Summary of Concepts in Variational Principles

Richard Chapling v1 14 October 2018

Part I. Finite-Dimensional Optimisation and Convexity

1. Stationary Points in $\mathbb{R}^n$

A function $f : \mathbb{R}^n \to \mathbb{R}$ is *differentiable* if there is a linear map L so that

$$f(\mathbf{x}+\mathbf{h}) = f(\mathbf{x}) + L(\mathbf{h}) + o(\|\mathbf{h}\|)$$

for $\|\mathbf{h}\| \to 0$. This linear map can be written as an inner product with a vector $\nabla f(\mathbf{x})$, the *gradient* of f.

A differentiable function $f : \mathbb{R}^n \to \mathbb{R}$ has a *stationary point* at $\mathbf{x}$ if $\nabla f(\mathbf{x}) = 0$. $\mathbf{x}$ is a *local minimum* if f is minimised by $\mathbf{x}$ in a neighbourhood, and a *global minimum* if f has its smallest value in its whole domain at $\mathbf{x}$. Exactly the same is said of maxima, where the function is locally or globally largest.

First order necessary condition for a local extremum If f is differentiable, we must have $\nabla f(\mathbf{x}) = 0$ for $\mathbf{x}$ to be an extremum (or we could find a vector $\mathbf{h}$ that made $\mathbf{h}\cdot\nabla f(\mathbf{x})$ have the wrong sign).

Second-order sufficient condition for a minimum Suppose that the Hessian $\nabla\nabla f(\mathbf{x})$ is nondegenerate. Then a sufficient condition for a minimum is

$$\mathbf{v}\cdot(\nabla\nabla f(\mathbf{x}))\cdot\mathbf{v} > 0 \tag{1}$$

for all $\mathbf{v} \neq 0$ ($\nabla\nabla f(\mathbf{x})$ is positive-definite).

If f has a maximum, $-f$ has a minimum, so the Hessian has to be negative-definite for a maximum.

One can prove this by using Taylor's theorem: for small $\|\mathbf{h}\| > 0$, we have

$$f(\mathbf{x}+\mathbf{h}) - f(\mathbf{x}) = \mathbf{h}\cdot\nabla f(\mathbf{x}) + \frac{1}{2}\mathbf{h}\cdot(\nabla\nabla f(\mathbf{x}))\cdot\mathbf{h} + o(\|\mathbf{h}\|^2),$$

and the first term has to be zero to avoid being negative. Then the right-hand side is positive for small enough $\|\mathbf{h}\| > 0$ provided that the Hessian is positive-definite.

1.1. Convexity

A set $A \subseteq \mathbb{R}^n$ is called *convex* if $(1-t)\mathbf{x} + t\mathbf{y} \in A$ for each $\mathbf{x}, \mathbf{y} \in A$, $t \in (0,1)$, i.e. the line segment joining $\mathbf{x}$ and $\mathbf{y}$ lies entirely in A.

$f : A \to \mathbb{R}$ is called *convex* if

$$f((1-t)\mathbf{x} + t\mathbf{y}) \geqslant (1-t)f(\mathbf{x}) + tf(\mathbf{y}). \tag{2}$$

for each $\mathbf{x}, \mathbf{y} \in A$, $t \in (0,1)$. If we can replace $\geqslant$ by $>$, f is called *strictly convex*. If $-f$ is (strictly) convex, f is (strictly) concave.

First-order convexity conditions Suppose f is differentiable. Then the following are equivalent:

1. $f(\mathbf{x})$ is convex,
2. $f(\mathbf{y}) \geqslant f(\mathbf{x}) + (\mathbf{y}-\mathbf{x})\cdot\nabla f(\mathbf{x})$,
3. $(\mathbf{y}-\mathbf{x})\cdot(\nabla f(\mathbf{y}) - \nabla f(\mathbf{x})) \geqslant 0$.

Proof. **1 $\Longrightarrow$ 2** Rearranging the definition of convexity, we have

$$f(\mathbf{y}) \geqslant f(\mathbf{x}) + \frac{f(\mathbf{x}+t(\mathbf{y}-\mathbf{x})) - f(\mathbf{x})}{t}$$

Taking $t \downarrow 0$, the second term on the right-hand side tends to $(\mathbf{y}-\mathbf{x})\cdot\nabla f(\mathbf{x})$, which is 2.

2 $\Longrightarrow$ 1 Let $\mathbf{z} = (1-t)\mathbf{x} + t\mathbf{y}$. Then applying 2 twice, we have

$$f(\mathbf{x}) \geqslant f(\mathbf{z}) + (\mathbf{x}-\mathbf{z})\cdot\nabla f(\mathbf{z})$$
$$f(\mathbf{y}) \geqslant f(\mathbf{z}) + (\mathbf{y}-\mathbf{z})\cdot\nabla f(\mathbf{z}),$$

and adding $(1-t)$ of the first to t of the second gives the convexity condition (the gradient terms cancel).

2 $\Longrightarrow$ 3 We have, by swapping x and y, the inequalities

$$f(\mathbf{y}) - f(\mathbf{x}) \geqslant (\mathbf{y}-\mathbf{x})\cdot\nabla f(\mathbf{x})$$
$$f(\mathbf{x}) - f(\mathbf{y}) \geqslant (\mathbf{x}-\mathbf{y})\cdot\nabla f(\mathbf{y})$$

Adding and rearranging gives 3.

3 $\Longrightarrow$ 2 The Fundamental Theorem of Calculus gives

$$\begin{aligned} f(\mathbf{x}+\mathbf{v}) - f(\mathbf{x}) &= \int_0^1 \frac{d}{dt} f(\mathbf{x}+t\mathbf{v})\,dt \\ &= \int_0^1 \mathbf{v}\cdot\nabla f(\mathbf{x}+t\mathbf{v})\,dt \\ &= \int_0^1 \mathbf{v}\cdot\big(\nabla f(\mathbf{x}+t\mathbf{v}) - \nabla f(\mathbf{x})\big)\,dt + \mathbf{v}\cdot\nabla f(\mathbf{x}) \\ &\geqslant \mathbf{v}\cdot\nabla f(\mathbf{x}), \end{aligned}$$

by applying 3 with $\mathbf{y} = \mathbf{x} + t\mathbf{v}$; this is (2). □

2 implies that f lies above its tangent plane, (given on the right-hand side). This also implies that any local minimum of a convex function is a *global* minimum.

In 1D, 3 states that the derivative is nondecreasing.

Second-order convexity condition If f is twice-differentiable, it is convex $\Longleftrightarrow$ $\nabla\nabla f(\mathbf{x})$ is nonnegative-definite.

Proof. $\Longleftarrow$ Starting from the LHS of 3 of the previous result and applying the FToC again,

$$\begin{aligned} \mathbf{v}\cdot(\nabla f(\mathbf{x}+\mathbf{v}) - \nabla f(\mathbf{x})) &= \int_0^1 \frac{d}{dt}\mathbf{v}\cdot\nabla f(\mathbf{x}+t\mathbf{v})\,dt \\ &= \int_0^1 \mathbf{v}\cdot\nabla\nabla f(\mathbf{x}+t\mathbf{v})\cdot\mathbf{v}\,dt \geqslant 0 \end{aligned}$$

if the condition on $\nabla\nabla f$ holds, and hence we have 3, and f is convex.

$\Longrightarrow$ Taking $t > 0$ and setting $y = \mathbf{x} + t\mathbf{v}$ in 3 again,

$$0 \leqslant \frac{\mathbf{v}\cdot(\nabla f(\mathbf{x}+t\mathbf{v}) - \nabla f(\mathbf{x}))}{t},$$

and then taking the limit as $t \downarrow 0$ gives the result, since the right-hand side tends to $\mathbf{v}\cdot\nabla\nabla f(\mathbf{x})\cdot\mathbf{v}$. □

Strict convexity does not imply that $\nabla\nabla f(\mathbf{x}) > 0$: e.g. x^4.

1.2. Variation with Constraints

Suppose we want to minimise or maximise $f(\mathbf{x})$ subject to some constraints on the values of $\mathbf{x}$, $\mathbf{g}(\mathbf{x}) = 0$: a typical example would be to demand that $\|\mathbf{x}\| = 1$. We start with one condition, $g(\mathbf{x}) = 0$.

If there were no constraints, we would look for $\nabla f(\mathbf{x}) = 0$, so that the function does not change to first order no matter which direction we go. Adding a constraint generally restricts us to working on a hypersurface represented by $g(\mathbf{x}) = 0$. This limits the directions that we can move to those that are perpendicular to the normal of this hypersurface, *viz.* those for which $\mathbf{v} \cdot \nabla g(\mathbf{x}) = 0$. Therefore, while we may have $\nabla f(\mathbf{x}) = 0$, this is no longer the only possibility: also available is $\nabla f(\mathbf{x})$ being parallel to $\nabla g(\mathbf{x})$. Therefore the equations we have to solve are

$$\nabla f(\mathbf{x}) = \lambda \nabla g(\mathbf{x}) \tag{3}$$

$$g(\mathbf{x}) = 0. \tag{4}$$

We can put these together by adding an extra part to the function we are trying to extremise:

$$L(\mathbf{x}, \lambda) = f(\mathbf{x}) - \lambda g(\mathbf{x}), \tag{5}$$

the *Lagrangian*, and now we have an unconstrained minimisation problem with one extra variable. λ is called the *Lagrange multiplier*; whether we use $+\lambda g(\mathbf{x})$ or $-\lambda g(\mathbf{x})$ is unimportant theoretically. Now, differentiating with respect to $\mathbf{x}$ gives the gradient equation, and differentiating wrt λ gives the constraint equation; we then solve the equations to find $\mathbf{x}$, and then determine what sort of stationary point we have. λ appears to become superfluous, but often has a physical interpretation, like tension holding the point to the surface. Also, when the constraints are satisfied, we obviously have $L(\mathbf{x}, \lambda) = f(\mathbf{x})$: the point of the Lagrangian is that it allows us to express the gradient condition easily.

If there are more constraints, we simply augment f to include one Lagrange multiplier for each constraint, turning it into $L(\mathbf{x}, \boldsymbol{\lambda}) = f(\mathbf{x}) - \boldsymbol{\lambda} \cdot \mathbf{g}(\mathbf{x})$: the idea then becomes that $\nabla f(\mathbf{x})$ lives in the subspace spanned by the normals.

2. Legendre Transform

We assume in this section that $f : A \to \mathbb{R}$ is differentiable. The *Legendre transform* allows us to convert a function of x into a function of p, its gradient. It is given by

$$f^*(p) = \sup_{x \in A} (px - f(x)). \tag{6}$$

We write A^* for the set of p for which this is finite. It is easy to see that A^* is convex, and on A^*, f^* is convex, since

$$\begin{aligned} f^*(tp + (1-t)q) &= \sup_{x \in A} \big((tp + (1-t)q)x - f(x)\big) \\ &\leqslant t \sup_{x \in A} (px - f(x)) + (1-t) \sup_{x \in A} (qx - f(x)) \\ &= t f^*(p) + (1-t) f^*(q). \end{aligned}$$

(Notice also that the inequality also implies that the set where f^* is finite is convex, since it shows that $f^*(tp + (1-t)q) < \infty$ whenever the endpoints are.)

If f is differentiable, we actually find f^* by differentiating,

$$\frac{d}{dx}(px - f(x)) = p - f'(x).$$

For this to be a maximum, we must have $p = f'(x)$. This relationship is invertible if and only if f' is increasing, i.e. f is convex, and we then find

$$f^*(p) = p(f')^{-1}(p) - f\left((f')^{-1}(p)\right).$$

For example, if $f(x) = ax^2/2$, we have $p - ax = 0$ for the minimum, so the closest point has $x = p/a$, and substituting this in gives $f^*(p) = \frac{1}{2a}p^2$.

Geometric interpretation Draw the graph of f and a line of gradient p that touches the graph of f:

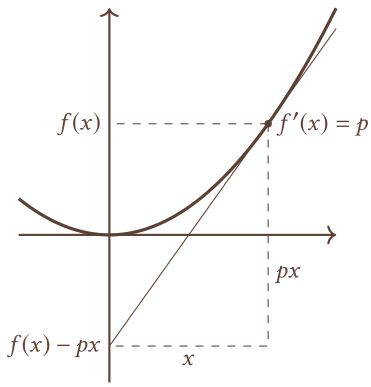

So $-f^*(p)$ is the smallest y-intercept of a line that intersects the graph of f and has gradient p. (Or equally, the largest the y-intercept can be so that the graph of f has no part below the line.)[1]

Inversion It is easy to show that $f^{**}(x) \leqslant f(x)$, but this can be strengthened to:

Theorem. *$f^{**}(x) = f(x)$ if and only if f is convex.*

Proof. Since f^{**} is convex, it is clear that if f itself is not convex, it can't be equal to f^{**}. It remains to check the equality when f is convex.

We have

$$f^{**}(x) = \sup_{p \in A^*} (px - f^*(p)) = \sup_{p \in A^*, \alpha \leqslant -f^*(p)} (px + \alpha)$$

The latter is a supremum over the linear functions, and we saw above that $-f^*(p)$ is the largest value of the y-intercept for which the graph of f has no part below the line. So $f^{**}(x)$ is the largest value assumed by a linear function of this kind at x. If f is convex, this is $f(x)$, because we can just choose the tangent.[2] □

In fact, f^{**} is the *convex hull* of f: this is exactly the supremum of all the linear functions that lie below the graph of f; it is also the largest convex function that is pointwise smaller than f.

Young's inequality By the definition of the Legendre transform, it is clear that for any x and p,

$$f(x) + f^*(p) \geqslant p \cdot x. \tag{7}$$

As a special case, we have

$$\frac{a}{2}x^2 + \frac{1}{2a}p^2 \geqslant p \cdot x, \tag{8}$$

which is sometimes called the *Peter–Paul inequality.*

[1] This interpretation can be adapted to non-differentiable f: we consider the set lines of gradient p, and find the lowest one that intersects f (or equivalently, the highest one with no part of f's graph below it). The y-intercept of this line is the Legendre transform of f.

[2] Were f not differentiable, we instead need to take a *support line* at $(x, f(x))$, which is exactly a line so that the graph of f lies wholly to one side of it, apart from the point $(x, f(x))$, which lies on it. We don't do enough convex analysis in this course to show that these always exist for convex functions, though, hence the restriction noted at the beginning of the section.

Thermodynamics Thermodynamics is frequently concerned with finding properties of a system based on varying some quantities while keeping others constant. What makes it exceptional is that most variables depend on others in complicated ways that make it difficult to talk about things consistently. The mathematical way around this is to define different *thermodynamic potentials*, with different explicit dependence on variables: these are chosen based on the type of equilibrium the system relaxes into when we hold such things constant.

- We start with the *internal energy* $U(S, V)$ of the system, which is given as a function of the entropy S (a measure of the disorder of the system), and the volume V. We define the *temperature* and *pressure* using this function, as

$$T = \left(\frac{\partial U}{\partial S}\right)_V, \quad p = -\left(\frac{\partial U}{\partial V}\right)_S; \tag{9}$$

  these are essentially equivalent to the differential definition, that expresses the change in internal energy as a change in heat and the amount of mechanical work the system does,

$$dU = T\,dS - p\,dV, \tag{10}$$

  (the *First Law of Thermodynamics*) and in both cases the minus sign is due to the energy decreasing if the system pushes its volume bigger.
- The entropy is generally rather hard to measure, and many processes we care about operate at a constant temperature, so we would like to change to a function that depends on temperature instead of entropy. (9) tells us what to do: the differential relation implies that we can do a Legendre transform, to define the *Helmholtz free energy*

$$F(T, V) = \inf_S (U - TS). \tag{11}$$

  The differential relation for this potential is

$$dF = dU - S\,dT - T\,dS = -S\,dT - p\,dV.$$

- F is normally the most useful thermodynamic potential, but we can equally obtain functions of other combinations by Legendre transformations: for processes at constant pressure, the *enthalpy*,

$$H(S, p) = \inf_V (U + pV)$$

  and at constant pressure and temperature, the *Gibbs free energy*

$$G(T, p) = \inf_{S,V} (U - TS + pV).$$

  These satisfy the differential relations

$$dH = T\,dS + V\,dp,$$
$$dG = -S\,dT + V\,dp.$$

Higher dimensions It is straightforward to extend the definition of the Legendre transform to functions from a set of vectors $A \in \mathbb{R}^n$:

$$f^*(\mathbf{p}) = \sup_{\mathbf{x}\in A}(\mathbf{p}\cdot\mathbf{x} - f(\mathbf{x})), \tag{12}$$

and the rest of the theory follows in much the same way.

For example, if $f(\mathbf{x}) = a\|\mathbf{x}\|^2/2$, we take the gradient to find that

$$\mathbf{p} - a\mathbf{x} = 0,$$

for the minimum, so the closest point has $\mathbf{x} = \mathbf{p}/a$, and substituting this in gives $f^*(\mathbf{p}) = \frac{1}{2a}\|\mathbf{p}\|^2$.

Part II. The Calculus of Variations

3. Functionals

A *functional* is a map from a space of functions (e.g. continuous functions, or smooth functions) to $\mathbb{R}$. We normally write a functional as $F[y]$, with the argument in square brackets. In this course, functionals are generally given by integrals of the form

$$F[y] = \int_a^b f(x, y(x), y'(x), \dots)\,dx. \tag{13}$$

The *variation*, or *first variation*, of a functional is given by

$$DF[y](\phi) = \lim_{t\to 0}\frac{F[y + t\phi] - F[y]}{t} = \frac{d}{dt}F[y + t\phi]\Big|_{t=0} \tag{14}$$

Conceptually, this is a directional derivative, in the direction of ϕ; it is also known as the *Gâteaux derivative*. It is a linear map on the space of ϕs. If we can separate it into the form $\int_a^b \phi g(x, y, y', \dots)\,dx$, which is like an inner product with ϕ, then we call g the *functional derivative* of F,

$$\frac{\delta F}{\delta y} = g, \tag{15}$$

or conversely, $DF[y](\phi) = \int_a^b \frac{\delta F}{\delta y}(x)\phi(x)\,dx$.

As with the finite-dimensional case, F is said to have a *stationary point* at y if $DF[y](\phi) = 0$ for all appropriate ϕ. The study of the stationary points of functionals is called the *Calculus of Variations*.

3.1. Euler–Lagrange Equations

Lemma (Fundamental Lemma of the Calculus of Variations). *Suppose that g is continuous, and*

$$\int_a^b \phi(x)g(x)\,dx = 0$$

for every smooth ϕ. Then $g(x) \equiv 0$.

Sketch of proof. Suppose g is not identically zero. Then we can find an interval where g has one sign, and construct a ϕ that is zero outside this interval and has one sign inside, and it follows that the integral cannot be zero. □

We compute a sufficient condition for the integral (13) to have a stationary point. Suppose that the integrand is $f(x, y, y')$, and we seek functions with known values at the endpoints. Thus we should only consider ϕ which are zero at the endpoints. We have

$$\begin{aligned}\frac{F[y + t\phi] - F[y]}{t} &= \frac{1}{t}\int_a^b [f(x, y + t\phi, y' + t\phi') - f(x, y, y')]\,dx \\ &= \int_a^b \left(\frac{\partial f}{\partial y}\phi + \frac{\partial f}{\partial y'}\phi' + o(t)\right)dx,\end{aligned}$$

by using Taylor's Theorem. Taking the $t \to 0$ gives $DF[y](\phi)$; we want to express this as a function of ϕ only, so integrating the ϕ' term by parts, we find

$$DF[y](\phi) = \left[\phi\frac{\partial f}{\partial y'}\right]_a^b + \int_a^b \left(\frac{\partial f}{\partial y} - \frac{d}{dx}\frac{\partial f}{\partial y'}\right)\phi\,dx. \tag{16}$$

The first term vanishes because we chose $\phi(a) = \phi(b) = 0$. Hence

$$\frac{\delta F}{\delta y} = \frac{\partial f}{\partial y} - \frac{d}{dx}\frac{\partial f}{\partial y'}. \tag{17}$$

We conclude that:

A necessary and sufficient condition for y to be a stationary point of $F[y] = \int_a^b f(x, y, y')\,dx$ is that

$$\frac{\partial f}{\partial y} - \frac{d}{dx}\frac{\partial f}{\partial y'} = 0 \tag{18}$$

in the interval (a, b).

This type of calculation is half of the work in this course, the other half being to solve the equations.

N.B. We are treating y and y' as independent variables, because we are actually applying Taylor's theorem in t: what we are really doing is differentiating f with respect to whatever happens to be in the nth "slot" in its list of arguments.

Independent of y Suppose that f does not depend on y. Then the Euler–Lagrange equation reduces to

$$\frac{d}{dx}\frac{\partial f}{\partial y'} = 0 \implies \frac{\partial f}{\partial y'} = A, \tag{19}$$

a constant.

Independent of x: Beltrami identity We can use the Euler-Lagrange equation to prove:

$$\begin{aligned}\frac{d}{dx}\left(y'\frac{\partial f}{\partial y'} - f\right) &= y''\frac{\partial f}{\partial y'} + y'\frac{d}{dx}\frac{\partial f}{\partial y'} - \frac{\partial f}{\partial x} - y'\frac{\partial f}{\partial y} - y''\frac{\partial f}{\partial y'} \\ &= -\frac{\partial f}{\partial x} + y'\left(\frac{d}{dx}\frac{\partial f}{\partial y'} - \frac{\partial f}{\partial y}\right) = -\frac{\partial f}{\partial x},\end{aligned}$$

so if f does not depend on x,

$$y'\frac{\partial f}{\partial y'} - f = A, \tag{20}$$

a constant.

Lagrange Multipliers Of course, we can also ask about *constrained* minimisation problems, such as the shape of the *catenary*, where the length of the chain is fixed. If we have an integral condition such as this on the whole curve, of the form $G[y] = \int_a^b g(x, y, y')\,dx = 0$, we can insert a constant Lagrange multiplier, so the total functional becomes

$$F[y] - \lambda G[y] = \int_a^b \big(f(x, y, y') - \lambda g(x, y, y')\big)\,dx. \tag{21}$$

On the other hand, if the condition must be enforced *pointwise*, such as finding geodesics on a sphere, where the point has to be told continually to stay on the sphere, we need to consider a *Lagrange multiplier function*, so the total Lagrangian has the form

$$\int_a^b \big(f(x, y, y') - \lambda(x) g(x, y, y')\big)\,dx. \tag{22}$$

Variable endpoints We may not have both ends of the curve fixed: in this case, we can't just use the Euler–Lagrange equation. We also need a condition on the endpoint so that we can still have a stationary point when we let the variation ϕ not take the value 0 at a and b. This we get out of (16): an endpoint term vanishes if

$$\frac{\partial f}{\partial y'}(a) = 0; \tag{23}$$

effectively a boundary condition to be used when solving the equations.

More functions If we have a vector of functions, $\mathbf{y}(x)$, then the directional derivative involves a vector function, $\boldsymbol{\phi}(x)$. Then the expansion of $f(x, \mathbf{y} + \boldsymbol{\phi}, \mathbf{y}' + \boldsymbol{\phi}')$ is

$$f(x, \mathbf{y}+\boldsymbol{\phi}, \mathbf{y}'+\boldsymbol{\phi}') = f(x, \mathbf{y}, \mathbf{y}') + \sum_i \frac{\partial f}{\partial y_i}\phi_i + \sum_i \frac{\partial f}{\partial y_i'}\phi_i' + o(\|\boldsymbol{\phi}\|, \|\boldsymbol{\phi}'\|), \tag{24}$$

and the usual integration by parts leads to one Euler–Lagrange equation for each function, or if you prefer, one for each independent variation ϕ_i:

$$\frac{\partial f}{\partial y_i} - \frac{d}{dx}\left(\frac{\partial f}{\partial y_i'}\right) = 0. \tag{25}$$

If $\frac{\partial f}{\partial y_i} = 0$, we again find that $\frac{\partial f}{\partial y_i'}$ is constant. There is still one Beltrami identity, arising from the total derivative

$$\frac{d}{dx}\left(\sum_i y_i'\frac{\partial f}{\partial y_i'} - f\right) = \cdots = -\frac{\partial f}{\partial x}. \tag{26}$$

More independent variables Suppose that the function is now a function of a vector $\mathbf{x}$. Then the integrand is a function of y and its various partial derivatives, $f(\mathbf{x}, y, \nabla y)$ or $f(x_1, \dots, x_n, y, \partial_1 y, \dots, \partial_n y)$, where we abbreviate $\frac{\partial}{\partial x_i} = \partial_i$; the functional is something like

$$\int_V f(x_1, \dots, x_n, y, \partial_1 y, \dots, \partial_n y)\,dx_1 \cdots dx_n. \tag{27}$$

Carrying out the usual variation gives a boundary term $\int_{\partial V} \phi \nabla y\, d\mathbf{S}$, and the Euler–Lagrange equation

$$\frac{\partial f}{\partial y} - \sum_i \frac{\partial}{\partial x_i}\frac{\partial f}{\partial(\partial_i y)} = 0; \tag{28}$$

or $\frac{\partial f}{\partial y} - \nabla \cdot \frac{\partial f}{\partial(\nabla y)} = 0$ if you think that's clearer. There is now a Beltrami identity for each independent variable, from the total derivatives

$$\frac{d}{dx_i}\left(y'\frac{\partial f}{\partial y'} - f\right) = \cdots = -\frac{\partial f}{\partial x_i}. \tag{29}$$

More derivatives The essential proceedure for $f(x, y, y', y'', \dots, y^{(n)})$ is the same as (16): integrate by parts enough to move the derivatives off the ϕ. We end up with some endpoint conditions that we won't worry about here, and

$$\sum_{k=0}^{n} (-1)^k \frac{d^k}{dx^k}\frac{\partial f}{\partial y^{(k)}} = 0. \tag{30}$$

While there is a Beltrami identity for this generalisation, it still contains total time derivatives, so is far less useful than the special case where f only depends on y and y'.

3.2. Examples

Shortest distance between two points We all know it's a straight line, but how to prove it using CoV? We may choose our axes so that the points are $(0, 0)$ and $(1, 0)$. We assume that the path is a function of $y(x)$ (can be more general and take a parametrised path $(x(t), y(t))$, which gives more equations, but trying to give a simple example). The total length of the curve is

$$L[y] = \int_0^1 \sqrt{1 + y'^2}\,dx. \tag{31}$$

Staring at this should suggest to you that $y' = 0$ is best. Let's check: Euler–Lagrange equation is

$$0 - \frac{d}{dx}\frac{y'}{\sqrt{1 + y'^2}} = 0.$$

Integrating once,

$$y' = A\sqrt{1+y'^2},$$

or

$$(1-A^2)y'^2 = A^2.$$

It follows that y' is constant. To satisfy the boundary conditions $y(0) = y(1) = 0$, we must have $y' = 0$, and hence the shortest distance is $y = 0$, as expected.

Brachistochrone Now we do a two-variable problem. Given $B = (0,0)$ and $C = (c,-d)$ lower than B, what shape is the curve joining B and C along which a particle will descend in least time? It's not a straight line in general.

To set up the problem, let the curve be $(x(t), y(t))$. The total time is given by

$$\int_0^T dt = \int_0^c \frac{dt}{dx}\,dx = \int_0^c \frac{dx}{\dot{x}},$$

providing that $\dot{x} \neq 0$ and we can write the curve as $y(x)$. Now, we also know that energy is conserved, so taking the GPE to be zero at B,

$$0 = mgy + \frac{1}{2}m(\dot{x}^2+\dot{y}^2) = mgy + \frac{1}{2}m\dot{x}^2\left(1+\left(\frac{dy}{dx}\right)^2\right)$$

Hence, writing $dy/dx = y'$, the integral can be written

$$\int_0^c \frac{\sqrt{1+y'^2}}{\sqrt{-2gy}}\,dx.$$

This is independent of x, so it is sensible to use the Beltrami identity, which gives

$$\frac{-1}{A\sqrt{2g}} = \frac{y'^2}{\sqrt{-2gy}\sqrt{1+y'^2}} - \frac{\sqrt{1+y'^2}}{\sqrt{-2gy}} = \frac{-1}{\sqrt{-2gy}\sqrt{1+y'^2}},$$

A a constant determined by the boundary conditions. Since we expect $y < 0$, we suppose $A > 0$, and rearranging gives

$$1 = \frac{\sqrt{y}}{\sqrt{A+y}}\frac{dy}{dx},$$

offering us

$$x = \int_0^y \frac{\sqrt{Y}}{\sqrt{A+Y}}\,dY,$$

since when $x = 0, y = 0$. Setting $Y = -A(1-u^2)$, so $dY = -2uA\,du$, we find

$$x = \int_{\sqrt{1-y/A}}^1 \frac{\sqrt{A}\sqrt{1-u^2}}{\sqrt{Au^2}} 2uA\,du = 2A\int_{\sqrt{1-y/A}}^1 \sqrt{1-u^2}\,du.$$

But this is the integral to calculate the area under a circle of radius A up to a point a distance $\sqrt{1-y/A} = U$ from the origin. Some geometry allows us to compute the area as the difference of a sector and a triangle, giving

$$x = A\left(\arccos U - U\sqrt{1-U^2}\right) = \tfrac{1}{2}A\left(\theta - \sin\theta\right),$$

where $\theta = 2\arccos U$ is double the angle from the vertical axis. Inverting the y equation gives

$$y = -A(1-\cos^2\tfrac{1}{2}\theta) = \tfrac{1}{2}A(\cos\theta - 1);$$

these are the equations for a cycloid. Making it pass through C is not analytically straightforward, but drawing a picture shows that any straight line with negative gradient passing through the origin intersects the cycloid once, so there is in fact a unique choice of A so that the cycloid passes through C.

Dido's problem Also known as the isoperimetric problem. Want to maximise the area we can enclose in a given length, L, of fence, say. We need a Lagrange multiplier to enforce the length condition. We can again approach this by assuming that y is a function of x, but this requires we assume symmetry about the x-axis, and it is better to consider a parametrised curve $(x(t), y(t))$. Then the Lagrangian is

$$L = \tfrac{1}{2}(yx' - xy') + \lambda\sqrt{x'^2+y'^2},$$

using the area formula we found from Green's Theorem in IA Vector Calculus. This looks like we should try the Beltrami identity, but alas, the expression vanishes! We have to use the Euler–Lagrange equations. There is one for each of x and y, but we can see that the Lagrangian is invariant under replacing $x \mapsto y$, $y \mapsto -x$, (i.e. a rotation by $\pi/2$) so they will actually both contain the same information. The x equation is

$$\begin{aligned} 0 &= \frac{1}{2}y' - \frac{d}{dt}\left(-\frac{1}{2}y + \lambda\frac{x'}{(x'^2+y'^2)^{1/2}}\right) \\ &\vdots \\ &= y'\left(1 - \lambda\frac{y'x''-x'y''}{(x'^2+y'^2)^{3/2}}\right). \end{aligned} \tag{32}$$

Repeating this with y gives the same equation with the leading y' replaced by $-x'$, and since we want to assume that the curve is sensibly parametrised so that the derivative never totally vanishes, we conclude that

$$\frac{y'x''-x'y''}{(x'^2+y'^2)^{3/2}} = \frac{1}{\lambda}$$

This seems rather a mess, but recall from Vector Calculus that the expression on the left is actually the curvature! Therefore the Euler–Lagrange equation just tells us that to make the area stationary, the curvature is constant. And of course, a curve of constant curvature is a circle. It follows that the maximum area is $L^2/(4\pi)$.[3]

Catenary What is the shape assumed by a chain of length L suspended between points $A = (-a, 0)$ and $B = (a, 0)$, where $2a < L$? (The general case of points at different height follows by rescaling and chopping the curve off after the right length, since the local equations are the same.) This time we will rely on our intuition that the curve is a single-valued function of $y(x)$. What do we need to minimise? The gravitational potential energy, which is given by $-\rho gy\,ds$ for a small chunk of curve. Hence the integral is

$$\int_{-a}^a (-\lambda - \rho gy)\sqrt{1+y'^2}\,dx.$$

Yet again the Beltrami identity's the simplest way: we have

$$\frac{1}{A} = \frac{(-\lambda-\rho gy)}{\sqrt{1+y'^2}}y'^2 - (-\lambda-\rho gy)\sqrt{1+y'^2} = \frac{(\lambda+\rho gy)}{\sqrt{1+y'^2}}.$$

Writing $A\lambda = \beta$ and $(A\rho g) = 1/\alpha$, the equation rearranges into

$$1 = \frac{y'^2}{(y/\alpha+\beta)^2 - 1}.$$

Integrating this equation gives

$$y = -\beta + \alpha\cosh\left((x-x_0)/\alpha\right),$$

and we need this to be symmetrical to fit through $(\pm a, 0)$, so $x_0 = 0$. We also have $\beta = \alpha\cosh(a/\alpha)$. The length is given by

$$L = \int_{-a}^a \sqrt{1+y'^2}\,dx = \int_{-a}^a \cosh(x/\alpha)\,dx = 2\alpha\sinh(a/\alpha).$$

One can check that there is only one possible value of α, since the right-hand side is a decreasing function of α that tends to $2a$ as $\alpha \to 0$. So the solution is

$$y = \alpha\big(\cosh(x/\alpha) - \cosh(a/\alpha)\big).$$

[3] Alternatively, (32) is actually a total derivative, so we see that $y + \lambda x'/(\sqrt{x'^2+y'^2}) = y_0$, and $-x + \lambda y'/\sqrt{x'^2+y'^2} = -x_0$. Rearranging, squaring and adding, we find that $(x-x_0)^2 + (y-y_0)^2 = \lambda^2$, the equation of a circle.

Part III. Applications

4. Fermat's Principle

Fermat's principle states that light rays take paths that make stationary the time taken; this is of course

$$T = \int_a^b \frac{ds}{v} = \frac{1}{c}\int_a^b n\,ds = \frac{1}{c}\int_a^b n\|\dot{\mathbf{x}}(t)\|dt, \tag{33}$$

where v is the speed of light at x, c is the speed of light in a vacuum, and $n = c/v$ is the refractive index at the point x.

If n is constant, our result on shortest distances shows that light travels in a straight line.

Snell's Law Suppose we have a 2D medium in which the speed of light varies with the height y, so the integral is

$$\int_a^b n(y)\sqrt{x'^2 + y'^2}\,dt.$$

Then the Euler–Lagrange equation for x is

$$\frac{d}{dt}\left(n(y)\frac{x'}{\sqrt{x'^2+y'^2}}\right) = 0.$$

If we let θ be the angle of the path to the vertical, this equation implies that $n(y)\sin\theta(y)$ is constant. This is *Snell's Law*, which is probably more familiar in the form when n suddenly jumps from n_1 to n_2,

$$n_1 \sin\theta_1 = n_2 \sin\theta_2. \tag{34}$$

5. Geodesics

A *geodesic* is a curve that minimises length (at least locally); we know that the infinitesimal length is given by $ds = \|\gamma'(t)\|\,dt$ for a curve γ. When we are looking for shortest distances on a surface $g(\mathbf{x}) = 0$, we could minimise the length of the curve $\int \|\mathbf{x}'(t)\|\,dt$ subject to the constraint that the curve lies in the surface (i.e minimise $\int \big(\|\mathbf{x}'(t)\| - \lambda(t)g(\mathbf{t})\big)\,dt$), but it is mathematically more fruitful to introduce coordinates $\mathbf{q}$ parametrising the surface, and minimise directly in these. This requires two new concepts, which are properly introduced in IB Geometry: the idea of a *tangent space* at a point on a surface, and with the *Riemannian metric*, which is a device for measuring length in the tangent space.

The tangent space at a point p is the set of vectors that are possible as tangent vectors of curves through p on the surface. It is a vector space, but there is no obvious way to measure the angle between two curves, or the length of a vector: there is no natural inner product on this space, so we appeal to the inner product that exists on Euclidean space. Suppose the parametrisation is $\mathbf{x} = \sigma(\mathbf{q})$. Then if $\mathbf{x}$ lies in the surface, $\mathbf{x}'(t) = \sum_i \sigma_{,i} q'^i(t)$, where we write the coordinates with indices upstairs, and $\sigma_{,i} = \partial\sigma/\partial q^i$ (essentially all you need to know at this point is that upstairs indices can only contract with downstairs and vice versa). Then

$$\|\mathbf{q}'\|^2 = \sum_{ij}(\sigma_{,i}\cdot\sigma_{,j})q'^i q'^j = \sum_{ij} g_{ij} q'^i q'^j, \tag{35}$$

where $g_{ij} = \sigma_{,i}\cdot\sigma_{,j}$ is called the *metric*; it is clear that it is positive-definite since the LHS is positive.[4] Because the surface bends, the basis $\sigma_{,i}$, and hence the tangent space, varies from point to point.

For a 2D surface, the line length can be written in the form

$$ds = \sqrt{Eu'^2 + 2Fu'v' + Gv'^2}\,dt, \tag{36}$$

and then the Euler–Lagrange equations are

$$\frac{E_u u'^2 + 2F_u u'v' + G_u v'^2}{2\sqrt{Eu'^2 + 2Fu'v' + Gv'^2}} - \frac{d}{dt}\frac{Eu' + Fv'}{\sqrt{Eu'^2 + 2Fu'v' + Gv'^2}} = 0. \tag{37}$$

For the general metric g_{ij}, we end up with equations for each $q^i(t)$, in the form

$$0 = \frac{d}{dt}(g_{ij}\dot{q}^j + g_{ji}\dot{q}^j) - \partial_i g_{jk}\dot{q}^j\dot{q}^k = 2g_{ij}\ddot{q}^j + 2\partial_k g_{ij}\dot{q}^j\dot{q}^k - \partial_i g_{jk}\dot{q}^j\dot{q}^k,$$

where we have assumed that t is the arc length, so the square roots all disappear. Contracting with the inverse of the metric gives the *geodesic equation*,

$$\ddot{q}^i + \Gamma^i_{jk}\dot{q}^j\dot{q}^k = 0, \tag{38}$$

where

$$\Gamma^i_{jk} = \frac{1}{2}g^{im}(\partial_j g_{mk} + \partial_k g_{mj} - \partial_m g_{jk}) \tag{39}$$

are the *Christoffel symbols*, which are essential in understanding derivatives in curved spaces (see Part II Differential Geometry and General Relativity).

5.1. Example: On the sphere

The sphere is simple enough to be treated both ways.

With constraints The Lagrangian is

$$\|\mathbf{x}'\| - \lambda(t)(\|\mathbf{x}\| - 1) \tag{40}$$

Differentiating with respect to λ gives the condition $\|\mathbf{x}\| = 1$. On the other hand, the Euler–Lagrange equations are

$$0 = -\lambda(t)\frac{\mathbf{x}}{\|\mathbf{x}\|} - \frac{d}{dt}\frac{\mathbf{x}'}{\|\mathbf{x}'\|} = -\lambda(t)\mathbf{x} - \frac{\mathbf{x}''}{\|\mathbf{x}'\|} + \frac{\mathbf{x}'(\mathbf{x}'\cdot\mathbf{x}'')}{\|\mathbf{x}'\|^3}$$

Now take $\mathbf{a}$ to be a constant vector perpendicular to $\mathbf{x}(0)$ and $\mathbf{x}'(0)$. Then this equation implies that $\mathbf{x}''(0)$ is also perpendicular to $\mathbf{a}$. Hence $\mathbf{x}$ is always perpendicular to $\mathbf{a}$, so the curve lies in a plane through the origin. (Recall from the Fresnet–Serret equations that $\mathbf{n}' = -\kappa\mathbf{t} + \tau\mathbf{b}$, and if $\tau = 0$, the curve is planar.) Such a circle is called a *great circle*.

With parametrisation The obvious parametrisation for the sphere is spherical coordinates, $\mathbf{x} = (\cos\varphi\sin\theta, \sin\varphi\sin\theta, \cos\theta)$. The tangent vectors are then

$$\sigma_{,\theta} = (\cos\varphi\cos\theta, \cos\varphi\sin\theta, -\sin\theta)$$
$$\sigma_{,\varphi} = (-\sin\varphi\sin\theta, \cos\varphi\sin\theta, 0).$$

From these we find the line element is

$$ds^2 = d\theta^2 + \sin^2\theta\,d\varphi^2, \tag{41}$$

and therefore the Lagrangian is

$$\ell[\theta,\varphi] = \int_a^b \sqrt{\theta'^2 + \sin^2\theta\,\varphi'^2}\,dt \tag{42}$$

The lack of φ-dependence implies that

$$\frac{\sin^2\theta\,\varphi'}{\sqrt{\theta'^2 + \sin^2\theta\,\varphi'^2}} = h,$$

a constant, and the Euler–Lagrange equation for θ is

$$\frac{\sin\theta\cos\theta\,\varphi'^2}{\sqrt{\theta'^2 + \sin^2\theta\,\varphi'^2}} - \frac{d}{dt}\frac{\theta'}{\sqrt{\theta'^2 + \sin^2\theta\,\varphi'^2}} = 0$$

Taking the arc length parametrisation, $\sqrt{\theta'^2 + \sin^2\theta\,\varphi'^2} = 1$, so

$$\sin^2\theta\,\varphi' = h, \qquad \sin\theta\cos\theta\,\varphi'^2 - \theta'' = 0.$$

Now, the sphere is symmetric enough that we can choose that $\varphi'(0) = 0$, $\theta'(0) = 1$, and so $h = 0$, and then $\sin^2\theta\,\varphi' = 0$. Assuming that $\sin\theta \neq 0$, we then have $\varphi' = 0$, and so $\theta = \theta(0) + t$, again a circle that lies in a plane through the origin.

[4] In more abstract spaces, we normally *define* the metric as a positive-definite symmetric tensor that does this measurement on the tangent space.

6. Lagrangian Mechanics

Newtonian mechanics is essentially based on the equation $m\ddot{\mathbf{x}} = \mathbf{F}$. Hamilton and others found another way to talk about mechanics, based on a variational principle, that makes it much easier to understand the significance of symmetries.

If the kinetic energy is written as T and the potential energy as V, then the *Lagrangian* is defined as $L = T - V$.[5] (For simple systems, $T = \frac{1}{2}m\|\dot{\mathbf{x}}\|^2$, while $V(\mathbf{x}, t)$ is the potential.) The *action* is

$$S[\mathbf{x}] = \int_{t_0}^{t_1} L(\mathbf{x}, \dot{\mathbf{x}}, t)\, dt. \tag{43}$$

Hamilton's Principle states that *classical paths make stationary the action.*

Taking the simple form of the Lagrangian, we find via the Euler–Lagrange equations that

$$0 = \frac{\partial L}{\partial \mathbf{x}} - \frac{d}{dt}\frac{\partial L}{\partial \dot{\mathbf{x}}} = -\nabla V - \frac{d}{dt}(m\dot{\mathbf{x}}) = -\nabla V - m\ddot{\mathbf{x}},$$

which we recognise as Newton's equations.

Example: central force Suppose we have central force, as in IA Dynamics and Relativity. Then $\|\dot{\mathbf{x}}\|^2 = \dot{r}^2 + r^2(\dot{\theta}^2 + \sin^2\theta\,\dot{\varphi}^2)$, and $V = V(r)$. Then L does not depend on φ, so

$$\frac{\partial L}{\partial \dot{\varphi}} = mr^2 \sin^2\theta\,\dot{\varphi} = hm$$

is constant. The other equations are

$$m\ddot{r} = mr(\dot{\theta}^2 + \sin^2\theta\,\dot{\varphi}^2) - V'(r)$$

for the radius, and

$$\frac{d}{dt}(r^2\dot{\theta}) = r^2\ddot{\theta} + 2r\dot{\theta} = r^2\sin\theta\cos\theta\,\dot{\varphi}^2.$$

This last shows that if we choose our axes so that $\theta(0) = \pi/2$ and $\dot{\theta}(0) = 0$, θ remains zero, so the path lives in the equatorial plane. Then $\sin\theta = 1$, so we have

$$m\ddot{r} = mr\dot{\varphi}^2 - V'(r) = \frac{mh^2}{r^3} - V'(r),$$

as usual; the RHS is the negative derivative of what is usually called the *effective potential*, $V_{\text{eff}}(r) = \frac{mh^2}{2r^2} + V(r)$.

7. Hamiltonian Mechanics

Lagrange's equations are the n second-order equations

$$\frac{\partial L}{\partial \mathbf{x}} - \frac{d}{dt}\frac{\partial L}{\partial \dot{\mathbf{x}}} = 0,$$

derived from the Lagrangian $L(\mathbf{x}, \dot{\mathbf{x}}, t)$. We can formulate mechanics in another way, using first-order equations. We do this by using the Legendre transform: define the *Hamiltonian*

$$H(\mathbf{p}, \mathbf{x}, t) = \sup_{\dot{\mathbf{x}} \in \mathbb{R}^n} (\mathbf{p}\cdot\dot{\mathbf{x}} - L(\mathbf{x}, \dot{\mathbf{x}}, t)). \tag{44}$$

L is a convex function of $\dot{\mathbf{x}}$, so this makes sense. We have, of course, $\mathbf{p} = m\dot{\mathbf{x}}$ if the Lagrangian has the simple kinetic term; either way, $\mathbf{p} = \partial L/\partial \dot{\mathbf{x}}$ is called the *conjugate momentum* of $\mathbf{x}$. Then we consider the integral

$$I[\mathbf{p}, \mathbf{x}] = \int (\mathbf{p}\cdot\dot{\mathbf{x}} - H(\mathbf{p}, \mathbf{x}, t))\, dt. \tag{45}$$

The Euler–Lagrange equations of this are first-order:

$$\dot{\mathbf{x}} = \frac{\partial H}{\partial \mathbf{p}}, \qquad \dot{\mathbf{p}} = -\frac{\partial H}{\partial \mathbf{x}}; \tag{46}$$

these are *Hamilton's equations.* There are $2n$ first-order equations. What about H? We can compute

$$\frac{dH}{dt} = \frac{\partial H}{\partial \mathbf{x}}\cdot\dot{\mathbf{x}} + \frac{\partial H}{\partial \mathbf{p}}\cdot\dot{\mathbf{p}} + \frac{\partial H}{\partial t} = \frac{\partial H}{\partial t}, \tag{47}$$

using Hamilton's equations, so H only changes by its direct dependence on t. In fact, H represents the *energy* of the system. (We see that in the old notation, $H = T + V$, although the variables are different.)

8. Noether's Theorem and Conservation Laws

$t \mapsto t'(t)$, $\mathbf{x} \mapsto \mathbf{x}'(t')$ is called a *symmetry* if the action S is unchanged by this substitution,

$$S'[\mathbf{x}'] = \int_{a'}^{b'} L(t', \mathbf{x}', \dot{\mathbf{x}}')\, dt' = S[\mathbf{x}]. \tag{48}$$

Theorem (Noether's Theorem). *For each continuous symmetry of S there is a corresponding quantity Q, which is conserved by solutions to the Euler–Lagrange equations.*

Proof. Parametrise the symmetry as $t \mapsto t_\alpha$, $x \mapsto x_\alpha$. If the symmetry is continuous, we can expand to first order in α about $\alpha = 0$, where to first order in α we have

$$t_\alpha = t + \alpha\tau(t) + o(\alpha), \qquad x_\alpha(t_\alpha) = x(t) + \alpha\xi(t) + o(\alpha).$$

For now suppose that α is a constant. From now on we discard all terms beyond order α without comment. Then the transformed time-derivative is

$$\frac{dx_\alpha}{dt_\alpha}(t_\alpha) = \frac{dt}{dt_\alpha}\frac{d}{dt}(x(t) + \alpha\xi(t)) = \frac{\dot{x}(t) + \alpha\dot{\xi}(t)}{1 + \alpha\dot{\tau}(t)} = \dot{x} + \alpha(\dot{\xi} - \dot{\tau}\dot{x}), \tag{49}$$

where all terms on the RHS are evaluated at t. Then the new action is

$$S_\alpha[x_\alpha] = \int_{a_\alpha}^{b_\alpha} L\left(t_\alpha, x_\alpha(t_\alpha), \frac{dx_\alpha}{dt_\alpha}(t_\alpha)\right) dt_\alpha,$$

Changing variables back to t, we have $dt_\alpha = (1 + \alpha\dot{\tau})\, dt$, so

$$S_\alpha[x_\alpha] = \int_a^b L\left(t_\alpha(t), x_\alpha(t_\alpha(t)), \frac{dx_\alpha}{dt_\alpha}(t_\alpha(t))\right)(1 + \alpha\dot{\tau})\, dt, \tag{50}$$

and now we can look directly at the difference $S_\alpha[x_\alpha] - S[x]$ by considering the integrands. Expanding L for small α gives

$$\begin{aligned} L&\left(t_\alpha, x_\alpha(t_\alpha), \frac{dx_\alpha}{dt_\alpha}(t_\alpha)\right)(1 + \alpha\dot{\tau}) - L(t, x, \dot{x}) \\ &= \alpha\dot{\tau}L + \alpha\tau\frac{\partial L}{\partial t} + \alpha\xi\frac{\partial L}{\partial x} + \alpha(\dot{\xi} - \dot{\tau}\dot{x})\frac{\partial L}{\partial \dot{x}} \\ &= \alpha\left(\tau\frac{\partial L}{\partial t} + \xi\frac{\partial L}{\partial x} + \dot{\tau}\left(L - \dot{x}\frac{\partial L}{\partial \dot{x}}\right) + \dot{\xi}\frac{\partial L}{\partial \dot{x}}\right). \end{aligned}$$

Therefore for the transformation to be a symmetry, we must have

$$\tau\frac{\partial L}{\partial t} + \dot{\tau}\left(L - \dot{x}\frac{\partial L}{\partial \dot{x}}\right) + \xi\frac{\partial L}{\partial x} + \dot{\xi}\frac{\partial L}{\partial \dot{x}} = 0 \tag{51}$$

Now we need the conserved current. It is not obvious that (51) is the total derivative of something (indeed, we'll find it isn't). We also need to use that x solves the Euler–Lagrange equation

$$E(L) := \frac{\partial L}{\partial x} - \frac{d}{dt}\frac{\partial L}{\partial \dot{x}} = 0.$$

We notice that the first two terms of (51) look like fragments of a derivative of $\tau(L - \dot{x}\partial L/\partial\dot{x})$, so it seems sensible to compute the total derivative of the bracket and see what happens. We find that

$$\frac{d}{dt}\left(L - \dot{x}\frac{\partial L}{\partial \dot{x}}\right) = \frac{\partial L}{\partial t} + \dot{x}\frac{\partial L}{\partial x} + \ddot{x}\frac{\partial L}{\partial x} - \ddot{x}\frac{\partial L}{\partial x} - \dot{x}\frac{d}{dt}\frac{\partial L}{\partial \dot{x}} = \frac{\partial L}{\partial t} + \dot{x}E(L).$$

[5] Why? Because this works.

Therefore in fact we have that

$$\tau \frac{\partial L}{\partial t} + \dot{\tau}\left(L - \dot{x}\frac{\partial L}{\partial \dot{x}}\right) = \frac{d}{dt}\left(\tau\left(L - \dot{x}\frac{\partial L}{\partial \dot{x}}\right)\right) - \tau\dot{x}E(L).$$

The other two terms are somewhat simpler: adding and subtracting the same term gives

$$\xi\frac{\partial L}{\partial x} + \dot{\xi}\frac{\partial L}{\partial \dot{x}} = \xi\left(\frac{\partial L}{\partial x} - \frac{d}{dt}\frac{\partial L}{\partial \dot{x}}\right) + \xi\frac{d}{dt}\frac{\partial L}{\partial \dot{x}} + \dot{\xi}\frac{\partial L}{\partial \dot{x}} = \xi E(L) + \frac{d}{dt}\left(\xi\frac{dL}{d\dot{x}}\right)$$

Hence the expression on the left of (51) becomes

$$(\xi - \tau\dot{x})E(L) + \frac{d}{dt}\left(\tau\left(L - \dot{x}\frac{\partial L}{\partial \dot{x}}\right) + \xi\frac{\partial L}{\partial \dot{x}}\right)$$

If x is a solution to $E(L) = 0$, the first term vanishes, and we see that

$$Q := \tau\left(L - \dot{x}\frac{\partial L}{\partial \dot{x}}\right) + \xi\frac{\partial L}{\partial \dot{x}} \qquad (52)$$

is the required conserved quantity. □

The calculation in the proof is fairly tortuous. There is actually a slightly puzzling method we can use to obtain Q directly from the difference of the integrands: suppose that α is a function of t. Then instead of (51) we find the difference of the integrands to first order in α and $\dot{\alpha}$ is

$$\begin{aligned} L\left(t_\alpha, x_\alpha(t_\alpha), \frac{dx_\alpha}{dt_\alpha}(t_\alpha)\right)&\left(1 + \dot{(\alpha\tau)}\right) - L(t, x, \dot{x}) \\ &= \alpha\tau\frac{\partial L}{\partial t} + \alpha\xi\frac{\partial L}{\partial x} + \dot{(\alpha\tau)}\left(L - \dot{x}\frac{\partial L}{\partial \dot{x}}\right) + \dot{(\alpha\xi)}\frac{\partial L}{\partial \dot{x}} \\ &\vdots \\ &= \alpha\big(\dot{Q} + (\xi - \tau\dot{x})E(L)\big) + \dot{\alpha}Q, \end{aligned}$$

so if we let α vary with t, we can find Q by simply computing this expansion and reading off the coefficient of $\dot{\alpha}$.

It is simple to generalise this analysis to more variables and more complicated Lagrangians: we simply obtain extra terms, with corresponding extra terms in the Euler–Lagrange equations, which simplify into a form similar to the above.

Spacial Translations A simple continuous transformation of space only is given by $\tau = 0, \xi = 1$, and then it's clear that this satisfies (51) if $\partial L/\partial x = 0$. Then $Q = \partial L/\partial \dot{x}$ is conserved, AKA the momentum p.

Time Translation The corresponding transformation for time is $\tau = 1, \xi = 0$. This is a symmetry if $\partial L/\partial t = 0$, and then the current is $Q = L - \dot{x}\frac{\partial L}{\partial \dot{x}}$, i.e. the energy.

Angular Momentum A small rotation about axis $\mathbf{n}$ is given by $\mathbf{x}_\alpha = \mathbf{x} + \alpha\mathbf{n} \times \mathbf{x}$. If L is invariant under this rotation, the conserved quantity is

$$\frac{\partial L}{\partial \dot{\mathbf{q}}} \cdot \boldsymbol{\xi} = \mathbf{p} \cdot (\mathbf{n} \times \mathbf{x}) = \mathbf{n} \cdot (\mathbf{x} \times \mathbf{p}) = \mathbf{n} \cdot \mathbf{L},$$

the component of the angular momentum in the direction of $\mathbf{n}$. If L is unchanged for any $\mathbf{n}$, every component of the angular momentum is conserved.

9. The Second Variation

We want to understand when the stationary point is a minimum, so we want an equivalent of the Hessian. This is the *second variation*, defined by

$$F[y + t\phi] - F[y] - tDF[y](\phi) = t^2 D^2F[y](\phi) + o(t^2), \qquad (53)$$

i.e. the coefficient of t^2 in the expansion.

We would like conditions on this that give us a minimum, say. First, we find an expression for $D^2F[y](\phi)$:

$$\begin{aligned} &F[y + t\phi] - F[y] - tDF[y](\phi) \\ &\quad = \int_a^b \left(L(x, y + t\phi, y' + t\phi') - L - t\phi\frac{\delta F}{\delta y}\right) dx \\ &\quad = \frac{1}{2}t^2 \int_a^b \left(\phi^2\frac{\partial^2 L}{\partial y^2} + 2\phi\phi'\frac{\partial^2 L}{\partial y'\partial y} + \phi'^2\frac{\partial^2 L}{\partial y'^2}\right) dx + o(t^2). \end{aligned}$$

Integrating by parts,

$$D^2F[y](\phi) = \frac{1}{2}\left[\phi^2\frac{\partial^2 L}{\partial y'\partial y}\right]_a^b + \frac{1}{2}\int_a^b \left(\phi^2\left(\frac{\partial^2 L}{\partial y^2} - \frac{d}{dt}\frac{\partial^2 L}{\partial y'\partial y}\right) + \phi'^2\frac{\partial^2 L}{\partial y'^2}\right) dx. \qquad (54)$$

For a minimum, we need this to be nonnegative at the point: if it is everywhere nonnegative, the functional is convex, and any minimum we can find will be a global one. We have to restrict ϕ to be in accordance with the boundary conditions, and differentiable enough for the integral to make sense. From now on, we take fixed y at the ends, so $\phi = 0$ at the endpoints and the boundary term always vanishes.

Legendre condition A simple necessary condition that $y = y_0$ should be a local minimum is that

$$\frac{\partial^2 L}{\partial y'^2}(y_0) \geqslant 0; \qquad (55)$$

this is easy to see if we imagine a very wiggly ϕ that remains quite small in magnitude: then if this coefficient is negative, we can make $D^2F[y_0](\phi)$ as negative as we like.

We can rename the coefficients to write the second variation more compactly as

$$D^2F[y](\phi) = \frac{1}{2}\int_a^b \left(P_y(x)\phi'^2 + Q_y(x)\phi^2\right) dx. \qquad (56)$$

Legendre's condition shows $P_{y_0} \geqslant 0$ is necessary, but this is not sufficient. A sufficient condition for a local minimum at y_0 is $Q_{y_0}(x) \geqslant 0$ and $P_{y_0}(x) > 0$ for $a < x < b$. In particular, the condition on P ensures that the first term is positive.

Jacobi condition To find a better set of conditions, we subtract 0 from the second variation, by using

$$0 = \frac{1}{2}\int_a^b (w\phi^2)'\, dx = \frac{1}{2}\int_a^b \left(w'\phi^2 + 2\phi\phi' w\right) dx. \qquad (57)$$

Then the second variation becomes

$$D^2F[y](\phi) = \frac{1}{2}\int_a^b \left(P\phi'^2 - 2w\phi\phi' + (Q - w')\phi^2\right) dx,$$

and assuming $P > 0$ and completing the square,

$$D^2F[y](\phi) = \frac{1}{2}\int_a^b \left(P\left(\phi' - \frac{w\phi}{P}\right)^2 + \left(Q - w' - \frac{w^2}{P}\right)\phi^2\right) dx. \qquad (58)$$

For the integrand to be a perfect square, we need to have

$$Q - w' - w^2/P = 0. \qquad (59)$$

If this is the case, the integral is nonnegative. In fact it is positive, since to be zero, we would have to have $\phi' = w\phi/P$, of which the only solution with a zero is identically zero. Hence the second variation is then strictly positive.

We therefore have to solve (59). Note first that w has no specified boundary conditions. (59) is nonlinear, but can be made linear at the cost of turning it into a second-order equation, by setting $w = -Pu'/u$; it becomes

$$-(Pu')' + Qu = 0. \qquad (60)$$

We just need to find a solution of this that has no zeros on $[a, b]$ (so w still makes sense), and if we can it follows that the second variation is positive. This is a *Sturm–Liouville equation*, which you will learn all about in IB Methods.