CHAPTER 11

The Fourier Transform

It is often useful to choose a convenient basis when you want to understand a collection of vectors. The machinery is quite straightforward – choose a transformation (usually, Euclidean or affine) that puts the vectors into a more convenient form, work with that form, then transform back if useful. It turns out that you can think of a function as being like a vector, so you can change the basis in which a function is expressed. Doing so is very useful. In this chapter, I will sketch a collection of related procedures that express various different functions on a basis of sinusoids of differing frequencies.

In the following chapter, I will show various applications of the idea, which is pervasive in signal processing, image processing and computer vision. For example, Figure 3.4 implies that sampling errors are related to fast changes in a signal. What is lost can be accounted for using a basis of sinusoids. As another example, when you smooth a function, you suppress some frequencies. Different smoothing procedures have quite different results, and these differences can be explained in terms of what was suppressed. As a final example, a sinusoidal basis shows how much smoothing is required to get good outcomes from sampling.

11.1 PRELIMINARIES

There are important and delicate matters I won't deal with in any detail here. These have to do with whether particular integrals exist; whether various series converge; and so on. For the functions of interest to us, the integrals will exist and the series will converge.

11.1.1 Fourier Series

Let $f(x)$ be some periodic function on the range $[0, 1]$. It is periodic, because $f(0) = f(1)$. You can represent $f(x)$ by a series

$$\sum_{u=-\infty}^{\infty} a_u e^{i2\pi ux}$$

where $i^2 = -1$, u is an integer and $e^{i2\pi ux} = \cos 2\pi ux + i \sin 2\pi ux$ is a complex exponential. This is a sinusoid, with frequency given by u. This representation works for any function of interest to us.

The coefficients of the series representing a particular function are easily recovered with a trick based on the properties of complex exponentials. Recall

$$\int_0^1 e^{i2\pi ux} dx = \begin{cases} 1 & \text{if } u = 0 \\ 0 & \text{otherwise} \end{cases}$$

which means that

$$\int_0^1 e^{i2\pi ux}e^{i2\pi vx}dx = \begin{cases} 1 & \text{if } u+v=0 \\ 0 & \text{otherwise} \end{cases}$$

All this means that the coefficients of the series are easy to recover, and

$$a_u = \int_0^1 f(x)e^{-i2\pi ux}dx.$$

An analogy may help understand this expression. Think of a periodic function as a vector with an infinite number of coefficients; then an integral is very like a dot product. This is a useful analogy because dot products measure the amount of one vector in the direction of another. If you think of a_u as the amount of $f(x)$ in the direction of $e^{i2\pi ux}$, then the series is an expression of $f(x)$ on a novel basis.

This representation works for any periodic function we care to work with. Different functions yield different series. There are a variety of interesting conditions under which a series a_u will yield a function; one is that

$$\sum_{-\infty}^{\infty} | a_u | < \infty.$$

There are interesting relationships between spaces of periodic functions and the rate of convergence of the series. Assuming no problems under any of these headings, this procedure defines a 1-1 change of basis – for any periodic function, you can construct a unique series that represents the function, and for any (reasonable) series, you have a periodic function.

11.1.2 Using a Fourier Series in 1D

Fourier series can be used to simplify all sorts of interesting applied mathematical problem. If you have a linear system and can deal with a complex exponential without too much trouble, then a Fourier series may simplify your problem. Here is a simple example, related to where the whole idea started, and worth a little trouble to understand. You wish to obtain $f(x,t)$, defined on $[0,1] \times [0,\infty)$ that solves

$$\frac{\partial^2 f}{\partial x^2} = \frac{\partial f}{\partial t}$$

subject to the boundary conditions that

$$f(x,0) = g(x) \text{ and } f(0,t) = f(1,t)$$

Now model $f(x,t)$ as a Fourier series in x. The coefficients vary with time, so

$$f(x,t) = \sum_{u=-\infty}^{\infty} a_u(t)e^{i2\pi ux}.$$

Now plug this expression into the equation, and get

$$\sum_{u=-\infty}^{\infty} -u^2 a_u e^{i2\pi ux} = \sum_{u=-\infty}^{\infty} \frac{da_u}{dt} e^{i2\pi ux}$$

which is true for each u individually, meaning

$$-(2\pi u)^2 a_u = \frac{da_u}{dt} \text{ so } a_u(t) = a_u(0)e^{-(2\pi u)^2 t}.$$

Furthermore, $f(x, 0) = g(x)$, so

$$a_u(0) = \int_0^1 g(x)e^{-i2\pi ux}dx.$$

What looked like a quite nasty PDE dissolved into a collection of independent and easy ODE's and was easily dealt with. You should think of this example as being rather like a case where one solves a linear system by changing basis so that the matrix is diagonal, and then exploiting the easy linear system that a diagonal matrix gives.

11.1.3 Fourier Series in 2D

Fourier series can be constructed in 2D as well. If $f(x, y)$ is a doubly periodic function on $[0, 1] \times [0, 1]$ (so $f(x, 0) = f(x, 1)$ and $f(0, y) = f(1, y)$), you can represent it by a series

$$\sum_{u=-\infty}^{\infty} a_{uv}e^{i2\pi(ux+vy)}$$

for u, v integers. Now $e^{i2\pi(ux+vy)}$ is a complex exponential in 2 dimensions; u gives its frequency in the x direction and v gives its frequency in the y direction. The coefficients are now

$$a_{uv} = \int_0^1 \int_0^1 f(x, y)e^{i2\pi(ux+vy)}dx.$$

Questions of convergence, uniqueness, etc. become interesting and delicate here, but there is nothing to concern us.

11.1.4 Using a Fourier Series in 2D

The previous example works in 2D, too. Now you wish to obtain $f(x, y, t)$, defined on $[0, 1] \times [0, 1] \times [0, \infty)$ that solves

$$\frac{\partial^2 f}{\partial x^2} + \frac{\partial^2 f}{\partial y^2} = \frac{\partial f}{\partial t}$$

subject to the boundary conditions that

$$f(x, y, 0) = g(x, y) \text{ and } f(x, 0, t) = f(x, 1, t) \text{ and } f(0, y, t) = f(1, y, t)$$

Now model $f(x, y, t)$ as a Fourier series in x as above, and discover that

$$a_{uv}(t) = a_{uv}(0)e^{-(2\pi)^2(u^2+v^2)t}$$

following the argument above. Furthermore, $f(x, y, 0) = g(x, y)$, so

$$a_{uv}(0) = \int_0^1 \int_0^1 g(x, y)e^{-i2\pi(ux+vy)}dxdy.$$

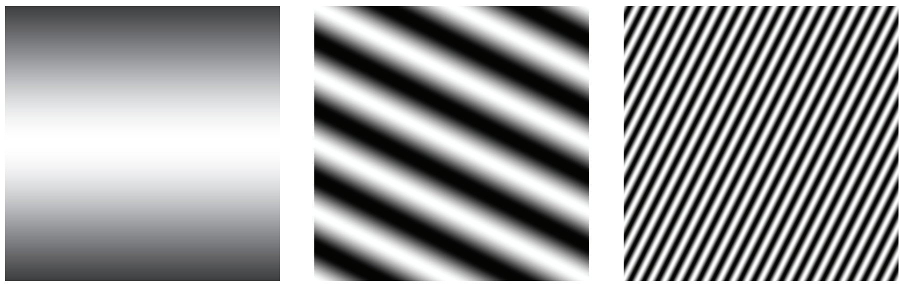

FIGURE 11.1: *The real component of Fourier basis elements shown as intensity images. The brightest point has value one, and the darkest point has value zero. The domain is* $[-1, 1] \times [-1, 1]$*, with the origin at the center of the image. On the* **left**, $(u, v) = (0, 0.4)$*; in the* **center**, $(u, v) = (1, 2)$*; and on the* **right** $(u, v) = (10, -5)$*. These are sinusoids of various frequencies and orientations described in the text.*

Again, something that looked like a very nasty PDE dissolved into a collection of independent and easy ODE's and was easily dealt with. Again, think of this example as being rather like a case where one solves a linear system by changing basis so that the matrix is diagonal, and then discovering that a diagonal matrix gives an easy system to solve.

> **Remember this:** *Periodic functions can be represented with a series of sinusoids using a Fourier series. This representation is a change of basis that makes solving some problems quite straightforward.*

11.2 FOURIER TRANSFORMS

A periodic function on the unit interval in 1D is very different from a general function in 1D. The values of the periodic function outside the unit interval can be ignored, but the general function can have non-zero values of significance for $|x|$ very large. It can "wiggle" at or near infinity. Because the general function can display much richer variation than a periodic function can, you should expect that a larger basis is required to represent its behavior. A similar argument applies to 2D as well.

11.2.1 A Basis for General Functions in 2D

For a periodic function in 2D, I used $e^{i2\pi(ux+vy)}$ as a basis, where u and v were integers that identified the basis function. For a general function in 2D, it turns out that the appropriate basis is

$$e^{i2\pi(ux+vy)}$$

but now u and v are arbitrary real numbers rather than integers. This means that there are many more basis elements than in the case of a Fourier series, and various complications of analysis occur. These can mostly be ignored except by specialists.

Any particular basis element is identified by a pair u and v of real numbers. For the moment, fix u and v, and consider the meaning of the value of the transform at that point. The exponential can be rewritten

$$e^{-i2\pi(ux+vy)} = \cos(2\pi(ux+vy)) + i\sin(2\pi(ux+vy)).$$

These terms are sinusoids on the x, y plane, whose orientation and frequency are chosen by u, v. Each term is constant when $ux+vy$ is constant (i.e., along a straight line in the x, y plane whose orientation is given by $\tan\theta = v/u$). The gradient of each term is perpendicular to lines where $ux+vy$ is constant, and the frequency of each sinusoid is $\sqrt{u^2+v^2}$. These sinusoids are often referred to as *spatial frequency components* or *spatial frequencies*; a variety are illustrated in Figure 11.1.

11.2.2 The Fourier Transform

As in the case of Fourier series, the change in basis is effected by integration. There is now one coefficient for every (u, v). The procedure for obtaining the coefficient is now called a *Fourier transform*.

Definition: 11.1 *The Fourier Transform*

The Fourier transform of a 2D signal $g(x, y)$ is

$$\mathcal{F}(g)(u,v) = \int\int_{-\infty}^{\infty} g(x,y)e^{-i2\pi(ux+vy)}dxdy$$

Everything we do here can be done in arbitrary dimension, but there is no need; those who care are likely to be able to fill in the details themselves. Be aware that there are a variety of definitions in the literature, which differ by constants (a $\sqrt{2\pi}$ term moves around from definition to definition, and engineers tend to prefer to write j for $\sqrt{-1}$).

Assume that appropriate technical conditions are true to make this integral exist. It is sufficient for all moments of g to be finite; a variety of other possible conditions are available [**?**]. The Fourier transform takes a complex valued function of x, y and returns a complex valued function of u, v. It is useful to recover a signal $g(x, y)$ from its Fourier transform $\mathcal{F}(g)(u, v)$.

Definition: 11.2 *The Inverse Fourier Transform*

The inverse Fourier transform is another change of basis with the form

$$g(x,y) = \int\int_{-\infty}^{\infty} \mathcal{F}(g)(u,v) e^{i2\pi(ux+vy)} du dv.$$

Proving that this inverse works requires a fair amount of ducking and weaving to do with limits and function spaces, and I will omit a proof (you could look one up in []).

Remember this: *Each type of geometric transformation is important. You should memorize the definitions.*

Remember this: *The Fourier transform is linear.*

$$\mathcal{F}(g+h) = \mathcal{F}(g) + \mathcal{F}(h)$$
$$\textit{and for } k \textit{ any constant}$$
$$\mathcal{F}(kg) = k\mathcal{F}(g).$$

To interpret the Fourier transform of an image, think of an image as a complex valued function with zero imaginary component. For fixed u and v, think about the value of the integral as the dot product between a sinusoid in x and y and the original function. Equivalently, the value of the transform at a particular u and v measures the amount of the sinusoid with given frequency and orientation in the signal, so the Fourier transform is a change of basis.

11.2.3 Filtering with a Fourier Transform

One obvious use of a Fourier transform is to change the amount of different spatial frequencies in an image. Do this by multiplying the Fourier transform by some set of weights, then applying an inverse Fourier transform to the result. The easiest case – which will prove fruitful later – is to use weights that are large around $(u,v) = (0,0)$ and fall off as the frequency increases. A natural choice is a gaussian *in spatial frequency space.* Write

$$g_\sigma(u,v) = \frac{1}{2\pi\sigma^2} e^{\left(-\frac{(u^2+v^2)}{2\sigma^2}\right)}.$$

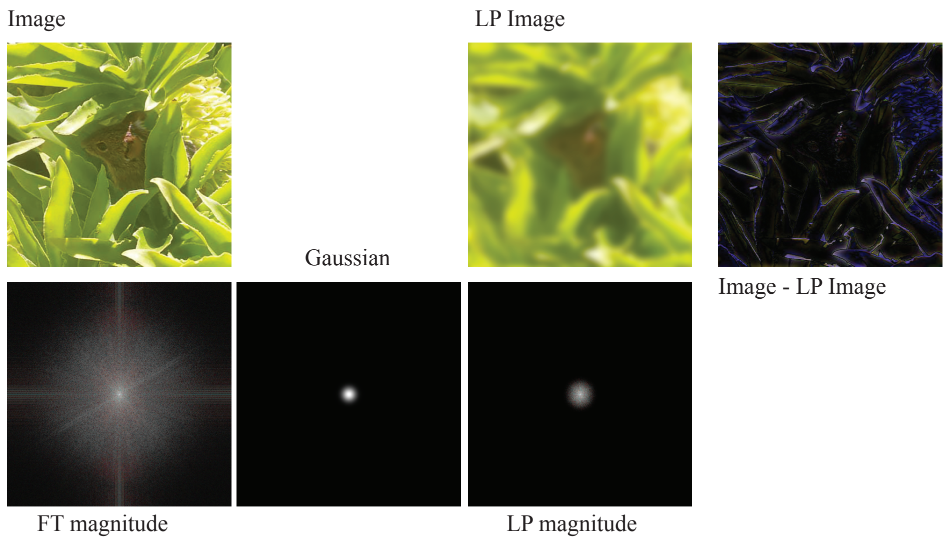

FIGURE 11.2: *On the* **top left***, the image of a four striped grass mouse with the log magnitude of its Fourier transform on the* **bottom left***.* **Center left***, the gaussian with* $\sigma = 10$ *in* u*,* v *space. This is multiplied by the weights, and the log magnitude of the result appears* **center right***.* **Above** *this is the image obtained by inverting the Fourier transform – equivalently, the low pass filtered image.* **Far left** *shows the high pass filtered image, obtained by subtracting the low pass filtered image from the original. I have* not *shown the log magnitude of the high pass filtered image, because scaling makes the result quite difficult to interpret (it doesn't look filtered). The low pass filtered version is heavily blurred, because only the lowest spatial frequencies appear in the result. Note the high pass filtered version contains what is missing from the low pass version, so has few large values which appear at edges.* Image credit: *Figure shows my photograph, taken at Kirstenbosch and Long Beach respectively.*

If σ is small, then the result of this process should have only low spatial frequencies, which will make it look blurry. The image has had a *low pass filter* applied. An alternative is to multiply the Fourier transform by $(1 - g_{\sigma}(u, v))$, which will yield an image of only high spatial frequencies (a *high pass filter*). Figure 11.2 and 11.3 show the results. Your suspicion of a strong relationship between multiplying the Fourier transform with a gaussian and convolving the image with a gaussian is correct.

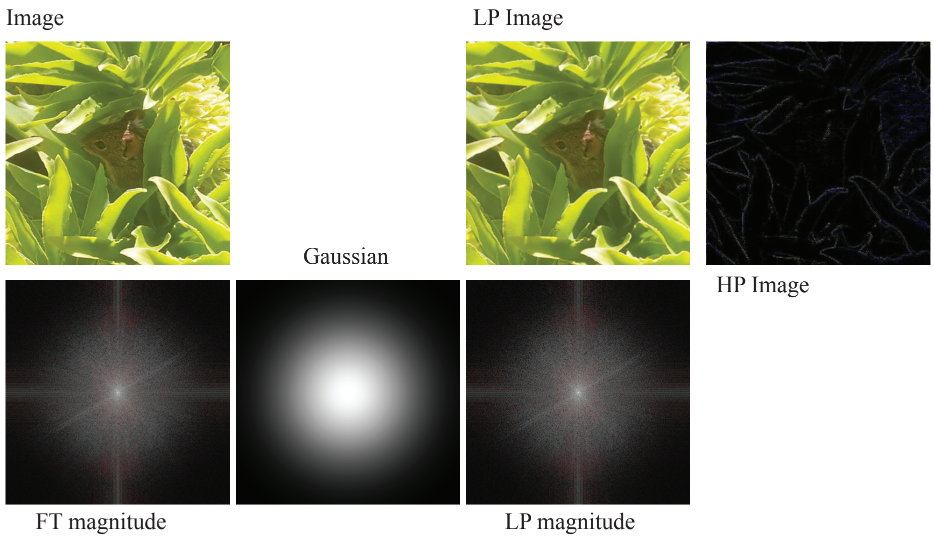

FIGURE 11.3: *On the* **top left**, *the image of a four striped grass mouse with the log magnitude of its Fourier transform on the* **bottom left**. **Center left**, *the gaussian with* $\sigma = 100$ *in* u, v *space. This is multiplied by the weights, and the log magnitude of the result appears* **center right**. **Above** *this is the image obtained by inverting the Fourier transform – equivalently, the low pass filtered image.* **Far left** *shows the high pass filtered image, obtained by subtracting the low pass filtered image from the original. I have* not *shown the log magnitude of the high pass filtered image, because scaling makes the result quite difficult to interpret (it doesn't look filtered). The low pass filtered version is less heavily blurred than that in Figure 11.2, because only the lowest spatial frequencies appear in the result. Note the high pass filtered version contains what is missing from the low pass version, so has very few large values which appear at edges.* Image credit: *Figure shows my photograph, taken at Kirstenbosch and Long Beach respectively.*

11.2.4 Phase and Magnitude

The Fourier transform consists of a real and a complex component:

$$\begin{aligned}\mathcal{F}(g(x,y))(u,v) &= \int\int_{-\infty}^{\infty} g(x,y)\cos(2\pi(ux+vy))dxdy + \\ &\quad i\int\int_{-\infty}^{\infty} g(x,y)\sin(2\pi(ux+vy))dxdy \\ &= \Re(\mathcal{F}(g)) + i * \Im(\mathcal{F}(g)) \\ &= \mathcal{F}_R(g) + i * \mathcal{F}_I(g).\end{aligned}$$

It is usually inconvenient to draw complex functions of the plane. One solution is to plot $\mathcal{F}_R(g)$ and $\mathcal{F}_I(g)$ separately; another is to consider the *magnitude* and *phase* of the complex functions, and to plot these instead. These are then called

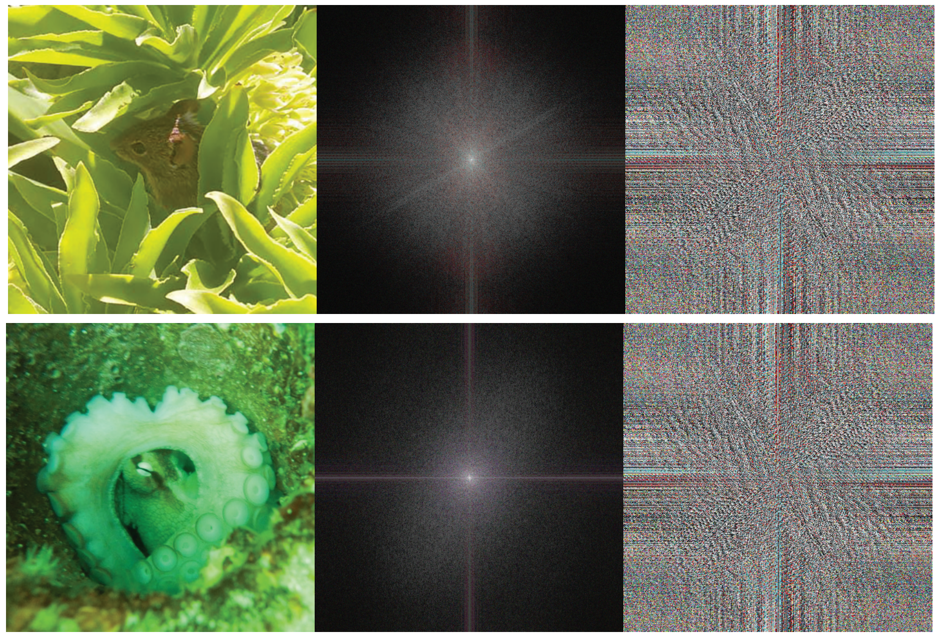

FIGURE 11.4: *On the* **left**, *images of a four striped grass mouse and an octopus;* **center**, *the log magnitude of the Fourier coefficients of the corresponding image, shown in a coordinate system where* $(0, 0)$ *is at the center of the image;* **right**, *the phase of the Fourier coefficients. The magnitude image appears monochrome because magnitudes in each color channel tend to be very similar. The phase appears colored because the phases in the color channels tend to be different. Notice that the magnitude images look quite similar, and that the phases are hard to interpret.* Image credit: *Figure shows my photographs, taken at Kirstenbosch and Long Beach respectively.*

the *magnitude spectrum* and *phase spectrum*, respectively. The magnitude spectra of images tend to be similar (look at Figure 11.4). This is a fact of nature (rather than something that can be proven axiomatically). It is related to the property that pixels mostly look like their neighbors, which is very important for denoising (Section 6.2). As a result, the magnitude spectrum of an image is surprisingly uninformative (see Figure 11.5 for an example).

Fourier transforms are known in closed form for a variety of useful cases; a large set of examples appears in **?**. I list a few in Table 11.1 for reference.

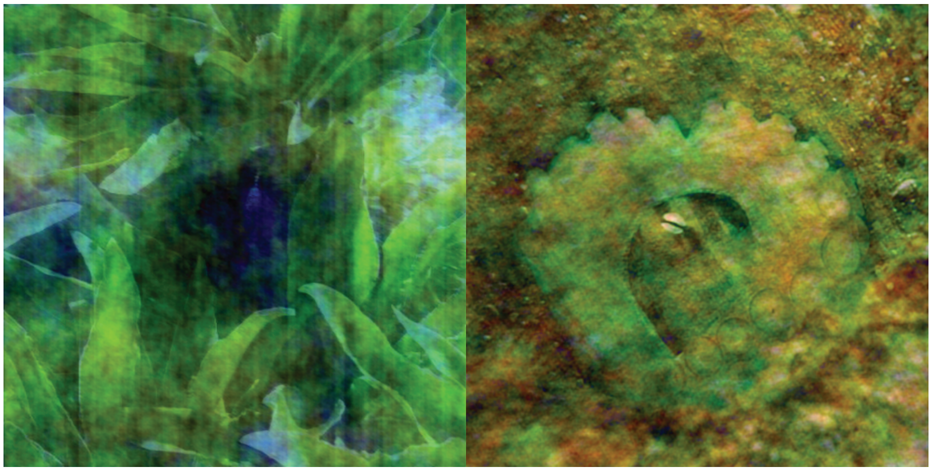

FIGURE 11.5: *Magnitudes of images tend to be the same, and most information is conveyed by phase. This is easily shown by swapping phase and magnitude for two images, applying an inverse, and looking at the result. This figure uses the images of Figure* **??**. *On the* **left**, *the phase comes from the mouse and the magnitude from the octopus; on the* **right**, *the phase comes from the octopus and the magnitude from the mouse. Although this swap leads to substantial image noise, it doesn't substantially affect the interpretation of the image, suggesting that the phase spectrum is more important for perception than the magnitude spectrum.*

Definition: 11.3 *The δ function in 2D*

The δ function in 2D is $\delta(x, y)$ and is defined by the properties that

$$\begin{aligned} \delta(x, y) &= 0 \text{if } x \neq 0 \text{ or } y \neq 0 \\ \int \delta(x, y) dx dy &= 1. \end{aligned}$$

Definition: 11.4 *The box function in 2D*

The box function in 2D is $\text{box}_2(x, y)$ and is defined by

$$\text{box}_2(x, y) = \begin{cases} 1 & |x| \leq 1/2 \text{ and } |x| \leq 1/2 \\ 0 & \text{otherwise} \end{cases} .$$

Table 11.1 contains mostly easy statements, made for reference and to save time. A few lines (2, 4, 5, 9, 12) require some care, and should be assumed true.

TABLE 11.1: *Some useful Fourier transform pairs.*

Function	Fourier transform	Tag
$f(x,y)$	$\int\int_{-\infty}^{\infty} f(x,y)e^{-i2\pi(ux+vy)}dxdy = \mathcal{F}(f)(u,v)$	1
$\int\int_{-\infty}^{\infty} \mathcal{F}(f)(u,v)e^{i2\pi(ux+vy)}dudv = f(x,y)$	$\mathcal{F}(f)(u,v)$	2
$\frac{\partial f}{\partial x}(x,y)$	$u\mathcal{F}(f)(u,v)$	3
$0.5\delta(x+a,y)+0.5\delta(x-a,y)$	$\cos 2\pi au$	4
$\cos 2\pi ax$	$0.5\delta(u+a,v)+0.5\delta(u-a,v)$	5
$e^{-\pi(x^2+y^2)}$	$e^{-\pi(u^2+v^2)}$	6
$box_2(x,y)$	$\frac{\sin u}{u}\frac{\sin v}{v}$	7
$f(ax,by)$	$\frac{\mathcal{F}(f)(u/a,v/b)}{ab}$	8
$\sum_{i=-\infty}^{\infty}\sum_{j=-\infty}^{\infty}\delta(x-i,y-j)$	$\sum_{i=-\infty}^{\infty}\sum_{j=-\infty}^{\infty}\delta(u-i,v-j)$	9
$f(x-a,y-b)$	$e^{-i2\pi(au+bv)}\mathcal{F}(f)$	10
$f(x\cos\theta - y\sin\theta, x\sin\theta + y\cos\theta)$	$\mathcal{F}(f)(u\cos\theta + v\sin\theta, -u\sin\theta + v\cos\theta)$	11
$(f*g)(x,y)$	$\mathcal{F}(f)\mathcal{F}(g)(u,v)$	12

Others are easy to derive assuming the form of the transform, that the integral exists, and so on (**exercises**).

There are a number of facts below the surface in this table. Write `swap` for the operation that swaps first and second arguments. Then

$$\mathcal{F}(f \circ \texttt{swap}) = \mathcal{F}(f) \circ \texttt{swap}$$

(use line 11 and line 2, or just do a change of variables in the integral). This means that

$$\mathcal{F}(\frac{\partial f}{\partial y}) = v\mathcal{F}(u, v)$$

(use line 3 in addition to the result about swapping).

Use a simple change of variables to get

$$\mathcal{F}(\mathcal{F}(f)) = f(-x, -y)$$

and notice that this means that, in principle, an inverse Fourier transform isn't really required (you could just Fourier transform twice, then rotate the resulting function). More interesting,

$$\mathcal{F}(\mathcal{F}(\mathcal{F}(\mathcal{F}(f)))) = f(x, y).$$

11.2.5 Fourier Transforms of Images

An image is a discrete signal. There is a version of the Fourier transform that maps discrete signals to discrete signals. In 1D, this version applies to a discrete signal where only the values at the sample points $[1, 2, \ldots, N]$ are non-zero. The Fourier transform is linear, and so is the discrete version. Viewing the Fourier transform as a change of basis should suggest that the discrete Fourier transform in 1D can be represented as multiplication by an $N \times N$ complex matrix; this is correct. However, discrete Fourier transforms can be computed very much faster than by routine matrix multiplication by careful management of intermediate values, justifying the name *fast Fourier transform* or, almost always, *FFT*. Details are out of scope.

All of this applies to 2D signals as well. Mostly, the FFT can be treated as a Fourier transform, but there are some important details to keep track of. The change of basis description should suggest to you that an $N \times N$ image will have an $N \times N$ Fourier transform, and this is the case. For most people, it is "natural" to think of the spatial frequency where $(u, v) = (0, 0)$ as lying at the center of the image, with u and v running from negative to positive values from left to right and bottom to top. For computational reasons, most API's report the FFT of an image in a rather odd coordinate system where the highest spatial frequencies are at the center and the lowest ones are at the corners. If your API does this, it will also have some form of shift command that changes the coordinate system.

11.2.6 Warnings

Fourier transforms are an extremely helpful conceptual device, and can be very powerful computational tools, but need to be approached with caution because a given Fourier transform coefficient *depends on the entire image.* Changing one pixel in an image will change some, but not all, results of a convolution with that image because convolution is local – only a window of pixels affects the results. But change one pixel in an image, and you change the whole Fourier transform. This isn't intuitive. It also means that processing images in Fourier domain is unusual.

It is much more usual to think in terms of magnitudes and phases rather than real and imaginary components of the complex values of the transform. This is mostly because the magnitude of the FFT at u, v can be interpreted as "how much" of that spatial frequency is present. Finally, the magnitude of a Fourier transform tends to have quite large dynamic range, and it is usual to show pictures of log magnitude (actually $\log(\mathsf{abs}\,(z)+1)$, to avoid problems with small numbers) rather than magnitude.

Remember this: *The functions of interest to us can be represented using a Fourier transform. The Fourier transform represents the function on a basis of sinusoids. The Fourier transform is a complex function. It is usual to look at magnitude and phase rather than real and imaginary components. Images tend to have about the same magnitude spectrum, and most image information is in the phase. Tables or change of variables can be used to obtain the Fourier transform for most cases of interest. There are fast, efficient methods to compute the Fourier transform of a discrete signal. The Fourier transform depends on every pixel in an image which makes it difficult to have accurate intuitions about what will result if you change a Fourier transformed version of an image. It is unusual to process images by taking a Fourier transform, manipulating the result, then applying an inverse Fourier transform.*

11.3 YOU SHOULD

11.3.1 remember these definitions:

11.3.2 remember these facts:

11.3.3 be able to:

- Visualize a Fourier series as a change of basis.
- Visualize a Fourier transform as a change of basis.
- Derive or remember simple Fourier transform pairs.
- Derive useful results from the table of Fourier transforms with a change of variable.

EXERCISES

QUICK CHECKS

11.1. What is the Fourier series for the function of one dimension with constant value 1?

11.2. What is the Fourier series for the function of one dimension $\sin 2\pi x$?

11.3. Look at the example of Section 11.1.2. Check that, in the limit $t \to \infty$, the solution is a constant. What is the constant?

11.4. What is the Fourier series for the function $\sin 2\pi(x + y)$?

11.5. What is the Fourier series for the function $\sin 2\pi(x + y)$?

11.6. Section 11.2.1 describes the real and complex components as $e^{i2\pi(ux+vy)}$ sinusoids on the x, y plane. Check that each term is constant when $ux + vy$ is constant.

11.7. Section 11.2.1 describes the real and complex components as $e^{i2\pi(ux+vy)}$ sinusoids on the x, y plane. Check that the frequency of each sinusoid is $\sqrt{u^2 + v^2}$.

11.8. In Section 11.2.4, I say the fact that the magnitude spectra of images tends to be similar is related to the property that pixels mostly look like their neighbors. Explain this relationship briefly.

11.9. In Section 11.2.6, I say: "change one pixel in an image, and you change the whole Fourier transform". Explain.

LONGER PROBLEMS

11.10. Some identities for Fourier transforms are (basically) just a change of variables in the integral. Here $f(x, y)$ is a function, and $\mathcal{F}(f)(u, v)$ is its Fourier transform, which you should assume exists.

(a) Show that the Fourier transform of $f(x - a, y - b)$ is

$$e^{-i2\pi(au+bv)}\mathcal{F}(f).$$

(b) Show that the Fourier transform of $f(ax, by)$ is

$$\frac{\mathcal{F}(f)(u/a, v/b)}{ab}.$$

(c) Show that the Fourier transform of $f(x\cos\theta - y\sin\theta, x\sin\theta + y\cos\theta)$ is

$$\mathcal{F}(f)(u\cos\theta + v\sin\theta, -u\sin\theta + v\cos\theta).$$

(d) Write `swap` for the operation that swaps first and second arguments. Show

$$\mathcal{F}(f \circ \texttt{swap}) = \mathcal{F}(f) \circ \texttt{swap}$$

using a change of variables in the integral.

(e) Use the result of the previous subexercise to show that

$$\mathcal{F}(\frac{\partial f}{\partial y}) = v\mathcal{F}(u, v).$$

11.11. The Fourier transform has some extraordinary properties.

(a) Use a simple change of variables to get

$$\mathcal{F}(\mathcal{F}(f)) = f(-x, -y).$$

(b) Now show

$$\mathcal{F}(\mathcal{F}(\mathcal{F}(\mathcal{F}(f)))) = f(x, y).$$

11.12. Assume that $f(x, y)$ is a real valued function, and write $\overline{z}$ for the complex conjugate of the complex number z.

(a) Show that $\overline{\mathcal{F}(f)}(u, v) = \mathcal{F}(f)(-u, -v)$.

(b) The *cross-correlation* of two real 2D functions $f(x, y)$ and $g(x, y)$ is given by

$$\text{xcorr}(f, g) = \int_{s,t} f(s, t)g(x + s, y + t)dsdt.$$

Show that

$$\mathcal{F}(\text{xcorr}(f, g)) = \overline{\mathcal{F}(f)}\mathcal{F}(g).$$

(c) The autocorrelation of $f(x, y)$ is xcorr(f, f). Show that

$$\mathcal{F}(\text{xcorr}(f, f)) = \overline{\mathcal{F}(f)}\mathcal{F}(f).$$

PROGRAMMING EXERCISES

11.13. Obtain 10 $M \times N$ color images $\mathcal{R}^{(i)}$, all of the same size (the internet might be useful; life is simpler if you choose $M = N$, and M should be at least 256).

(a) Write $\mathcal{R}_{ijk}$ for the i, j, k' th pixel value (there are three indexes because it is a color image) and $3MN$ for the total number of pixel values. Choose one image, and compute the value of

$$s(\mathcal{R}^{(1)}) = \sqrt{\frac{\sum_{ijk}\left(\mathcal{R}^{(1)}_{ijk}\right)^2}{3MN}}.$$

Now scale the other images so that $s(\mathcal{R}^{(i)}) = s(\mathcal{R}^{(1)})$. Write $\mathcal{I}^{(i)}$ for the scaled images.

(b) Use whatever API appeals (I used numpy) to compute an FFT of each scaled image. Compute the magnitude and phase for each.

(c) Write $\pi(i)$ for the i'th element of a permutation of the indices $1, \dots 10$, and IFFT$(\text{mag}_i, \text{pha}_j)$ for the image you get from an inverse FFT using the magnitude of image i and the phase of image j. Now compute

$$\sum_i \left[\| \mathcal{I}^{(i)} - \text{IFFT}(\text{mag}_i, \text{pha}_{\pi(i)}) \|_2{}^2\right]$$

(which is what happens when you swap magnitudes and keep phases) and also

$$\sum_i \left[\| \mathcal{I}^{(i)} - \text{IFFT}(\text{mag}_{\pi(i)}, \text{pha}_i) \|_2{}^2\right]$$

(which is what happens when you keep magnitudes and swap phases).

(d) Do the results confirm the claim that phases contain most of the information in the image?

(e) Write $\overline{\text{mag}}$ for the average of the magnitudes over all 10 images. Now compute

$$(1/10)\sum_i \left[\| \mathcal{I}^{(i)} - \text{IFFT}(\overline{\text{mag}}, \text{pha}_{\pi(i)}) \|_2{}^2\right].$$

Compare this to

$$(1/10)\sum_i \left[\| \mathcal{I}^{(i)} - \mathcal{I}^{(\pi(i))} \|_2^{\,2}\right]$$

(which is an estimate of the average distance between two images). What conclusions can you draw?

(f) Can you conclude that images (mostly) have the same autocorrelation function? Why?

11.14. Gaussian and Laplacian pyramids admit a natural interpretation in terms of Fourier transforms. This is most easily seen if you ignore the downsampling step. Write D_σ for the operation that smoothes an image with a gaussian of scale σ (but doesn't downsample) and G_k for the k'th layer of a gaussian pyramid. Recall an N level gaussian pyramid then can be written as:

$$\begin{aligned} G_1 &= \mathcal{I} \\ &\ldots \\ G_k &= D_\sigma(G_{k-1}) \\ &\ldots \\ G_N &= D_\sigma(G_{N-1}). \end{aligned}$$

(a) You wish to construct a this pyramid using a Fourier transform. To obtain a layer in the pyramid, you will: FFT an image; multiply the FFT by some function; then apply an inverse FFT. What function do you use to obtain the k'th layer?

(b) The previous subexercise interprets the layers of a Gaussian pyramid as the results of low-pass filtering an image, using a different low-pass filter for each layer. You now wish to construct a this pyramid using a Fourier transform. To obtain a layer in the pyramid, you will: FFT an image; multiply the FFT by some function; then apply an inverse FFT. What function do you use to obtain the k'th layer? Visualize the functions for $k = 1, 2, 3, 4$ for a 4 layer pyramid.

(c) The previous subexercise interprets the layers of a Laplacian pyramid as the results of bandpass filtering an image, using a different bandpass filter for each layer. If the origin of the FFT is at the center of the image, then the FFTs of these filters look like rather roughly like rings around the origin. Break each layer of the Laplacian pyramid into four components. Do this by multiplying the bandpass filter by a function that is 1 for one quadrant and zero for the other three quadrants. Construct the resulting pyramid for an image, and interpret its elements.

11.15. The *cross-correlation* of two real 2D functions $f(x, y)$ and $g(x, y)$ is given by

$$\text{xcorr}(f, g) = \int_{s,t} f(s, t)g(x + s, y + t)dsdt.$$

(a) Show that $\mathcal{F}(\text{xcorr}(f, g)) = \overline{\mathcal{F}(f)}$