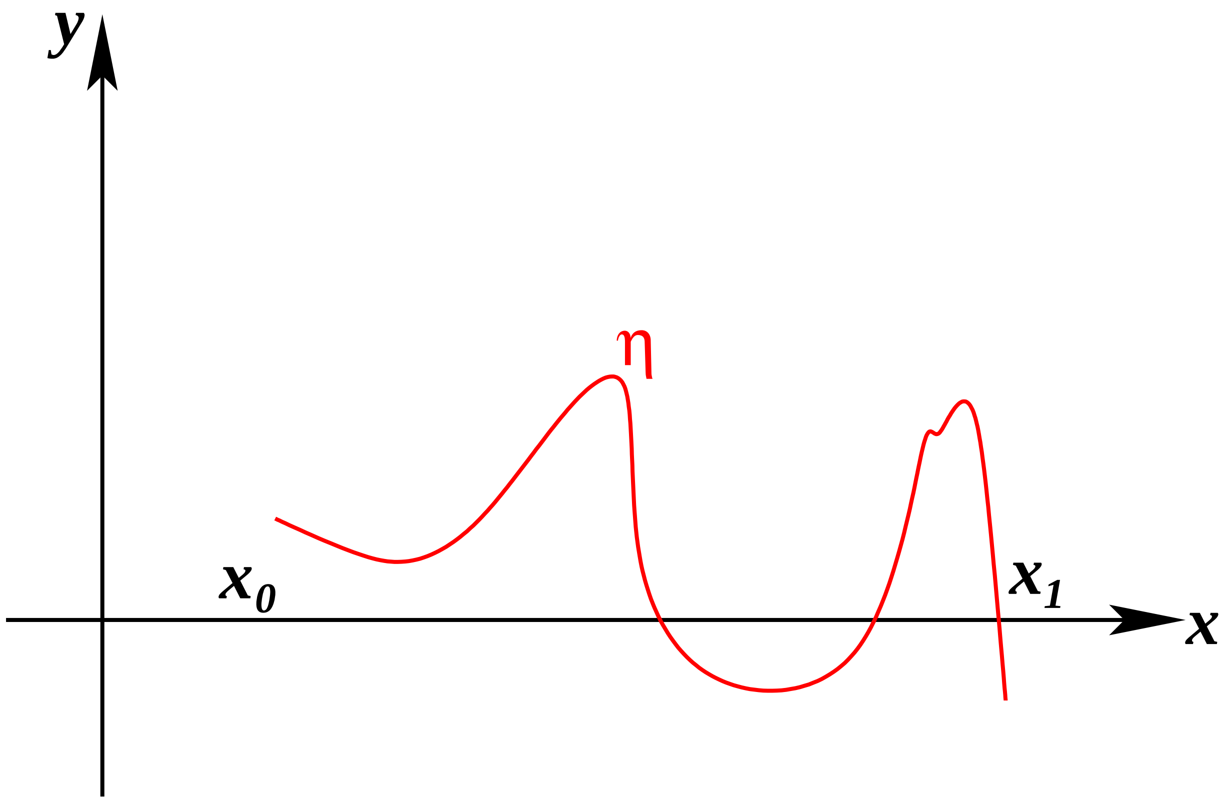

Same derivation of the first variation

Simple case where $F\{y\} = \int_{x_0}^{x_1} f(x,y,y')\,dx$

$$
\begin{aligned}
f(x,\hat{y},\hat{y}') &= f(x,y,y') + \varepsilon\left[\eta\frac{\partial f}{\partial y} + \eta'\frac{\partial f}{\partial y'}\right] + O(\varepsilon^2) \\
F\{\hat{y}\} - F\{y\} &= \int_{x_0}^{x_1} f(x,\hat{y},\hat{y}')dx - \int_{x_0}^{x_1} f(x,y,y')dx \\
&= \varepsilon\int_{x_0}^{x_1}\left[\eta\frac{\partial f}{\partial y} + \eta'\frac{\partial f}{\partial y'}\right]dx + O(\varepsilon^2) \\
\delta F(\eta,y) &= \lim_{\varepsilon\to 0}\frac{F\{y+\varepsilon\eta\} - F\{y\}}{\varepsilon} \\
&= \int_{x_0}^{x_1}\left[\eta\frac{\partial f}{\partial y} + \eta'\frac{\partial f}{\partial y'}\right]dx
\end{aligned}
$$

The first variation

As before, we can vary the sign of ε, so for $F\{y\}$ to be a local minima it must be the case that

$$\delta F(\eta, y) = 0, \qquad \forall \eta \in \mathcal{H}$$

however, now $\mathcal{H}$ allows curves with arbitrary end-points, so that $\eta(x_0) \neq 0$, and $\eta(x_1) \neq 0$ are possible.

Hence when we integrate by parts we get

$$\delta F(\eta, y) \quad = \quad \left[\eta \frac{\partial f}{\partial y'}\right]_{x_0}^{x_1} + \int_{x_0}^{x_1} \eta \left[\frac{\partial f}{\partial y} - \frac{d}{dx}\left(\frac{\partial f}{\partial y'}\right)\right] dx$$

But now the first term $\left[\eta \frac{\partial f}{\partial y'}\right]_{x_0}^{x_1}$ is not necessarily zero.

The first variation

However, $\delta F(\eta, y) = 0$ for all η, which includes cases where $\eta(x_0) = \eta(x_1) = 0$, and so the Euler-Lagrange equation must still be satisfied for such and extremal.

Given the E-L equation is satisfied by an extremal, the condition $\delta F(\eta, y) = 0$ next implies that

$$\left[\eta \frac{\partial f}{\partial y'} \right]_{x_0}^{x_1} = 0$$

and we can likewise choose curves η such that $\eta(x_0) \neq 0$ and $\eta(x_1) = 0$, or visa versa, so that we must have

$$\left. \frac{\partial f}{\partial y'} \right|_{x_0} = 0 \qquad \text{and} \qquad \left. \frac{\partial f}{\partial y'} \right|_{x_1} = 0$$

Euler-Lagrange again

Hence, as before, the extremal must satisfy the E-L equations

$$\frac{d}{dx}\left(\frac{\partial f}{\partial y'}\right) - \frac{\partial f}{\partial y} = 0$$

but now that the boundary conditions were not specified as part of the problem, we get natural boundary conditions

$$\left.\frac{\partial f}{\partial y'}\right|_{x_0} = 0 \quad \text{and} \quad \left.\frac{\partial f}{\partial y'}\right|_{x_1} = 0$$

which specify that the derivative at the end-points will be zero.

Extensions (i)

What happens if we fix one end point, e.g. $y(x_0) = y_0$.

The result is we cannot vary this end-point when perturbing, so $\eta(x_0) = 0$, and therefore the condition

$$\left[\eta \frac{\partial f}{\partial y'}\right]_{x_0}^{x_1} = 0$$

collapses to give just one extra condition

$$\left.\frac{\partial f}{\partial y'}\right|_{x_1} = 0$$

Hence the boundary conditions are **modular** in the sense that when we remove one, we replace it automatically with the natural boundary condition.

Extensions (ii)

The above results can be extended as before, in particular, consider a functional containing higher order derivatives:

$$F\{y\} = \int_{x_0}^{x_1} f(x, y, y', y'')\, dx,$$

$$\begin{aligned}
\delta F(\eta, y) \quad = \quad & \left[\eta\left(\frac{\partial f}{\partial y'} - \frac{d}{dx}\frac{\partial f}{\partial y''}\right)\right]_{x_0}^{x_1} + \left[\eta' \frac{\partial f}{\partial y''}\right]_{x_0}^{x_1} \\
& + \int_{x_0}^{x_1} \left[\eta \frac{\partial f}{\partial y} - \eta \frac{d}{dx}\frac{\partial f}{\partial y'} + \eta \frac{d^2}{dx^2}\frac{\partial f}{\partial y''}\right] dx
\end{aligned}$$

where we see integration by parts introduces terms including η and η'.

Extensions (ii)

The Euler-Lagrange equations are

$$\frac{\partial f}{\partial y} - \frac{d}{dx}\frac{\partial f}{\partial y'} + \frac{d^2}{dx^2}\frac{\partial f}{\partial y''} = 0$$

where the natural boundary conditions are

$$\left.\frac{\partial f}{\partial y'} - \frac{d}{dx}\frac{\partial f}{\partial y''}\right|_{x_0} = 0 \qquad \text{and} \qquad \left.\frac{\partial f}{\partial y'} - \frac{d}{dx}\frac{\partial f}{\partial y''}\right|_{x_1} = 0$$

$$\left.\frac{\partial f}{\partial y''}\right|_{x_0} = 0 \qquad \text{and} \qquad \left.\frac{\partial f}{\partial y''}\right|_{x_1} = 0$$

where the first two replace absent conditions on the value of y at the end-points, and the second two replace absent conditions on y' at the end points.

Bent beam

Let $y : [0, d] \to I\!R$ describe the shape of the beam, and $\rho : [0, d] \to I\!R$ be the load per unit length on the beam.
For a bent elastic beam the potential energy from elastic forces is

$$V_1 = \frac{\kappa}{2} \int_0^d y''^2 \, dx, \qquad \kappa = \text{flexural rigidity}$$

The potential energy is

$$V_2 = - \int_0^d \rho(x) y(x) \, dx$$

Thus the total potential energy is

$$V = \int_0^d \frac{\kappa y''^2}{2} - \rho(x) y(x) \, dx$$

Bent Beam: see earlier

The Euler-Lagrange equation is

$$\begin{aligned}
\frac{\partial f}{\partial y} - \frac{d}{dx}\frac{\partial f}{\partial y'} + \frac{d^2}{dx^2}\frac{\partial f}{\partial y''} &= 0 \\
y^{(4)} &= \frac{\rho(x)}{\kappa}
\end{aligned}$$

This DE has solution

$$y(x) = P(x) + c_3 x^3 + c_2 x^2 + c_1 x + c_0$$

where the c_k's are the constants of integration, and $P(x)$ is a particular solution to $P^{(4)}(x) = \rho(x)/\kappa$.

Bent Beam: Example 1

Doubly clamped: see earlier lectures.

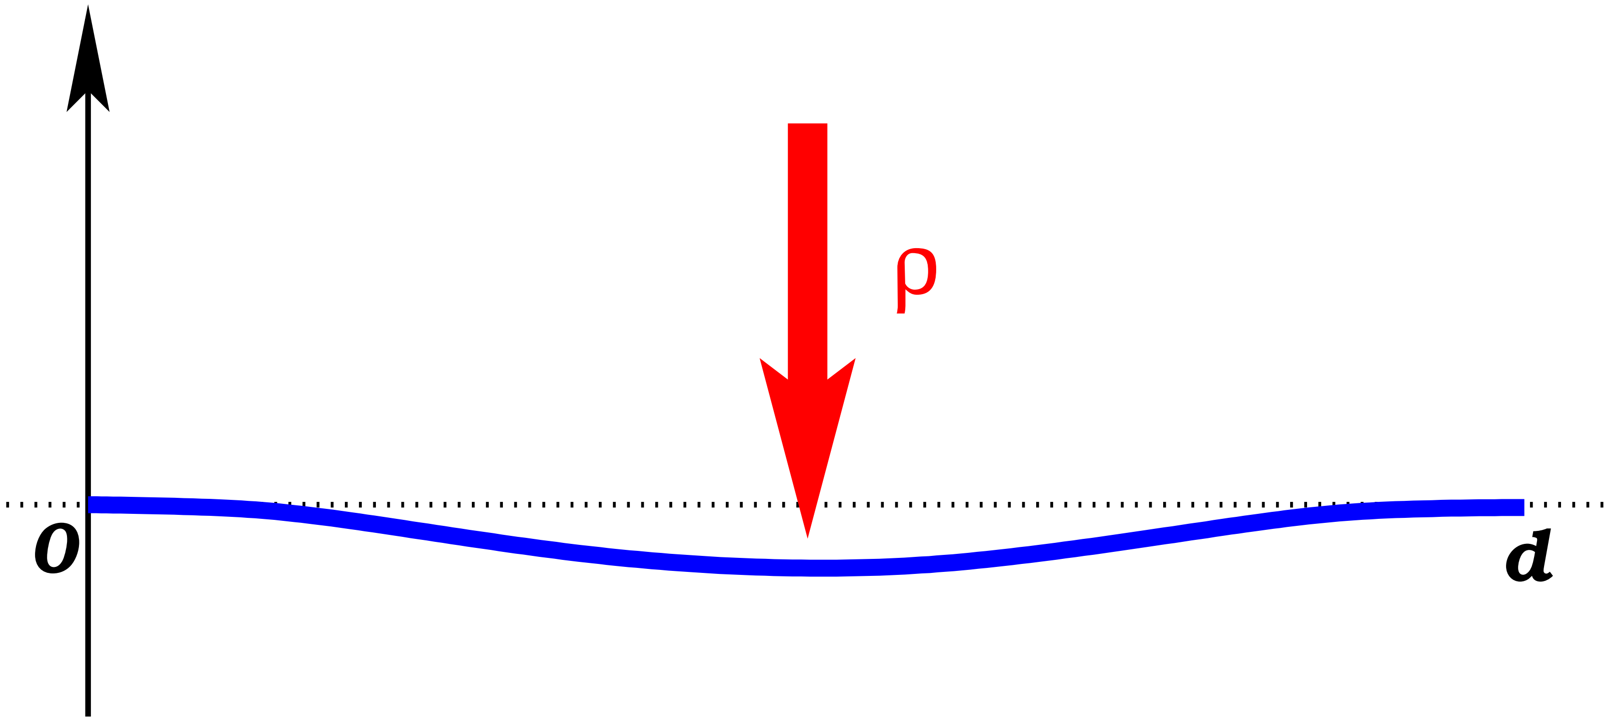

Two end-points are fixed, and clamped so that they are level, e.g. $y(0) = 0$, $y'(0) = 0$, and $y(d) = 0$ and $y'(d) = 0$.

Bent Beam: Example 1

Doubly clamped: see earlier lectures.
Choose a solution of the form

$$y(x) = \frac{\rho(d-x)^2 x^2}{24\kappa}$$

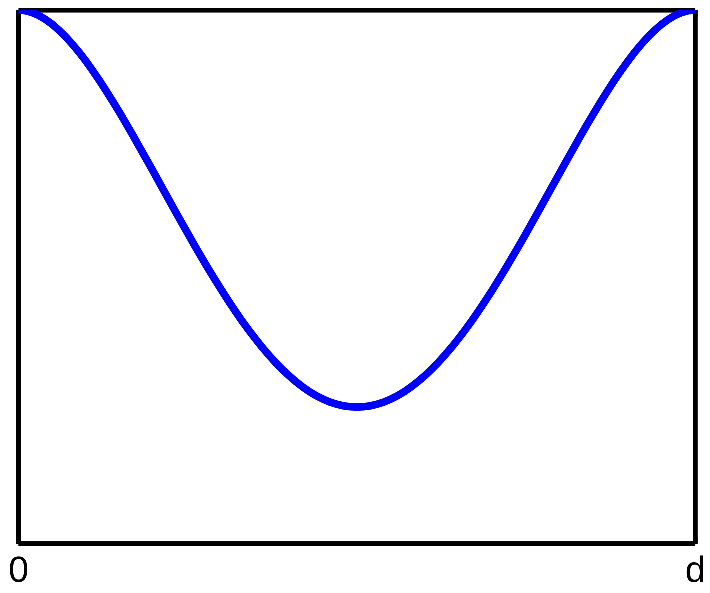

Then the derivative

$$y'(x) = \frac{2\rho(d-x)x^2}{12\kappa} + \frac{\rho(d-x)^2 x}{12\kappa}$$

We can see that the constraints are satisfied

$$y(0) = 0 \qquad \text{and} \qquad y(d) = 0$$
$$y'(0) = 0 \qquad \text{and} \qquad y'(d) = 0$$

Bent Beam Example 2

Freely supported, uniform load
The natural constraints are

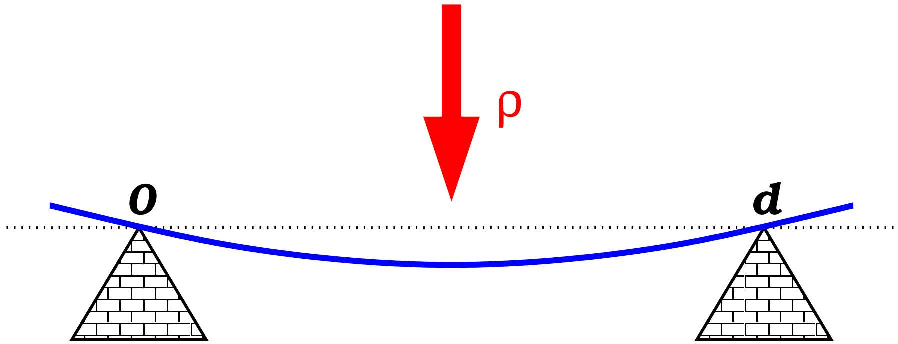

$$
\begin{aligned}
\left.\frac{\partial f}{\partial y''}\right|_{x_0} &= \kappa y''(x_0) = 0 \\
\left.\frac{\partial f}{\partial y''}\right|_{x_1} &= \kappa y''(x_1) = 0
\end{aligned}
$$

The fixed end-points are $y(0) = y(d) = 0$, so uniform load solution looks like

$$
y(x) = \frac{\rho x \left(d^3 - 2dx^2 + x^3\right)}{24\kappa}
$$

Bent Beam Example 3

One end-point fixed, and clamped.
Called a **Cantilever**
The natural constraints are

$$
\begin{aligned}
\left.\frac{\partial f}{\partial y''}\right|_{x_1} &= \kappa y''(x_1) = 0 \\
\left.\frac{\partial f}{\partial y'} - \frac{d}{dx}\frac{\partial f}{\partial y''}\right|_{x_1} &= -\frac{d}{dx}\kappa y''\Big|_{x_1} \\
&= \kappa y'''(x_1) = 0
\end{aligned}
$$

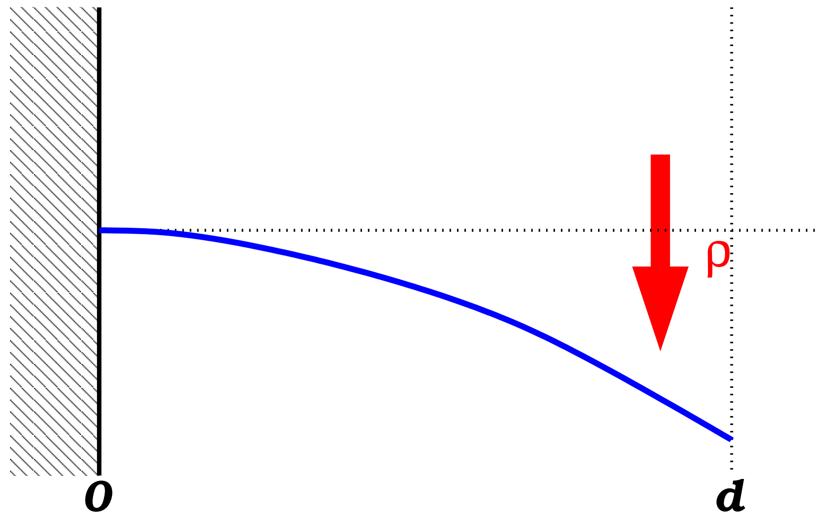

The clamped end-point introduces constraints $y(0) = 0$ and $y'(0) = 0$ so the solution for uniform load is

$$
y(x) = \frac{\rho x^2(6d^2 - 4dx + x^2)}{24\kappa} \quad \text{and} \quad y(d) = \frac{\rho d^4}{8\kappa}
$$

Bent beam, end-points conditions

End-point conditions are modular: i.e., we can use different end-point conditions at each end of the beam.

- **clamped:** specifies the position, and the derivative.
- **freely supported:** specifies the position. Natural boundary condition is that the second derivative is zero at the end point.
- **no condition:** neither position, nor end-point are specified, so the natural boundary conditions fix the second and third derivatives at the end point to be zero.

Bent Beam Example 4

One end-point fixed, but not clamped.

In this case the beam just collapses, and lies vertical.

The approach doesn't work, but this is a failure of the **model**, not the **method**.

In this case, the cantilever approximation (that x_1 is fixed) no longer works, and we need to consider a more general model that allows x_1 to vary as well.

Intro to Optimal Control (part II)

Often in optimal control problems we may specify the initial state, but not the final state. However, there may be a cost associated with the final state, and we include this in the functional to be minimized (or maximized). We call this a terminal cost.

Optimal control with terminal costs

In an optimal control problem we again have a non-holonomic constraint

$$\dot{\mathbf{x}}(t) = \mathbf{g}(t, \mathbf{x}, \mathbf{u})$$

given initial state $x(t_0)$, but now the final state will be free and we wish to minimize a functional

$$F\{\mathbf{u}\} = \phi(t_1, \mathbf{x}(t_1)) + \int_{t_0}^{t_1} f(t, \mathbf{x}, \mathbf{u})\, dt$$

the term $\phi(t_1, \mathbf{x}(t_1))$ is called the **terminal cost**.

Terminal costs reformulation

Note that

$$\phi(t_1, \mathbf{x}(t_1)) = \phi(t_0, \mathbf{x}(t_0)) + \int_{t_0}^{t_1} \frac{d}{dt}\phi(t, \mathbf{x})\, dt$$

so we can rewrite

$$\begin{aligned} F\{\mathbf{u}\} &= \phi(t_1, \mathbf{x}(t_1)) + \int_{t_0}^{t_1} f(t, \mathbf{x}, \mathbf{u})\, dt \\ &= \phi(t_0, \mathbf{x}(t_0)) + \int_{t_0}^{t_1} \left[f(t, \mathbf{x}, \mathbf{u}) + \frac{d}{dt}\phi(t, \mathbf{x}) \right] dt \end{aligned}$$

where note that the first term is fixed by the starting point, and so we can drop it from the problem.

Terminal costs: example

Imagine the problem we wish to solve is to minimize the time, i.e. t_1. ,We could write this as a terminal cost problem, e.g. minimize

$$F\{\mathbf{u}\} = t_1$$

So $\phi(t) = t$, and $\frac{d}{dt}\phi = 1$ and therefore, we can write the minimum time problem in the form

$$F\{\mathbf{u}\} = \int_{t_0}^{t_1} 1\, dt$$

Terminal costs and E-L equations

Given a problem like

$$F\{\mathbf{u}\} = \int_{t_0}^{t_1} \left[f(t,\mathbf{x},\mathbf{u}) + \frac{d}{dt}\phi(t,\mathbf{x}) \right] dt$$

Note that

$$\frac{d}{dt}\phi(t,\mathbf{x}) = \frac{\partial \phi}{\partial t} + \sum_{i=1}^{n} \frac{\partial \phi}{\partial x_i}\dot{x}_i$$

E-L equations:

$$\begin{aligned} \frac{d}{dt}\frac{\partial f}{\partial \dot{x}_k} - \frac{\partial f}{\partial x_k} + \frac{d}{dt}\frac{\partial \phi}{\partial x_k} - \frac{\partial^2 \phi}{\partial x_k \partial t} - \sum_{i=1}^{n} \frac{\partial^2 \phi}{\partial x_k \partial x_i}\dot{x}_i &= 0 \\ \frac{d}{dt}\frac{\partial f}{\partial \dot{x}_k} - \frac{\partial f}{\partial x_k} &= 0 \end{aligned}$$

Terminal costs and E-L equations

Hence terminal costs play no part in the Euler-Lagrange equations, which makes sense

- fixed end-point problem
  - terminal cost is fixed (by the end-point)
  - so Euler-Lagrange equations unchanged
- free end-point problem
  - Euler-Lagrange equations aren't effected by freeing up the end-points

Terminal costs and boundary conditions

Terminal costs play no part in the Euler-Lagrange equations, but for free-end points we get a new natural boundary condition:

- Take a functional written in the form:

$$F\{\mathbf{x}\} = \int_{t_0}^{t_1} \left[f(t,\mathbf{x},\dot{\mathbf{x}}) + \frac{d}{dt}\phi(t,\mathbf{x}) \right] dt = \int_{t_0}^{t_1} h(t,\mathbf{x},\dot{\mathbf{x}})\, dt$$

- Natural boundary condition

$$\left.\frac{\partial h}{\partial \dot{x}_k}\right|_{t=t_1} = \left.\frac{\partial f}{\partial \dot{x}_k} + \frac{\partial}{\partial \dot{x}_k}\frac{d\phi}{dt}\right|_{t=t_1} = \left.\frac{\partial f}{\partial \dot{x}_k} + \frac{\partial \phi}{\partial x_k}\right|_{t=t_1} = 0$$

where we use
$$\frac{d}{dt}\phi(t,\mathbf{x}) = \frac{\partial \phi}{\partial t} + \sum_{i=1}^{n} \frac{\partial \phi}{\partial x_i}\dot{x}_i$$

Example: stimulated plant growth

Plant growth problem:

- market gardener wants to plants to grow as much as possible within a fixed window of time $[t_0, t_1] = [0, 1]$
- supplement natural growth with lights as before
- growth rate dictates $\dot{x} = 1 + u$
- cost of lights

$$F\{u\} = \int_0^1 \frac{1}{2} u(t)^2 \, dt$$

- value of crop is proportional to the height at $t_1 = 1$

$$\phi(t_1, x(t_1)) = kx(1)$$

Plant growth problem statement

Minimize (equivalent to maximizing the profit)

$$F\{u,x\} = -kx(1) + \int_0^1 \frac{1}{2}u^2\,dt = \int_0^1 \frac{1}{2}u^2 - k\,dt$$

Subject to $x(0) = 0$,

$$\dot{x} = 1 + u$$

- note that the extra constant in F will not effect the E-L equations, so the solution must still have the same form, i.e., $u = const$
- but the end conditions have changed

Plant growth

Including the Lagrange multiplier $\lambda(t)\left[\dot{x}-1-u\right]$

$$H\{u,x\} = \int_0^1 h(t,u,\dot{x}) + \frac{d}{dt}\phi(x)\,dt$$

where

$$\begin{aligned} h(t,u,\dot{x}) &= \frac{1}{2}u^2 + \lambda(t)\left[\dot{x}-1-u\right] \\ \phi(x) &= -kx \end{aligned}$$

Now the independent variable is t, and there are three dependent variables x, u, λ.

Plant growth: E-L equations

Three dependent variables, so three E-L equations

$$
\begin{aligned}
\frac{d}{dt}\frac{\partial h}{\partial \dot{x}} + \frac{\partial h}{\partial x} &= 0 && (1)\\
\frac{d}{dt}\frac{\partial h}{\partial \dot{u}} + \frac{\partial h}{\partial u} &= 0 && (2)\\
\frac{d}{dt}\frac{\partial h}{\partial \dot{\lambda}} + \frac{\partial h}{\partial \lambda} &= 0 && (3)
\end{aligned}
$$

Notice that $d\phi/dt$ is a constant, so it plays no part.

- h is linear in $\dot{x}$ so equation (1) is degenerate
- equation (2) gives us the E-L equation we had before
- equation (3) just gives us back the constraint

Plant growth: natural boundary cond.

Natural boundary conditions at $t_1 = 1$.

$$
\begin{aligned}
\left.\frac{\partial h}{\partial \dot{x}} + \frac{\partial \phi}{\partial x}\right|_{t_1} &= 0 \\
\left.\frac{\partial h}{\partial \dot{u}} + \frac{\partial \phi}{\partial u}\right|_{t_1} &= 0
\end{aligned}
$$

The second is trivial, i.e., $0 = 0$, so consider the first:

$$
\frac{\partial h}{\partial \dot{x}} + \frac{\partial \phi}{\partial x} = \lambda - k = 0
$$

We already know from the E-L equations that $\lambda = u$, and $u = const$, so the end result is that $u = k$.

Plant growth solution

The solution is $u = k$, and so

$$x(1) = 1 + k$$

When $k = 1$ we get the same solution we got before, but that isn't a general rule.

Also the optimization objective will be

$$F\{u, x\} = -1 - k + 1.5k^2$$

written in terms of profit we get

$$\text{profit} = 1 + k - 1.5k^2$$

Plant growth solution

Another way to see how the end-point conditions work

- The E-L equations still apply
- So u is still a constant
- $x(1) = 1 + u$ is the solution to the system DE

The height at $t_1 = 1$ would be $1 + u$ and so the profit would be

$$F\{u,x\} = 1 + ku - \int_0^1 \frac{1}{2}u^2\,dt = 1 + ku - \frac{1}{2}u^2$$

Clearly, the maximum here occurs for $u = k$.

Optimal Control

We will continue with optimal control later in the course when we have considered a bit more theory, but consider the following problem:

> Replace the previous plant growth problem by a similar problem, but instead of a terminal cost (related to value of plant), we aim to get the plants to height 2 in time that minimizes the cost.

Now t_1 is also a free variable – how can we deal with this?

Freeing up the independent variable

We can deal with both the optimal control problem and the collapsing beam by freeing up the value of the dependent variable.