Lecture 1

1.1 The basic optimal control problem

System:

$$\dot{x} = f(t, x, u), \qquad x(t_0) = x_0$$

where $x \in \mathbb{R}^n$ is the state, $u \in \mathcal{U} \subset \mathbb{R}^m$ is the control ($\mathcal{U}$ is a control set).

Cost:

$$J(u) := \int_{t_0}^{t_1} L(t, x(t), u(t))dt + K(x(t_1))$$

where:

- L and K are given functions (*running cost* and *terminal cost*);
- Final time t_1 is either free or fixed;
- Final state $x(t_1) = x_1$ is either free or fixed or belongs to some target set.

Problem: find u which minimizes $J(u)$ over all admissible controls.

Later we'll come back to this and fill in some technical details (regularity of f, admissible functions u, precise formulation of versions of the above problem, etc.)

Examples of applications:

- Send a rocket to the moon, minimize fuel consumption;
- Produce a given amount of a chemical, minimize time and/or amount of catalyst used—or maximize amount of produced chemical in given time, etc.;
- Bring sales of a new product to a desired level through advertising, minimize the amount of money spent on advertising;
- Communication: maximize throughput/accuracy for given bandwidth/capacity;
- More examples?

$\longrightarrow$ This course: mathematical theory.

Note: this is *dynamic* optimization, in the sense that it involves a dynamical system and time.

Can view it as choosing the best *path* among all paths feasible for the system, with respect to the given cost function. In this sense, the problem is *infinite-dimensional*, because the space of paths is an infinite-dimensional function space.

Compare with static, finite-dimensional optimization problem: find a minimum of a given function $f : \mathbb{R}^n \to \mathbb{R}$.

It is useful to have some knowledge of finite-dimensional optimization, as there are similarities (but don't need to know very much).

Rough outline of the course (see course homepage for more details):

- Brief review of finite-dimensional optimization
- Calculus of variations
- The maximum principle (MP) of optimal control
- The Hamilton-Jacobi-Bellman (HJB) equation and dynamic programming
- Linear Quadratic Regulator (LQR) problem
- Other topics (maximum principle on manifolds; further connections between the Maximum Principle and the HJB equation; robust control)

This subject has a rich history. We'll make some historical remarks along the way, and will try to follow the chronological development (especially from variational calculus to the MP).

1.2 Some background on finite-dimensional optimization

References: Luenberger, Linear and Nonlinear Programming; also A-F, Sections 5.2-5.4.

Consider a function

$$f : D \subset \mathbb{R}^n \to \mathbb{R}$$

where D is a region, domain of f.

Definition 1 A point x^* is a *local minimum* of f if $\exists\, \varepsilon > 0$ such that $\forall\, x \in D$ satisfying $|x - x^*| < \varepsilon$ we have

$$f(x^*) \leq f(x). \tag{1.1}$$

If the inequality is strict, then we have a *strict* local minimum.

If (1.1) holds $\forall\, x \in D$, then the minimum is *global*.

Maxima are defined similarly.

1.2.1 Unconstrained optimization

This is when all x near x^* in R^n are in D, i.e., x^* belongs to D together with its $\mathbb{R}^n$-neighborhood. The simplest case is $D = \mathbb{R}^n$, which is sometimes called *completely unconstrained*. However, as far as *local* minimization is concerned, it's enough to assume that D is an open subset of $\mathbb{R}^n$, or at least that x^* is an interior point of D.

The first-order necessary condition

Let x^* be a local min.

Suppose $f \in \mathcal{C}^1$ (continuously differentiable).

Take any direction $d \in \mathbb{R}^n$. Then $x^* + \alpha d \in D$ for all $\alpha > 0$ small enough.

(Note: we are in the unconstrained case.)

Consider $f(x^* + \alpha d)$ as a function of $\alpha \in \mathbb{R}$, say $g(\alpha)$, for a fixed d. This reduces the vector case to the scalar case (ordinary calculus). We can write (think Taylor series):

$$g(\alpha) = g(0) + g'(0)\alpha + o(\alpha)$$

or, expanding in terms of f,

$$f(x^* + \alpha d) = f(x^*) + \nabla f(x^*) \cdot \alpha d + o(\alpha) \tag{1.2}$$

where

$$\nabla f := (f_{x_1}, \dots, f_{x_n})$$

is the *gradient* of f. (f_{x_i} is a shorthand for the partial derivative $\frac{\partial f}{\partial x_i}$.)

Recall: $o(\alpha)$ goes to 0 faster than α:

$$\lim_{\alpha \to 0} \frac{o(\alpha)}{\alpha} = 0$$

(it represents "higher-order terms").

To have $f(x^* + \alpha d) \geq f(x^*)$ for small α, we must have $\nabla f(x^*) \cdot d \geq 0$, and this must be true for all d.

To prove this: start with (1.2), move $f(x^*)$ to the left, then divide both sides by α and take the limit as $\alpha \searrow 0$.

Actually, since we can replace d by $-d$, we see that $\nabla f(x^*) \cdot d$ must be 0. And since d was arbitrary, this implies

$$\boxed{\nabla f(x^*) = 0}$$

This is the *first-order necessary condition for optimality.*

A point x^* satisfying this condition is called a *stationary point.*

"First-order" – because it's derived using the first-order expansion (1.2).

Note: we assumed that $f \in \mathcal{C}^1$ and that we're dealing with the unconstrained problem ($x^* \in \text{int} D$).

Exercise 1 (due Jan 24) Justify the validity of the expansion (1.2) rigorously, i.e., verify the $o(\alpha)$ property. Remember that f is only assumed to be $\mathcal{C}^1$.

Remark 1 The term $o(\alpha)$ above depends not only on α but also on d. With no loss of generality we can restrict attention to $|d| = 1$. From the Mean Value Theorem we easily see that the convergence of $o(\alpha)$ to 0 as $\alpha \to 0$ is actually uniform with respect to d, in the sense that $\forall\, \varepsilon > 0\ \exists\, \delta > 0$ such that $\alpha < \delta \Rightarrow |o(\alpha)| < \varepsilon\ \forall\, d$ with $|d| = 1$. $\square$

Second-order conditions

Now, let's expand $g(\alpha) = f(x^* + \alpha d)$ again, but this time include *second-order terms. Assume* $f \in \mathcal{C}^2$.

$$g(\alpha) = g(0) + g'(0)\alpha + \frac{1}{2}g''(0)\alpha^2 + o(\alpha^2).$$

Assume $g'(0) = 0$ as above. What is $g''(0)$?

We have: $g'(\alpha) = \nabla f(x^* + \alpha d) \cdot d$. In coordinates, this is

$$\sum_{i=1}^{n} f_{x_i}(x^* + \alpha d)d_i.$$

Differentiating again:

$$g''(\alpha) = \sum_{i,j=1}^{n} f_{x_i x_j}(x^* + \alpha d)d_i d_j$$

So

$$g''(0) = \sum_{i,j=1}^{n} f_{x_i x_j}(x^*)d_i d_j$$

or in matrix notation:

$$g''(0) = d^T \nabla^2 f(x^*)d$$

where

$$\nabla^2 f := \begin{pmatrix} f_{x_1x_1} & \cdots & f_{x_1x_n} \\ & \ddots & \\ f_{x_nx_1} & \cdots & f_{x_nx_n} \end{pmatrix}$$

is the *Hessian* matrix.

Thus our second-order expansion is (assuming first-order necessary condition holds):

$$f(x^* + \alpha d) = f(x^*) + \frac{1}{2}d^T \nabla^2 f(x^*)d\alpha^2 + o(\alpha^2). \tag{1.3}$$

Divide by α^2 and take the limit as $\alpha \searrow 0$:

$$\lim_{\alpha \searrow 0} \frac{f(x^* + \alpha d) - f(x^*)}{\alpha^2} = \frac{1}{2}d^T \nabla^2 f(x^*)d + \lim_{\alpha \searrow 0} \frac{o(\alpha^2)}{\alpha^2}$$

The last term is zero, and if x^* is a local minimum then the whole thing must be ≥ 0. Since d was arbitrary, the matrix $\nabla^2 f(x^*)$ must be positive semidefinite:

$$\boxed{\nabla^2 f(x^*) \geq 0} \quad \text{(positive semidefinite)}$$

This is the *second-order necessary condition for optimality.*

Note that this condition:

- Must be used in combination with the first-order condition;
- Assumes that $f \in \mathcal{C}^2$;
- Applies only to the unconstrained case ;
- Distinguishes minima from maxima (by inequality sign) while the first-order one does not.

How about a sufficient condition?

Suppose that we strengthen the necessary conditions to:

1) $\nabla f(x^*) = 0$

2) $\nabla^2 f(x^*) > 0$ (positive *definite*)

Consider again the expansion (1.3), assuming $f \in \mathcal{C}^2$. Since $\nabla^2 f(x^*) > 0$, we have

$$d^T \nabla^2 f(x^*) d \geq \lambda_{\min}(\nabla^2 f(x^*)) |d|^2$$

where $\lambda_{\min}$ is the smallest eigenvalue and in our case it is > 0. Since $\dfrac{o(\alpha^2)}{\alpha^2} \to 0$ as $\alpha \to 0$, we can pick $\varepsilon > 0$ small enough so that

$$0 < \alpha < \varepsilon \quad \Rightarrow \quad |o(\alpha^2)| < \frac{1}{2} \lambda_{\min}(\nabla^2 f(x^*)) |d|^2 \alpha^2$$

and we get $f(x^* + \alpha d) > f(x^*)$. Thus x^* is actually a *strict* local minimum.

We have derived the *second-order sufficient conditions for optimality.*

Feasible directions, global minima, and convex problems

Constrained case: What if x^* is a boundary point of D? (E.g., D closed.)

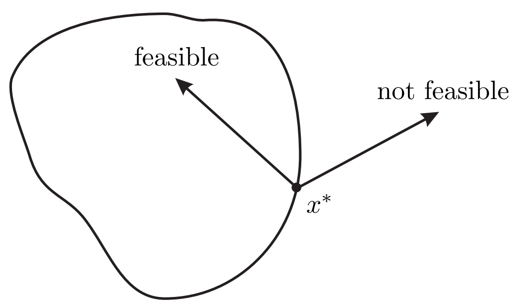

Figure 1.1: Feasible directions

Call $d \in \mathbb{R}^n$ a *feasible direction* (at x^*) if $x^* + \alpha d \in D$ for small enough $\alpha > 0$. Then $\nabla f(x^*) \cdot d \geq 0$ still true for every feasible direction d. Since not all d are feasible, $\nabla f(x^*) = 0$ is no longer true.

If D is *convex*, then the feasible direction approach is suitable: can connect x^* to any other $x \in D$ by a ray d . If D is not convex, then the approach is conservative

What if $D \subset \mathbb{R}^n$ is a lower-dimensional surface?

Feasible directions no longer make sense (in the sense that the condition in terms of them is vacuous). Need to refine the approach—more on this later.

What about global minima?

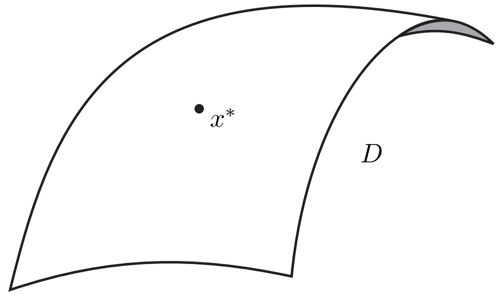

Figure 1.2: Surface

Weierstrass Theorem: If f is a continuous function and D is a compact set, then the global min

$$\min_{x \in D} f(x)$$

exists.

Compact sets can be defined, in $\mathbb{R}^n$, in three equivalent ways:

1) Closed and bounded
2) Open covers have finite subcovers
3) Every sequence has a subsequence converging to some point in D (sequential compactness)

Weierstrass Theorem will be useful later when discussing existence of optimal controls.

In practice, to find a global minimum:

- Find all interior points satisfying $\nabla f(x^*) = 0$ (stationary points);
- If f is not differentiable everywhere, include also points where ∇f does not exist (together with stationary points these give *critical points*);
- Find all *boundary points* satisfying $\nabla f(x^*) \cdot d \geq 0$ for all feasible d;
- Compare values at all these points and choose the smallest one.

Can also use the second-order necessary condition, but harder computationally.

Convex problems: [Luenberger, Sect. 6.4 and 6.5]

If D is a convex set and f is a convex function, then:

- A local minimum is automatically a global one;
- The first-order necessary condition is also a sufficient condition (when $f \in \mathcal{C}^1$). Thus if $\nabla f(x^*) \cdot d \geq 0$ for all feasible directions d, or in particular $\nabla f(x^*) = 0$ if x^* is an interior point, then x^* is a minimum (local hence also global).

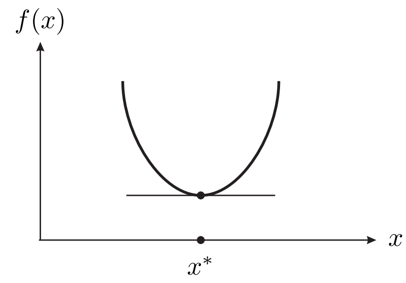

Figure 1.3: Convex function

Convex function: the graph lies above the linear approximation $f(x^*) + \nabla f(x^*) \cdot (x - x^*)$.

Efficient numerical algorithms exist for converging to points satisfying $\nabla f(x^*) = 0$ (stationary points), and so for convex problems they give a global minimum.

$\longrightarrow$ So, we see that local results are relevant for global optimization too, provided these global problems have some nice features.

1.2.2 Constrained optimization

Suppose that D is a surface in $\mathbb{R}^n$ given by the equations

$$h_1(x) = h_2(x) = \cdots = h_m(x) = 0 \tag{1.4}$$

where h_i are $\mathcal{C}^1$ functions from $\mathbb{R}^n$ to $\mathbb{R}$.

$\longrightarrow$ Here $f \in \mathcal{C}^1$ is tacitly assumed whenever needed.

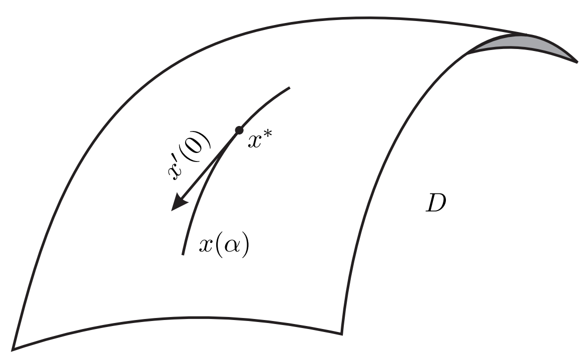

Figure 1.4: Tangent vector

The first-order necessary condition (Lagrange multipliers)

Assume that $x^* \in D$ is a *regular point* in the sense that the gradients ∇h_i, $i = 1, \ldots, m$ are linearly independent at x^*.

Instead of lines/rays $x^*+\alpha d$, $\alpha > 0$, we need to consider *curves* $x(\alpha) \in D$, $x(0) = x^*$. The previously considered $x(\alpha) = x^*+\alpha d$ becomes a special case. For any such curve $x(\alpha)$ we can consider the function $g(\alpha) = f(x(\alpha))$. Let's still keep $\alpha > 0$ as earlier.

Similarly to before, we see that x^* is a minimum $\Rightarrow g'(0) \geq 0$. But

$$g'(\alpha) = \nabla f(x(\alpha)) \cdot x'(\alpha) \;\Rightarrow\; g'(0) = \nabla f(x^*) \cdot x'(0)$$

so we must have $\nabla f(x^*) \cdot x'(0) \geq 0$. Here $x'(0)$ is a tangent vector to D at x^* (see figure). It lives in the tangent space to D at x^*, which is denoted by $T_{x^*}D$.

Since we don't have inequality constraints, D has no boundary $\Rightarrow$ we can reverse the direction of the curve to get $-x'(0)$ instead of $x'(0)$. So, must actually have $\nabla f(x^*) \cdot x'(0) = 0$. (Cf. replacing d by $-d$ earlier.)

What do we know about $x'(0)$?

$h_i(x(\alpha)) = 0$ for all α (with small enough α), for all $i = 1, \dots, m$. Differentiating:

$$0 = \left.\frac{d}{d\alpha}\right|_{\alpha=0} h_i(x(\alpha)) = \nabla h_i(x^*) \cdot x'(0)$$

which is a condition on $x'(0)$.

Actually, one can show that all tangent vectors to D at x^* are exactly vectors d satisfying

$$\nabla h_i(x^*) \cdot d = 0 \qquad \forall\, i$$

(we proved in one direction; see [Luenberger, Sect. 10.2] for the converse). So:

$$\nabla f(x^*) \cdot d = 0 \qquad \forall\, d \;\text{ such that } \nabla h_i(x^*) \cdot d = 0 \quad \forall\, i \tag{1.5}$$

What relation between $\nabla f(x^*)$ and $\nabla h_i(x^*)$ does this imply? (We want to get rid of d in the condition.)

Claim:

$$\nabla f(x^*) \in \operatorname{span}\{\nabla h_i(x^*),\, i = 1, \ldots, m\}. \tag{1.6}$$

Proof. Otherwise $\nabla f(x^*)$ has a component orthogonal to $\operatorname{span}\{h_i(x^*)\}$, i.e, $\exists\, d \neq 0$ such that

$$\nabla h_i(x^*) \cdot d = 0 \qquad \& \qquad \nabla f(x^*) = \sum_{i=1}^{m} \lambda_i \nabla h_i(x^*) + d$$

Take inner product with d:

$$\nabla f(x^*) \cdot d = d \cdot d \neq 0$$

which contradicts the previous fact. □

Geometrically: $\nabla f(x^*) \perp D$ at x^*.

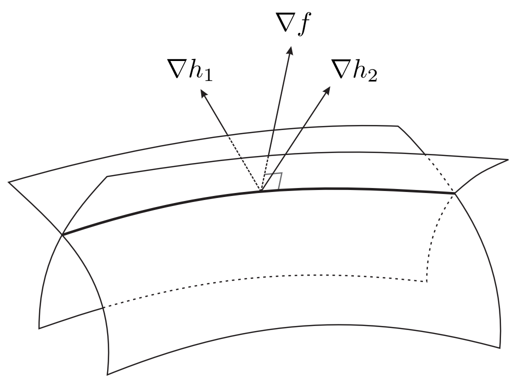

Figure 1.5: Gradients

Intuition: Unless ∇f is normal to D, there are curves in D whose tangent vectors make both positive and negative inner products with ∇f, hence in particular f can decrease.

The above proof relies on geometry. It is a bit involved; there was the gap that we didn't prove the inverse implication as in Luenberger. If you don't like geometry, there's another proof of (1.6) which is purely analytic (although the geometric intuition is lost). This will actually be useful when we get to calculus of variations problems with constraints.

Alternative proof: Assume for simplicity that there is only one constraint ($m = 1$); the extension to $m > 1$ is straightforward. Suppose that $\nabla f(x^*)$ and $\nabla h_1(x^*)$ are linearly independent, which means that the matrix

$$\begin{pmatrix} f_{x_1} & f_{x_2} & \cdots & f_{x_n} \\ (h_1)_{x_1} & (h_1)_{x_2} & \cdots & (h_1)_{x_n} \end{pmatrix}(x^*)$$

has rank 2. Then it has a nonsingular 2×2 submatrix, and we may assume with no loss of generality (reordering the coordinates if necessary) that this is

$$\begin{pmatrix} f_{x_1} & f_{x_2} \\ (h_1)_{x_1} & (h_1)_{x_2} \end{pmatrix}(x^*) \tag{1.7}$$

Consider the mapping

$$F : (x_1, x_2, x_3, \ldots, x_n)^T \mapsto (f(x), h_1(x), x_3, \ldots, x_n)^T$$

In view of the nonsingularity of (1.7), we can apply the Inverse Function Theorem (see, e.g., [Rudin, p. 221]) and conclude that there are neighborhoods of x^* and $F(x^*)$ on which the mapping F is a bijection (has an inverse). But this means that there are points x near x^* such that $h_1(x) = h_1(x^*)$ but $f(x) < f(x^*)$, a contradiction. □

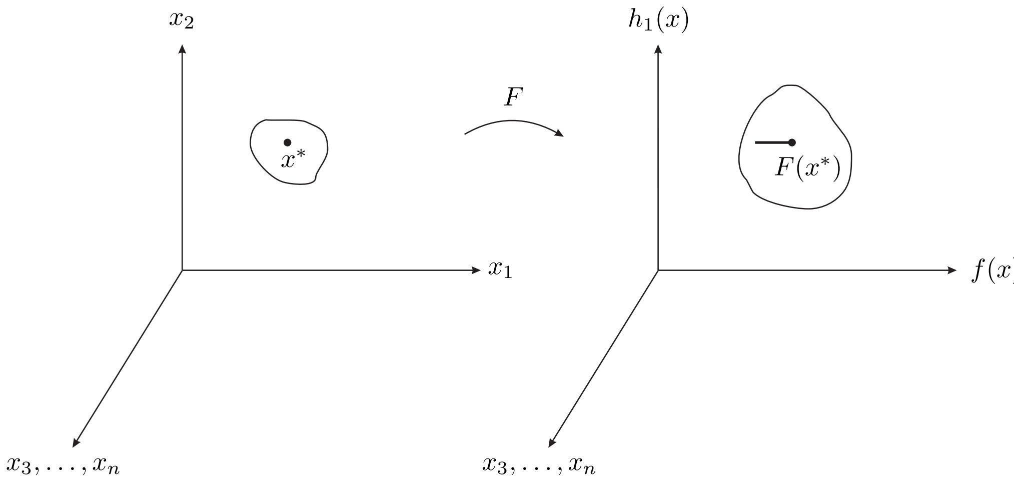

Figure 1.6: Illustrating the alternative proof

The condition (1.6) implies $\exists\, \lambda_1^*, \dots, \lambda_m^*$ such that

$$\boxed{\nabla f(x^*) + \lambda_1^* \nabla h_1(x^*) + \cdots + \lambda_m^* \nabla h_m(x^*) = 0} \tag{1.8}$$

The λ_i^*'s are called *Lagrange multipliers*. This is the *first-order necessary condition for constrained optimality* (generalizes the unconstrained case).

In the unconstrained case, $\nabla f(x^*) = 0$ is n equations in n unknowns. Now, (1.4) and (1.8) is $n+m$ equations in $n+m$ unknowns (x^*, λ^*).

Exercise 2 (due Jan 31) Give an example where a local minimum x^* is not a regular point and the above necessary condition is false (justify both these claims).

Interpreting the above result:

Define the function $l : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}$ by

$$l(x, \lambda) = f(x) + \sum_{i=1}^{m} \lambda_i h_i(x)$$

(*Lagrangian*, or *augmented cost*).

Then at (x^*, λ^*) we have

$$\nabla_x l(x^*, \lambda^*) = \nabla f(x^*) + \sum_{i=1}^{m} \lambda_i \nabla h_i(x) = 0$$

(first-order necessary condition) and

$$\nabla_\lambda l(x^*, \lambda^*) = h(x^*) = 0$$

(constraints). In other words, (x^*, λ^*) is an unconstrained stationary point.

$\longrightarrow$ So, adding Lagrange multipliers basically converts a constrained problem into an unconstrained one.

Lagrange's original reasoning was to replace the problem of minimizing $f(x)$ with respect to x by that of minimizing $l(x, \lambda)$ with respect to (x, λ). If we solve this latter problem, then we must have $h_i(x) = 0 \; \forall \, i$ (because otherwise playing with λ's can decrease l), and subject to these constraints x should minimize f (because otherwise it wouldn't minimize l). But this is not a rigorous argument; for one thing, it doesn't make a proper distinction between necessary and sufficient conditions.

Second-order conditions

1. Necessary condition:

Suppose x^* is a regular point and local minimum (over D), $f \in \mathcal{C}^2$, and λ^* is from the first-order necessary condition (which we assume holds). Then the Hessian matrix

$$\nabla_x^2 l(x^*, \lambda^*) := \nabla^2 f(x^*) + \sum_{i=1}^{m} \lambda_i^* \nabla_x^2 h_i(x^*)$$

is positive semidefinite on the tangent space to D at x^*. (Meaning that $d^T \nabla_x^2 l(x^*, \lambda^*) d \geq 0 \; \forall \, d \in T_{x^*} D$.)

2. Sufficient condition:

First order condition, plus

$$d^T \nabla_x^2 l(x^*, \lambda^*) d > 0 \qquad \forall \, d \;\text{ such that } \nabla h_i(x^*) \cdot d = 0 \quad \forall \, i \tag{1.9}$$

imply that x^* is a (strict) local minimum (over D).

Regularity of x^* is not needed here. However, if x^* is a regular point, then (1.9) is equivalent to saying that $\nabla_x^2 l(x^*, \lambda^*)$ is *positive definite* on $T_{x^*} D$.

$\nabla_x^2 l > 0$ (positive definite on the whole $\mathbb{R}^n$) is certainly sufficient for this.

For proofs, see [Luenberger, Sect. 10.5].

1.3 A preview of the infinite-dimensional case

References: A-F, Sect. 5.5; G-F, Chapter 1. (A-F more or less follows G-F.)

1.3.1 Functional minimization

Instead of minimizing a function $f : \mathbb{R}^n \to \mathbb{R}$, we want to minimize a *functional* $J : V \to \mathbb{R}$ where V is an (infinite-dimensional) vector space.

Our vector space V will usually be a space of functions $f(x)$ or $y(x)$.

("Functional" reflects this difference.)

There does not exist a "universal" function space, like $\mathbb{R}^n$.

Examples:

- $\mathcal{C}^k([a,b],\mathbb{R}^n)$ – the space of k-times continuously differentiable ($k \geq 0$) functions from $[a,b]$ to $\mathbb{R}^n$.
- Bounded functions, piecewise continuous functions, analytic functions, etc.
- other?

Why are these infinite-dimensional?

Polynomials $1, x, x^2, x^3, \ldots$ are linearly independent. (Another example: Fourier harmonics.)

We need to define a *topology* on V, so that we can talk about neighborhoods, convergence, etc. For our purposes, the easiest thing to do is to introduce a *norm*: $\|f\|\ \forall\, f \in V$. This norm must satisfy the usual axioms :

Positive definiteness: $\|f\| \geq 0$, $\|f\| > 0$ if $f \neq 0$;

Homogeneity: $\|\lambda f\| = |\lambda| \cdot \|f\|$ for $\lambda \in \mathbb{R}$, $f \in V$;

Triangle inequality: $\|f + g\| \leq \|f\| + \|g\|$.

Norm gives a *distance* (*metric*): $d(f,g) := \|f - g\|$.

Unlike in $\mathbb{R}^n$, different norms are not equivalent.

Examples of norms:

- Sup-norm:
$$\|f\| := \sup_{[a,b]} |f(x)|$$
Actually, it's a max if f is at least $\mathcal{C}^0$ (assuming that the interval is finite). Sup is needed for bounded not necessarily continuous functions.

- On $\mathcal{C}^1$, can use
$$\|f\|_1 := \sup |f| + \sup |f'|$$
- L_p-norm:
$$\|f\|_{L_p} := \left(\int_a^b |f(x)|^p dx \right)^{1/p}$$

Sup-norm is L_∞-norm, special case of L_p-norm.

The L_p-norms are used to define the spaces
$$\mathcal{L}_p := \{ f : \|f\|_{L_p} < \infty \}, \qquad p = 1, 2, \ldots, \infty.$$

Once we have a metric, we can define local minima:

Definition 2 A function $f^* \in A \subset V$ is a *local minimum* of J over A if $\exists\, \varepsilon > 0$ such that $\forall\, f \in A$ satisfying $\|f - f^*\| < \varepsilon$ we have
$$J(f^*) \leq J(f).$$

Here we switch notation from x^* to f^* to emphasize the infinite-dimensional case, but could also keep x^*. Also, A plays the role of D earlier, J plays the role of f, V plays the role of $\mathbb{R}^n$.

Note: the choice of space V and norm $\|\cdot\|$ is important (cf. strong vs. weak minima—more on this later).

Conditions for minima—?

1.3.2 The first variation and first-order necessary condition

A *linear* functional $\delta J|_{f^*}$ on V is called the *first variation* of J around f^* if for all f and all α (small enough),

$$J(f^* + \alpha f) = J(f^*) + \delta J|_{f^*}(f)\alpha + o(\alpha)$$

Linearity is understood in the usual sense.

This corresponds to the standard notion of (Gateaux) derivative

$$\delta J|_{f^*}(f) = \lim_{\alpha \to 0} \frac{J(f^* + \alpha f) - J(f^*)}{\alpha}.$$

The first variation may not exist (J may not be differentiable).

First order necessary condition for optimality:

$$\boxed{\delta J|_{f^*} = 0} \quad \text{(zero functional)}$$

This works both for minima and maxima. Proof—completely analogous to the finite-dimensional case (leaving this for you to check). But: need to be able to compute δJ for a given functional J—we'll see how later (for now, look at the example in A-F, Sect. 5.5).

1.3.3 The second variation and second-order conditions

Recall [G-F, p. 98; A-F, Sect. 2.11] that a functional $B(f,g)$ on $V \times V$ is called *bilinear* if it's linear separately in each argument (the other one being fixed). $Q(f) := B(f,f)$ then defines a *quadratic functional*, or *quadratic form*, on V (sometimes the latter term is reserved for finite dimensions, as in G-F, but A-F use it in infinite dimensions too).

A quadratic form $\delta^2 J\big|_{f^*}$ on V is called the *second variation* of J around f^* if $\forall f \in V$ and $\forall \alpha$ we have

$$J(f^* + \alpha f) = J(f^*) + \delta J|_{f^*}(f)\alpha + \delta^2 J\big|_{f^*}(f)\alpha^2 + o(\alpha^2)$$

Second-order necessary condition for a local minimum: first-order necessary condition holds, plus

$$\boxed{\delta^2 J\big|_{f^*} \geq 0} \quad \text{(positive semidefinite)}$$

i.e., $\delta^2 J\big|_{f^*}(f) \geq 0 \;\forall f$.

Again, the proof is as before, but the main question is how to find $\delta^2 J$.

What about strengthening this to obtain a sufficient condition?

$\delta^2 J\big|_{f^*} > 0$ is not enough—recall the finite-dimensional proof.

How about:

1) $\delta J|_{f^*} = 0$

2) $\delta^2 J\big|_{f^*}(f) \geq \lambda_{\min}\|f\|^2 \; \forall f \in V$, for some $\lambda_{\min} > 0$.

Condition 2 is not automatic from positive definiteness. It's *strong*, or *uniform*, positive definiteness. To show that the second-order term indeed dominates the higher-order term, we also need to assume that the $o(\alpha^2)$ term decays uniformly with respect to perturbation f (cf. Remark 1).

1.3.4 Global minima and convex problems

What about global minima?

$A \subset V$, want $\min_{f \in A} J(f)$.

Weierstrass' Theorem is still valid, but A must be compact in the sense of 2nd or 3rd definitions given earlier (sequential compactness or finite subcover). Recall the three definitions of compactness (Lecture 2):

1) Closed and bounded

2) Open covers have finite subcovers

3) Every sequence has a subsequence converging to some point in D (sequential compactness)

2 and 3 are also OK in the infinite-dimensional case, but 1 is not. Actually, 2 $\Leftrightarrow$ 3 for metric spaces [Sutherland].

Exercise 3 (due Feb 7) Give an example of $A \subset V$ and J such that J is continuous, A is closed and bounded (justify all these properties), but global $\min_{f \in A} J(f)$ does not exist.

Begin optional:____

Convex problems:

- Convexity of a functional and of a subset of an infinite-dimensional linear vector space is defined as in the finite-dimensional case;
- The result "local $\Rightarrow$ global" is still true;
- In optimal control, under certain (restrictive) conditions we may get a convex infinite-dimensional optimization problem. We will not directly use results from convex functional optimization in this course.
- Concepts from finite-dimensional convex analysis will be useful for MP;
- There is large literature on convex optimization in Hilbert spaces. See Luenberger, "Optimization by vector space methods" (ECE 580 book), Chapter 7 for more on this.

End optional ____

Chapter 2

Calculus of variations

2.1 Examples of variational problems

References: Sussmann's notes, Chap. 1 (see class homepage); G-F, Sect. 1; MacCluer; Sussmann-Willems CSM '97 paper (on class homepage).

2.1.1 Dido's isoperimetric problem

According to the legend about the foundation of Carthage (~850 B.C.), Dido purchased from a local king all the land that could be enclosed by the hide of an ox along the Northern African coastline.

She sliced the hide into very thin strips and was able to tie them together to enclose a sizable area, which became the city of Carthage.

Let's formulate this mathematically. Assume the coast is a straight line (segment).

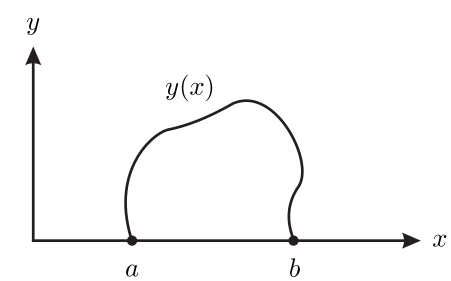

Figure 2.1: Dido's problem

Assume that there is a physical limit on how thin the stripes can be. Then we have a fixed length.

Problem: maximize the enclosed area.

Let $y(x)$ be a solution curve, $y(a) = y(b) = 0$.

Area: $J = \int_a^b y(x)dx \to \max$.

Arclength constraint:

$$\int_a^b \sqrt{1 + (y'(x))^2}dx = C \quad \text{(fixed constant)} \tag{2.1}$$

(This is $\int_0^C ds$ where s is arclength. Use Pythagoras Theorem.)

Can you guess the optimal solution?

Answer: Arc of a circle.

This was known to Zenodorus, ~150 B.C. We'll derive this answer later in the course.

2.1.2 Light reflection and refraction

[Sussmann; also Feynman, I-26-2]

First, in free space: light travels along a path of shortest distance, which is a straight line. This is already a curve minimization principle!

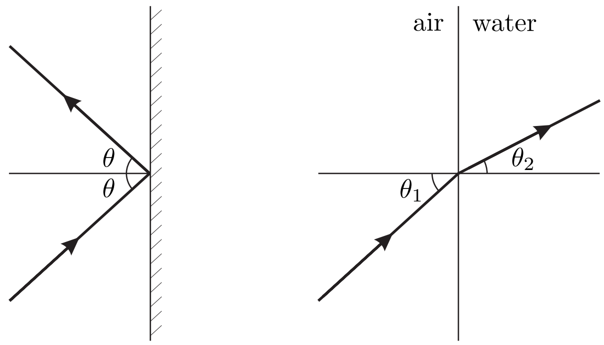

Figure 2.2: Light reflection and refraction

Reflection:

Hero of Alexandria suggested that light still takes the path of shortest distance. Can show with some geometry that then the angles must be the same. Can generalize this to curved reflecting surfaces—not trivial, try it!

Refraction:

This is harder.

Ptolemy (=Ptolemeus?) made a list of pairs (θ_1, θ_2) containing quite a few values (140 A.C.). This is ancient Greek experimental physics!

Snell's Law (1621): $\sin\theta_1 = n\sin\theta_2$, where n is the refraction coefficient ($n \approx 1.33$ for air to water). He found the pattern in Ptolemy's results (but didn't have an explanation).

Fermat's principle (1650): this is not shortest distance any more, but it's the path of *shortest time.*

Snell's law can be derived from Fermat's principle by differential calculus. In fact, this was one of the examples that Leibniz gave to illustrate the power of calculus in his original calculus monograph (1684).

2.1.3 Catenary and brachistochrone

[Sussmann; MacCluer, CVP 3 and 4, solved on p. 49–; catenary also solved in G-F, p. 20]

Catenary:

(Posed and solved incorrectly by Galileo in 1630s, correctly by Johann Bernoulli in 1691.)

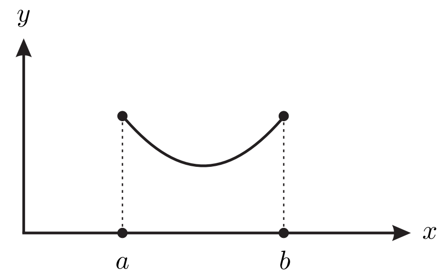

Figure 2.3: Catenary

Find the shape of a chain of given length, with uniform mass density, suspended between two points A and B.

$y(x)$ as before, $y(a) = A$, $y(b) = B$. Can take these to be equal.

Center of mass (y-coordinate):

$$J = \int_a^b y\sqrt{1 + (y'(x))^2}dx \;\rightarrow\; \min$$

(minimizing the y-coordinate of the CM $\equiv$ minimizing potential energy).

Length constraint is again (2.1).

Solution is $y = c\cosh(x/c)$ – "catenaries". Galileo thought that it was a parabola.

Brachistochrone: ($\beta\rho\alpha\chi\iota\sigma\tau o\chi\rho o\nu o s$ means "shortest-time")

Also posed and solved incorrectly by Galileo (1638), solved correctly by Johann Bernoulli, Leibniz, and others during a competition called by Johann Bernoulli in 1696 [Sussmann-Willems].

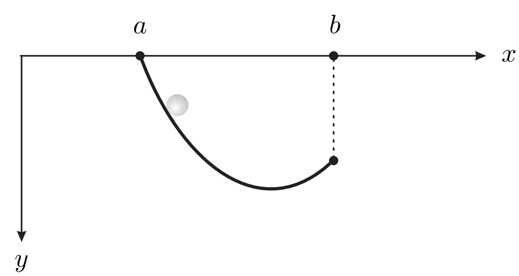

Figure 2.4: The brachistochrone problem

Let a particle (ball) slide without friction along a path from a to b. (The path defines the surface of a "well.") Which path gives the shortest travel time?

$$\text{Time} = \int \frac{\text{Arclength}}{\text{Velocity}}$$

Arclength is given by $\int_a^b \sqrt{1 + (y'(x))^2}dx$ as before.

To calculate velocity, let's use conservation of energy:

$$\frac{mv^2}{2} - mgy = E.$$

Normalizing (choosing suitable units), we can take $m = 1$ and $g = 1/2$. Also, by our choice of the y-axis, $E = E_0 = 0$. Thus, $v = \sqrt{y}$ and we get

$$J = \int_a^b \frac{\sqrt{1 + (y'(x))^2}}{\sqrt{y(x)}} dx.$$

Bicycle wheel interpretation [S-W, Fig. 4]:

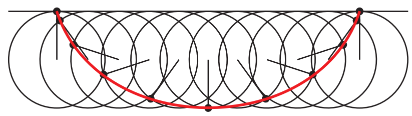

Figure 2.5: A cycloid

Bernoulli's original solution was actually based on the law of light refraction! [Sussmann-Willems]. Treat the particle as light, speed of light is proportional to the square root of height. Discretizing, we partition the plane into strips, so that we get a straight line in each strip and refraction on the boundaries.

Also useful for pilots!

Cannot always be a straight line [Sussmann-Willems]. And Galileo thought it was a circle arc.

2.2 Weak and strong extrema

The above problems are all more or less of this form (*if we ignore the constraints for now*):

Let $L(x, y, z)$ be a $\mathcal{C}^2$ function (this is certainly enough later, can be relaxed). Among all functions $y(x)$, $x \in [a, b]$, $y \in \mathcal{C}^1$, $y(a) = A$, $y(b) = B$ (boundary conditions), find the (local) minima of the cost

$$J(y) = \int_a^b L(x, y(x), y'(x))dx.$$

($L =$"Lagrangian", or running cost.)

Here $y \in \mathbb{R}^1$ (single planar curve) – *single degree of freedom*, will start with this. Or can have $y \in \mathbb{R}^n$ – *multiple degrees of freedom*. Useful for spatial curves or the motion of many particles (the latter was proposed by Lagrange, Mécanique Analytique, 1788 [Sussmann]).

Important: y and y' are position and velocity, but $L(\cdot, \cdot, \cdot)$ is to be viewed as a function of three independent variables.

Recall: to discuss local extrema, we must first fix the norm. Our curves are $\mathcal{C}^1$, so there are two natural candidates for the norm:

$$\|y\|_0 := \max_{a \le x \le b} |y(x)|$$

(can relax the $\mathcal{C}^1$ condition then) and

$$\|y\|_1 := \max_{a \le x \le b} |y(x)| + \max_{a \le x \le b} |y'(x)|$$

Extrema (minima and maxima) with respect to $\| \cdot \|_0$ are called *strong extrema*, and those with respect to $\| \cdot \|_1$ are called *weak extrema*.

<u>***Lecture 5***</u>

A strong minimum is automatically a weak minimum if it is $\mathcal{C}^1$, but not vice versa. Reason: ε-neighborhood with respect to $\|\cdot\|_0$ includes more functions, hence it's harder to satisfy $J(y^*) \leq J(y)$ for all y in the ε-neighborhood.

$\longrightarrow$ Also, we'll see that finding weak minima (or conditions for them) is easier, for the same reason.

Caveat: we'll see that in optimal control, we want strong minima. We also want to relax $y \in \mathcal{C}^1$ to $y \in \mathcal{C}^1$ a.e. (absolutely continuous), which is OK if we use the 0-norm.

Example:

Figure 2.6: Closeness in weak and strong sense

Also imagine y_3 which is like y_2 but slightly out of phase.

Then y_1 is close to y_0 both in strong and weak sense, i.e., with respect to both 0-norm and 1-norm. But y_2 is close to y_0 only with respect to 0-norm, but not 1-norm. Also, y_2 and y_3 are close with respect to 0-norm but not 1-norm.

(Technically speaking, we need to smooth out the corners to take 1-norm.)

If y_2 and y_3 are obtained by control:

$$\frac{dy}{dx} = u \in \{-1, 1\}$$

then the difference is shifting the switching times a little. This is a small perturbation of control, we want to include it—hence we want to work with 0-norm and strong extrema.

Exercise 4 (due Feb 7) Consider the problem of minimizing

$$J = \int_0^1 (y'(x))^2(1 - (y'(x))^2)dx$$

subject to $y(0) = y(1) = 0$. Is the curve $y(x) \equiv 0$ a weak minimum (over $\mathcal{C}^1$ curves)? Is it a strong minimum (over a.e. $\mathcal{C}^1$ curves)? Is there another curve that is a strong minimum?

2.3 First-order necessary conditions for weak extrema

Strong extremum $\Rightarrow$ weak extremum $\Rightarrow$ necessary condition, so it works for both. But the way it's derived, it's designed to test for weak extrema.

2.3.1 The Euler-Lagrange equation

$$J = \int_a^b L(x, y(x), y'(x))dx \;\rightarrow\; \min \tag{2.2}$$

(or write $J(y)$ to be more explicit), $L = L(x, y, z) \in \mathcal{C}^2$, $y : [a, b] \to \mathbb{R}$ (one degree of freedom), $y \in \mathcal{C}^1$,

$$y(a) = A, \;\; y(b) = B \;\; \text{(fixed points)} \tag{2.3}$$

Note: From now on we'll be using the more compact notation L_x, L_y, $L_{y'}$, L_{xx}, etc. for partial derivatives.

Let $y(x)$ be a given test curve.

Consider its perturbation

$$y + \alpha\eta$$

where η is another $\mathcal{C}^1$ curve and $\alpha \in \mathbb{R}$ (small).

Recall the discussion of weak vs. strong minima: this only works for a weak minimum. Both $\alpha\eta$ and $\alpha\eta'$ are small for small α.

$$\|\alpha\eta\|_1 = |\alpha|(\sup|\eta| + \sup|\eta'|)$$

$\longrightarrow$ To get a stronger necessary condition for a strong minimum, we'll need to use a different technique.

To satisfy the boundary conditions for this new curve, we must have

$$\eta(a) = \eta(b) = 0. \tag{2.4}$$

Recall the first-order necessary condition in terms of the first variation: $\delta J|_y = 0$, where

$$J(y + \alpha\eta) = J(y) + \delta J|_y(\eta)\alpha + o(\alpha)$$

We have

$$J(y + \alpha\eta) = \int_a^b L(x, y + \alpha\eta, y' + \alpha\eta')dx$$

Expanding L (inside the integral) and using chain rule:

$$J(y + \alpha\eta) = \int_a^b \big[L(x, y, y') + \big(L_y(x, y, y')\eta(x) + L_z(x, y, y')\eta'(x)\big)\alpha + o(\alpha)\big]dx$$

so the first variation is

$$\delta J|_y(\eta) = \int_a^b \big[L_y(x, y, y')\eta + L_z(x, y, y')\eta'\big]dx$$

This depends not just on η but also on η' – not surprising since L depends on y'. But let's convert to η only, integrating by parts:

$$= \int_a^b \left[L_y\eta - \frac{d}{dx}L_z\eta\right]dx + L_z\eta|_a^b$$

where the last term is 0 due to the boundary conditions (2.4). So, we must have

$$\int_a^b \left[L_y - \frac{d}{dx}L_z\right]\eta dx = 0$$

for all $\mathcal{C}^1$ curves $\eta(x)$ vanishing at $x = a$ and $x = b$.

Lemma 1 *If a continuous function $\xi : [a, b] \to \mathbb{R}$ is such that*

$$\int_a^b \xi(x)\eta(x)dx = 0$$

for all C^1 functions $\eta : [a, b] \to \mathbb{R}$ with $\eta(a) = \eta(b) = 0$, then $\xi(x) \equiv 0$ on $[a, b]$.

This lemma is sometimes attributed to DuBois-Reymond, a German (despite the name!) mathematician from the 19th century. It was known before, but he generalized it; a more proper name is Euler-Lagrange lemma.

Proof. Suppose $\xi(x) \neq 0$ for some x, say, $\xi(x) > 0$. By continuity, $\xi > 0$ on some interval (c, d) containing x.

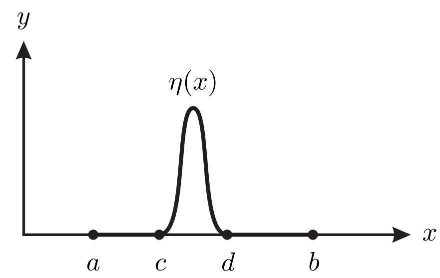

Figure 2.7: The graph of η

For η as shown, $\int \xi\eta dx > 0$—a contradiction. □

We proved that if $y(\cdot)$ is a weak minimum (or maximum! so can just say "extremum") then

$$L_y(x, y(x), y'(x)) = \frac{d}{dx}L_z(x, y(x), y'(x)) \qquad \forall\, x \in [a, b].$$

This is the *Euler-Lagrange equation*.

Sometimes it's written in shorter form:

$$\boxed{L_y = \frac{d}{dx}L_{y'}} \tag{2.5}$$

However, keep in mind that we must take derivatives treating y' as an independent variable.

We'll call trajectories satisfying (2.5) *extremals*.

Example 1 (the answer is obvious but a good illustration of Euler-Lagrange) Shortest path between two points. $L = \sqrt{1 + (y'(x))^2}$, or $L(x, y, z) = \sqrt{1 + z^2}$ (cf. earlier examples).

$$L_y = 0, \qquad L_z = \frac{z}{\sqrt{1 + z^2}}, \qquad L_z(x, y(x), y'(x)) = \frac{y'(x)}{\sqrt{1 + (y'(x))^2}}.$$

No need to compute $\frac{d}{dx}$:

$$\frac{d}{dx}L_z = 0 \;\Leftrightarrow\; L_z = \text{const} \;\Leftrightarrow\; y'(x) = \text{const}$$

so solutions are straight lines. □

Note that two given points uniquely determine the extremal.

2.3.2 Some technical remarks

- Multiple degrees of freedom: $y = (y_1, \dots, y_n)^T \in \mathbb{R}^n$. Same derivation (need to use inner products now) gives the condition

$$L_{y_i} = \frac{d}{dx} L_{y_i'}, \qquad i = 1, \dots, n$$

- Differentiability assumptions:

  Expanded form of Euler-Lagrange :

$$L_y = L_{xy'} + L_{yy'}y' + L_{y'y'}y''$$

  We see that $\frac{d}{dx}L_{y'}$, when expanded, contains the term $L_{y'y'}y''$, which indicates that we must assume $y \in \mathcal{C}^2$. However, a more careful derivation shows that $y \in \mathcal{C}^1$ is enough, and the existence of $L_{y'y'}y''$ can be shown as part of the necessary condition [G-F, pp. 16-17 and the derivation of Euler-Lagrange equation itself]. Some other assumptions, such as existence of L_x (which doesn't appear in the Euler-Lagrange equation), can be relaxed [Sussmann, Handout 2].

- The Euler-Lagrange equation is invariant under arbitrary changes of coordinates—which allows us to pick convenient ones.

Euler-Lagrange is a *necessary* condition. Can give an example where the extremal it provides is not optimal. (This issue will arise later in optimal control.)

$$\min \int_0^1 y(x)(y'(x))^2 dx, \qquad y(0) = 0,\ y(1) = 0$$

Euler-Lagrange :

$$2\frac{d}{dx}(yy') = (y')^2$$

$y \equiv 0$ is a solution—actually, can show that it is a unique solution. But it's obviously neither a minimum nor a maximum.

2.3.3 Some historical remarks

- Leonhard Euler (1707-1783) , Swiss, lived in St. Petersburg for many years.

Derived the equation in 1740 using polygonal approximations.

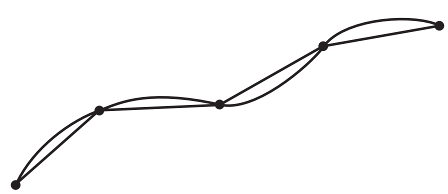

Figure 2.8: Euler's derivation

Finite-dimensional problem: replace a curve by the approximation and optimize the locations of the n points. The "Euler-Lagrange " equation results as $N \to \infty$.

- Joseph-Louis Lagrange (1736-1813), Turin - Italy (France).

  In 1755 (19 years old!) he derived the equation using variational calculus. He wrote a letter to Euler, who was very impressed and called Lagrange's method "calculus of variations."

2.3.4 Special cases

The Euler-Lagrange equation in expanded form (already gave this before):

$$L_{xy'} + L_{yy'}y' + L_{y'y'}y'' = L_y$$

(second-order differential equation for $y(x)$). Boundary conditions: (2.3).

These two conditions should uniquely define $y(x)$ (cf. Example 1).

The second-order equation is hard to solve in general, but in some cases it reduces to a first-order differential equation.

Case 1: "no y". $L = L(x, y')$.

Euler-Lagrange becomes

$$\frac{d}{dx}L_{y'} = 0 \;\Leftrightarrow\; L_{y'} = \text{const}$$

(*integral of motion*). This is first-order. In Example 1, we just had $L = L(y')$ – even simpler.

$L_{y'}$ is called *momentum*.

Case 2: "no x". $L = L(y, y')$.

Euler-Lagrange :

$$0 = \frac{d}{dx}L_{y'} - L_y = L_{yy'}y' + L_{y'y'}y'' - L_y$$

Multiply by y':

$$0 = L_{yy'}(y')^2 + L_{y'y'}y'y'' - L_y y' = \frac{d}{dx}\left(L_{y'}y' - L\right)$$

(easily verified: $L_{y'}y''$ terms cancel). So in this case,

$$L_{y'}y' - L = \text{const}$$

and again we have an integral of motion and Euler-Lagrange reduces to a first-order DE.

$L_{y'}y' - L$ is the *Hamiltonian*—to be discussed more next. (Also appears in the general formula for δJ [G-F]).

Exercise 5 (due Feb 14) Derive the solution of the brachistochrone problem, using the above result. (This is classical and can be found in many places. Go on your own as far as you can. If you look up the final answer, you should still demonstrate that it indeed satisfies the conditions.)

2.3.5 Variable terminal point problems

The above derivation of the Euler-Lagrange equation relied on deriving the first variation of J with respect to a class of perturbations $\eta(x)$ satisfying $\eta(a) = \eta(b) = 0$ (these are the admissible perturbations in the case of both endpoints being fixed).

The class of admissible perturbations depends on the problem. In other problems, other types of perturbations may be allowed.

Example of a variable terminal point problem:

Same cost, $y(a) = A$, but $y(b)$ is arbitrary.

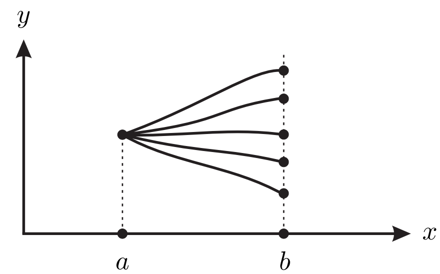

Figure 2.9: Variable terminal point

([G-F] also have the case of both points free on vertical lines—all these cases are similar.)

Perturbed curve is $y + \alpha\eta$ where $\eta(b)$ is no longer 0.

Repeating same arguments:

$$\delta J|_y(\eta) = \int_a^b \left[L_y - \frac{d}{dx} L_z \right] \eta dx + L_z \eta(b) = 0.$$

Perturbations such that $\eta(b) = 0$ are still allowed. Hence, the Euler-Lagrange equation must still hold. Thus the integral is 0, and we have an additional condition:

$$L_z(b, y(b), y'(b))\eta(b) = 0$$

or, since $\eta(b)$ is arbitrary,

$$L_z(b, y(b), y'(b)) = 0.$$

This replaces the condition $\eta(b) = 0$. We still want two boundary conditions to uniquely specify the extremal. Comparing with the previous case, where both endpoints were fixed, here we have one fewer constraints, but on the other hand we have a richer perturbation family which let us obtain one extra condition.

Example 2 For $L = \sqrt{1 + z^2}$, we have

$$L_{y'} = \left. \frac{z}{\sqrt{1 + z^2}} \right|_{z=y'} = \frac{y'}{\sqrt{1 + (y')^2}} = 0$$

and this implies $y'(b) = 0$ (orthogonality). □

Exercise 6 (due Feb 21) Consider a more general version of the variable terminal point problem, with vertical line replaced by a curve:

$$J := \int_a^{x_f} L(x, y(x), y'(x))dx$$

$y(a) = A$ (fixed), x_f unspecified, $y(x_f) = \varphi(x_f)$ where $\varphi : \mathbb{R} \to \mathbb{R}$ is a $\mathcal{C}^1$ curve. Derive a necessary condition for a weak minimum. Your answer should contain, besides the Euler-Lagrange equation, an additional condition ("transversality condition") which accounts for variations in x_f and which explicitly involves φ'.

This is similar to what we'll encounter later in the Maximum Principle.

2.4 Hamiltonian formalism and mechanics

References: G-F, Chap. 4 (skipping Chapters 2 and 3 for now, will go back to them); Sussmann, Handout 3; Arnol'd, MMCM.

2.4.1 Hamilton's canonical equations

When analyzing the Euler-Lagrange equation for some special cases ("no x", "no y"), we came across two quantities conserved in these cases:

("no y") *Momentum*:

$$p := L_{y'}$$

("no x") *Hamiltonian*:

$$H(x, y, y', p) = p \cdot y' - L(x, y, y')$$

(if $y, y' \in \mathbb{R}^n$ then also $p \in \mathbb{R}^n$ and we need inner product).

The variables y and p are called *canonical variables*. Euler-Lagrange implies the following ODEs for them (expressing everything in terms of H):

$$\frac{dy}{dx} = y' = H_p$$
$$\frac{dp}{dx} = p' = \frac{d}{dx}L_{y'} \underset{\text{(E-L)}}{=} L_y = -H_y$$

$$\boxed{y' = H_p, \quad p' = -H_y}$$

These are *Hamilton's canonical equations.*

This reformulation of the Euler-Lagrange equation was proposed by Hamilton, ~1835.

In addition, if we differentiate H with respect to y':

$$H_{y'} = p - L_{y'} = 0$$

by definition of p. This suggests that along the optimal trajectory, H has an extremum as a function of y'. Later, we will see that:

- This is true;
- This extremum is in fact a maximum ;
- This is true even if H is not differentiable with respect to y', or if y' is a boundary point and $H_{y'} \neq 0$ —what matters is the maximum.

This makes more sense when we treat y' as an independent variable (*control Hamiltonian*).

19th century mathematicians (Legendre, Hamilton, Weierstrass) didn't write it this way. Instead, they wrote $H(x, y, p)$ and treated y' as a dependent variable, given by the implicit relation $p = L_{y'}(x, y, y')$. This is why it wasn't until the 1950s that the Maximum Principle was discovered (see the nice discussion in Sussmann's notes).

Mathematically, the Lagrangian L and the Hamiltonian H are related via a transformation which is very classical and useful in many areas (duality in optimization, geometry, etc.), called *Legendre transformation*. This will also clarify the relation between p and y' that we just mentioned.

2.4.2 Legendre transformation

Consider a function $f(\xi)$. (Doing this for scalar case, but easy to extend to $f : \mathbb{R}^n \to \mathbb{R}$ [Boyd], [Arnold, p. 64].)

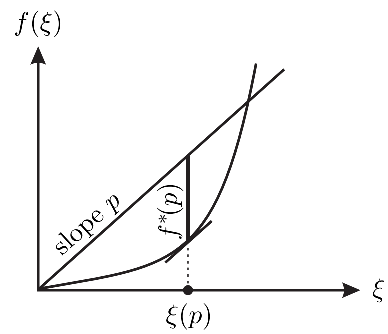

Figure 2.10: Legendre transformation

In the picture f is convex, but it doesn't have to be.

Legendre transform of f will be a new function, f^*, of a new variable, p.

Let p be given. Draw the line through 0 with slope p ($y = p\xi$). Take the point $\xi = \xi(p)$ at which the graph of f is the farthest from the line, measured vertically. That is,

$$\xi(p) = \arg\max_{\xi}\{p\xi - f(\xi)\} \tag{2.6}$$

(If the line is below the graph then the formulas are still the same but the picture is no longer correct.) Then define

$$f^*(p) = p\xi(p) - f(\xi(p)) = \max_{\xi}\{p\xi - f(\xi)\}$$

(This is also called the *conjugate function* of f [Boyd]. $f^*(p)$ is the "maximal gap" between $f(\xi)$ and $p\xi$.)

Can also write this more symmetrically: $f^*(p) + f(\xi) = p\xi$, where p and ξ are related as above.

The maximization condition (2.6) implies that the derivative of $p\xi - f(\xi)$ at $\xi(p)$ with respect to ξ should be 0:

$$p - f'(\xi(p)) = 0$$

(geometrically: tangent slope is the same, p). Note that $\xi(p)$ may not exist, so the domain of f^* is not known a priori.

Legendre transform has nice properties, e.g.:

- f^* is a convex function, even if f is not convex (reason: f^* is a pointwise maximum of functions that are affine in p). Useful in optimization (dual methods [Boyd, p. 221]);
- If f is convex, then Legendre transform is involutive: $f^{**} = f$.

Claim: $H(x, y, p)$, as a function of p (with x, y fixed), is the Legendre transform of $L(x, y, y')$ as a function of $\xi = y'$ (with x, y fixed).

PROOF.

$$L^*(p) = \left(py' - L(x, y, y')\right)\Big|_{p=L_{y'}} = H(p).$$

□

This is to be understood as follows: y' is no longer an argument of H, but it's a dependent variable expressed in terms of x, y, p by the implicit relation

$$p = L_{y'}(x, y, y'). \tag{2.7}$$

Note that the above argument is formal. We're ignoring the issue of whether or not y' is uniquely determined by (2.7). Anyway, from the point of view of the Maximum Principle this is the wrong idea, and we actually won't use it (it's given here just for historical reasons).

2.4.3 Hamilton's principle of least action and conservation laws

Euler-Lagrange :

$$\frac{d}{dx}\left(L_{y'}\right) = L_y$$

Newton's second law:

$$\frac{d}{dt}(m\dot{q}) = -U_q$$

where q are (generalized, if n particles) coordinates, and U is potential, so $m\dot{q}$ is momentum and U_q is force (we're assuming the existence of a potential, from which the force is obtained).

Look similar! Let's make them even more similar by switching notation.

$x \mapsto t$ (wasn't good for brachistochrone, or in calculus of variations in general, but OK in mechanics and control)

$y \mapsto q$

$y' \mapsto q' = \dot{q}$

Then Euler-Lagrange becomes

$$\frac{d}{dt}\left(L_{\dot{q}}\right) = L_q$$

This is exactly Newton's second law if we define the Lagrangian by

$$L := \frac{m\dot{q}^2}{2} - U(q) = T - U$$

where T is kinetic energy and U is potential energy.

(Check: $L_{\dot{q}} = m\dot{q}$, $L_q = -U_q$, assuming U doesn't depend on $\dot{q}$.)

So, the laws of Newtonian mechanics can be recovered from a path optimization condition.

Hamilton's principle of least action: trajectories of mechanical systems are extremals (in many cases, minima) of the functional

$$\int_{t_0}^{t_1} L dt$$

where $L = T - U$. This is called the *action integral*.

This is a necessary condition—but so is Newton's second law.

This suggests an (alternative) view of mechanics in which paths made by objects are the ones along which the action integral is minimized. In other words, these are "straight lines" (geodesics) in the curved space whose metric is determined by forces (mass $\Rightarrow$ gravity). No forces $\Rightarrow$ straight lines. Indeed, the extremals of $\int_{t_0}^{t_1} \frac{m\dot{q}^2}{2} dt$ are straight lines. These are the same extremals as in Example 1, but there the Lagrangian was arclength while here it's the action functional (the two give the same extremals).

What is the meaning of the Hamiltonian?

If $L = T - U = \frac{m\dot{q}^2}{2} - U$, then

$$H = L_{\dot{q}}\dot{q} - L = \frac{m\dot{q}^2}{2} + U = T + U = E$$

which is the total energy!

In case of many particles, T is a more general quadratic form (of velocity variables), but $H = E$ is still true.