

8 Filtering

The estimated spectrum of a time series gives the distribution of variance as a function of frequency. Depending on the purpose of analysis, some frequencies may be of greater interest than others, and it may be helpful to reduce the amplitude of variations at other frequencies by statistically filtering them out before viewing and analyzing the series. For example, the high-frequency (year-to-year) variations in a gauged discharge record of a watershed may be relatively unimportant to water supply in a basin with large reservoirs that can store several years of mean annual runoff. Where low-frequency variations are of main interest, it is desirable to *smooth* the discharge record to eliminate or reduce the short-period fluctuations before using the discharge record to study the importance of climatic variations to water supply. *Smoothing* is a form of filtering which produces a time series in which the importance of the spectral components at high frequencies is reduced. Electrical engineers call this type of filter a *low-pass* filter, because the low-frequency variations are allowed to “pass through” the filter. In a low-pass filter, the low frequency (long-period) waves are barely affected by the smoothing.

It is also possible to filter a series such that the low-frequency variations are reduced and the high-frequency variations unaffected. This type of filter is called a *high-pass* filter. Detrending is a form of high-pass filtering: the fitted trend line tracks the lowest frequencies, and the residuals from the trend line are what remains after the lowest frequencies have been removed. A third type of filtering, called *band-pass* filtering, reduces or filters out both high and low frequencies, and leaves some intermediate frequency band relatively unaffected.

In this lesson, we cover several methods of smoothing, or *low-pass* filtering. We have already discussed how the cubic smoothing spline might be useful for this purpose. Four other types of filters are discussed here: 1) simple moving average, 2) binomial, 3) Gaussian, and 4) windowed (Hamming method). Considerations in choosing a type of low-pass filter are the filter length and desired frequency response of the filter.

8.1 Mathematical operation

A smoothed time series value is “merely an estimate of what the value in the series would be if undesired high frequencies were not present” (Panofsky and Brier 1958, p. 147). A statistical filter, or digital filter, is a series of weights that when cumulatively multiplied by consecutive values of a time series gives the filtered series. The series of weights is sometimes called the

filtering function, or simply the *filter*. The operation of filtering is illustrated in Table 1.

Year	Filter	Time Series	Filtered Values
1		12	
2	.25 x	17	14.00
3	.50 x	10	14.75
4	.25 x	22	17.25
5		15	15.75
6		11	13.75
7		18	18.50
8		27	21.50
9		14	

Assume that the numbers 12, 17,...,14 in column three of the table are a time series, and that the filter has weights 0.25, 0.50, 0.25. The filtered values are the cumulative products of the weights and the original time series. Filtering proceeds by sliding the filter alongside the time series one value at a time, each time computing a cumulative product. For example, in Table 1, the filter is centered on year 3, such that the filtered value for year 3 is computed from the series at times 2, 3 and 4 as follows:

$$(0.25)(17) + (0.50)(10) + (0.25)(22) = 14.75$$

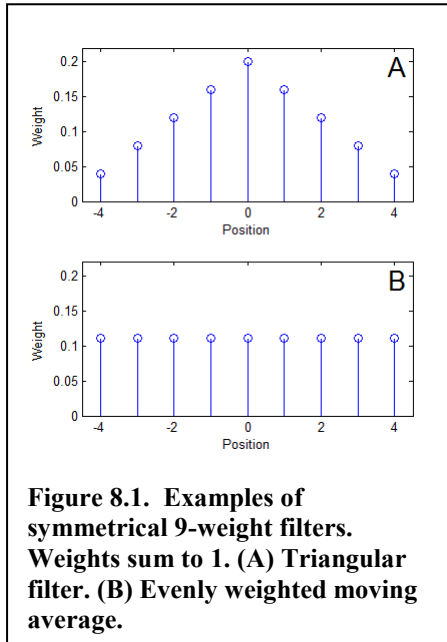


Figure 8.1. Examples of symmetrical 9-weight filters. Weights sum to 1. (A) Triangular filter. (B) Evenly weighted moving average.

The filtering can be described by the equation

$$s_t = \sum_{i=-n}^n w_i x_{t+i} \quad (1)$$

where x_t is the original time series,

$w_i, i = -n, (-n+1), \dots, 0, 1, \dots, n$ are the weights, with central weight w_0 , and s_t is the smoothed, or filtered, series. The filtered value is assigned to the year corresponding to the central value of the sliding weights, so that features in the smoothed series are not shifted relative to their position in the original series. Usually the weights are fractional values whose sum is one: this guarantees that the mean of the filtered series approximately equals the mean of the original series. The *filter length* is the total number of weights. A filter is *symmetrical* if the weights to left of the central weight are the same as those to the right of the central weight (Figure 8.1). For example, the filter used in Table 1 is symmetrical because the same weight (0.25) flanks the central weight on either side. Symmetry of the filter weights is important to avoid phase shifts (see

frequency response) in filtering. For a filter with a central weight and n weights to either side, the *filter length* is

$$N = 2n + 1 \quad (2)$$


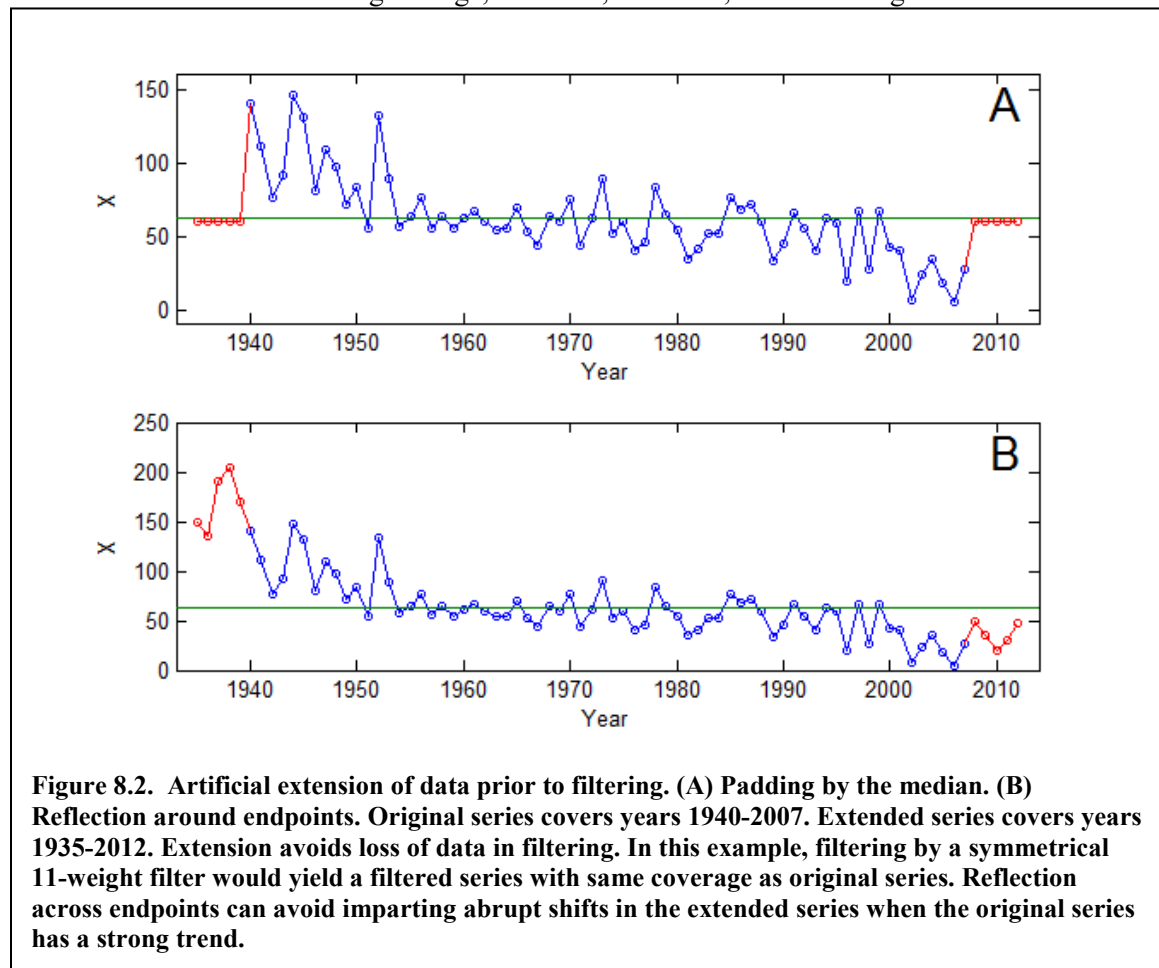
The filtered series in Table 1 is shorter than the original time series because of the loss of starting and ending values. For the example, the first available filtered value is centered on the *second* observation in the original time series, and the last filtered value is centered on the *next-to-last* original time series observation. For a filter length (odd) of N , a total of $(N-1)/2$ values are lost off the front and back of the series because of the requirement for startup values. For example, $(N-1)/2 = (3-1)/2 = 1$ value is lost from each end in applying the three-weight filter in Table 1.

Sometimes the original time series is extended forward and backward artificially before filtering so that the filtered series covers the same observations as the original series. Because no “real” data exist outside the ends of the original time series, this procedure can lead to disagreeable *end effects* in the filtered series. Two commonly used extension methods are 1) *substituting the long-term mean or median*, and 2) *reflecting the data across the end points* (Figure 8.2).

8.2 Frequency response

The frequency response of a filter describes the effect of the filter on sinusoidal inputs at different frequencies. The frequency response has two components – amplitude and phase. The phase at a given frequency describes the shift in the position of a wave at that frequency along the time axis. For some filtering applications, it is desirable that the phase is zero, so that peaks and troughs representing waves in the original data are not shifted in the filtered data. Filters for which the phase of the frequency response is zero at all frequencies are called *zero-phase*. Symmetrical (example in Table 1) filters are zero-phase. Other types of symmetrical filters to be

discussed below are the moving average, binomial, Gaussian, and Hamming filters.




The other component of the frequency response of a filter is the amplitude. The amplitude of the frequency response at a given frequency is the ratio of the amplitude of an output sine wave to an input sine wave at that frequency (Figure 8.3). Smoothing filters have little effect on the gradual, or low-frequency, variations while damping or more-or-less removing the high-frequency variations. For this reason, smoothing filters are also called “low pass” filters. The amplitude of frequency response as a function of frequency is sometimes called the *frequency response function* or just the *response function* (Figure 8.4).

Symmetrical digital filters such as the filter in Table 1 are examples of *finite impulse response (FIR)* filters. The *impulse response* of a system is defined as the output for a unit

pulse of input restricted to a single time step. An FIR filter has the property that if the input series

has just a unit departure at one specific time, the response in the filtered series is restricted to a finite number of times. For example, the response to a unit impulse for the filter in Table 1 would be distributed over three time points.



The frequency response function of an FIR filter can be computed as the Fourier transform of the filter weights. For a symmetrical FIR filter, the frequency response function can be written as


$$u(f) = w_0 + 2 \sum_{k=1}^n w_k \cos(2k\pi f \Delta t) \quad (3)$$

where $u(f)$ is the frequency response, f is frequency, w_k is the k^{th} weight numbered outward from the central weight w_0 , and Δt is the data interval, the time between successive observations in the time series (Panofsky and Brier 1958, p. 149).

8.3 Simple moving average

An example of a symmetrical low-pass filter is the *simple moving average* filter of length N , where N is an odd integer. The individual weights of the moving average are equal to $1/N$, so that the sum of the weights is $N(\frac{1}{N}) = 1$. An example of a simple moving average filter is the 9-weight moving average $\left\{ \frac{1}{9} \quad \frac{1}{9} \quad \frac{1}{9} \quad \frac{1}{9} \quad \frac{1}{9} \quad \frac{1}{9} \quad \frac{1}{9} \quad \frac{1}{9} \quad \frac{1}{9} \right\}$, whose weights are plotted in Figure 8.1B. Application of the N -weight moving-average filter is equivalent to computing a sample mean for each subset of N values. The simple moving-average filter is therefore also called the *running mean*. The running mean has the practical advantage of simplicity over some other types of filter.

The frequency response of the running mean of length N is 1.0 at the lowest frequency $f = 0$, corresponding to infinite wavelength, and decreases to 0 at $f = 1/N$, corresponding to a wavelength the same as the filter length. Thus, for example, the frequency response of a nine-year running mean decreases from 1.0 at a wavelength of infinity to 0 at a wavelength of 9 years



(Figure 8.5). But the frequency response of the running mean has an undesirable property: for wavelengths shorter than the filter length, N – or frequencies higher than $1/N$ – the frequency response oscillates rather than stays close to zero. While the frequency response is zero at frequencies that are multiples of $1/N$ (e.g., $f=1/N$, $f=2/N$, ...), it is nonzero at intermediate frequencies (Figure 8.5). Consequently, high-frequency variations in the original series at those intermediate frequencies are *not* removed. For example, a series smoothed with a 9-year running mean filter might retain considerable variance at $f=0.15$ (wavelength ~ 6.7 yr) if the original series contains much variance near $f=0.15$. Essentially, this low-pass filter is letting high-frequency variation pass through at selected frequencies (e.g., at the high lobes in Figure 8.5). Because some higher-

frequency fluctuations may not be damped out, the interpretation of the filtered series is made more complicated.

Ideally, a low-pass filter has a smoothly declining frequency response that remains low at frequencies higher than specified threshold frequency. Several classes of symmetrical filters whose filter weights decrease in size away from the central weight (unlike the moving average) have this desired trait. Such “better” low-pass filters include the binomial, Gaussian and Hamming. Another drawback of the moving average as a “smoothing” filter is that the smoothed series can jump abruptly in response to large single-year anomalies in the time series; the smoothed series can therefore appear jagged. Nevertheless, the moving average is often used because of the ease of interpreting each smoothed value as an arithmetic average over N observations.

8.4 Binomial filter

For the binomial filter, the weights are set proportional to the binomial coefficients (Panofsky and Brier, 1958; Mitchell et al. 1966). The binomial filter can be computed by repeated convolution of the sequence of weights $[0.5 \ 0.5]$, corresponding to equal probabilities of success or failure for a binomial distribution. If we let $b_0 = [0.5 \ 0.5]$, the three-weight binomial filter is given by the convolution of b_0 with itself

$$b_1 = \text{conv}(b_0, b_0) = [0.25 \ 0.50 \ 0.25] \quad (4)$$

The four-weight binomial filter, say b_2 , is formed by convoluting b_1 with b_0 . The five-weight binomial is formed by convoluting b_2 with b_0 , and so forth. The weights of an $N + 1$ weight binomial filter can be computed conveniently as follows

$$c_k = \frac{N!}{k!(N-k)!} \quad k = 0, 1, \dots, N$$

$$b_k = c_k / \sum_{k=0}^N c_k \quad (5)$$

Following Mitchell et al. (1966), the appropriate value of filter length, $N + 1$, can be computed for any desired period of 50% frequency response. The standard deviation of the binomial

distribution is $\sigma_B = \sqrt{N} / 2$, and the 50% response period occurs approximately at six standard deviations. Thus, considering an annual time series, if the 50% response period in years is p , the relationship

$$6\sigma_B = 3\sqrt{N} = p \quad (6)$$

yields

$$N = \left(\frac{p}{3}\right)^2 \quad (7)$$

To ensure that the filter length $N + 1$ is odd, N is rounded to the nearest **even** integer before being substituted into (5) to compute the filter weights. Weights much smaller than the maximum weight (say, less than 5% of it) are dropped, and then the filter is normalized such that the weights sum to 1 (Mitchell et al. 1966).

As an example, say the desired 50% response period is 10 years. The computed value of N is

$$N = \left(\frac{10}{3}\right)^2 = 11.111 \quad (8)$$

which is rounded to 12. The coefficients, computed from (5), truncated to remove excessively small values, and normalized to sum to 1 are

[0.0162 0.0541 0.1216 0.1946 0.2270 0.1946 0.1216 0.0541 0.0162]

(Figure 8.6).

As N becomes large, the weights for the binomial filter approximate the ordinates of the Gaussian, or normal, distribution. An alternative to the binomial filter is to set the weights proportional to the probability points of a normal distribution.

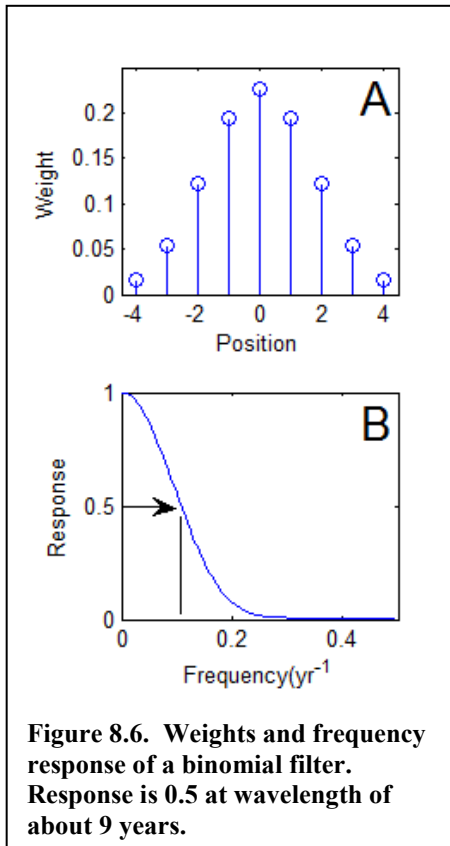


Figure 8.6. Weights and frequency response of a binomial filter. Response is 0.5 at wavelength of about 9 years.

8.5 Gaussian filter

Weights of the Gaussian filter are proportional to ordinates of a normal probability density function with a specified standard deviation (Mitchell et al. 1966). The Gaussian filter is convenient because the standard deviation of the appropriate Gaussian distribution can be related directly to the 50% frequency response of the filter. According to Mitchell et al. (1966), “the response ... drops below 50 per cent at wavelengths equal to about 6 standard deviations of the Gaussian curve.”

The appropriate Gaussian distribution therefore has standard deviation

$$\sigma_G = \frac{\lambda_{0.5}}{6} \quad (9)$$

where $\lambda_{0.5}$ is the desired wavelength at which the amplitude of frequency response is 0.5. The filter weights are obtained by sampling the pdf of the standard normal distribution at

t - values $0, \pm 1/\sigma_G, \pm 2/\sigma_G, \pm 3/\sigma_G, \dots$. These weights are next truncated to exclude values less than 5

percent of the maximum weight, and then are scaled so that the weights sum to 1.0.

For example, say we want a Gaussian filter with frequency response of 0.5 at a wavelength of 10 years. The appropriate Gaussian filter has standard deviation

$$\sigma_G = 10/6 = 1.66667$$

The t-distribution is sampled at t-values $0, \pm 0.6, \pm 1.2, \pm 2.4, \dots$, where the sampling is continued out to a large number of points – say as many points as observations in the series to be filtered. For 21 sample points, the pdf values are (to 4 digits)

0.0000	0.0000	0.0000	0.0001	0.0006	0.0044	0.0224	0.0790	0.1942	0.3332
0.3989	0.3332	0.1942	0.0790	0.0224	0.0044	0.0006	0.0001	0.0000	0.0000
0.0000									

Truncating to exclude all values less than 5 percent of 0.3989 yields

0.0224	0.0790	0.1942	0.3332	0.3989	0.3332	0.1942	0.0790	0.0224.
--------	--------	--------	--------	--------	--------	--------	--------	---------

These weights sum to about 1.6565. Dividing the weights by the sum yields the final weights


0.0134	0.0474	0.1165	0.1999	0.2394	0.1999	0.1165	0.0474	0.0134.
--------	--------	--------	--------	--------	--------	--------	--------	---------

These weights sum to 1 (ignoring rounding error).

8.6 Hamming-window filter

The binomial filter approaches a “bell shape” as filter length, N , increases, and the Gaussian filter is by definition bell shaped. Other ‘bell-shaped’ filters have the desired trait for low-pass filtering of a frequency response that drops steadily from 1.0 at low frequencies to zero at some frequency and remains at zero at higher frequencies.

A different approach to filter design consists of applying a smoothing window, or smoothing filter, to a mathematically derived *ideal* digital filter. The ideal filter is specified by a *cutoff frequency*, f_0 , defined such that the amplitude of the frequency response is 1 for all frequencies less than f_0 and is 0 for all frequencies greater than or equal to f_0 . Such a frequency response is called a *brick-wall* response. Recall that the frequency response of a filter is the Fourier transform



of the impulse response of the filter, and that the impulse response of a symmetrical digital filter is proportional to the filter itself. The ideal filter is accordingly computed as the inverse Fourier transform of the brick-wall frequency response (Figure 8.7). The ideal filter as so defined is not implementable because its impulse response is infinite and non-causal (The MathWorks, 1998, p. 2-19). To create a finite-duration impulse response, the ideal filter is truncated by applying a “window.”

A useful window for this purpose is the *Hamming window*, or *raised cosine window* (Karl 1989, The MathWorks 1999). The Hamming window weights are

computed as a function of a cosine

$$w_i = (0.54 - 0.46 \cos[2\pi i / (N - 1)]) / \sum_{i=0}^{N-1} w_i \quad 0 \leq i \leq (N - 1) \quad (10)$$

where N is the length of the window, or filter. N includes the central weight and the weights on either side of it. For example, a 5-weight Hamming-window filter¹, with $N = 5$ and central weight w_2 , has weights $[0.0357 \ 0.2411 \ 0.4464 \ 0.2411 \ 0.0357]$. The Hamming window applied to the ideal low-pass filter yields an implementable filter that in a sense is ideal given the specified constraint on the filter length.

The filter design problem in the windowed method is reduced to 1) specifying a desired cutoff frequency, and 2) specifying a desired filter length. As the filter length is increased, the algorithm comes closer to the objective of an “ideal” filter in terms of frequency response, but more data is lost off the ends of the series because of the large number of weights. The filter weights sum to 1, but for longer filter lengths some weights can be negative (Figure 8.8). This is a necessary consequence of the mathematics, but can be disturbing for practical interpretation.

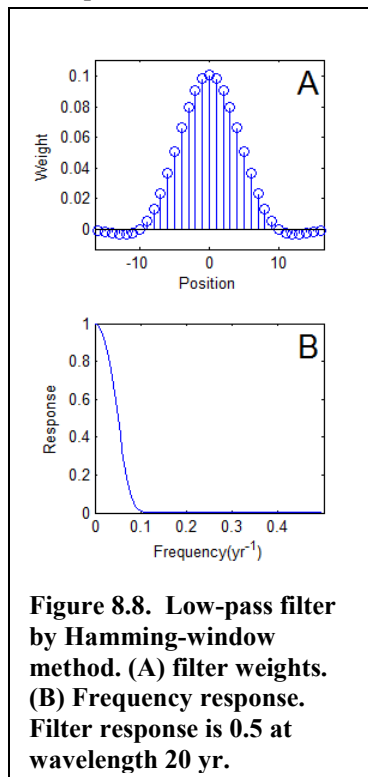


Figure 8.8. Low-pass filter by Hamming-window method. (A) filter weights. (B) Frequency response. Filter response is 0.5 at wavelength 20 yr.

The windowing method of filter design can be useful for band-pass and high-pass as well as low-pass windows. The method is probably most applicable when well-defined frequency ranges are of interest. For example, a tree-ring series might be filtered with a band-pass filter targeted on the frequencies that dominate the variance of the 11-year sunspot cycle. In most dendroclimatological studies, however, the precise cutoff frequency of variations of interest is difficult to specify, and the complexity of the windowing method might be overkill. If so, a simpler filter (e.g., binomial, Gaussian) with a more gradual transition between the frequencies retained and eliminated may suffice.

8.7 Effects of filtering on the time series and its spectrum

The effect of low-pass filtering on a pure sine wave is to “damp”, or reduce the amplitude, of the wave to some degree, depending on the wavelength of the sine wave and the frequency response of the filter (e.g., see Figure 8.3).

Geophysical time series are generally mixtures of variance at a wide range of frequencies, such that application of a low-pass

filter will smooth out the high frequencies and leave the lower frequencies relatively unaltered.


Because high-frequency variance is removed, the filtered series can be expected to have a lower variance than the original series. How much lower depends on the spectrum of the original series and the frequency response of the filter. If a time series has no variance at high frequencies, smoothing with a low-pass filter will have relatively little effect on the total variance. Conversely, if the time series is dominated by high-frequencies variations, low-pass filtering will greatly reduce the total variance of the series.

As an example of the effects of low-pass filtering, consider a time series of an annual tree-ring index plotted (Figure 8.9). The series does not appear to be cyclic, but does have some large-amplitude fluctuations with wavelength exceeding a decade, as well as considerable high-frequency variability.



As expected, filtering with a 9-weight binomial filter yields a smoother time plot and reduced range in the time series (Figure 8.10). The frequency response for the binomial filter (see figure

¹ Matlab function **hamming** yields a hamming window. Function **fir1** returns the weights of a low-pass filter designed by the Hamming-window method.

8.6b) indicates that variance at frequencies higher than about $f=0.2$ will be severely reduced in filtering.² Indeed, the spectra plotted in Figure 8.11 show that variance has essentially been completely eliminated at those high frequencies. The ratio of the areas under the two spectra is about 1/3, meaning that filtering has removed 2/3 of the original variance.



² Theoretically, for a linear system with the original time series as input and the filtered time series as output, the spectrum of the input multiplied by the square of the frequency response at a given frequency will give the spectrum of the output at that frequency (see Chatfield 2004).



8.8 References

- Chatfield, C., 2004. The analysis of time series, an introduction, sixth edition: New York, Chapman & Hall/CRC.
- Karl, J.H., 1989. An introduction to digital signal processing, Academic Press, Inc., San Diego, California 92101.
- Mitchell, J.M., Jr., Dzerdzeevskii, B., Flohn, H., Hofmeyr, W.L., Lamb, H.H., Rao, K.N., and Wallen, C.C., 1966. Climatic change, Technical Note 79: Geneva, World Meteorological Organization.
- Panofsky, H.A., and Brier, G.W., 1958. Some applications of statistics to meteorology: The Pennsylvania State University Press, 224 p.
- The MathWorks, Inc., 1998. Signal processing toolbox for use with Matlab, User's Guide, version 4. Apple Hill Drive, Natick, MA 01760-2098