## 2D and 3D Fourier transforms

## The 2D Fourier transform

The reason we were able to spend so much effort on the 1D transform in the previous chapter is that the 2D transform is very similar to it. The integrals are over two variables this time (and they're always from $-\infty$ to $\infty$ so I have left off the limits). The FT is defined as

$$
\begin{equation*}
G(u, v)=\iint g(x, y) e^{-i 2 \pi(x u+y v)} d x d y \tag{1}
\end{equation*}
$$

and the inverse FT is

$$
\begin{equation*}
g(x, y)=\iint G(u, v) e^{i 2 \pi(x u+y v)} d u d v \tag{2}
\end{equation*}
$$

The Gaussian function is special in this case too: its transform is a Gaussian.

$$
\begin{equation*}
e^{-\pi\left(x^{2}+y^{2}\right)} \xrightarrow{F T} e^{-\pi\left(u^{2}+v^{2}\right)} \tag{3}
\end{equation*}
$$

The Fourier transform of a 2D delta function is a constant

$$
\begin{equation*}
\delta(x) \delta(y) \xrightarrow{F T} 1 \tag{4}
\end{equation*}
$$

and the product of two rect functions (which defines a square region in the $x, y$ plane) yields a $2 \mathrm{D} \operatorname{sinc}$ function:

$$
\begin{equation*}
\operatorname{rect}(x) \operatorname{rect}(y) \xrightarrow{F T} \operatorname{sinc}(u) \operatorname{sinc}(v) . \tag{5}
\end{equation*}
$$



The function rect( $\mathbf{x}) \operatorname{rect}(\mathbf{y})$ is shown on the left. Its transform is the function $\sin (u) \operatorname{sinc}(\mathrm{v})$ shown on the right. (Ignore the units in the axes, they are the units of the discrete FT used to make the figure.)


One special 2D function is the circ function, which describes a disc of unit radius. Its transform is a Bessel function,

$$
\begin{equation*}
\operatorname{circ}(r) \xrightarrow{F T} \frac{J_{1}(2 \pi \rho)}{\rho} \tag{6}
\end{equation*}
$$

here the variables $r$ and $\rho$ represent $\left(x^{2}+y^{2}\right)^{1 / 2}$ and $\left(u^{2}+v^{2}\right)^{1 / 2}$, respectively. Again, as in the case of the rect function, something with "sharp edges" in one domain transforms into something with ripples in the other.



The circ function (shown on the left) has the transform on the right (a Bessel function, also known as the Airy function.) The Airy function is circularly symmetric, but doesn't quite look like that here because of aliasing artifacts from the discrete FFT on a computer (more about that later).

The 2D FT has a set of properties just like the 1D transform.

1. Linearity

$$
g+h \xrightarrow{F T} G+H
$$

2. Scale

$$
g(a x, b y) \xrightarrow{F T} \frac{1}{|a b|} G\left(\frac{u}{a}, \frac{v}{b}\right)
$$

3. Shift

$$
g(x-a, y-b) \xrightarrow{F T} G(u, v) e^{-i 2 \pi(a u+b v)}
$$

4. Convolution

$$
g * h \xrightarrow{F T} G \cdot H
$$

5. Rotation

$$
g\left(x^{\prime}, y^{\prime}\right) \xrightarrow{F T} G\left(u^{\prime}, v^{\prime}\right),
$$

where $\left(x^{\prime}, y^{\prime}\right)$ and $\left(u^{\prime}, v^{\prime}\right)$ are rotated about the origin through the same angle.

The rotation property is the only one we haven't seen before. You can understand it this way: if we define the vectors $\mathbf{x}=(x, y)$ and $\mathbf{u}=(u, v)$ then we can rewrite the definition of the FT (eqn. 13) as

$$
\begin{equation*}
G(\mathbf{u})=\iint g(\mathbf{x}) e^{-i 2 \pi \mathbf{x} \mathbf{u}} d \mathbf{x} \tag{7}
\end{equation*}
$$

where the vectors appear only through a dot product in the exponential function. Thus if you rotate the coordinate system for $\mathbf{x}$, a corresponding rotation of $\mathbf{u}$ will give the same result for the dot product.

More formally, let $\mathbf{R}$ be a rotation matrix. Then the FT of a rotated function $g(\mathbf{R x})$ can be gotten through the substitution $\mathbf{x}^{\prime}=\mathbf{R x}$ in this way:

$$
\begin{aligned}
G_{\mathbf{R}}(\mathbf{u}) & =\iint g(\mathbf{R x}) e^{-i 2 \pi \mathbf{x} \mathbf{u}} d \mathbf{x} \\
& =\iint g\left(\mathbf{x}^{\prime}\right) e^{-i 2 \pi\left(\mathbf{R}^{-1} \mathbf{x}^{\prime}\right) \mathbf{u}} d \mathbf{x}^{\prime} \\
& =\iint g\left(\mathbf{x}^{\prime}\right) e^{-i 2 \pi \mathbf{x}^{\prime}(\mathbf{R u})} d \mathbf{x}^{\prime} \\
& =G(\mathbf{R u})
\end{aligned}
$$

Like the 1D lattice, a 2D lattice with unit spacing transforms into the same function in the other domain. Making use of the scaling rule, it is then easy to show that the general 2D lattice transforms this way:

$$
\begin{equation*}
\left[\sum_{n=-\infty}^{\infty} \delta(a x-n)\right]\left[\sum_{m=-\infty}^{\infty} \delta(b y-m)\right] \xrightarrow{F T} \frac{1}{a b}\left[\sum_{n=-\infty}^{\infty} \delta\left(\frac{u}{a}-n\right)\right]\left[\sum_{m=-\infty}^{\infty} \delta\left(\frac{v}{b}-m\right)\right] \tag{8}
\end{equation*}
$$

This means that a lattice with spacings $1 / a$ and $1 / b$ transforms to a lattice with spacings $a$ and $b$, respectively. This is the origin of the term "reciprocal space" for the Fourier transform space.

## 2D Power spectrum

The 2D power spectrum is an important tool in electron microscopy. Recall from last time that the power spectrum is the magnitude squared of the FT, so in 2 D it is

$$
\begin{equation*}
S(u, v)=|F(u, v)|^{2} . \tag{9}
\end{equation*}
$$

If the random signal has been filtered by some filter function $H$ such that $F=G \cdot H$, then the expectation value of the spectrum is given by

$$
\begin{equation*}
S(u, v)=|G(u, v)|^{2}|H(u, v)|^{2} \tag{10}
\end{equation*}
$$

Let's let $H$ be the contrast transfer function. In two dimensions the CTF should be circularly symmetric (why do you think?) We'll define the magnitude $s$ of the spatial frequency as

$$
s=\sqrt{u^{2}+v^{2}}
$$

so the CTF has the form-ignoring spherical aberration but including the envelope function

$$
\begin{equation*}
H(s)=\sin \left(-\pi \delta s^{2}-\alpha\right) e^{-B s^{2} / 4} \tag{11}
\end{equation*}
$$

Below left is the theoretical CTF, the middle is a micrograph, and on the right is the power spectrum of the micrograph. Note that the power spectrum looks a lot like the CTF (actually according to (10) it is proportional to the CTF squared).


## The 2D FT and diffraction

The diffraction pattern is the Fourier transform of the amplitude pattern of a source of radiation. Consider the following system. A plane wave is propagating in the $+z$ direction, passing through a scattering object at $\mathrm{z}=0$, where its amplitude becomes $A_{o}(x, y)$. We want to know the amplitude of the wave at the detector in the $u, v$ plane, which is a distance $z$ from the $x, y$ plane.


Classical wave theory says that we can solve this problem by treating each point in the $x, y$ plane as the source of a spherical wavefront. Here is the wave function of a spherical wave of magnitude $A_{0}$ that starts from the origin,

$$
\Psi(r)=\frac{A_{0}}{r} e^{i 2 \pi r / \lambda}
$$

and the wave function at the detector plane is gotten by adding up a spherical wave emanating from each point in the $\mathrm{x}, \mathrm{y}$ plane. Using the weak-phase approximation we'll scale the amplitude of each diffracted wave by $1 / i \lambda$ and get:
$\psi_{d}(u, v)=\iint_{x, y} \frac{A_{0}(x, y)}{i \lambda \sqrt{(u-x)^{2}+(v-y)^{2}+z^{2}}} e^{i 2 \pi \sqrt{(u-x)^{2}+(v-y)^{2}} / \lambda} d x d y$.
Making various approximations, specifically assuming $u$ and $v$ are large compared to $x$ and $y$ and that $z$ is much larger than all other dimensions, the Fraunhofer formula for the wave amplitude at the detector plane is

$$
\begin{equation*}
\Psi_{d}(u, v)=\frac{e^{i 2 \pi z / \lambda}}{i \lambda z} e^{i 2 \pi\left(u^{2}+v^{2}\right) / 2 \lambda z} \iint_{(x, y)} A_{o}(x, y) e^{-i 2 \pi(x u+y v) / \lambda z} d x d y \tag{12}
\end{equation*}
$$

or, letting $G$ be the Fourier transform of the amplitude distribution $A(x, y)$,

$$
\begin{equation*}
\Psi_{d}(u, v)=\frac{e^{i 2 \pi z / \lambda}}{i \lambda z} e^{i \pi\left(u^{2}+v^{2}\right) / \lambda z} G\left(\frac{u}{\lambda z}, \frac{v}{\lambda z}\right) \tag{13}
\end{equation*}
$$

Where $G$ is the Fourier transform of the amplitude distribution $A_{0}(x, y)$. The intensity of X-rays or photons at the detector will be the squared magnitude of $\Psi_{d}$, in which the various phase factors drop out and we have

$$
\begin{equation*}
I(u, v)=\frac{1}{\lambda z}\left|G\left(\frac{u}{\lambda z}, \frac{v}{\lambda z}\right)\right|^{2} \tag{14}
\end{equation*}
$$

The intensity at the detector is proportional to the magnitude squared of the FT of the scattering pattern. This is the basis of electron crystallography. Working from the measured intensity at the detector, the first step is to find what the amplitude is at the detector (which you get by taking the square root of the intensity and assigning phases). Then the IFT give you the original pattern of amplitudes $A_{o}$ at the object.

## The 3D Fourier transform

In the same way, there exists a 3D Fourier transform as well. It is defined as a triple integral, and it has all the properties of the 2D FT, including rotations.

$$
\begin{equation*}
G(u, v, w)=\iiint g(x, y, z) e^{-i 2 \pi(x u+y v+w z)} d x d y d z \tag{15}
\end{equation*}
$$

Central to the theory of 3D reconstruction is the "central slice theorem". It is based on the fact that for any 3D distribution of density $g(x, y, z)$ there is a 3D Fourier transform volume $G(u, v, w)$. If you take a projection through $g$ to obtain a 2D image, it turns out that the Fourier transform of that image has the same values as slice through $G$. If you have data to fill up the Fourier volume with slices, then you can do an inverse transform to obtain the density map $g$. We'll talk more about this next time.

## Discrete Fourier transform and terminology

In this course we will be talking about computer processing of images and volumes involving Fourier transforms. The computer operates on data that have been sampled at regular, finite intervals and produces results that we view as individual pixels or voxels. There is an alternative Fourier transform in which the integrals are replaced by sums for working on such finite data sets. This is called the discrete Fourier transform (DFT). Its theory and use are very similar to those of the continuous FT that we have discussed here. There are some considerations about sampling intervals that have to be made, but otherwise things work just fine.

In the 1960s people discovered a beautiful algorithm for computing the DFT that is extremely efficient. It is called the fast Fourier transform. Its results are exactly the same as the DFT, only obtained much
more quickly. Here is an example. Suppose you want to do the 2D DFT of a $1000 \times 1000$ pixel image. To get the $1000 \times 1000$ element DFT, you have to do $10^{12}$ arithmetic operations (just think, you have to use all values of $x, y, u$ and $v$ in the calculation). It is spectacular that this calculation can be done more than $10^{4}$ times faster with the FFT. The FFT of a $1000 \times 1000$ image takes less than 0.1 second on my laptop computer, requiring about $10^{8}$ operations!

So, whenever someone talks about computing an FT on a computer, they usually say they are "taking the FFT". They are actually taking the discrete FT using the FFT algorithm.

More about the DFT can be found in the Appendix below.

## Appendix. The discrete Fourier transform

The Fourier transform is defined as an integral over all of space. How could we evaluate this on a computer? We will have to take a finite number of discrete samples both in the real $(x, y)$ space and in the $(u, v)$ frequency space.

Let's first do this in one dimension, and we'll model discrete samples by multiplying functions by series of delta functions. Recall the Fourier transform pair

$$
\begin{equation*}
\sum_{n=-\infty}^{\infty} \delta(a x-n) \xrightarrow{F T} \frac{1}{a} \sum_{m=-\infty}^{\infty} \delta\left(\frac{u}{a}-m\right) \tag{1.1}
\end{equation*}
$$

We can use this to understand the properties of the discrete Fourier transform (DFT). Start with our favorite FT pair,

$$
e^{-\pi x^{2}} \xrightarrow{F T} e^{-\pi u^{2}}
$$

which we'll call $g_{0}$ and $G_{0}$. But now we'll multiply $g_{0}$ by a series of delta functions (a scaled Shah function) to yield

$$
\begin{equation*}
g_{1}(x)=e^{-\pi x^{2}} \sum_{n=-\infty}^{\infty} \delta(a x-n) \tag{1.2}
\end{equation*}
$$

where $a$ is set to 4 , that is we sample $x$ in units of $1 / 4$. We know what the FT of (1.2) will be: it will be the convolution of $G(u)$ with a series of delta functions, so we'll get copies of $G$ spaced every 4 units on the $u$ axis. See the middle row of Figure A1. This spacing looks to be sufficient that if we were just to cut out the central copy of the Gaussian function, between $u=$ -2 to $u=2$, we'd have the original function back. This demonstrates a very important theorem, the sampling theorem. If the FT of some function $g_{0}$ is zero outside a limited interval of $u$, it can be reconstructed perfectly from discrete samples. The only requirement is that the sampling frequency must be at least twice the highest frequencies $\pm u$ for which $G(u)$ is nonzero.

You will understand what will happen if this criterion, called the Nyquist criterion, is violated. Imagine that we had sampled $g(x)$ more coarsely, say in $x$ steps of $1 / 2$. Then the periodicity of the repeats in $G_{1}$ would be only two units and there would be overlap of the copies of the Gaussians. This overlap is called aliasing and it reflects the loss of information-we can no longer reconstruct $g_{1}$ perfectly from $G_{1}$.

Now, what happens if we sample $G_{1}$, say in steps of $1 / 4$ in $u$, to yield the function $G_{2}$. If we evaluate its inverse transform $g_{2}$ we see that it also consists of copies of the original function


Figure A1. The continuous functions g 0 and $\mathbf{G 0}$ are sampled, resulting in periodic copies.
(see the bottom row of Fig. A1). Now we have discrete samples in both $x$ and $u$, which is exactly what we'd need to process data on a computer. Discrete samples mean that both $g$ and $G$ have become periodic functions, but as long as we make the periods long enough (by sampling finely enough) we can work with these discrete samples just fine, by chopping out one period to work with.

So, if we restrict both $x$ and $u$ to integer values, we can define the discrete Fourier transform and its inverse (DFT and IDFT, respectively) as

$$
\begin{align*}
& G(u)=\sum_{x=0}^{N-1} g(x) e^{-i 2 \pi u x / N}  \tag{1.3}\\
& g(x)=\frac{1}{N} \sum_{u=0}^{N-1} G(u) e^{i 2 \pi u x / N} \tag{1.4}
\end{align*}
$$

The values of $x$ range from 0 to $N-1$, as do the values of $u$. This can be confusing, as demonstrated in Fig. A2. Part A of the figure shows how I'd like to visualize a Gaussian function in an array with $N=16$ for use with the DFT, with index values of $x=-8,-7, \ldots-2,-1,0$, 1, 2...7. Part B shows however the conventional representation for using the DFT on a computer. The successive points correspond to my $x$ values of $0,1, \ldots 6,7,-8,-7, \ldots-1$. Remember that the DFT treats its input as if it were periodic, so it will "splice together" the end points of the interval; this means that the representations in A and B are equivalent. Note also that some programming languages have array indices starting with 1 instead of 0 . (Matlab and Fortran are like this.) As I'm using Matlab you'll see that in panel B the point corresponding to the peak of the Gaussian is at position 1 on the $x$ axis, which is actually showing array indices.


Figure A2. Two discrete representations of a Gaussian function using $\boldsymbol{N}=\mathbf{1 6}$ points.
Two commonly-used Matlab functions for "unwrapping" functions for use in the DFT are fftshift and ifftshift. To convert part B of the figure to part A, one uses fftshift; to convert A to $B$ you use ifftshift. If the number of points $N$ is even, fftshift and ifftshift do exactly the same thing. So if you have an array $g$ that is indexed like A, you can compute the DFT and get a centered Gaussian back by typing

```
G=fftshift(fft(ifftshift(g)));
```

In Matlab you have the functions fft and ifft (one-dimensional forward and inverse transforms) and fftn and ifftn ( n -dimensional transforms).

The FFT algorithm happens to work best when the number of points $N$ is a power of 2 or a product of small factors. That is one reason why people often use array sizes like 1024 instead of, say, 1020.

## Windowing

Figure A3 demonstrates another problem that arises from the finite sum in the DFT, that we never had with the FT because the latter uses an infinite integral. Function $g_{1}$ is a cosine function that has 8 cycles in the $N=80$ points, and its DFT looks a lot like the two delta functions we expect from the FT of a cosine function. (Right?) The cosine function $g_{2}$ has a slightly higher frequency, so 8.5 cycles fit in the 80 points. Its DFT looks horrible. Why all the ripples if we're taking the transform of a nice smooth cosine function? Remember that the DFT thinks its input is a periodic function. If we shift the display of $g_{2}$ by half of the window (bottom row in the figure) you can see that $g_{2}$ actually contains a discontinuity which is what causes the trouble.


Figure A3. DFTs of a cosine function. $\boldsymbol{N}=\mathbf{8 0}$.

One can get rid of the problems with a discontinuity by using a windowing function $w(x)$ that tapers gently toward zero at the edges. Here in Fig. A4 are the same functions except that $g_{1}^{\prime}$ and $g_{2}^{\prime}$ are the products of $g_{1}$ and $g_{2}$ with a smooth window function. It gives better looking transforms.


Figure A4. A window function has been applied to the cosine.
Note that in the upper two panels of this figure I used Matlab's $f f t s h i f t$ function to move the center of the waveforms to the center of each plot; but I changed the labeling of the $x$-axis to run from -40 to +39 , rather than from 1 to 80 .

## Scaling the DFT

Once we start using integer values of $x$ and $u$ for the DFT, we have to keep in mind the relationship between the integer values and real-world quantities. For example, let $x$ index the pixels of a micrograph. Then one unit in $x$ corresponds to a pixel size $d x$, e.g. 1.7 $\AA$. When we take the DFT to get the result in terms of the spatial frequency $u$, one unit in $u$ corresponds to a frequency of $d u=1 /(N d x)$, that is 1 / (size of the micrograph). So for example if $N=1024$ then $d u=5.7 \times 10^{-4} \AA^{-1}$. We can choose the integer $u$ values to range from $-N / 2$ to $N / 2-1$ and therefore correspond to real-world spatial frequencies of about -0.3 to 0.3 inverse angstroms, that is $\pm 1 /(2 d x)$. The maximum accessible frequency, $1 /(2 d x)$ is called the Nyquist frequency.

