Legendre transforms

Mark Alford, Jan 2015

1 Legendre transforms in undergraduate physics

If you know basic thermodynamics or classical mechanics, then you are already familiar with the Legendre transformation, perhaps without realizing it. The Legendre transformation connects two ways of specifying the same physics, via functions of two related ("conjugate") variables.

Context	Relationship	Conjugate variables
Classical particle mechanics	$H(p) = p\dot{x} - L(\dot{x})$ $L(\dot{x}) = p\dot{x} - H(p)$	$p = \partial L/\partial \dot{x}$ $\dot{x} = \partial H/\partial p$
Gibbs free energy	$G(T,\ldots) = TS - U(S,\ldots)$ $U(S,\ldots) = TS - G(T,\ldots)$	$T = \partial U/\partial S$ $S = \partial G/\partial T$
Enthalpy	$H(P,\ldots) = PV + U(V,\ldots)$ $U(V,\ldots) = -PV + H(P,\ldots)$	$P = -\partial U/\partial V$ $V = \partial H/\partial P$
Grand potential	$\Omega(\mu,\ldots) = -\mu N + U(n,\ldots)$ $U(n,\ldots) = \mu N + \Omega(\mu,\ldots)$	$\mu = \partial U/\partial N$ $N = -\partial \Omega/\partial \mu$

Table 1: Examples of the Legendre transform relationship in physics. In classical mechanics, the Lagrangian L and Hamiltonian H are Legendre transforms of each other, depending on conjugate variables $\dot{x}$ (velocity) and p (momentum) respectively. In thermodynamics, the Gibbs free energy G and the internal energy U are Legendre transforms of each other, depending on conjugate variables T (temperature) and S (entropy) respectively. Enthalpy H and internal energy are also Legendre transforms, with conjugate variables P (pressure) and V (volume), and with a plus sign instead of a minus in the relationship between them. The grand potential Ω is another Legendre transform of the internal energy, with conjugate variables μ (chemical potential) and N (number of particles).

Table 1 shows some examples, expressed in the standard way. However, these equations conceal one subtlety. Take the classical mechanics of a single particle, specified by the Legendre transform pair $L(\dot{x})$ and $H(p)$ (we suppress the x dependence of each of them). On the one hand it is crucial to keep $\dot{x}$ and p distinct, since the Lagrangian depends only on $\dot{x}$ and the Hamiltonian depends

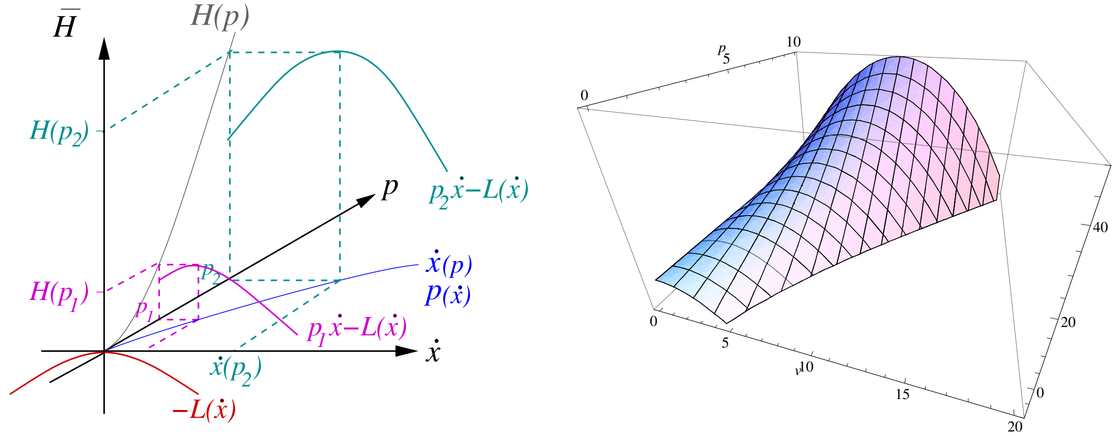

Figure 1: The function $\tilde{H}(p, \dot{x}) = p\dot{x} - L(\dot{x})$ for a free nonrelativistic particle of unit mass. The right panel shows a three-dimensional visualization of the function. The left panel shows how $H(p)$ is the value of $\tilde{H}$ along the "top of the ridge", i.e. for each p it is the value at the stationary point with respect to $\dot{x}$.

only on p:

$$\begin{aligned} \frac{\partial L}{\partial p} &= 0 \quad \Rightarrow \quad \dot{x} - \frac{\partial H}{\partial p} = 0 \ , \\ \frac{\partial H}{\partial \dot{x}} &= 0 \quad \Rightarrow \quad p - \frac{\partial L}{\partial \dot{x}} = 0 \ . \end{aligned} \tag{1}$$

On the other hand, we all know that velocity $\dot{x}$ and momentum p are just functions of each other, so what does it mean for something to be a function of one and not a function of the other?

2 Understanding the Legendre transform

Continuing with classical mechanics as our canonical example, let us clarify the meaning of (1) by constructing two functions that have no subtleties in their definitions:

$$\begin{aligned} &\text{untrained Hamiltonian} \ \ \tilde{H}(p, \dot{x}) \equiv p\dot{x} - L(\dot{x}) \\ &\text{untrained Lagrangian} \ \ \ \ \tilde{L}(\dot{x}, p) \equiv p\dot{x} - H(p) \end{aligned} \tag{2}$$

If we know the Lagrangian $L(\dot{x})$ (for example $L(\dot{x}) = -\frac{1}{2}m\dot{x}^2$ for a free nonrelativistic particle) then we can construct the untrained Hamiltonian $\tilde{H}(p, \dot{x})$, which is a function of $\dot{x}$ and of a completely independent variable p. In Fig. 1 we show how the untrained Hamiltonian depends on momentum and velocity. In the $p = 0$ plane, it is just $-L(\dot{x})$, which has its maximum at $\dot{x} = 0$. This maximum is the stationary point of $\tilde{H}$ with respect to $\dot{x}$ at fixed p. As p rises, the maximum of $\tilde{H}(p, \dot{x})$ with respect to $\dot{x}$ shifts to larger values of $\dot{x}$, and rises upwards. The

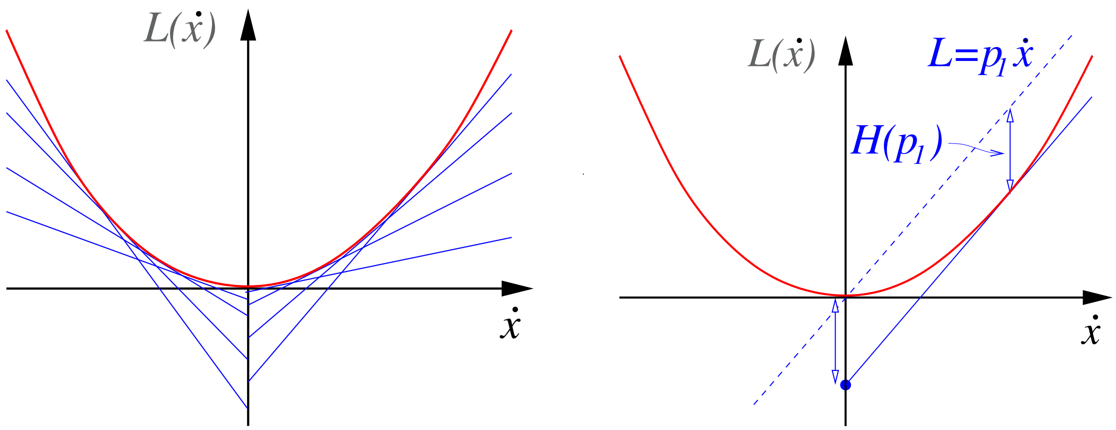

Figure 2: Left panel: the curve $L(\dot{x})$ can be reconstructed just by knowing the y-intercept of each tangent. Tangents are characterized by their slope p. Right panel: if we take a tangent (solid straight line) and shift it (solid dashed line) to pass through the origin, we can see that $H(p_1)$, defined as the maximum of $p\dot{x} - L(\dot{x})$, is also the negative of the y-intercept of the tangent.

Hamiltonian $H(p)$ is defined to be the value of the untrained Hamiltonian along this rising ridge, i.e.

$$H(p) \equiv \max_{\dot{x}} \tilde{H}(p, \dot{x}) = \tilde{H}(p, \dot{x}(p)) \tag{3}$$

where $\dot{x}(p)$ tracks the position of the maximum with respect to $\dot{x}$ of the untrained Hamiltonian $\tilde{H}(p, \dot{x})$. For a convex function $L(\dot{x})$ the maximum is the unique stationary point, so $\dot{x}(p)$ is defined as the solution of

$$\frac{\partial \tilde{H}}{\partial \dot{x}}(p, \dot{x}(p)) = 0 \ . \tag{4}$$

Now we can understand (1) more clearly: The statement that "the Hamiltonian only depends on momentum, not velocity" is really the statement that the Hamiltonian at a given momentum is equal to the untrained Hamiltonian at that momentum and with the velocity set to the value where the untrained Hamiltonian doesn't depend (to first order) on velocity.

The untrained Hamiltonian is a function that can be varied independently with respect to p and $\dot{x}$. The Hamiltonian H is the value of that function along the line in the $\dot{x}$-p plane that tracks the "ridge" (maximum with respect to variation of $\dot{x}$) in the function $\tilde{H}$.

3 Legendre transform and convex functions

The Legendre transform exploits a special feature of a convex (or concave) function $f(x)$: its slope $f'(x)$ is monotonic and hence is a single-valued and invertible function of x. This means that the function can be specified in the conventional way, by giving the value of $f(x)$ for each x, or it can be specified indirectly by giving the y-intercept of each tangent line to the function.

This is illustrated in the left panel of Fig. 2. If we know the vertical positioning (i.e. the y-intercept) of each tangent line to $L(\dot{x})$, then we can draw them all in the correct places, and $L(\dot{x})$ can be reconstructed as the envelope of all the tangents. In the right panel of Fig. 2 we show how this relates to the standard definition of the Legendre transform. If we draw the line of slope p, i.e. $L = p\dot{x}$ (dashed straight line) then (see (3)) the Legendre transform $H(p_1)$ is the maximum difference between this line and the curve $L(\dot{x})$,

$$H(p_1) \equiv \max_{\dot{x}} \Big(p_1 \dot{x} - L(\dot{x}) \Big) \ . \tag{5}$$

By shifting the line so that it is tangent to $L(\dot{x})$, we see that the y-intercept is $-H(p_1)$. So specifying $H(p)$ is equivalent to specifying the y-intercepts of all the tangents to $L(\dot{x})$, which is just an alternative way of specifying $L(\dot{x})$.

4 Multiple Legendre transforms

We can simultaneously Legendre transform with respect to many variables. For a multiparticle classical system, the Lagrangian and Hamiltonian are related by

$$H(p_i, x_i) = \sum_i p_i \dot{x}_i - L(\dot{x}_i, x_i) \ . \tag{6}$$

We can naturally generalize this construction to a classical field theory where the discrete index i becomes a continuous index $\vec{x}$ labelling a different degree of freedom at each point in space, and the Lagrangian and Hamiltonian become functionals:

$$H[\Pi, \varphi] = \int \Pi(x) \dot{\varphi}(x)\, dx - L[\dot{\varphi}, \varphi] \ . \tag{7}$$

Finally, in quantum field theory, we find a Legendre transform relationship between $W[J]$, the generator of connected Greens functions, and $\Gamma[\varphi]$, the generator of one-particle-irreducible Greens functions:

$$\Gamma[\varphi] = -\int J(x)\varphi(x)\, dx + W[J] \ . \tag{8}$$

Note that the signs here are opposite to the conventional Legendre transform, as in the case of enthalpy and internal energy (Table 1). This just means that we will get some minus signs in the resultant relationships between conjugate variables. Another way to say it is that we should have defined Γ differently, as the negative of how it is conventionally defined.