# CHAPTER 6

# 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 change the basis in which a function is expressed, and that 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. Figure 2.8 implies that sampling errors are related to fast changes in a signal. A basis of sinusoids allows exact accounting for what is lost. Different smoothing procedures have quite different results, and a basis of sinusoids explains where Further, when you smooth a function, you suppress some frequencies. results, and this basis can be used to explain this effect. Finally, this basis shows how much smoothing is required to get good outcomes from sampling.

# 6.1 PRELIMINARIES

### 6.1.1 Fourier Series

Let f(x) be some periodic function on the range [0,1]. It is periodic, because  $f(0) = f(2\pi)$ . 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. Recall

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

which means that

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

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

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

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.

There are important and delicate matters I won't deal with in any detail here. This representation works for any periodic function we care for. 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 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. This series construction 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.

## 6.1.2 Using a Fourier Series in 1D

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

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

$$-u^{2}a_{u} = \frac{da_{u}}{dt}$$
 so  $a_{u}(t) = a_{u}(0)e^{-u^{2}t}$ .

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

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

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 discovering that a diagonal matrix gives an easy system to solve.

# 6.1.3 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

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

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

### 6.1.4 Using a Fourier Series in 1D

The previous example works here, 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

$$f(x, y, 0) = g(x, y)$$
 and  $f(x, 0, t) = f(x, 1, t)$  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

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

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

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

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.



FIGURE 6.1: 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.

## 6.2 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. This argument works for 2D as well.

#### A Basis for General Functions in 2D 6.2.1

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 complex numbers. 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 complex 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. For example, consider the real term, which is constant when ux+vyis constant (i.e., along a straight line in the x, y plane whose orientation is given by  $\tan \theta = v/u$ ). The gradient of this term is perpendicular to lines where ux + vy is constant, and the frequency of the 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 6.1.

#### 6.2.2 The Fourier Transform

As in the case of Fourier series, the change in basis is effected by integration. The procedure is now called a Fourier transform. Recall i for  $\sqrt{-1}$ , and define the Fourier transform of a 2D signal g(x, y) to be

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

(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. For all this to make sense, 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.

The Fourier transform is linear:

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

It is useful to recover a signal g(x,y) from its Fourier transform  $\mathcal{F}(g)(u,v)$ . This is another change of basis with the form

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

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 []).

## 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



FT magnitude LP magnitude

FIGURE 6.2: 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 qaussian 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.

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

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

If  $\sigma$  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 6.2 and 6.3 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.

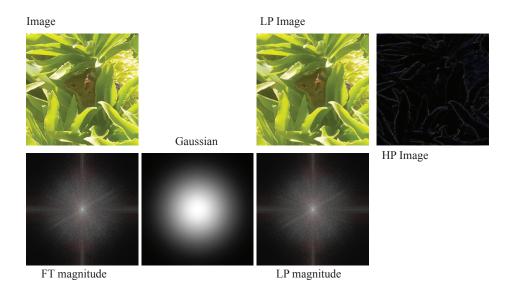


FIGURE 6.3: On the top left, the image of a four striped grass mouse with the log magnitude of its Fourier transform on the bottom left. Center left, the gaussian with  $\sigma = 100$  in u, v space. This is multiplied by the weights, and the log magnitude of the result appears center right. Above this 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 6.2, 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.

#### 6.2.4 Phase and Magnitude

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

$$\mathcal{F}(g(x,y))(u,v) = \int \int_{-\infty}^{\infty} g(x,y) \cos(2\pi(ux+vy)) dx dy + i \int \int_{-\infty}^{\infty} g(x,y) \sin(2\pi(ux+vy)) dx dy$$
$$= \Re(\mathcal{F}(g)) + i * \Im(\mathcal{F}(g))$$
$$= \mathcal{F}_{B}(g) + i * \mathcal{F}_{I}(g).$$

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

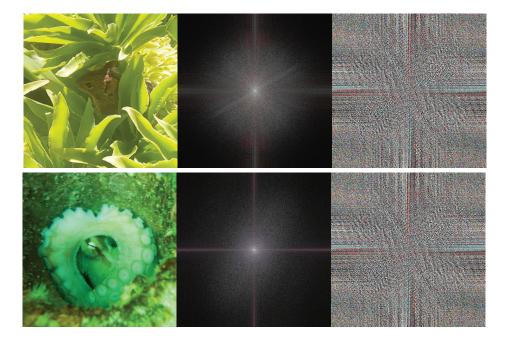


FIGURE 6.4: On the left, images of a four striped grass mouse and an octopus; center, the log magnitude of the Fourier coefficients of the corresponding image, shown in a coordinate system where (0,0) is at the center of the image; right, the phase of the Fourier coefficients. The magnitude image appears 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 tends to be similar (look at Figure 6.4). This appears to be a fact of nature, rather than something that can be proven axiomatically. As a result, the magnitude spectrum of an image is surprisingly uninformative (see Figure 6.5 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 6.1 for reference. Recall that  $\delta(x,y)$  is defined by the properties that

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

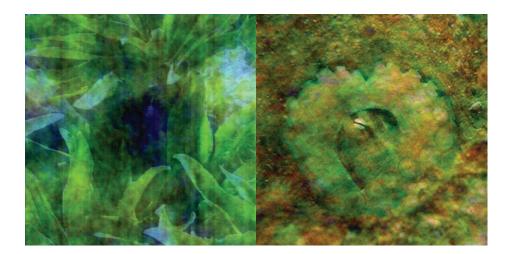


FIGURE 6.5: Magnitudes of images tend to be the same, and most information is conveyed by phase. This is easily shown by swapping phase and magnitude for two images, applying an inverse, and looking at the result. This figure uses the images of Figure??. On the left, the phase comes from the mouse and 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.

The function  $box_2(x, y)$  is defined by

$$box_2(x,y) = \begin{cases} 1 & |x| \le 1/2 \text{ and } |x| \le 1/2 \\ 0 & \text{otherwise} \end{cases}.$$

Table 6.1 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. 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. Write swap for the operation that swaps first and second arguments. Then

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

(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

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

TABLE 6.1: Some useful Fourier transform pairs.

Function	Fourier transform	Tag
f(x,y)	$\int_{-\infty}^{\infty} f(x,y)e^{-i2\pi(ux+vy)}dxdy = \mathcal{F}(f)(u,v)$	1
$\iint_{-\infty}^{\infty} \mathcal{F}(f)(u,v)e^{i2\pi(ux+vy)}dudv = f(x,y)$	$\mathcal{F}(f)(u,v)$	2
$\frac{\partial f}{\partial x}(x,y)$	$u\mathcal{F}(f)(u,v)$	3
$0.5\delta(x+a,y) + 0.5\delta(x-a,y)$	$\cos 2\pi a u$	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

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,

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

### 6.2.5 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, ..., 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.

#### 6.2.6 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  $\log(abs(z) + 1)$ , 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.

# 6.3 YOU SHOULD

# 6.3.1 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.