G2S3 :: July 2011 :: Waves and Imaging :: Fourier day(s), July 4 version

Preliminaries: from continuous to discrete Fourier transforms.

Fourier transform in $\mathbb{R}^n$ (physicist's convention):

$$\hat{f}(k) = (2\pi)^{-n/2} \int e^{-ix\cdot k} f(x)\, dx, \qquad f(x) = (2\pi)^{-n/2} \int e^{ix\cdot k} \hat{f}(k)\, dk.$$

Note that there are other commonplace conventions[1].

Consider the one-dimensional case ($n = 1$), with $x \in [0, 2\pi)$. Use N equally spaced points to discretize space as

$$x_j = jh, \qquad 0 \le j \le N-1, \qquad h = 2\pi/N.$$

Assume N is even, and discretize the frequency space as

$$k \in \mathbb{Z}, \qquad -N/2 \le k \le N/2 - 1.$$

The discrete Fourier transform (DFT) is obtained by sampling $f_j = f(x_j)$, replacing the integral over x by $h\sum_j$, and the integral over k by $\sum_k$. The constant in front of $\sum_j$ is further put to 1 (at the expense of a constant $1/N$ in front of $\sum_k$):

$$\hat{f}_k = \sum_{j=0}^{N-1} e^{-2\pi ijk/N} f_j, \qquad f_j = \frac{1}{N} \sum_{k=-N/2}^{N/2-1} e^{2\pi ijk/N} \hat{f}_k.$$

Show that $\hat{f}_k$ is periodic with period N. As a result of periodicity, it makes no difference whether $-N/2 \le k \le N/2 - 1$ or $0 \le k \le N-1$.

Matlab further complicates things (or not?) by starting indices at 1 instead of zero, so that $1 \le j, k \le N$. Matlab's formulas for the FFT compensate for this fact by subtracting 1 form both j and k:

$$\hat{f}_k = \sum_{j=1}^{N} e^{-2\pi i(j-1)(k-1)/N} f_j, \qquad f_j = \frac{1}{N} \sum_{k=1}^{N} e^{2\pi i(j-1)(k-1)/N} \hat{f}_k.$$

This way, when you call the k-th component of $\hat{f}_k$, the "physical" frequency is really $k-1$.

[1] Engineer's convention:

$$\hat{f}(k) = \int e^{-ix\cdot k} f(x)\, dx, \qquad f(x) = (2\pi)^{-n} \int e^{ix\cdot k} f(x)\, dx.$$

For time domain transforms, some electrical engineers and geophysicists sometimes put a plus sign in the direct transform:

$$\hat{f}(\omega) = \int e^{it\omega} f(t)\, dt, \qquad f(t) = (2\pi)^{-1} \int e^{-it\omega} \hat{f}(\omega)\, d\omega.$$

Mathematician's convention:

$$\hat{f}(k) = \int e^{-2\pi ix\cdot k} f(x)\, dx, \qquad f(x) = \int e^{2\pi ix\cdot k} f(x)\, dx.$$

1. **Indexing for the FFT**

   Fix N = 128. Consider the cosine wave

$$f_j = \cos(k_0 x_j),$$

   with x_j as defined earlier, and $k_0 = 20$. *Question: compute its DFT with the fft command. Plot the result.*

   - For which values of k is the DFT nonzero?
   - In view of $\cos a = (e^{ia} + e^{-ia})/2$, what frequencies k do you expect to be present in $\cos(k_0 x_j)$?
   - How do you explain the discrepancy between the answers to the 2 questions above?

   The zero frequency appears at $k = 1$ in the result of Matlab's fft, so it is the first component. For displaying a DFT it is usually preferred to apply a cyclic permutlation so that the zero frequency appears centered (negative frequencies to the left, positive frequencies to the right.) This is done with the fftshift command. *Question: apply fftshift to the DFT and re-plot it. Undo this shift with ifftshift, apply the inverse DFT with ifft, and plot the result to check that you recover f_j.* How accurate is the reconstruction?

   Consider now the function

$$g_\pi(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-\pi)^2}{2\sigma^2}},$$

   with $\sigma = 1/15$. Its (continuous) Fourier transform over $\mathbb{R}$ is

$$\hat{g}_\pi(k) = \frac{1}{\sqrt{2\pi}} e^{-\frac{k^2\sigma^2}{2}} e^{-ik\pi}.$$

   Sample g_π over $N = 128$ points x_j in $[0, 2\pi)$, perform a DFT, and find the proper normalization constant so that the result approximates the gaussian. *Question: show graphically that the normalized DFT agrees with the explicit gaussian formula,* for instance by plotting the error vector. How large is the error?

2. **Symmetries**

   - **Real-even symmetry**. Show (by pen and paper) that the Fourier transform of a real function $f(x)$ obeys the symmetry relation $\hat{f}(-k) = \overline{\hat{f}(k)}$. Show that if furthermore $f(x)$ is even ($f(-x) = f(x)$), then it has a Fourier transform which is also real and even. Similar statements hold for the DFT.

     Question: test these statements on a discretization of

$$g_0(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{x^2}{2\sigma^2}}.$$

     (This is an exercise in centering g_0 and using periodicity in x. The symmetry relation $\hat{g}_0(-k) = \overline{\hat{g}_0(k)} = \hat{g}_0(k)$ should be satisfied to within 14 digits.)

- **Translation covariance**. Show that the Fourier transform of $f(x-a)$ is $e^{-ik\cdot a}\hat{f}(k)$. A similar property holds for the DFT.

  Question: compute $Im \log(\hat{g}_\pi/\hat{g}_0)$ *and relate the numerical answer to your prediction.*

- **Dilation covariance**. Show that the Fourier transform of $f(x/a)$ is $a^n\hat{f}(ak)$ for $a>0$. An L^∞-preserving dilation in x becomes an L^1-preserving dilation in k, and vice-versa.

3. **Sampling**

Some information is usually lost when a function is sampled on a regular (Cartesian) grid. Two complex exponentials e^{ik_1x} and e^{ik_2x} are called aliases when they are indistinguishable on the grid $x_j = jh$; this happens when $k_1 - k_2$ is a multiple of $2\pi/h$. The k_1-th and k_2-th Fourier coefficients are also identical under the same condition; this is the periodicity property alluded to earlier. (Remember, $N = 2\pi/h = 2\pi/(2\pi/N)$.)

This can be made precise by the statement that sampling a function on a grid with spacing h periodizes its Fourier transform with a period $2\pi/h$. Let

$$\hat{f}(k) = \frac{1}{\sqrt{2\pi}}\int e^{-ixk} f(x)\,dx, \qquad \hat{g}(k) = \frac{h}{\sqrt{2\pi}}\sum_{j\in Z} e^{-ix_jk} f(x_j).$$

Then it can be shown that

$$\hat{g}(k) = \sum_{m\in Z} \hat{f}(k + \frac{2\pi}{h}m). \qquad \text{(Poisson summation formula)}$$

Question: using the Poisson summation formula, refine your guess for the DFT of $g_\pi(x_j)$*. Repeat your numerical evaluation of the error vector.*

Question: Put to zero every other sample of $g_\pi(x_j)$*. Compute the DFT of the result. Interpret your result with the Poisson summation formula.*

Note that if $\hat{f}(k)$ is compactly supported in k, with support in $[-\pi/h, \pi/h]$, then the periodization does not give rise to overlap, and $\hat{g}(k) = \hat{f}(k)$. We call such functions bandlimited.

4. **Convolution, filtering**

The convolution of two functions is

$$f * g(x) = \int f(y)g(x-y)\,dy.$$

Convolutions become multiplications in the Fourier domain, a property called the convolution theorem:

$$\widehat{(f*g)}(k) = (2\pi)^{n/2}\,\hat{f}(k)\,\hat{g}(k).$$

Conversely, a multiplication in x corresponds to a convolution in k:

$$(2\pi)^{n/2}\widehat{fg}(k) = \hat{f} * \hat{g}(k). \tag{1}$$

The circular convolution of two sequences is

$$(f \circ g)_j = \sum_{j'=0}^{N-1} f_{j'} g_{(j-j') (\text{mod } N)}, \qquad 0 \le j \le N-1.$$

This is the proper notion of convolution that gives rise to the discrete convolution theorem:

$$\widehat{(f \circ g)}_k = \hat{f}_k \, \hat{g}_k.$$

The regular, non-circular convolution is

$$(f * g)_j = \sum_{j'=0}^{N-1} f_{j'} g_{j-j'}, \qquad -N+1 \le j \le N-1.$$

The support of $f * g$ now has length $2N-1$: it is larger because it contains every j for which there exists $0 \le j' \le N-1$ such that we simultaneously have $0 \le j - j' \le N-1$.

Question: implement a non-circular discrete convolution directly, and via FFT. Check that your results agree on vectors such as randn(N,1). The non-circular convolution can be viewed as a circular convolution by padding each of the vectors to be convolved by $N-1$ zeros.

Convolution is the basic operation behind filtering: low-pass filtering for enhancing the low frequencies in a signal, high-pass filtering for enhancing the high frequencies, bandpass filtering, etc. All these are realized by convolving an input signal by a well-chosen function/vector called filter. The Fourier transform of the filter is called transfer function in EE. The property that filtering can be realized by multiplication by the transfer function is often useful to speed up computation.

5. **Quadrature**

When $k = 0$, the Poisson summation formula yields

$$h \sum_{j \in \mathbb{Z}} f(x_j) = \sqrt{2\pi} \sum_{m \in Z} \hat{f}(\frac{2\pi}{h} m).$$

The term $m = 0$ is the integral of which the left-hand-side is the trapezoidal rule:

$$\sqrt{2\pi} \hat{f}(0) = \int f(x) \, dx.$$

The terms $m \neq 0$ quantify the error made in applying the trapezoidal rule for computing the integral of f. If f is bandlimited with band limit π/h, i.e. $\hat{f}(k) = 0$ for $|k| > 2\pi/h$, then the trapezoidal rule will be exact. The same conclusion is true over the interval $[0, 2\pi)$: the trapezoidal rule will be exact for bandlimited functions provided the band limit is less than π/h, and provided the function is periodic (this is important).

Consider now inner products, for which

$$h \sum_{j \in \mathbb{Z}} f(x_j) g(x_j) = \sqrt{2\pi} \sum_{m \in Z} \widehat{fg}(\frac{2\pi}{h} m).$$

From equation (1), we see that $\widehat{fg}$ is obtained form convolving $\hat{f}$ and $\hat{g}$. If $\hat{f}$ and $\hat{g}$ have the same support, then the convolution will double this support. In particular, if both f and g are bandlimited with band limit π/h, then fg is bandlimited with band limit $2\pi/h$. In that case, the quadrature on a grid with spacing h will be exact again.

The results that the quadrature of $\int f$, resp. $\int fg$, are exact provided f and g are bandlimited with sufficiently small band limit ($2\pi/h$, resp. π/h), are together known as the Shannon sampling theorem. On the other hand, the situation where periodization leads to overlap of the Fourier transform is called aliasing.

Question: compute $h\sum_j g_\pi(x_j)$ *and show that it is a very good approximation to* $\int g_\pi(x)\,dx = 1$ *– much better than a naive analysis of the trapezoidal rule would show. Why is that? If in the definition of* $g_\pi(x)$ *we take* $\sigma = 1/50$ *instead, why has the accuracy worsened? If we take* $\sigma = 1/2$*, why is the accuracy poor as well?*

6. **Smoothness**

The previous question shows that two types of errors arise in the computation of inner products:

- the truncation error, because the interval over which the integral is taken might not contain the support of all functions.
- the sampling error, because using a Cartesian grid amounts to periodizing the Fourier transform, resulting in aliasing.

For integrals to be computed accurately, both the integrand and its Fourier transform should decay fast. These requirements are to some extent contradictory: it is impossible for a function and its Fourier transform to be simultaneously compactly supported. This is called the uncertainty principle: if a function is peaked in x, it has to be spread in k, and vice-versa.

Decay in k is linked to smoothness in x: the smoother a function, the faster the decay of its Fourier transform.

As an example, let $a > 0$, and consider the function

$$f(x) = \frac{1}{2a}\chi_{[-a,a]}(x-\pi)$$

It is easy to show that

$$\hat{f}(k) = \frac{\sin(ka)}{ka}e^{-ik\pi} \equiv \mathrm{sinc}(ka)e^{-ik\pi}.$$

Take $a = 1$ in the numerical example below.

Question: discretize $f_j = f(x_j)$*, and consider the successive autoconvolutions of* f_j*. Plot them – they re called B-splines. Plot the corresponding DFT and predict the rate of decay as* $k \to \infty$ *as a negative power of* k*.* Question for mathematicians: what do the B-splines converge to as an infinite number of convolutions are taken? Why?

7. **Interpolation**

   In a nutshell, spectral interpolation of a sequence at a higher rate, still on a Cartesian grid, is done by zeropadding then inverting the Fourier transform.

   Question: Consider $g_\pi(x_j)$ again, with $\sigma = 1/2$. Start from a grid with $N = 128$, and interpolate on a finer grid $N = 256$. If you follow the prescription above, your result will contain an imaginary part. Why is that? How should it be removed?

   Question: do the same for $g_\pi(x_j)$ with $\sigma = 1/15$. Compare the accuracy of your answer with that of linear interpolation.

8. **Parseval identity**

$$\int f(x)\overline{g(x)}\,dx = \int \hat{f}(k)\overline{\hat{g}(k)}\,dk.$$

   Question: find and implement the corresponding identity for the DFT. Check its accuracy on vectors with random entries.

9. **Unequispaced FFT**

   Implement the Dutt-Rokhlin algorithm, and tune its parameters for best accuracy.