An Intuitive Explanation of Fourier Theory

Based on an html page by Steven Lehars, slehar@cns.bu.edu

http://sharp.bu.edu/~slehar/fourier/fourier.html

Fourier's theorem is used fairly extensively to design and simplify psychophysical experiments. Consequently, It is useful to understand some of the basic ideas behind it.

The theory is complicated mathematically. But there are some simple holistic concepts behind Fourier theory which are relatively easy to explain intuitively.

Basic Principles

Fourier theory states that any visual stimulus in space or time can be expressed as a sum of a series of spatial and temporal sinusoids. In the spatial "domain", these are sinusoidal variations in brightness across the image. For example, the sinusoidal pattern shown below can be captured in a single Fourier term that encodes 1: the spatial frequency, 2: the magnitude (positive or negative), and 3: the phase.

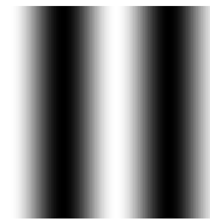

Steven Lehars

These three values capture all of the information in the sinusoidal image. The spatial frequency is the frequency across space (the x-axis in this case) along which the brightness modulates. The image on the right shows another sinusoid with a higher spatial frequency.

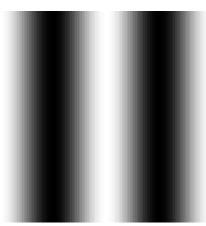

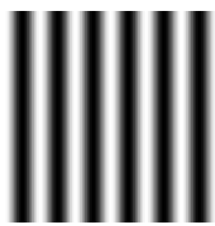

The magnitude of the sinusoid corresponds to its contrast, or the difference between the darkest and brightest peaks of the image. The phase represents how the wave is shifted relative to the origin.

Steven Lehars

A Fourier transform encodes not just a single sinusoid, but a whole series of sinusoids from high spatial frequencies (up to the "Nyquist frequency", i.e. the highest spatial frequency that can be resolved and encoded by the eye) to low spatial frequencies (down to zero, i.e. no modulation).
A signal containing only a single spatial frequency of frequency f is plotted as a single peak at point f along the spatial frequency axis, the height of that peak corresponding to the amplitude, or contrast of that sinusoidal signal.

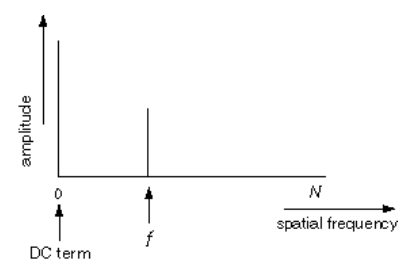

There is also a "DC term" corresponding to zero frequency, which represents the average brightness across the whole image. A zero DC term is impossible for visual stimuli because there is no such thing as a negative light. All real images have a positive DC term, as here.

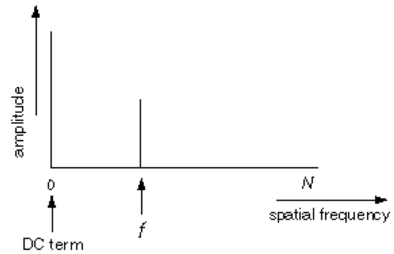

The Fourier transform also plots a mirror-image of the spatial frequency plot reflected across the origin, with spatial frequency increasing in both directions from the origin. These two plots are always mirror-image reflections of each other, with identical peaks at f and -f as shown below.

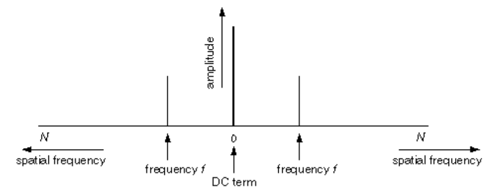

The Optical Fourier Transform

A simple lens can perform a Fourier transform in real-time. Place an image, for example a slide transparency, at the focal length of the lens, and illuminate that slide with coherent light, like a collimated laser beam. At the other focus of the lens place a frosted glass screen. That's it! The lens will automatically perform a Fourier transform on the input image, and project it onto the frosted glass screen. For example if the input image is a sinusoidal grating, as shown below, the resultant Fourier image will have a bright spot at the center, the DC term, with two flanking peaks on either side, whose distance from the center will vary with the spatial frequency of the sinusoid.

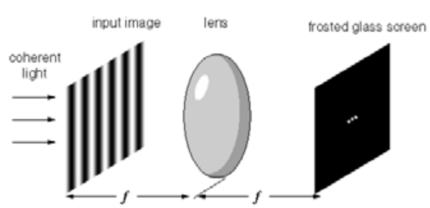

The image below shows a sinusoidal brightness image, and its two-dimensional Fourier transform, presented here also as a brightness image. Every pixel of the Fourier image is a spatial frequency value, the magnitude of that value is encoded by the brightness of the pixel. In this case there is a bright pixel at the very center - this is the DC term, flanked by two bright pixels either side of the center, that encode the sinusoidal pattern. The brighter the peaks in the Fourier image, the higher the contrast in the brightness image

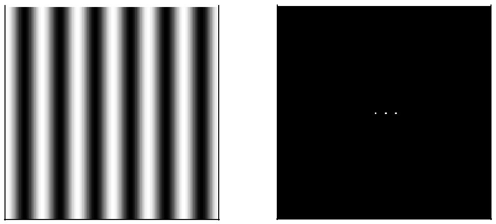

Steven Lehars

On the right is another sinusoidal brightness image, this time with a lower spatial frequency, together with it's two- dimensional Fourier transform showing three peaks as before, except this time the peaks representing the sinusoid are closer to the central DC term, indicating a lower spatial frequency.

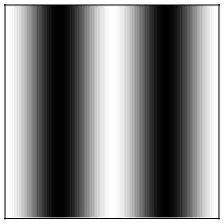

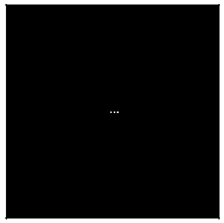

Steven Lehars

The significant point is that the Fourier image encodes exactly the same information as the brightness image, except expressed in terms of amplitude as a function of spatial frequency, rather than brightness as a function of spatial displacement. An inverse Fourier transform of the Fourier image produces an exact pixel-for-pixel replica of the original brightness image.

The orientation of the sinusoid correlates with the orientation of the peaks in the Fourier image relative to the central DC point. In this case a tilted sinusoidal pattern creates a tilted pair of peaks in the Fourier image.

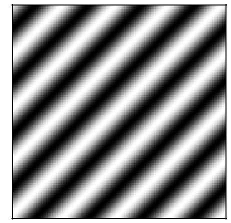

Steven Lehars

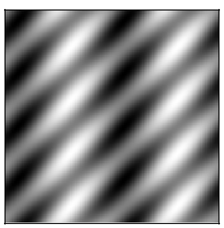

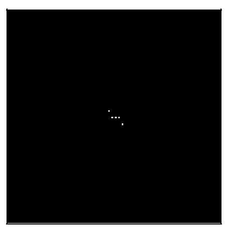

The brightness and the Fourier images are completely interchangeable, because they contain exactly the same information. The combined brightness image shown above could have been produced by a pixel-for- pixel adding of the two brightness images, or by a pixel-for-pixel addition of the corresponding Fourier transforms, followed by an inverse transform to go back to the brightness domain. Either way the result would be exactly identical

Steven Lehars

Higher Harmonics

The basis set for the Fourier transform is the smooth sinusoidal function, which is optimized for expressing smooth rounded shapes. But the Fourier transform can actually represent any shape, even harsh rectilinear shapes with sharp boundaries, which are the most difficult to express in the Fourier code, because they need so many higher order terms, or higher harmonics. How these "square wave" functions are expressed as smooth sinusoids will be demonstrated by example.

Higher Harmonics of a square wave

1 3 5 7

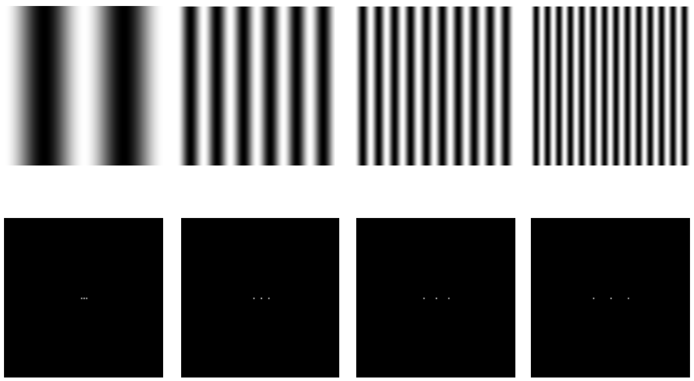

Steven Lehars

Higher Harmonics of a square wave

1 1+3 1+3+5 1+3+5+7

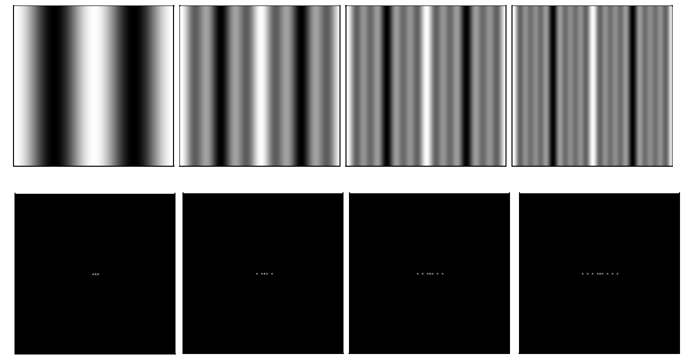

Steven Lehars

Fourier transform of a line

Brightness Image

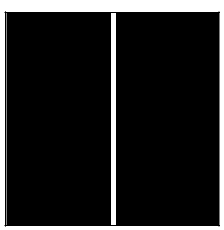

Fourier transform

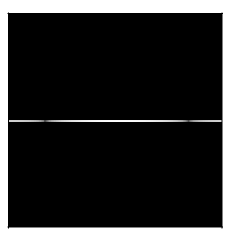

Steven Lehars

Fourier Filtering

The Fourier transform can be used to perform filtering operations to adjust the spatial frequency content of an image. We begin with an input image shown below, and perform a Fourier transform on it, then we do an inverse transform to reconstruct the original image. This reconstructed image is identical, pixel-for-pixel, with the original brightness image.

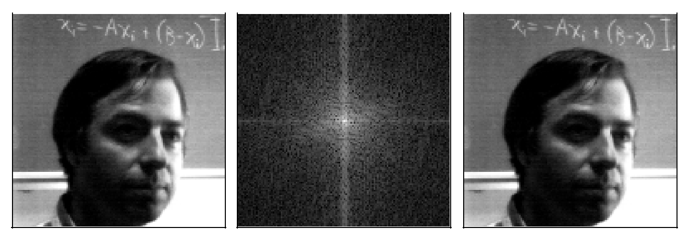

Steven Lehars

Low-Pass Filter

We begin with a low- pass filter, i.e., a filter that allows the low spatial-frequency components to pass through, but cuts off the high spatial frequencies. Since the low frequency components are found near the central DC point, all we have to do is define a radius around the DC point, and zero every point in the Fourier image that is beyond that radius. An inverse Fourier transform applied to this low-pass filtered image produces the inverse transformed image shown next.

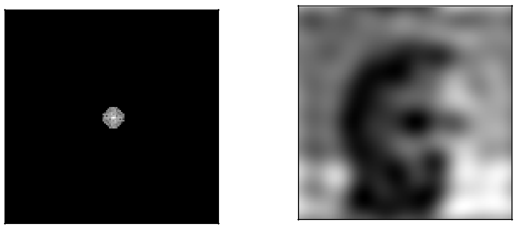

The low-pass filtered image is blurred, preserving the low-frequency, broad smooth regions of dark and bright, but losing the sharp contours and crisp edges. Low-pass filtering is equivalent to optical blurring.

Steven Lehars

High-Pass Filter

Next, high-pass filtering, where we use the same spatial frequency threshold to define a radius in the Fourier image. All spatial frequency components that fall within that radius are eliminated, preserving only the higher spatial frequency components.

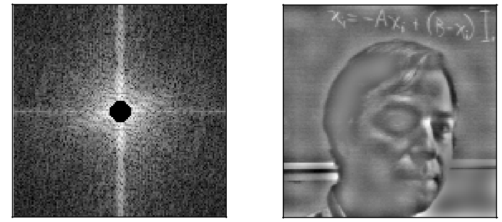

If the low-pass filtered inverse-transformed image is added pixel-for-pixel to the high-pass inverse- transformed image, this would exactly restore the original unfiltered image. These images are complementary therefore, each one encodes the information which is missing from the other.

Steven Lehars

Band-Pass Filters

Finally, band-pass filtering that preserves only those spatial frequencies that fall within a band, greater than a low cut-off, but less than a higher cut-off.

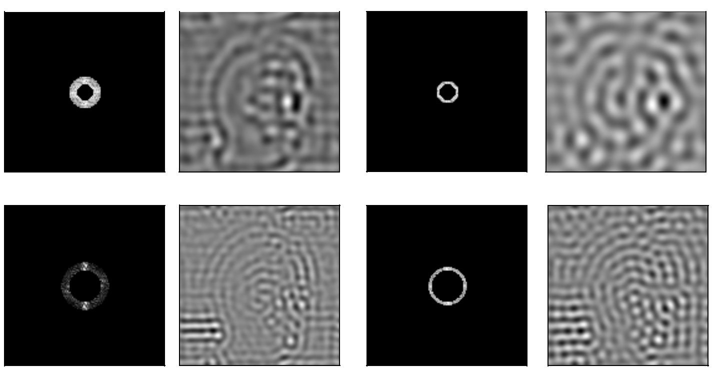

Steven Lehars