Legendre Transforms

Markus Deserno
Department of Physics, Carnegie Mellon University, Pittsburgh, PA 51213, USA
(Dated: December 26, 2012)

A Legendre transform is a procedure for expressing the information content of some function by using a different independent variable, namely, the derivative of this function with respect to (one of) its argument(s). These notes explain how this is done and why simply performing some sort of algebraic substitution instead would destroy information.

I. INTRODUCTION

A. Information content of functions

Functions are *mappings*, often from some set of real (or complex) numbers into another such set. They tell us about some relationship and thus contain *information*. In these notes we will be concerned with the question for how to *represent* that information content, without necessarily being able to quantify it. This sounds very vague, so let us be (marginally) more specific.

How much information is contained in a function? It turns out that this is a nontrivial question with many subtle ramifications. Consider for instance the simple function

$$f : \begin{cases} \mathbb{R} \;\to\; \mathbb{R}_0^+ \\ x \;\mapsto\; x^2 \end{cases} . \tag{1}$$

It takes little to write this down, so the information content appears small. And yet, we could decide to instead "implement" this function by a lookup table. Since `single` precision floating point numbers require 32 bits of information, a computer can store $2^{32} \approx 4.29 \times 10^9$ different single precision floating point numbers. A lookup table would thus contain twice as many numbers, each with 32 bits, requiring 32 GB of data (and note that we can't actually store all squares in single precision…). But then, how much information *is* there in that function?

It turns out that below we will be interested in ways to represent the same information content in different ways, hence we must only answer the question *"do two different functions contain the same information?"* This is a simpler question, because it only requires us to *compare* information content, not actually to *quantify* it. This is quite analogous to the situation that permits us to decide whether two sets have the same size ("magnitude") without actually counting elements: A sack of apples and a sack of oranges contain the same number of elements (and thus have the same magnitude) if we can *pair up* the apples and oranges without any of them remaining unpaired. And as you surely know, by such means Georg Cantor has first arrived at the quite unexpected and nontrivial conclusion that the two sets of natural numbers and rational numbers actually have the same size.

B. Transforms

We are well used to the fact that the information contained in a function can be represented in different ways. Let us make two examples. Take the function $y = f(x)$ and let us assume it's invertible. Then, clearly, the *inverse* function $x = f^{-1}(y)$ contains the same information. A simple way to convince yourself that this is true is to recall that the graph of the inverse function is just the mirror image of the graph of the original function, mirrored at the line $y = x$. And clearly, mirroring preserves all the information.

The second, maybe more interesting example is that we can transform functions into other functions. For instance, we can Fourier transform $f(x)$ into the function $\tilde{f}(k)$. For a suitable set of starting functions, Fourier transforms can be inverted, and hence we can recover $f(x)$ from $\tilde{f}(k)$. This would then let us surmise that the function and its Fourier transform contain the same information. Indeed, in functional analysis we would learn that functions are members of abstract Hilbert spaces and that they can be *represented* in different ways—meaning, using different basis sets—but that it's always the same "function" we're talking about and that a change of basis set doesn't change the information content of the function.

The topic of these notes, *Legendre transforms*, are yet another way to transform one function into another function while preserving information content.

C. Functions and variables

Physicists love streamlined notation. The detailed notation in equation (1) is rarely found in the physics literature. Physicists often don't even write $y = f(x)$, they just talk of the function $y(x)$. And when they refer to the inverse $f^{-1}(y)$ they call that the function $x(y)$. However, what seems awfully sleek in fact has the danger of confusing *three* distinct concepts. What, for instance, do we now mean by "x"? It could be either one of the following three:

1. the *independent variable* x.
2. the *(inverse) function* f^{-1}.
3. the *value* $f^{-1}(y)$ resulting from inserting the independent variable y into the inverse function f^{-1}.

Usually, we don't need to distinguish between these three things very carefully. However, when it comes to Legendre transforms, *the whole point* is to change a function by representing it through a different independent variable, which in turn is defined through (a derivative of) the function itself. If we are now sloppy, we might miss the whole concept and everything is somehow redundant and mysterious. Hence, please make sure that before you read on, you do understand that there *is* a difference between the three concepts above.

Now, the trick is to distinguish between them, *without* making the notation clumsy. In fact, we would like to avoid using the seemingly unnecessary extra symbol "f". I therefore suggest the following: When we speak of independent variables, we use their italic symbols. When we speak of functions, we use roman type, and if we speak of values of functions, we add the independent variable in parentheses. This way, the three concepts from above are distinguished as such: What we formerly all called "x" is now:

1. the *independent variable* "x".
2. the *(inverse) function* "x".
3. the *value* "$\mathrm{x}(y)$" resulting from inserting the independent variable y into the inverse function x.

This means, we now have $y = \mathrm{y}(x)$ and $x = \mathrm{x}(y)$. This might look unusual. If you feel it is unnecessary, then please go back and convince yourself that these three concepts are different things.

II. LEGENDRE TRANSFORM

A. Aim

Take a function $\mathrm{y} : x \mapsto \mathrm{y}(x)$. It contains a lot of information. For instance, it tells us the y-values for many chosen values of x, but it also tells us the *slope* at any given value of x (as long as the function is differentiable). Sometimes that derivative turns out to be of so much interest, that one is tempted to use it as a variable itself. So, if we define

$$p := \mathrm{y}'(x) \ , \tag{2}$$

we might want to find a way to express the information content of the function y using the derivative p as the independent variable. Of course, we would not want to lose information along the way. So how could we do that?

The obvious thing to try seems to be this: Solve Eqn. (2) for x as a function of p, and insert this back into $\mathrm{y}(x)$. This now gives a function of p:

$$\mathrm{y}(x) \to \tilde{\mathrm{y}}(p) = \mathrm{y}(x(p)) = \mathrm{y}\big(\mathrm{y}'^{-1}(p)\big) \ . \tag{3}$$

The crucial question is: Does this new function $\tilde{\mathrm{y}}$ contain the same information as the original function y? The quick answer is: No, it doesn't. And the best way to see this is by a simple example.

Take for instance the function $\mathrm{y} : x \mapsto \frac{1}{2}(x - x_0)^2$. The derivative is $p = \mathrm{y}'(x) = x - x_0$, and this can be uniquely solved for x as a function of p, giving $x = p + x_0$. If we insert this back into our original function $\mathrm{y}(x)$, we get

$$\mathrm{y}\big(x(p)\big) = \frac{1}{2}\big(x(p) - x_0\big)^2 = \frac{1}{2}p^2 \ . \tag{4}$$

Notice that the value of x_0 has dropped out! Functions with different values of x_0 would map to the same final function $\frac{1}{2}p^2$. Hence, we do not know, if someone tells us that they have obtained $\frac{1}{2}p^2$, which function they started with. Information is lost.

There is a subtlety here that is worth understanding fully. You might object that the value of x_0 is contained in the *transformation equation*, and of course I could transform $\frac{1}{2}p^2$ back, together with the right value of x_0, if I use the right transformation equation. I'd just have to memorize which transformation I did, and that might be different for different initial functions (here: different values of x_0). This is true, but this is not the point. The point is that I do not want to memorize the transformed function *together with the transformation equation.* I only want to know *by what general procedure* the transformation was accomplished.

There is a way to solve this problem, but not for all possible functions. It turns out that we will only be able to do things nicely, if our original function is of a special form, and this requires one more interlude:

B. Convex and concave functions

Definition (convexity/concavity): Take a function $\mathrm{y} : x \mapsto \mathrm{y}(x)$ defined over some interval $[a, b]$. It is called "convex" over $[a, b]$ if for every choice of numbers $\{x_1, x_2\} \in [a, b]$ and every $t \in [0, 1]$ we have

$$\mathrm{y}\big(t\,x_1 + (1-t)\,x_2\big) \ \le \ t\,\mathrm{y}(x_1) + (1-t)\,\mathrm{y}(x_2) \ . \tag{5}$$

Geometrically, this means that the graph of the function lies below the line segment joining any two points of the graph. A function y is called concave if its negative $-\mathrm{y}$ is convex.

Several remarks are in order:

1. Convex and concave functions are necessarily continuous. In fact, they have to be differentiable, except maybe at countably many points.
2. If a convex function $\mathrm{y}(x)$ is everywhere differentiable, then it lies above any of its tangents.
3. If a convex function $\mathrm{y}(x)$ is everywhere differentiable *twice*, then $\mathrm{y}''(x) \ge 0$.

C. Definition of Legendre transform

It turns out that a transformation from a function $\mathrm{y}(x)$ to a new function ${}^\star\mathrm{y}(p)$ where (i) $p = \mathrm{y}'(x)$ and (ii) no information is lost is possible if and only if the function y is either convex or concave. The corresponding transformation works as follows:

Definition (Legendre Transformation):

$$ {}^\star\mathrm{y}(p) \;:=\; \begin{cases} \max_x \big\{xp - \mathrm{y}(x)\big\} & \text{if y is convex} \\ \min_x \big\{xp - \mathrm{y}(x)\big\} & \text{if y is concave} \end{cases} . \tag{6} $$

Again, several remarks are in order:

1. Strictly speaking we should take the supremum (sup) instead of the maximum, and the infimum (inf) instead of the minimum.

2. Notice that this is well-defined even if $\mathrm{y}(x)$ contains isolated points at which it is not differentiable.

3. If $\mathrm{y}(x)$ is differentiable, finding the maximum or minimum just requires setting the derivative of the expression in curly brackets with respect to x to zero, leading to

   $$ 0 = \frac{\partial}{\partial x}\big\{xp - \mathrm{y}(x)\big\} = p - \mathrm{y}'(x) \ , \tag{7} $$

   and hence we find that the new variable is $p = \mathrm{y}'(x)$, as we aimed to. This is then solved for x, and inserting the (p-dependent) result will lead to the value of the maximum or minimum we look for.

4. If a function deviates from being convex or concave only in localized regions, the definition of the Legendre transform still makes sense.

5. There are different sign conventions in use. Some people swap the terms between the curly brackets together with the min/max out front:

   $$ \mathrm{y}^\star(p) \;:=\; \begin{cases} \min_x \big\{\mathrm{y}(x) - xp\big\} & \text{if y is convex} \\ \max_x \big\{\mathrm{y}(x) - xp\big\} & \text{if y is concave} \end{cases} . \tag{8} $$

   One may check that ${}^\star\mathrm{y}(p) = -\mathrm{y}^\star(p)$. The convention followed with ${}^\star\mathrm{y}$ has the advantage that the Legendre transform of a convex function is convex and that of a concave function is concave. For the $\mathrm{y}^\star$-convention convexity/concavity switches upon transformation—but it's the latter which is common in thermodynamics. Unfortunately, it will also turn out that what precisely is the inverse Legendre transform depends on the convention. We'll get to that soon.

D. Properties of the Legendre Transform

1. *Geometric interpretation*

The tangent $t_{x_0}(x)$ of some function $f(x)$ constructed at some point $x = x_0$ has the equation

$$ t_{x_0}(x) = f(x_0) + f'(x_0)\,(x - x_0) \ . \tag{9} $$

Its value at $x = 0$ thus satisfies

$$ t_{x_0}(0) = f(x_0) - x_0 f'(x_0) \ . \tag{10} $$

The right hand side can be viewed as the Legendre transform of $f(x)$, in the following sense: If we label the tangent by its slope $p = f'(x_0)$ instead of its contact x-value x_0, and if we solve this for $x_0 = x_0(p)$, we get

$$ f^\star(p) = f(x_0(p)) - x_0(p)\,p \ , \tag{11} $$

which is exactly of Legendre-transform type. Hence, the Legendre transform of a convex function can be viewed as the ordinate value of the tangent to $f(x)$ of slope p: $f^\star(p) = t_{x_0(p)}(0)$. Notice that if the function $f(x)$ is differentiable and convex, there exists at most one tangent to every value of the slope p, and hence exactly one corresponding ordinate value. Such a function can be described by the envelope of all its tangents, and since each such tangent has a unique ordinate value, we begin to see why Legendre transforms conserve information.

2. *Convexity/Concavity*

Let us take the definition of a Legendre transformation as applied to convex functions:

$$ {}^\star\mathrm{y}(p) \;:=\; \max_x \big\{xp - \mathrm{y}(x)\big\} \ . \tag{12} $$

Observe, again, that this often makes sense even if the function y is not convex to begin with (it might be "almost" convex, but somewhere have a small concave bump). However, the resulting function ${}^\star\mathrm{y}$ is *always* convex. This follows simply checking the convexity definition: Let p_1 and p_2 be two points in the domain where ${}^\star\mathrm{y}$ is defined and let t be within $[0,1]$. Then,

$$ \begin{aligned} {}^\star\mathrm{y}(tp_1 + (1-t)p_2) \;&=\; \max_x \big\{x[tp_1 + (1-t)p_2] - \mathrm{y}(x)\big\} \\ &=\; \max_x \Big\{t\big[xp_1 - \mathrm{y}(x)\big] + (1-t)\big[xp_2 - \mathrm{y}(x)\big]\Big\} \\ &\overset{*}{\leq}\; t\,\max_x \big\{xp_1 - \mathrm{y}(x)\big\} + (1-t)\,\max_x \big\{xp_2 - \mathrm{y}(x)\big\} \\ &=\; t\,{}^\star\mathrm{y}(p_1) + (1-t)\,{}^\star\mathrm{y}(p_2) \ , \end{aligned} \tag{13} $$

where at "$*$" we used the inequality $\max_x\{f(x)+g(x)\} \leq \max_x\{f(x)\} + \max_x\{g(x)\}$. □

For the same reason, a Legendre transform using the "min" procedure in Eqn. (6) leads to a concave function, even if the function to start with is not concave.

*y-convention		$\mathrm{y}^\star$-convention	
convex	concave	convex	concave
${}^\star\mathrm{y}(p) = \max_x\{xp - \mathrm{y}(x)\}$	${}^\star\mathrm{y}(p) = \min_x\{xp - \mathrm{y}(x)\}$	$\mathrm{y}^\star(p) = \min_x\{\mathrm{y}(x) - xp\}$	$\mathrm{y}^\star(p) = \max_x\{\mathrm{y}(x) - xp\}$
$\mathrm{y}(x) = \max_x\{xp - {}^\star\mathrm{y}(p)\}$	$\mathrm{y}(x) = \min_x\{xp - {}^\star\mathrm{y}(p)\}$	$\mathrm{y}(x) = \max_x\{\mathrm{y}^\star(x) + xp\}$	$\mathrm{y}(x) = \min_x\{\mathrm{y}^\star(x) + xp\}$
${}^\star\mathrm{y}' = \mathrm{y}'^{-1}$	${}^\star\mathrm{y}' = \mathrm{y}'^{-1}$	$\mathrm{y}^{\star\prime} = -\mathrm{y}'^{-1}$	$\mathrm{y}^{\star\prime} = -\mathrm{y}'^{-1}$
$\mathrm{dy} = p\,\mathrm{d}x$	$\mathrm{dy} = p\,\mathrm{d}x$	$\mathrm{dy} = p\,\mathrm{d}x$	$\mathrm{dy} = p\,\mathrm{d}x$
$\mathrm{d}^\star\mathrm{y} = x\,\mathrm{d}p$	$\mathrm{d}^\star\mathrm{y} = x\,\mathrm{d}p$	$\mathrm{dy}^\star = -x\,\mathrm{d}p$	$\mathrm{dy}^\star = -x\,\mathrm{d}p$

TABLE I: Legendre transform pairs and useful relations in the *y- and $\mathrm{y}^\star$-convention, both for convex and concave functions. Notice that for the *y-convention the Legendre transform is its own inverse, while in the $\mathrm{y}^\star$-convention there is an additional minus-to-plus switch in the xp-term. Also notice that in the *y-convention convexity and concavity does not change upon transformation, and hence the "min" and "max" is the same for both directions. In the $\mathrm{y}^\star$-convention convexity and concavity swap, and so do the "min" and "max". Notice that when one actually *computes* the Legendre transforms, one anyways ends up searching for the derivative in the min- or max-terms, and hence it doesn't greatly matter whether it actually is a min- or a max-procedure that is to be implemented. Notice also that the differential relations are rather easy to memorize, and indeed they are frequently used in thermodynamics.

If we use the alternative definition of the Legendre transform from Eqn. (8) then—since convexity and concavity swap upon transformation—the Legendre transform of an "almost" convex function is concave and that of an almost concave function is convex.

3. Inverse of derivatives

The new independent variable of a Legendre transform is the derivative of the original function:

$$\frac{\partial \mathrm{y}(x)}{\partial x} = p(x) \ . \qquad (14a)$$

What is the derivative of the Legendre transform $\mathrm{y}^\star$ with respect to its independent variable? Assuming differentiability and convexity of y, we know that $\mathrm{y}^\star(p) = \mathrm{y}(x(p)) - x(p)p$, where $x(p) = \mathrm{y}'^{-1}(p)$. Hence,

$$\frac{\partial \mathrm{y}^\star(p)}{\partial p} = \underbrace{\frac{\partial \mathrm{y}(x)}{\partial x}}_{p}\frac{\partial x}{\partial p} - x(p) - \frac{\partial x}{\partial p}p = -x(p) \ . \qquad (14b)$$

Since evidently $p(x)$ and $x(p)$ are inverses of each other, Eqns. (14a) and (14b) show that – up to a minus sign – y' and $\mathrm{y}^{\star\prime}$ are *also* inverse functions of each other:

$$\mathrm{y}^{\star\prime} = -\mathrm{y}'^{-1} \ . \qquad (15)$$

Using a slightly more sloppy notation (which, however, is popular in thermodynamics), we can state this as follows: If we have a function $\mathrm{y}(x)$ and its Legendre transform $\mathrm{q}(p)$ (using thermodynamics convention (8), then we have the following two matching pair of differentials:

$$\mathrm{dy} = p\,\mathrm{d}x \quad \longleftrightarrow \quad \mathrm{dq} = -x\,\mathrm{d}p \ . \qquad (16)$$

4. Inverse Legendre Transformation

For convex or concave functions y we have

$${}^{\star\star}\mathrm{y} = \mathrm{y} \ . \qquad (17)$$

This says that for convex or concave functions *the Legendre transform is its own inverse.* This, unfortunately, is only true for the *y convention. The inverse for the $\mathrm{y}^\star$ convention needs to be calculated differently.

Proof: We will be a bit lazy and only prove this for differentiable functions. Assume without loss of generality that y is convex. In that case we first have

$${}^\star\mathrm{y}(p) = \max_x \big\{xp - \mathrm{y}(x)\big\} = x(p)p - \mathrm{y}(x(p)) \ , \qquad (18)$$

where $p(x) = \mathrm{y}'(x)$ and hence $x(p) = \mathrm{y}'^{-1}(p)$. Notice that this is the point where we need convexity: If $\mathrm{y}(x)$ is convex, then $\mathrm{y}'(x)$ is monotonic and hence *we can uniquely solve for* $x(p)$.

Now, Legendre transforming one more time (and remembering that *y is also convex if y is convex) we get

$${}^{\star\star}\mathrm{y}(q) = \max_p \big\{pq - \mathrm{y}^\star(p)\big\} = p(q)q - \mathrm{y}^\star(p(q)) \ . \qquad (19)$$

Inserting $\mathrm{y}^\star(p)$, we get

$$\begin{aligned} \mathrm{y}^{\star\star}(q) &= p(q)q - \mathrm{y}^\star(p(q)) \\ &= p(q)q - \big[\underbrace{x(p(q))}_{q}p(q) - \mathrm{y}(\underbrace{x(p(q))}_{q})\big] \\ &= \mathrm{y}(q) \ , \end{aligned}$$

which is what we wanted to prove. □

As it turns out, in the $\mathrm{y}^\star$-convention it is also necessary to also swap the sign of the "xp"-term. This can be easily checked by the same type of calculation with which we proved what the inverse transform was for the *y-case.

Table I summarizes the Legendre transform rules in both conventions, for convex and concave functions.

5. Preservation of information

The obvious corollary of the previous section is that Legendre transformations preserve information: If it is possible to recreate y from y*, then no information can have gotten lost. Observe, though, that this does only hold if the function we started with was convex or concave.

Since the Legendre transform still makes sense if we have *local* deviations from convexity or concavity (say, an overall convex function with a local concave "bump"), we might ask what now happens after two Legendre transforms. The answer is that we recover the *convex (concave) envelope* of the original function. This finding plays an important role in the theory of phase transitions.

III. APPLICATION TO THERMODYNAMICS

A. Energy and free energy

In thermodynamics we find that the entropy S as a function of energy E, volume V, and particle number N is a *thermodynamic potential*, meaning that it contains all the thermodynamic information we can hope for. We also learn that certain derivatives of the entropy are of thermodynamic interest, for instance we know that

$$\frac{1}{T} = \left(\frac{\partial S}{\partial E}\right)_{V,N} \ , \tag{20}$$

where T is the temperature. It is then natural to ask whether we can construct other thermodynamic potentials that contain such derivatives as their natural independent variables, and whether we can find them starting from the entropy. Since we want to replace some variables by others which are derivatives of the function we start with, and since we certainly don't want to lose precious thermodynamic information along the way, a Legendre transform appears to be the winning ticket.

To look at the best know example, let us first of all suppress the volume and particle number dependence and only look at the energy. Next, instead of looking at $S = \mathrm{S}(E)$, let us look at the inverse thermodynamic potential $E = \mathrm{E}(S)$. (Warning: Some people use a roman "E" to denote the "exergy", a different thermodynamic potential. We don't.) Since the entropy is monotonic over the ranges over which (canonical) thermal equilibrium can be achieved, this inverse actually exists and we hence do not lose information. Moreover, we of course then also have

$$T = \left(\frac{\partial E}{\partial S}\right)_{(V,N)} \ . \tag{21}$$

For the reasons outlined above, simply solving this equation for S as a function of T and the inserting this into $\mathrm{E}(S)$ will not solve the problem. We do get a correct equation, namely the (caloric) equation of state $E = E(T)$, but we lost information along the way, and hence $E(T)$ is not a thermodynamic potential anymore.

What we need to do instead is to calculate the Legendre transform of $\mathrm{E}(S)$, where $T = \partial E/\partial S$ will be the new variable. Since $\mathrm{S}(E)$ is concave, $\mathrm{E}(S)$ is convex, and using the sign-reversed definition of the Legendre transform, we get

$$\mathrm{E}^\star(T) = \min_S \left\{ \mathrm{E}(S) - TS \right\} \ . \tag{22a}$$

Notice that in order to actually find the minimum, for every given value of T, we have to search through all pairs $\{S, \mathrm{E}(S)\}$. But obviously we could also look at all pairs $\{\mathrm{S}(E), E\}$, thus using the entropy $\mathrm{S}(E)$ and not its inverse $\mathrm{E}(S)$ to evaluate the expression in curly brackets. (If it helps, imagine performing this in a computer program, and changing the loop index needed to scan through all possible values of the expression in curly brackets from "S" to "E".) But this means that we could just as well write

$$\mathrm{E}^\star(T) = \min_E \left\{ E - T\mathrm{S}(E) \right\} \ . \tag{22b}$$

The difference between Eqn. (22a) and Eqn. (22b) is subtle, but well visible in our notation: In both formulas we have essentially "$E - TS$", but in the first we view the energy as a function of entropy and minimize over all values of the entropy, in the second we view the entropy as a function of energy and hence minimize over all values of the energy.

Of course, in thermodynamics we don't write the Legendre transform as $\mathrm{E}^\star(T)$ but instead give it a new symbol. Usually it's F, but sometimes one also finds A. So people write $\mathrm{F}(T)$ or simply $F(T)$ or simply F. And they call it the (Helmholtz) *free energy*. From what we have learned about Legendre transforms, the free energy contains the same thermodynamic information as the entropy. Our discussion above also shows that it is a concave function of temperature.

Other Legendre transforms are evidently possible, for instance replacing the volume by the pressure, or the particle number by the chemical potential, thus leading to all kinds of other equivalent thermodynamic potentials with all kinds of names, such as "enthalpy" or "grand potential".

B. Relation to Laplace transforms and partition functions

The entropy $S(E)$ is the logarithm of a function $\Omega(E)$, which is essentially the density of states and a multiplicity $N!$ divided out:

$$S(E) = k_\mathrm{B} \ln \Omega(E) \ . \tag{23}$$

The canonical partition function is the Laplace transform of $\Omega(E)$, and the free energy essentially the logarithm of the canonical partition function:

$$e^{-\beta F(T)} = Z(T) = \int dE\, \Omega(E)\, e^{-\beta E} \ . \tag{24}$$

It is easy to see that the Laplace transform relation between partition functions translates to a Legendre transform relation between the thermodynamic potentials *in the thermodynamic limit*. To see this, we need to Laplace-evaluate the Laplace transform.

First, we use (23) to rewrite (24):

$$e^{-\beta F(T)} = \int dE\, e^{-\beta[E - TS(E)]} \ . \tag{25}$$

We next need to make extensivity explicit. We will write $E = N\epsilon$ and $S = N s_N(\epsilon)$.

$$e^{-\beta F(T)} = \int d\epsilon\, N\, e^{-\beta N[\epsilon - T s_N(\epsilon)]} \ . \tag{26}$$

Notice that $s_N(\epsilon)$ still depends on N. All we know is that if the thermodynamic limit exists, it will converge against and N-independent value $s_\infty(\epsilon)$. We can thus write $s_N(\epsilon) = s_\infty(\epsilon) + \delta s_N(\epsilon)$, where the latter is a function that decays to zero in the thermodynamic limit. And thus we are in the position to perform a saddle-point (or "Laplace"-) evaluation of the integral:

$$\begin{aligned} e^{-\beta N f_N(T)} &= \int d\epsilon\, N\, e^{-\beta N[\epsilon - T s_\infty(\epsilon) - T\delta s_N(\epsilon)]} \\ &\overset{N\to\infty}{\sim} N\, e^{-\beta N \min_\epsilon[\epsilon - T s_\infty(\epsilon)]} \end{aligned} \tag{27}$$

Taking the logarithm and dividing by $-\beta N$, we get

$$f(T) = \min_\epsilon \left\{\epsilon - T s(\epsilon)\right\} \ , \tag{28}$$

which expresses the now well-known Legendre transform between the specific energies, entropies and free energies (for which the limit $N \to \infty$ exists).