

Detecting Regular Patterns Using Frequency Domain Self-filtering

Donald G. Bailey
Physics Department, Massey University
Palmerston North, New Zealand
D.G.Bailey@massey.ac.nz

Abstract

Filtering is often used in image processing to smooth noise, and to enhance or detect features within an image. Images which have regular patterns in the spatial domain have peaks in the frequency domain corresponding to the spatial frequencies of the regular patterns. When processing such images, it is often desirable to keep such peaks, enhancing the pattern and removing noise or irregularities. This is effectively a bandpass filtering operation. The problem with such filtering is that it requires a priori knowledge of the contents of the image so that the filter can be 'tuned' to select the appropriate frequencies. Self-filtering overcomes this by multiplying the frequency domain image with its own magnitude. This gives a bandpass filter that is automatically tuned to the frequency content of the image. Applications include: detecting and enhancing regular patterns; interpolating or extrapolating regular patterns; and smoothing or reducing noise.

1. Introduction

In image processing terms, a filter is any mathematical operation that, as its name implies, modifies the information content of an image in such a way that some specified information is retained, while other unwanted information is removed. Filters are used in a wide range of applications, with the type of filter depending on the nature of the information that is to be retained or discarded.

Filters may be implemented by replacing each pixel in an image by some mathematical operation applied to the pixel values within a window or neighbourhood centred on that pixel. For computational reasons, the neighbourhood is usually small or local, but in general it is not constrained to be so.

Linear filters use a weighted average or linear combination of the pixel values within the window as the filter function. Different types of filters may be constructed by using different weights within the window

[1]. This convolution in the spatial domain may be readily implemented as a multiplication in the frequency domain.

1.1 Filtering in the frequency domain

Let $g(x, y)$ be the image being filtered, and $G(u, v)$ be its Fourier transformation. Linear filtering may be represented as

$$F(u, v) = G(u, v)W(u, v) \quad (1)$$

where $W(u, v)$ is the frequency domain filter function, and $F(u, v)$ is the Fourier transform of the filtered image.

Filtering in the frequency domain therefore corresponds to weighting the relative importance of the different spatial frequency components of an image. The particular application of the filter governs what information is important, and should be kept, and what is not, and should be discarded. The selection of an appropriate filter function requires some knowledge of how the image information is represented in the frequency domain [1].

1.2 Properties of regular patterns

Since the Fourier transform decomposes an image in terms of sinusoids, an image with a periodic pattern in the spatial domain will have distinct peaks in the frequency domain. For real images, G is conjugate symmetric, so all of the peaks and other features occur in pairs about the origin. The features of the frequency domain representation that are of particular interest for filtering are the position, amplitude, phase, shape and distribution of the peaks.

The position of a peak in G indicates the spatial frequency in g associated with that peak. The frequency is proportional to the distance of the peak from the origin, and the direction of the peak from the origin corresponds to the direction of the sinusoidal variation. This means that if an object is rotated, the peaks in G corresponding to the object will also rotate by the same amount.

The magnitude of a peak indicates the amplitude of the corresponding sinusoidal component. The position of an object in g is conveyed in the phase of its associated peaks in G . Shifting the object to a different position within the image will affect only the phase of the peaks associated with that object.

The shape of a peak contains information about the shape and size of the object within the image. This can be thought of in the following way: if the object consists of a single sinusoidal variation which is localised in the image, it can be considered as the product of a non-localised sinusoid, s , and an envelope function, $e(x,y)$, which defines the shape and size. This product corresponds to a convolution in the frequency domain

$$g = s \times e(x,y) \Leftrightarrow G = S \otimes E(u,v) \quad (2)$$

where S , consisting of a pair of peaks in the frequency domain is convolved with the Fourier transform of the envelope function, E . Therefore E specifies the shape of the peaks, while S specifies their position.

As the Fourier transform is a linear operation, if there are several independent objects within g , then G is simply the sum of the frequency components of each of the objects. This has useful implications for filtering if the frequency components are all independent. However, in general there will be overlap, with the frequency components of one object interfering with the same components of another object.

If a regular pattern in g contains sharp edges, this results in peaks positioned at integer multiples of the fundamental peak in G . In general, these harmonic peaks have amplitudes inversely proportional to the spatial frequency. A general trend in the frequency domain representation of images is that most of the energy is concentrated in the lower spatial frequencies regardless of whether or not there are observable regular patterns in the image.

Random noise in an image has random pattern in frequency domain. White noise (uncorrelated additive Gaussian noise) has a uniform probability density function in the frequency domain. Any object peaks with amplitudes below this noise level are unable to be distinguished from the noise. Patterned noise, however, will also have distinct peaks corresponding to the regular pattern of the noise.

1.3 Filtering regular patterns

Examples of different filtering applications involving regular patterns are detecting the regular patterns, suppressing noise, interpolating missing data and extrapolating data from regular patterns.

Pattern detection involves detecting specific combinations of spatial frequencies corresponding to the patterns being detected. By amplifying the peaks associated with the pattern being detected and attenuating the remaining areas, the pattern is enhanced at the expense of other unwanted information. The problem with this approach is that it is necessary to know beforehand where the peaks of interest will be located. This requires knowing the exact scale and orientation of the object since both these factors influence the positions of the peaks in the frequency domain.

Another application is the suppression of random noise. Since most of the information in an image is concentrated in the lower spatial frequencies, the simplest approach is to attenuate the high frequency components. The disadvantage with this method is that the higher frequency harmonics associated with edges in the image are also attenuated. This has the effect of blurring the edges in the image. A better approach when filtering images containing regular patterns is to attenuate where there are no peaks.

Filters may also be used for interpolating or extrapolating regular patterns within an image. Since the shape of a peak contains the shape and size information, modifying the shape of the peak will modify the shape and size of the pattern in the image. More specifically, making the peak narrower will extend the regular pattern in the image, providing some degree of extrapolation.

2. Frequency domain self-filtering

All of these applications require information on the positions of the peaks in the frequency domain for the formulation of appropriate filters.

The problem becomes given G , how do we determine an appropriate filter function W ? One approach is to use a self-filter, where the image itself in some way defines the filter function. The simplest such self-filter is to use the magnitude of the Fourier transformed image as the filter function. That is

$$W(u,v) = |G(u,v)| \quad (3)$$

2.1 Properties

This filter has a number of useful properties. First, it is a zero phase filter. That is, when it is applied, it does not affect the phase of G . As a result, the position information contained in the phase is retained. Second, since G is conjugate symmetric, W will also be conjugate symmetric. This means that the filtered image $f(x,y)$ will also be real.

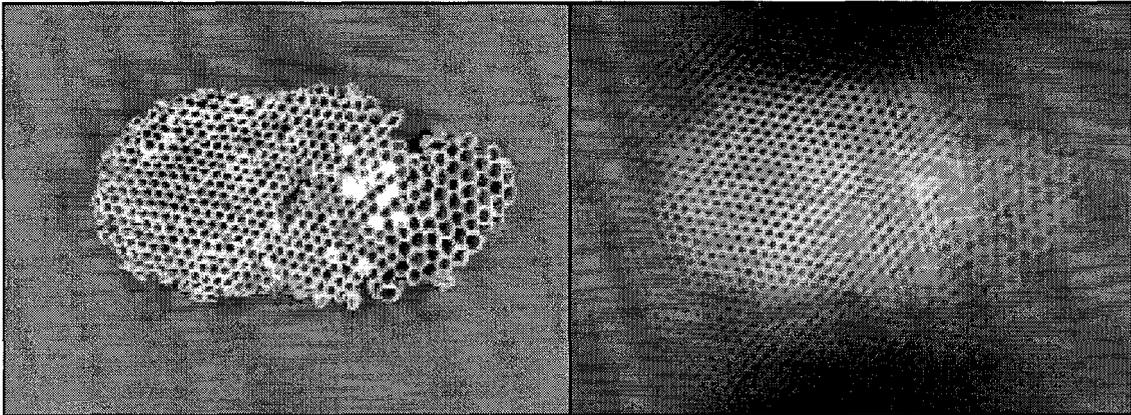


Figure 1: Detecting regular patterns. Before (left) and after (right) filtering.

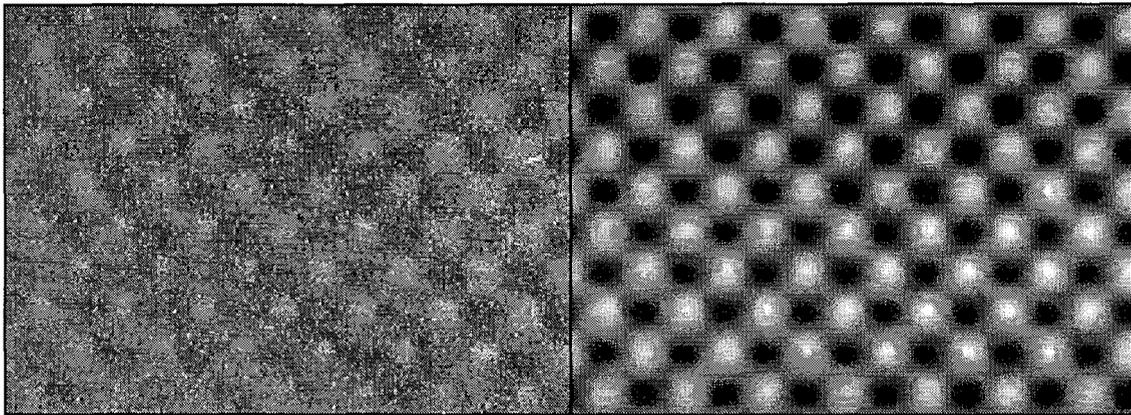


Figure 2: Suppression of noise. Before (left) and after (right) filtering.

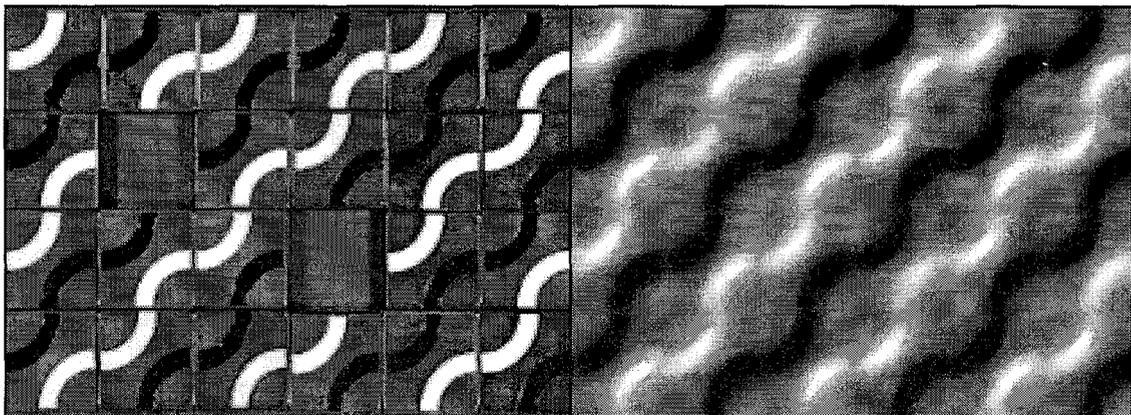


Figure 3: Interpolating missing features in regular patterns. Before (left) and after (right) filtering.

The filter weights peaks according to their strength. Therefore regular patterns in an image are amplified, and the components of the image which do not have a pattern are attenuated. Since higher weighting is given to the higher peaks, the filter enhances the stronger patterns more than weaker patterns. This effect is illustrated clearly in figure 1, where the cells in the on the left edge of the image are enhanced more than the cells on the right.

Since there is a larger area of cells on the left, the peaks corresponding to those cells have greater energy and are therefore amplified more.

Since the filter is amplifying the peaks and attenuating between the peaks, it is good at suppressing random noise. In figure 2, the noise in the image associated with the bubbles in the foam is almost completely removed.

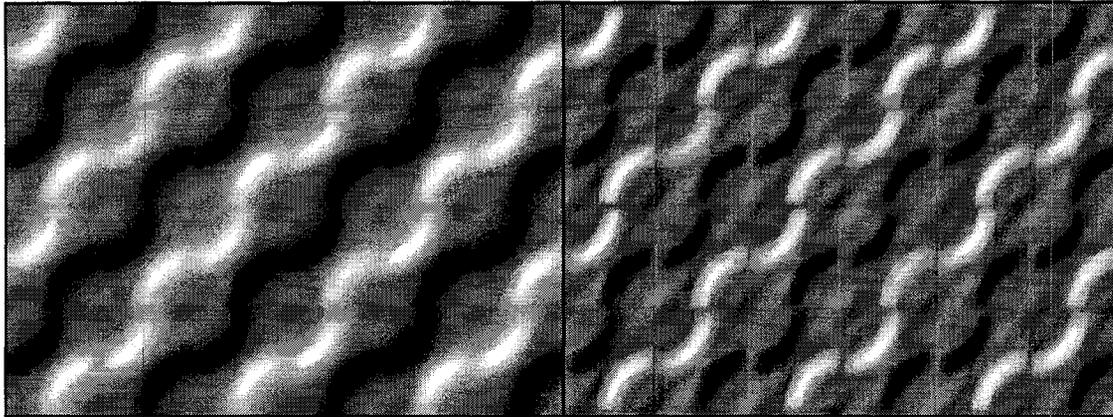


Figure 4: Interpolation without (left) and with (right) frequency compensation for sharp edges.

The frequency domain self-filter is also applicable to extrapolation and interpolation applications as illustrated in figure 3. Since the centre of the peak has the highest weight, it is amplified more than the flanks. The overall effect of this is for the peak to become narrower in relation to its height. Empirically, this will cause the size and shape of the object to grow in the image, providing extrapolation past the previous boundaries, or interpolation into holes.

2.2 Limitations and extensions

One disadvantage of the filter represented by equation (3) is that it will also blur sharp edges. Sharp edges result in a series of peaks weighted inversely proportional to the spatial frequency, therefore the peaks associated with the high frequency harmonics will be attenuated relative to the fundamental peak. This problem may be overcome by weighting the filter function by spatial frequency to keep the peaks in their correct relative proportions:

$$W(u, v) = \sqrt{u^2 + v^2} |G(u, v)| \quad (4)$$

Figure 4 shows clearly the effect of increasing the weight to the higher frequencies. Note that although the sharp edges are retained, there is not a significant increase in the noise within the image.

The most significant limitation of frequency domain self-filtering is that if there are two significant patterns in the image then the stronger pattern will be enhanced more. A related difficulty is that self-filtering is ineffective at removing pattern noise. Such noise is indistinguishable from the regular patterns which are of interest. If some information is known about the two patterns beforehand, then it is possible to weight the self filter to selectively amplify the pattern of interest or to remove the pattern noise.

A problem may be encountered when using the FFT to perform the self-filtering. Since the FFT is periodic, if there is an intensity gradient across the image, this will appear as a low frequency sawtooth wave. This regular pattern is amplified by filtering and can mask the effects of any other regular patterns in the image. These effects may be minimised by the appropriate use of preprocessing and windowing techniques [2].

Finally, the filter as presented here is that it distorts the shape of the object. In applications where the shape of the object must be retained, it is necessary to amplify all of a peak by the same value, rather than giving greater amplification at the centre. This may be accomplished to a limited degree by extending the maximum of each peak out a few pixels, retaining more of the shape of the flanks of the peak while filtering.

Although the frequency domain self-filter is described here in terms of linear filter theory, it is not actually a linear filter. Since the filter depends on the image being filtered, in general filtering two images separately and adding the results is not equivalent to filtering the sum of the two images.

3. Summary

It has been shown that self-filtering is effective at enhancing and detecting regular patterns within an image, suppressing noise, and interpolating missing features within the image. Overall, the frequency domain self-filter may be considered as a bandpass filter which is automatically tuned to the dominant frequency content of the image.

References

- [1] Gonzalez, R.C. and P. Wintz, P. *Digital Image Processing*. Addison-Wesley, 1987.
- [2] Brigham, E.O. *The Fast Fourier Transform*. Prentice-Hall, 1974.