Computing Fourier Transforms Numerically

1. Real Even Functions

Superficially it seems that the numerical calculation of a Fourier transform ought to be quite easy—you just take the FFT of the function, sampled evenly in t. Rather than treating the general case, let us discuss the simplest situation first. Suppose we want the FT of a real, even function $g(t)$, whose transform is also real and even. Then

$$\hat{g}(f) = \int_{-\infty}^{\infty} g(t)\, \mathrm{e}^{-2\pi i f t}\, dt \tag{1}$$

$$= 2\,\mathrm{Re} \int_{0}^{\infty} g(t)\, \mathrm{e}^{-2\pi i f t}\, dt\,. \tag{2}$$

This seems perfectly adapted for approximation by the FFT:

$$\hat{g}_m = \sum_{n=0}^{N-1} g_n\, \mathrm{e}^{-2\pi i m n/N}, \quad m = 0, 1, 2, \cdots N-1 \tag{3}$$

where we can replace the integral by the sum. There are some important details to get right, if this is to work out. Notice that here the index on a subscripted variable starts at zero, while in MATLAB it begins at unity; so be careful in computer code.

First, how rapidly should the function g be sampled in time? We will say

$$g_n = g(n\,\Delta t)\,. \tag{4}$$

Roughly we want the highest frequencies in g to be represented in the sampling. For simple smooth functions, like the Planck formula (Figure 1)

$$E(k) = \frac{\beta k^3}{\exp(\alpha k) - 1} \tag{5}$$

Figure 1.

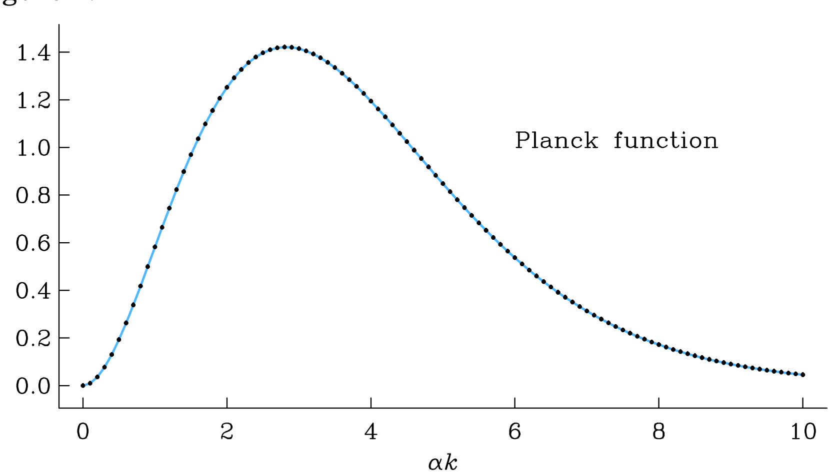

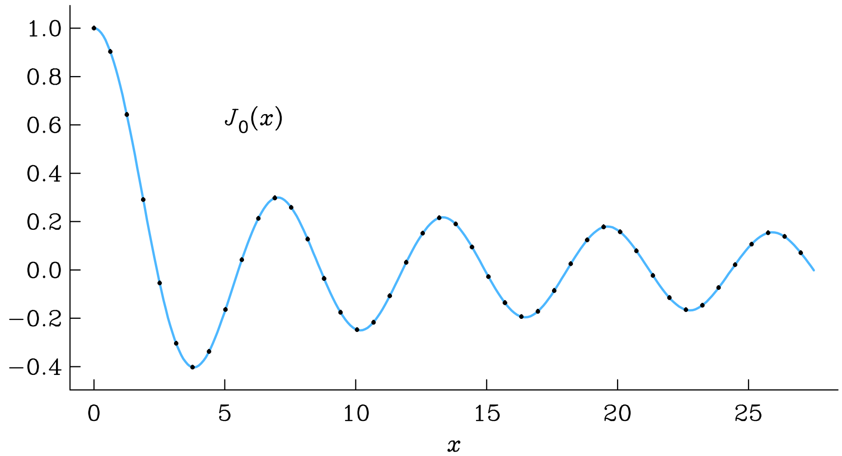

Figure 2.

we can see the graph of the function that there is probably no need to sample more densely than $\alpha k = 0.1$ or $\Delta k = 0.1/\alpha$.

In another common situation the function oscillates as it decays, like the Bessel function $J_0(x)$ we should sample with at least 10 samples in the dominant period, as shown in Figure 2. We conclude that $\Delta x = 0.63$. So if possible, plot the function you want to transform to get an idea of the frequency content.

The next thing we should estimate is N the number of samples of $g(t)$. Most functions we need to transform die away with large t. If the numerical FT is to be a reasonable approximation to $\hat{g}$ the function should be quite small at the end of the data series. Obviously if the function is still large at $t_{\max} = N\,\Delta t$ the answer will be inaccurate: if you want your result to be accurate to a part in 10^4, it is plausible to insist that $|\,g(N\,\Delta t)\,| < \max\,|\,g\,| \times 10^{-4}$. If all you want to do is plot $\hat{g}$ a part in 10^3 will certainly suffice.

Let us try this on the Planck formula. When we look only at the exponential we have $\exp(\alpha N\,\Delta k) = 10^3$, or $\alpha N\,\Delta k = \ln 1000 = 6.9$. But this obviously too small, just by looking at Figure 1. We must take the numerator k^3 into account. Try increasing by a factor of 2: with 13.8 we find $E = 2.7 \times 10^{-3}$, so let us say $k_{\max} = \alpha N\,\Delta k = 15$ and we conclude that $N = 15/(\alpha \Delta k) = 150$ a very modest number.

However when we try this on the Bessel function we find a different story. The magnitude of J_0 decays only as $\sqrt{\pi/2x}$ and so even for 1 part in 1000 $x_{\max} = N\,\Delta x = 6.4 \times 10^5$; we see that $N = 10^6$. This is certainly possible, but if we need a part in a million accuracy, things get out of hand since then $N = 10^{12}$. We will return to possible solutions to this problem later.

Let us substitute into (3) to see what we have so far: comparing (3) with (2) we see we need to write things in terms of $t = n\,\Delta t$.

$$\hat{g}_m = \sum_{n=0}^{N-1} g(n\Delta t) \exp \frac{-2\pi i m(n\Delta t)}{N\Delta t} . \tag{6}$$

Then if we say the frequency $f = m\,\Delta f$, we conclude from the exponential function that $\Delta f = 1/N\Delta t$. Hence the FFT can be written

$$\hat{g}_m = \sum_{n=0}^{N-1} g(n\Delta t)\ e^{-2\pi i(mDf)(n\Delta t)} . \tag{7}$$

The terms in the sum now correspond exactly to the factors in the integrand in (2). Next we need to convert the sum to an approximation of the integral. Recall the Trapezoid Rule

$$\int_a^b g(t)\ dt = ½\Delta t[g(a) + g(b)] + \Delta t \sum_{n=1}^{N-1} g(a + n\,\Delta t) + R . \tag{8}$$

This differs from (7) in two respects: first we must weight the sum in (7) by Δt; second, the first and last terms in the sum in (7) must be halved. If N is so large that $g(N\,\Delta t)$ may be neglected, we can ignore the discrepancy in the last term. So we conclude that the approximate Fourier integral for (1) is given by

$$\hat{g}(m\,\Delta f) = 2\,\Delta t\,\mathrm{Re}(\hat{g}_m - ½g(0)) \tag{9}$$

where $\Delta f = 1/(N\,\Delta t)$. Notice that the frequencies at which the transform is evaluated are not under our control if we use the FFT: they are dictated by the rigid form of the FFT and the fact that we must choose Δt and N to satisfy essential conditions for accuracy. Of course, one can always make N larger and Δt smaller provided that $N\,\Delta t$ exceeds the essential lower bound, but often Δf is found to be much too small, and then the FFT algorithm itself should be replaced with digital Fourier transform whose output frequencies are more flexible.

One last thing before the examples: the FFT of a real vector exhibits a symmetry about the middle frequency:

$$\hat{g}_m = (\hat{g}_{N-m})^* . \tag{10}$$

That symmetry does not hold for the Fourier integral, and so (9) cannot be expected to be even approximately correct if $f > ½N\,\Delta f$. By the halfway point in frequency, the sampled sines and cosines are being aliased into low-frequency versions, and thus the sum ceases to resemble the integral. So, asserting the symmetry of $\hat{g}(f)$ for even functions $g(t)$, we should write (9) as

$$\hat{g}(\pm m\,\Delta f) = 2\,\Delta t\,\mathrm{Re}(\hat{g}_m - ½g(0)),\ \ 0 \le m < ½N . \tag{11}$$

2. Examples

We start with the Planck formula. It's convenient to set $\alpha = \beta = 1$ in (5). We can recover the true units by scaling lengths by α and wavenumbers by $1/\alpha$; for the homework problem $\alpha = 2.62 \times 10^{-6}$ meters. We decided on the basis of Figure 1 to set $\Delta k = 0.1$ and $N = 150$. Then we generate the function E and take its FFT and apply (11). The result is a bit disappointing: see Figure 3a. The output vector is mostly close to zero. The interesting stuff is happening near $x = 0$. Since it is absolutely certain that the true $\hat{E}$ does not have sharp corners (why?) we must conclude that $\Delta x = 1/N\Delta k$, the sampling rate in the output function, is really about ten times too large! We can blow up the low m portion, so from $m = 0$ to 10. In *plotxy* we can smooth the result with the `smooth` command, and that's Figure 3b. But that's really not correct, even if it is smooth. The best way is to decrease Δx, by increasing N, but leaving Δk alone. In Figure 3c we see the correct answer. But observe how, in order to display the interesting portion in enough detail, we had to take the FFT on a series 1500 terms long, then throw away 650 of the 750 output values, because they were essentially all zeros.

Figure 3

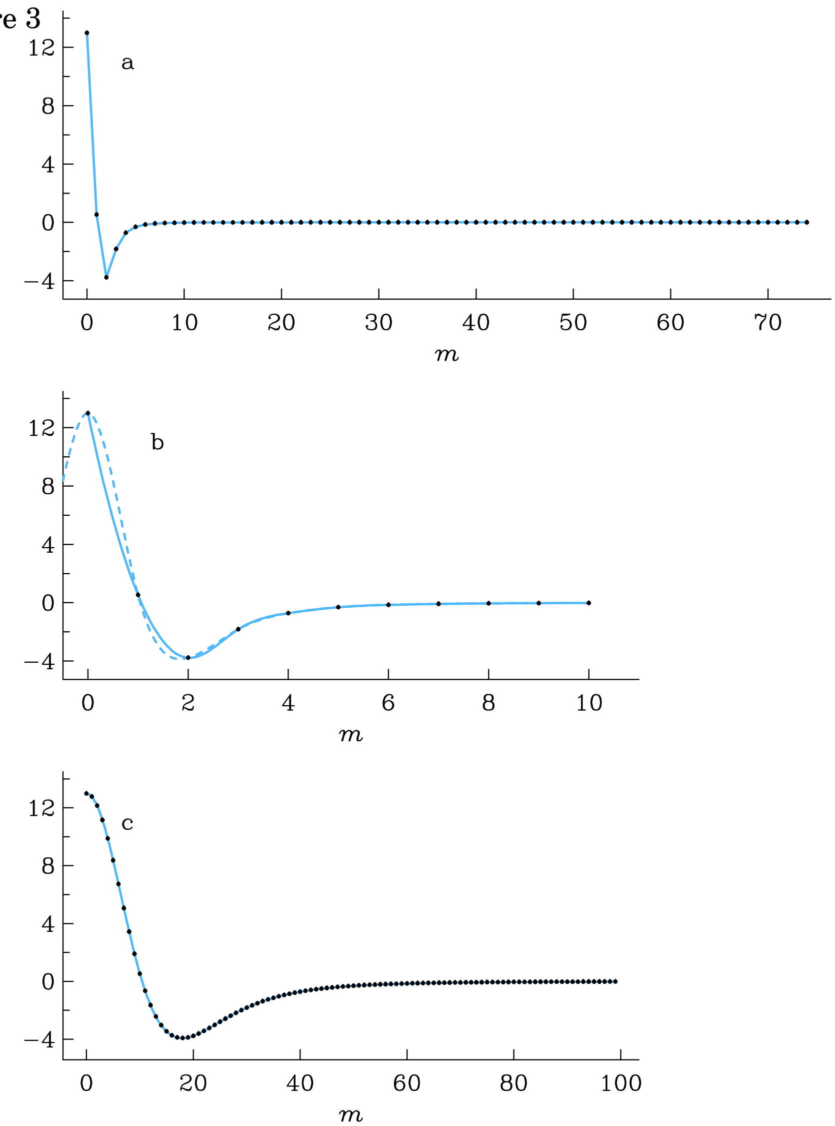

In fact we can get a better interpolation than Figure 3b, because we did not recognize that the transformed function $\hat{E}$, like E, itself is even. So we could spline smooth the rather sparse series if we can set the derivative at $x = 0$ to zero. That's shown dashed in 3b.

Next we turn to our Bessel function problem. Although we found that $N = 10^6$ would get us a part in a thousand, I decided against trying that. In Figure 4 we see how well we can do when we use (11) with $\Delta x = 0.63$ and $N = 1000$. Only the part near the origin has been plotted, because as before most of the output vector is near zero. In this case we know the exact transform; it is

$$\mathcal{F}\,[J_0] = \begin{cases} \dfrac{2}{(1-4\pi^2k^2)^{1/2}}\,, & |\,x\,| < 1/2\pi \\ 0, & |\,x\,| > 1/2\pi \end{cases} \tag{12}$$

This function is plotted in red in the Figure. The results are not too bad, except for the horrible undershoot near $x = 0.16$.

We can understand this phenomenon easily enough by recalling the analysis of spectral leakage in the PSD. The approximation to the true FT is obtained by setting to zero values outside the interval $(-X, +X)$, in other words multiplying the true g by a taper

$$\psi(x) = \mathrm{box}\left(\frac{x}{2X}\right) \tag{13}$$

Then

$$\tilde{g}(k) = \mathcal{F}\,[\psi\, g] = \hat{\psi} * \hat{g} \tag{14}$$

and we see that the approximate transform is the true one convolved with the FT of the box, namely

Figure 4.

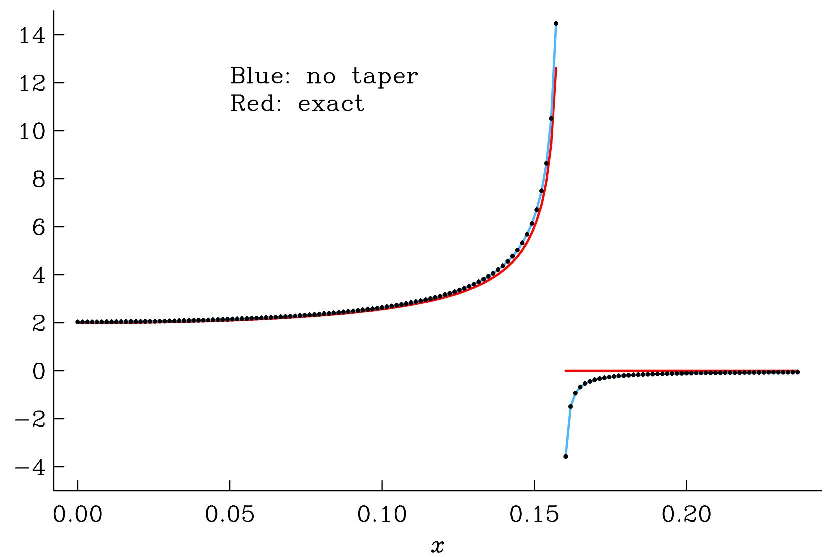

$$\hat{\psi}(k) = 2X \operatorname{sinc}(2kX) \tag{15}$$

The singular peak at $k = 1/2\pi$ in the true $\hat{g}$ of the Bessel function FT is being leaked into the region. This suggest that the solution to the problem is to multiply the original function with a taper that forces things to zero in a smooth manner for large x, and thus replaces the sinc in (15) with something that spreads peak values far less. I tried the following taper because it is simple to code, although a prolate would probably be better:

$$\psi(x) = \begin{cases} \cos^2 \dfrac{\pi x}{2X}, & |x| < X \\ 0, & |x| \geq X \end{cases} \tag{16}$$

for which the transform is

$$\hat{\psi}(k) = \frac{X \operatorname{sinc}(2kX)}{1 - 4k^2 X^2} \tag{17}$$

This transform decreases like k^{-3} rather than as k^{-1} for the untapered function.

Multiplying the Bessel function by the taper in (16) and then applying (11) yields remarkable improvement. The new approximation is so good that it is useless plotting the new result on the same graph as the correct transform because they are too similar. So in Figure 5 we plot the errors. In blue we see the absolute error in the FT of the tapered function, in light red that of the original approximation. There is thousand-fold improvement over most of the plotted range. Admittedly the absolute error rises near the singularity, and so if high accuracy were needed there, we would probably have to go to a prolate spheroidal taper, or perhaps some other completely different approach.

Figure 5.

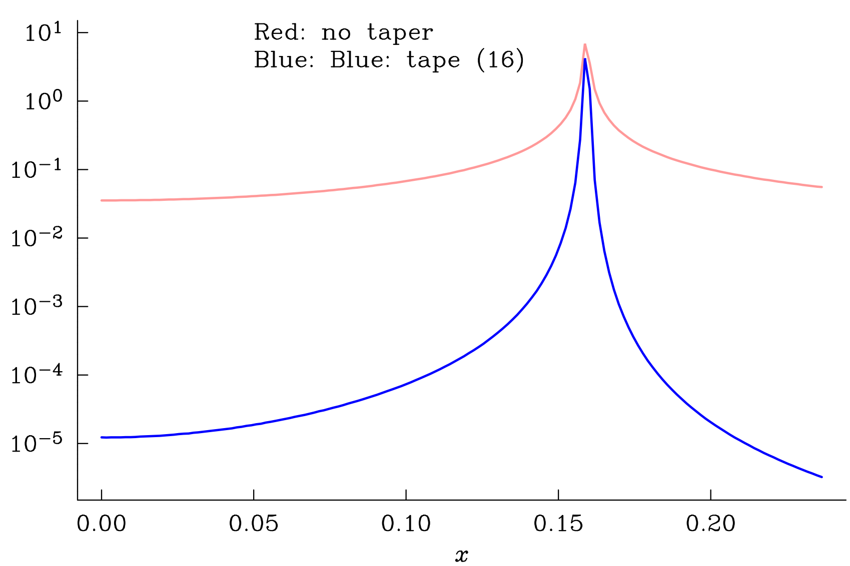

3. Slightly more General Functions

The analysis so far has treated real even functions. Little of consequence arises when we allow $g(t)$ to be a more general form. If $g(t)$ is real and odd, we can replace (11) with

$$\hat{g}(\pm m\,\Delta f) = \pm 2\mathrm{i}\,\Delta t\,\mathrm{Im}(\hat{g}_m),\;\; 0 \le m < \tfrac{1}{2}N\,. \tag{18}$$

The same taper should be used for slowly decaying odd functions as even ones.

Allowing a function that is neither even nor odd requires some reformulation, but can of course be handled by transforming the two functions:

$$g_e(t) = g(t) + g(-t);\quad g_o = g(t) - g(-t) \tag{19}$$

then reassembling after transforming with (11) and (18):

$$\hat{g}(f) = \tfrac{1}{2}(\hat{g}_e(f) + \hat{g}_o(f)) \tag{20}$$

We'll stop there, leaving the student to work out the extension to complex functions.