Legendre transforms

Mark Alford, 2019-02-15

1 Introduction to Legendre transforms

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. Table 1 shows some examples of Legendre transformations in basic mechanics and thermodynamics, expressed in the standard way.

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

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 internal energy U can be Legendre transformed into various thermodynamic potentials, with associated conjugate pairs of variables such as temperature-entropy, pressure-volume, and "chemical potential"-density.

The standard way of talking about Legendre transforms can lead to contradictorysounding statements. We can already see this in the simplest example, 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). The standard account of their relationship can be summarized like this:

$$
\begin{array}{rcc}
 & L(\dot{x}) = p\dot{x} - H(p) & H(p) = p\dot{x} - L(\dot{x}) \\
 & L \text{ only depends on } \dot{x}, & H \text{ only depends on } p, \\
 & \text{not on } p, & \text{not on } \dot{x}, \\
 & \text{so} \quad \left(\dfrac{\partial L}{\partial p}\right)_{\dot{x}} = 0 \ , & \text{so} \quad \left(\dfrac{\partial H}{\partial \dot{x}}\right)_{p} = 0 \ , \\
\text{giving the eqn of motion:} & \dot{x} - \dfrac{\partial H}{\partial p} = 0 \ . & p - \dfrac{\partial L}{\partial \dot{x}} = 0 \ . \\
\text{For a free particle,} & H = \dfrac{p^2}{2m} \ , & L = \dfrac{1}{2} m\dot{x}^2 \ , \\
\text{so} & \dot{x} = \dfrac{p}{m} \ . & p = m\dot{x} \ .
\end{array}
\tag{1}
$$

This derivation leads to the right equations of motion, but the logic sounds self-contradictory:

- On the one hand, we are supposed to treat $\dot{x}$ and p as independent variables, since when we take the derivative of the Lagrangian with respect to p we keep $\dot{x}$ constant.
- On the other hand, we used the assumption that $\dot{x}$ and p are independent variables to derive the equation of motion which says that velocity $\dot{x}$ and momentum p are functions of each other.

What is going on? Are velocity and momentum independent variables or not?

2 Understanding the Legendre transform

To properly understand the Legendre transform we need to distinguish the "ordinary" and "extended" versions of the functions we are dealing with. The "extended" version is defined by taking the standard Legendre transform relations, like those listed in Table 1, and allowing all the variables to vary independently with no constraints from the physics.

For the simple example of a free classical particle in one dimension the extended versions of the Hamiltonian and Lagrangian are

$$
\begin{array}{ll}
\text{extended Hamiltonian} & \tilde{H}(p, \dot{x}) \equiv p\dot{x} - L(\dot{x}) \ , \\
\text{extended Lagrangian} & \tilde{L}(\dot{x}, p) \equiv p\dot{x} - H(p) \ .
\end{array}
\tag{2}
$$

The essential distinction is that, while the ordinary Hamiltonian $H(p)$ depends only on the momentum, the extended Hamiltonian $\tilde{H}(p, \dot{x})$ is a function of two

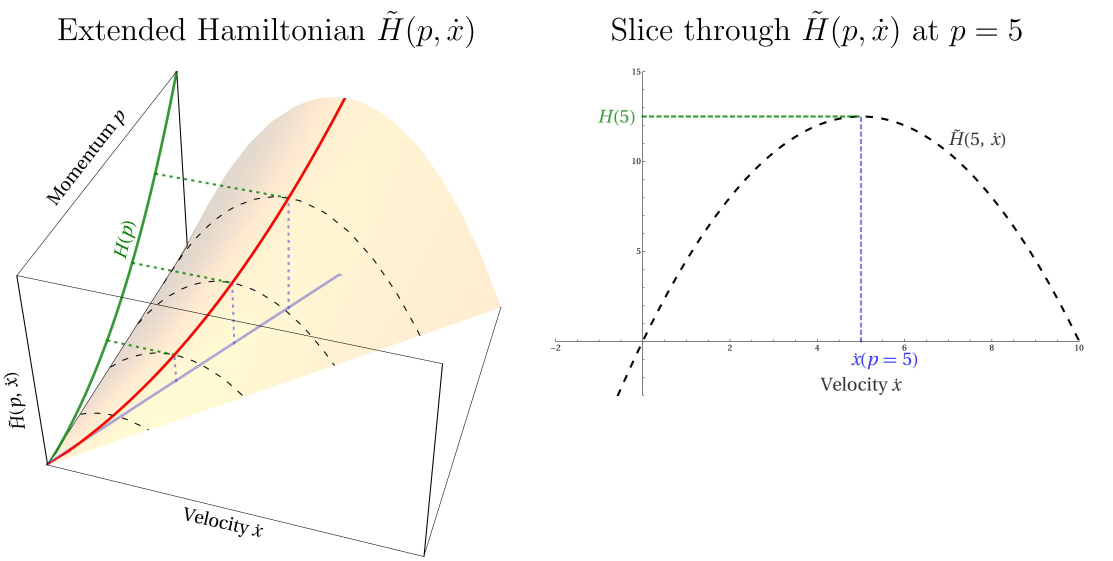

Figure 1: Left panel: 3D plot of the extended Hamiltonian $\tilde{H}(p,\dot{x}) = p\dot{x} - L(\dot{x})$ for a free nonrelativistic particle, $L = \frac{1}{2}m\dot{x}^2$. The ordinary Hamiltonian $H(p)$ (green curve) is the value of the extended Hamiltonian $\tilde{H}(p,\dot{x})$ on the red line that runs along the "top of the ridge". For each momentum p it is the maximum value of $\tilde{H}(p,\dot{x})$ when we vary the velocity $\dot{x}$ (e.g., along one of the dashed black lines), as shown in the right panel which is a slice through the 3D plot at a fixed momentum.

independent variables, the momentum and the velocity. This is illustrated in Fig. 1. The extended Hamiltonian lives on a two-dimensional $(\dot{x},p)$ space of states that are mostly unphysical: the momentum and velocity are not related by the equation of motion.

At each value of the momentum p we can vary the velocity and this traces out the dashed black lines as shown in the left panel of Fig. 1. As we vary $\dot{x}$ at fixed p, the extended Hamiltonian function rises to a maximum and then drops again, as shown in the right panel of Fig. 1. The value of $\dot{x}$ at which $\tilde{H}$ reaches its maximum is the physically valid value of the velocity for the momentum at which we are working. The value of the extended Hamiltonian at that maximum is the value of the ordinary Hamiltonian at that momentum.

As we scan upwards in momentum, the maximum of $\tilde{H}(p,\dot{x})$ with respect to $\dot{x}$ shifts to larger values of $\dot{x}$, and rises to higher values of the Hamiltonian function. The ordinary Hamiltonian $H(p)$ (the green curve in Fig. 1) is just the value of the extended Hamiltonian along this rising ridge of maxima (marked by the red line in Fig. 1),

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

where $\dot{x}(p)$ (the blue curve in Fig. 1) tracks the position of the maximum with respect to $\dot{x}$ of the extended Hamiltonian $\tilde{H}(p,\dot{x})$. For a convex function $L(\dot{x})$ the

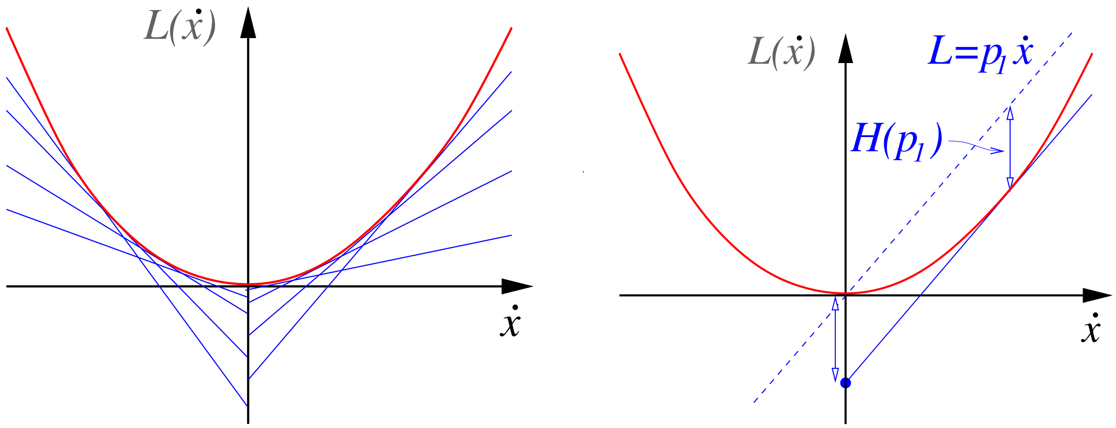

Figure 2: Left panel: the curve $L(\dot{x})$ can be reconstructed just by knowing the y-intercept of each tangent. Tangents are characterized by their slope p. Right panel: if we take a tangent (solid straight line) and shift it (solid dashed line) to pass through the origin, we can see that $H(p_1)$,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

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 a shorthand version of a more complicated statement: the ordinary Hamiltonian at a given momentum is the value of the extended Hamiltonian for that momentum, with the velocity set to the unique value where the extended Hamiltonian doesn't depend (to first order) on velocity.

If we think of nature as "looking for the minimum of the Lagrangian" then we can think of the momentum as the coefficient of a bias field that we subtract from the Lagrangian in order to move that minimum to higher velocities. The result of doing this is the extended Hamiltonian $\tilde{H}$, which depends on momentum and velocity. At a given momentum, nature favors the state with the maximum value of $\tilde{H}$, so the velocity adjusts itself to that value. The equation of motion just tells us which velocity will be favored for each value of the momentum.

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.