Frequency Domain Self-filtering for Pattern Detection

Donald G Bailey

Image Analysis Unit, Massey University
Palmerston North
Ph: +64 6 350 4063, Fax: +64 6 354 0207
E-mail: 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.

The approach proposed here to overcome this problem is frequency domain self-filtering. By multiplying the frequency domain image with its own magnitude, peaks in the frequency domain are enhanced based on the strength of those peaks. Regions of low activity in the frequency domain are attenuated relative to the peaks. This gives a bandpass filter which is automatically tuned to the frequency content of the image.

Applications of such a filter include: detecting and enhancing regular patterns; interpolating or extrapolating the regular pattern to regions in the image where it is not present; and smoothing or reducing noise in the image.

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. Examples include noise suppression, pattern detection, edge enhancement or detection, interpolation, and extrapolation.

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.

One class of filters, linear filters, uses a weighted average or linear combination of the pixel values within the window. Different types of filters may be constructed by using different weights within the image [1]. Mathematically, this may be represented as

$$g(x,y) = \sum_{i=-m/2}^{m/2} \sum_{j=-n/2}^{n/2} f(x-i, y-j) w(i,j) \qquad (1)$$

where $f(x,y)$ is the image before filtering, $g(x,y)$ is the result of filtering, and $w(i,j)$ are the weights within an mxn window. This operation is just a discrete convolution with a finite sized kernel (w). It can be shown that such a convolution in the spatial domain corresponds to a multiplication in the frequency domain. That is

$$f(x,y) = g(x,y) \otimes w(x,y) \quad \Leftrightarrow \quad F(u,v) = G(u,v)W(u,v) \tag{2}$$

where $\otimes$ denotes convolution and $F(u,v)$, $G(u,v)$, and $W(u,v)$ are two dimensional Fourier transforms of f, g, and w respectively. Equation (2) suggests an alternate implementation of linear filters: take the Fourier transform of an image, multiply by a filter function W, and inverse Fourier transform the result. The selection of an appropriate filter function requires some knowledge of how the image information is represented in the frequency domain [1].

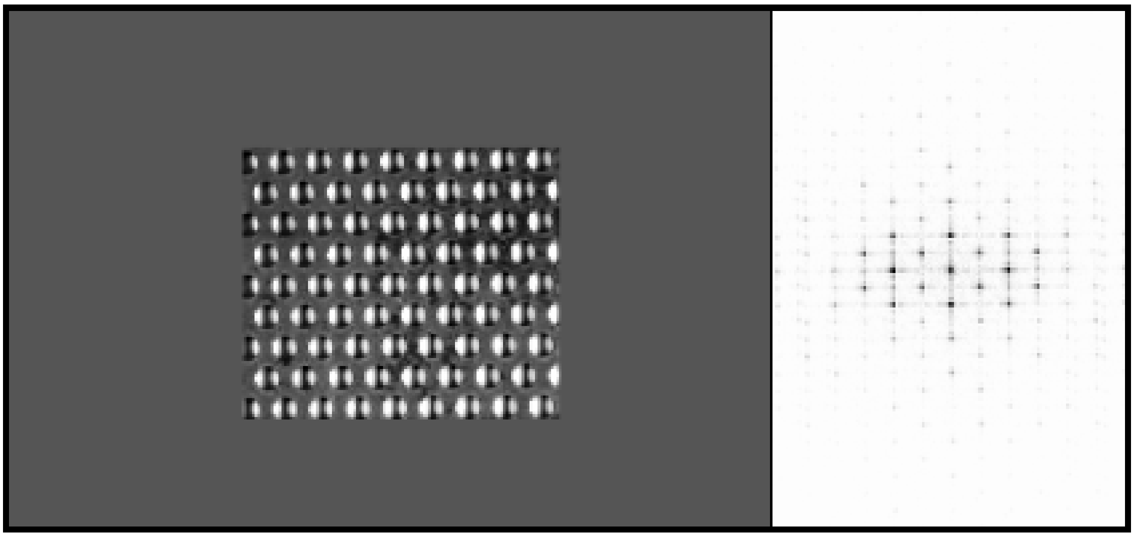

Figure 1: An image containing a localised regular pattern and its frequency domain representation.

1.1 Information in the frequency domain

In general, the frequency domain representation, $G(u,v)$, of an image, $g(x,y)$, will be complex. However, for real images, it can be shown that G is conjugate symmetric. That is

$$G(u,v) = G^*(-u,-v) \tag{3}$$

where * denotes complex conjugate. An image with a periodic pattern in the spatial domain will have distinct peaks in the frequency domain. Each peak represents a sinusoidal component of the periodic pattern. This is illustrated in figure 1. Note that, because of conjugate symmetry, 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

$$g(x,y) = s(x,y)e(x,y) \quad \Leftrightarrow \quad G(u,v) = S(u,v) \otimes E(u,v) \tag{4}$$

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 (or harmonics) of the fundamental peak in G. In general, these harmonic peaks have amplitudes inversely proportional to the spatial frequency. In fact, 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.

1.2 Frequency domain filtering

Filtering in the frequency domain 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 of importance, and should be kept, and what is not, and should be discarded. Examples of different filtering applications will be described.

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. Note however, that not all images have regular patterns, and therefore do not always have discernible peaks in the frequency domain.

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 above applications require information on the positions of the peaks in the frequency domain for the formulation of appropriate filters. One approach to this problem is to use the peaks of G to define the filter function. This introduces the concept of 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 in equation (2). That is

$$W(u,v) = |G(u,v)| \tag{5}$$

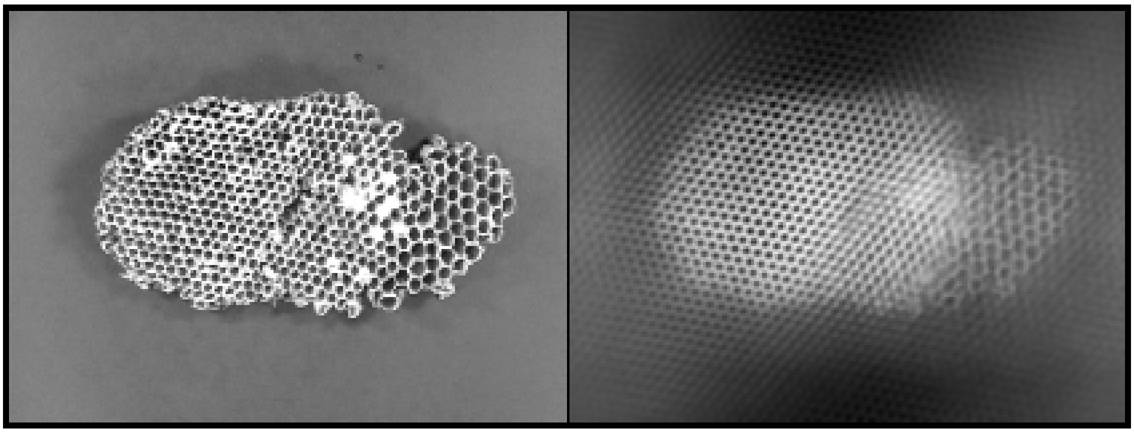

Figure 2: Detecting regular patterns. Before (left) and after (right) filtering.

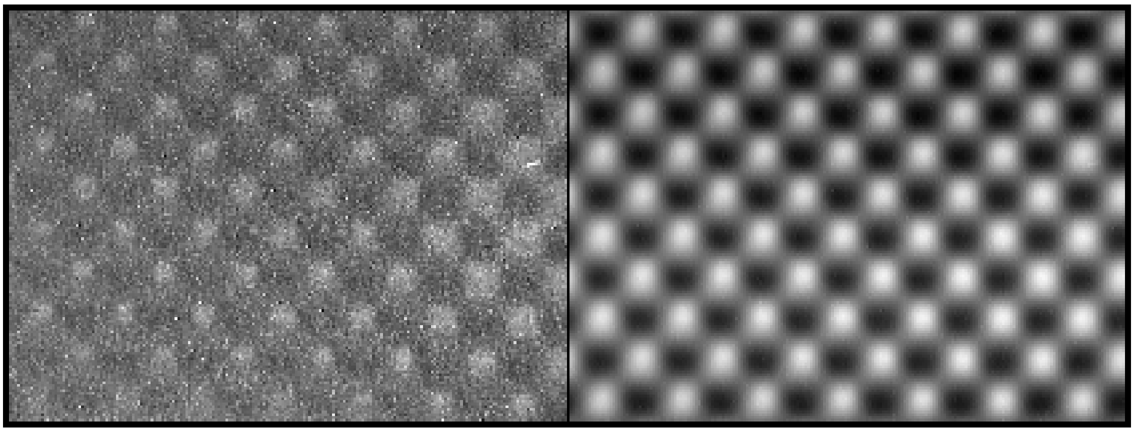

Figure 3: Suppression of noise. Before (left) and after (right) filtering.

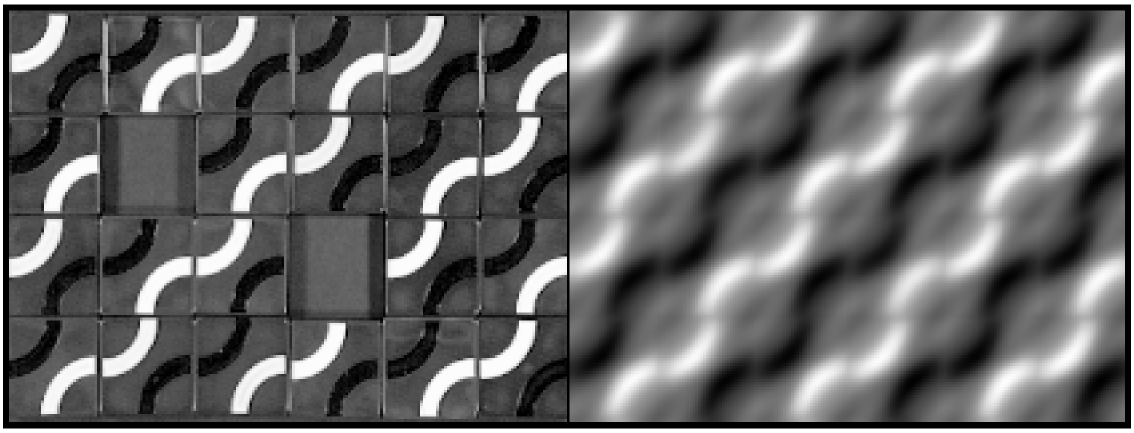

Figure 4: Interpolating missing features in regular patterns. Before (left) and after (right) filtering.

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.

The effect of frequency domain self-filtering is illustrated for three filtering applications in figures 2 to 4.

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 2, 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 amplified more.

Since the filter is amplifying the peaks and attenuating between the peaks, it is good at suppressing random noise. In figure 3, the noise in the image associated with the bubbles in the foam is almost completely removed.

The frequency domain self-filter is also applicable to extrapolation and interpolation applications as illustrated in figure 4. 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 Retention of sharp edges

One disadvantage of the filter represented by equation (5) is that it will also blur sharp edges. That is because sharp edges result in a series of peaks weighted inversely proportional to the spatial frequency. This means that 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. This modified self-filter may be represented by equation (6).

$$W(u,v) = \sqrt{u^2 + v^2}\,|G(u,v)| \qquad (6)$$

The effect of this modification is shown in figure 5. Note that although the sharp edges are retained, there is not a significant increase in the noise in the image.

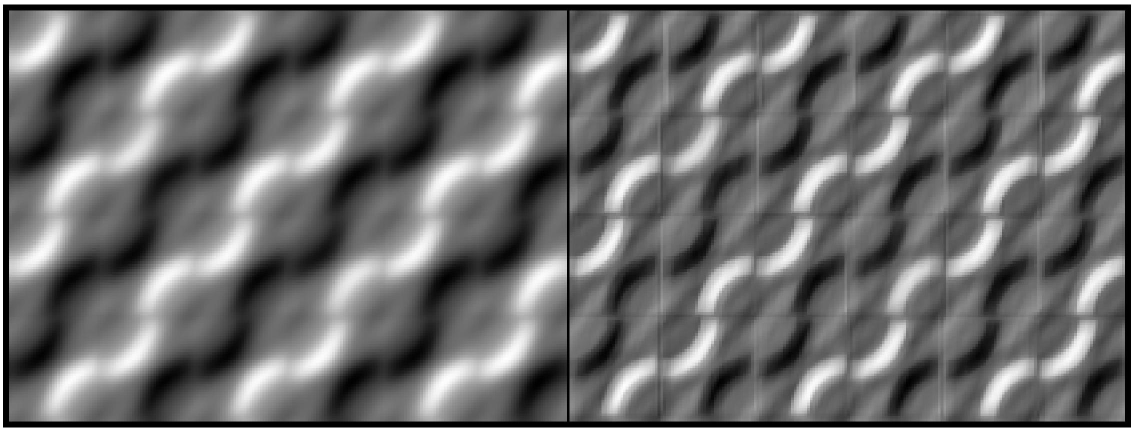

Figure 5: Trax tiles from figure 4 filtered with equation (5) (left) and equation (6) (right).

2.3 Relationship with autocorrelation

Although the frequency domain self-filter is related to autocorrelation, there are some important differences. Self-filtering uses a zero phase filter so the phase information in the frequency domain, or the positional information is not affected. Autocorrelation uses the complex conjugate of G as the filter function.

$$W(u,v) = G^*(u,v) \qquad (7)$$

When this filter is applied, the phase of the filter cancels exactly the phase of the image, causing the location information to be lost.

2.4 Equivalent convolution kernel

The self-filter has been represented so far in terms of its frequency domain representation. Equation (2) allows the filter to be represented as a convolution. The equivalent convolution kernels used in figures 2 and 3 are shown in figure 6. Note that if implemented using the Fast

Fourier Transform without padding, the convolution represented by equation (1) will need to be performed with wrapping from one side of the image to the other (that is the positions will need to be calculated modulo N where N is the image size) [2].

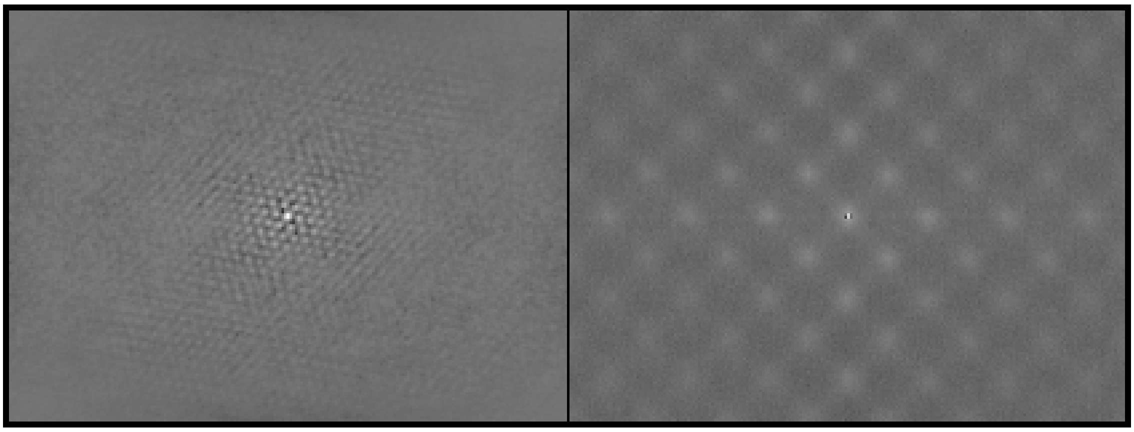

Figure 6: Equivalent convolution kernels used in figures 2 and 3.

2.5 Limitations

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. This can present a problem in a number of ways.

If the weaker pattern is of interest, it will not be amplified as much as the stronger pattern. If some information is known about the two patterns beforehand, then it is possible to weight the self filter to selectively amplify the pattern if interest. In general terms, this may be represented as

$$W(u,v) = H(u,v)\left|G(u,v)\right| \tag{8}$$

Equation (6) is a special case of this where more weight is given to the higher frequencies to prevent blurring of edges.

A related difficulty is that self-filtering is ineffective in removing pattern noise. Such noise is indistinguishable from the regular patterns which are of interest. Again, if the spatial frequency content is known beforehand, it is possible to weight against the noise as in equation (8).

A related problem mat 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].

Another limitation of the filter as presented here is that it distorts the shape of the object. To prevent this, filter function needs 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 presented here as a linear filter, the frequency domain self-filter is not a linear operation. Since the filter depends on the image being filtered, in general filtering two images separately and adding the results in not equivalent to filtering the sum of the two images. In other words

$$G_1(x,y)\left|G_1(x,y)\right| + G_2(x,y)\left|G_2(x,y)\right| \neq \left(G_1(x,y) + G_2(x,y)\right)\left|G_1(x,y) + G_2(x,y)\right| \tag{9}$$

This is because the images, when filtered individually, are actually filtered with different filters. When filtered in combination, a third filter is used.

3. Summary

Many filtering applications require information about the position of peaks in the frequency domain. If this information is not precisely known beforehand, then it is difficult to devise a filter tailored to the application. One approach around this problem is to use the image itself to specify the filter. As a result, a frequency domain self-filter was devised which uses the magnitude of the image in the frequency domain as the filter function.

It has been shown that such a filter is effective at enhancing and detecting regular patterns within an image, suppressing noise, and interpolating missing features within the image. Patterns are enhanced by selectively amplifying frequency domain features according to their strength. Random noise is suppressed since information between the peaks in the frequency domain is suppressed. As the shape of the peaks in the frequency domain contain information on the shape and size of the patterns, making the peaks narrower extends the size of the pattern in the image, making it useful for extrapolation and interpolation.

One disadvantage of the filter is that it tends to blur sharp edges. Its performance may be improved in this regard by compensating of the drop in amplitude of the harmonic peaks with increasing frequency. A limitation of the filter in applications where there are multiple patterns in an image, is that the patterns are amplified according to the amplitude of the peaks produced in the frequency domain. This means that weaker patterns will be suppressed relative to stronger patterns. In situations where the weaker patterns are those of interest, it is necessary to modify the filter to improve the gain of those patterns.

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 **R.C. Gonzalez and P. Wintz.** *Digital Image Processing.* Addison-Wesley, 1987.

2 **E.O. Brigham.** *The Fast Fourier Transform*. Prentice-Hall, 1974.