

CHAPTER 3

EMD EQUIVALENT FILTER BANKS, FROM INTERPRETATION TO APPLICATIONS

Patrick Flandrin, Paulo Gonçalves and Gabriel Rilling

Huang's data-driven technique of empirical mode decomposition (EMD) is given a filter bank interpretation from two complementary perspectives. First, a stochastic approach operating in the frequency domain shows the spontaneous emergence of an equivalent dyadic filter bank structure when EMD is applied to the versatile class of fractional Gaussian noise processes. Second, a similar structure is observed when EMD is operated in the time domain on a deterministic pulse. A detailed statistical analysis of the observed behavior is carried out involving extensive numerical simulations that suggest a number of applications. New EMD-based approaches are used to estimate the scaling exponents in the case of self-similar processes, to perform a fully data-driven spectral analysis, and to denoise-detrend signals that contain noise.

3.1. Introduction

Empirical mode decomposition (EMD) has been recently pioneered by Huang et al. (1998) for adaptively decomposing signals into a sum of “well-behaved” AM-FM components consisting of natural “intrinsic” building blocks that describe the complicated waveform. The technique has already been employed successfully in various applications (Coughlin and Tung 2004; Fournier 2002; Huang et al. 1998; Neto et al. 2004; Wu et al. 2001).

Although EMD is quite simple in principle, it still lacks a theoretical foundation. Indeed, it is presently defined only as the output of an iterative algorithm, with no analytical definition that could be used for performance evaluation. The only way to better understand this technique is to resort to extensive numerical simulations in well-controlled situations. Such an “input-output” approach is adopted here, with the objective of obtaining a detailed, yet empirical, statistical knowledge of the EMD behavior, just as we might do for some unknown “filter” in signal processing.

Because the EMD algorithm is not uniquely defined since it depends on a number of user-controlled tunings such as the particular interpolation scheme for envelope extraction, the stopping criterion used in the sifting process, and the manner in which border effects are treated, we assume that its principle and the manner in which it is implemented are known. More precisely, the algorithm used in this study was developed on the basis on algorithmic considerations described in earlier

publications (see Rilling et al. 2003) and is available as a MATLAB code on the Internet (<http://perso.ens-lyon.fr/patrick.flandrin/emd.html>).

3.2. A stochastic perspective in the frequency domain

Our first characterization of EMD is carried out in the frequency domain from a stochastic perspective. The idea is to apply EMD to some broadband noise in order to understand how a full spectrum process is split into its “intrinsic mode functions” (IMF). One versatile class of full spectrum processes is provided by *scaling processes* for which wavelets (unanimously considered as a naturally fitted analysis tool; see Abry et al. 2000) can be used as a benchmark for performance evaluation.

3.2.1. Model and simulations

Fractional Gaussian noise (fGn, see Embrechts and Maejima 2002; Mandelbrot and van Ness 1968) is a generalization of ordinary white noise. It is a versatile model for a homogeneously spreading broadband noise without any dominant frequency band, is an intrinsically discrete-time process, and may be described as the increment process of fractional Brownian motion (fBm) since fBm is the only self-similar Gaussian process with stationary increments. Consequently the statistical properties of fGn are entirely determined by its second-order structure, which depends solely upon one single scalar parameter, H , its Hurst exponent. More precisely, $\{x_H[n], n = \dots, -1, 0, 1, \dots\}$ is a fGn of index H (with $0 < H < 1$) if and only if it is a zero-mean Gaussian stationary process whose autocorrelation sequence $r_H[k] := \mathbb{E}\{x_H[n]x_H[n+k]\}$ is

$$r_H[k] = \frac{\sigma^2}{2} (|k-1|^{2H} - 2|k|^{2H} + |k+1|^{2H}). \quad (3.1)$$

It is well known that the special case $H = \frac{1}{2}$ reduces to (discrete-time, uncorrelated) white noise, whereas other values induce non-zero correlations, either negative when $0 < H < \frac{1}{2}$ or positive when $\frac{1}{2} < H < 1$ (long-range dependence). Taking the discrete Fourier transform of (3.1), we readily obtain the power spectrum density of fGn, or

$$S_H(f) = C \sigma^2 |e^{i2\pi f} - 1|^2 \sum_{k=-\infty}^{\infty} \frac{1}{|f+k|^{2H+1}}, \quad (3.2)$$

with $|f| \leq \frac{1}{2}$. If $H \neq \frac{1}{2}$, we have $S_H(f) \sim C \sigma^2 |f|^{1-2H}$ when $f \rightarrow 0$. It therefore follows that fGn is a convenient model for power-law spectra at low frequencies. From its spectral properties, the particular value $H = \frac{1}{2}$ delineates two domains with contrasting behaviors. In the regime $0 < H < \frac{1}{2}$, we have $S_H(0) = 0$, and the spectrum is high-pass (sometimes referred to as an “ultraviolet” situation). On the other hand, within the range $\frac{1}{2} < H < 1$, we have $S_H(0) = \infty$ with a “ $1/f$ ”-type spectral divergence (“infrared” catastrophe). In both situations, the power-law form

of the spectrum, although not exactly verified, is well approximated over most of the Nyquist frequency band. In other words, we have a quasi-linear relation in log-log coordinates,

$$\log \mathcal{S}_H(f) \approx (1 - 2H) \log |f| + C,$$

for most frequencies $-\frac{1}{2} \leq f \leq \frac{1}{2}$.

3.2.2. Equivalent transfer functions

Extensive simulations were carried out on fGn processes, with H values ranging from 0.1 to 0.9. The present study [whose results were first proposed in Flandrin et al. (2004) with further extensions in Flandrin and Gonçalves (2004)] generalizes the study conducted independently by Wu and Huang (2004) for white noise only ($H = \frac{1}{2}$) and consistently supports their findings.

In all of our simulations, the data length was taken to be $N = 512$, and, for each value of H , $J = 5000$ independent sample paths of fGn were generated via the Wood and Chan (1994) algorithm. EMDs were computed for all sample paths $\{x_H^{(j)}[n]; n = 1, \dots, N\}$ (with $j = 1, \dots, J$), resulting in a collection of IMFs referred to as $\{d_{k,H}^{(j)}[n]; n = 1, \dots, N; k = 1, \dots, K_j\}$. Although the number K_j of IMFs varied from one realization to the other, none of the realizations generated less than 7 modes. Therefore, $K = 7$ has been taken in this study as the common number of modes for all realizations.

Given this dataset, a spectral analysis was carried out mode by mode with the estimated power spectrum density (PSD) given by

$$\hat{\mathcal{S}}_{k,H}(f) := \sum_{m=-N+1}^{N-1} \hat{r}_{k,H}[m] w[m] e^{-i2\pi f m}, \quad |f| \leq \frac{1}{2},$$

where $w[n]$ is a Hamming taper, and

$$\hat{r}_{k,H}[m] = \frac{1}{J} \sum_{j=1}^J \left(\frac{1}{N} \sum_{n=1}^{N-|m|} d_{k,H}^{(j)}[n] d_{k,H}^{(j)}[n + |m|] \right), \quad |m| \leq N - 1$$

is the ensemble average (over the J realizations) of the empirical estimates of the auto-correlation function. The result of this spectral analysis is plotted in Fig. 3.1, whose graphs reveal a number of striking features:

- (1) Regardless of the value of the Hurst exponent H , the behavior of the first IMF (thick line) differs from that of the other modes. To a first approximation, it possesses the characteristics of a high-pass filter while higher order modes behave similarly to a band-pass filter. The (roughly half-band) high-pass character of the first mode must be tempered, however, by the fact that the maximum

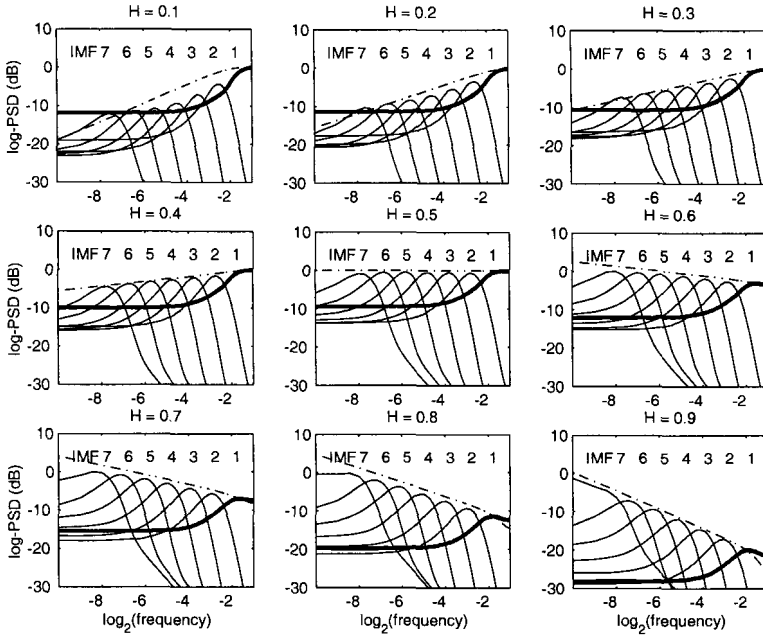


Figure 3.1: IMF power spectra in the case of fractional Gaussian noise. The logarithm of the estimated power spectrum densities (log-PSD) is plotted as a function of the logarithm of the normalized frequency for the first seven IMFs. For $H = 0.1, 0.2, \dots, 0.9$, the spectral estimates have been computed on the basis of 5000 independent sample paths of 512 data points. Theoretical PSDs of the full processes are superimposed as dashed-dotted curves. (Originally published in *Int. J. Wavelets Multiresolut. Inform. Process.*, **2**, 477–496, ©2004 World Scientific.)

attenuation in the stop-band is no more than 10 dB (as compared to the maximum which occurs at the Nyquist frequency $\frac{1}{2}$, and there is a non-negligible contribution in the lower half-band in “ultraviolet” situations $H < \frac{1}{2}$).

- (2) As H varies from 0.1 to 0.9, the spectrum of the last IMF ($k = 7$) progressively evolves from band-pass to increasingly low-pass, in accordance with the increasing predominance of low frequencies (“infrared catastrophe”).
- (3) In a similar, but more general manner, the energy balance among the different modes reflects the behavior of the global spectrum (superimposed dashed-dotted curve) described by (3.2), the flat spectrum when $H = \frac{1}{2}$ (the case of white noise), and the increasing (decreasing) power-law spectrum when $H < \frac{1}{2}$ ($H > \frac{1}{2}$).
- (4) For modes $k = 2$ to 6 (band-pass IMFs), all the spectra appear nearly the same, with some shifts in abscissa and ordinate, and this finding is surprisingly reminiscent of what is currently being observed in wavelet decompositions (see Flandrin 1999; Mallat 1998).




Figure 3.2: (a) IMF average number of zero-crossings *zero-crossing* in the case of fractional Gaussian noise. For clarity, only those curves corresponding to the extreme indices $H = 0.1$ (circles) and $H = 0.9$ (squares) have been plotted in the diagram; the remaining cases ($H = 0.2, 0.3, \dots, 0.8$) lead to regularly intertwined similar curves. The superimposed solid lines correspond to linear fits within the IMF range $k = 2$ to 6. (b) Corresponding decrease rate of zero-crossings (circles and dashed line), with the least-squares quadratic fit given by (3.4) superimposed as a solid line. (Originally published in *Int. J. Wavelets Multiresolut. Inform. Process.*, **2**, 477–496, ©2004 World Scientific.)

This last observation suggests that we should examine in greater detail how the different spectra are related to each other for a given H . To this end, we can use the unique structure of IMFs: all extrema appear as an alternation of local minima and maxima separated by only one zero-crossing. Finding the average number of zero-crossings in a mode is, therefore, a meaningful way of characterizing its mean frequency. The average number of zero-crossings $z_H[k]$ is plotted in Fig. 3.2a as a function of the IMF number k ; this figure suggests the functional relation

$$z_H[k] \propto \rho_H^{-k}, \quad (3.3)$$

where ρ_H very nearly equals 2.

A more precise check of (3.3) is shown in Fig. 3.2b, where the estimated scaling factor ρ_H is given by the slope from a linear fit of a semi-log diagram of $\log_2 z_H[k]$ vs. k for $k = 2$ to 6. The observed decrease in the number of zero-crossings as the order of the IMFs increases is nearly equal to 2 and may be approximated by the




Figure 3.3: Renormalized IMF spectra in the case of fractional Gaussian noise. For each value of H , the band-pass IMFs ($k = 2$ to 6) of Figure 3.1 are plotted according to the renormalization given by (3.5) with $\alpha = 2H - 1$, and the values of ρ_H are given in Fig. 3.2b. (Originally published in *Int. J. Wavelets Multiresolut. Inform. Process.*, **2**, 477–496, ©2004 World Scientific.)

quadratic expression:*

$$\rho_H \approx 2.01 + 0.2 \left(H - \frac{1}{2} \right) + 0.12 \left(H - \frac{1}{2} \right)^2. \quad (3.4)$$

Using (7.4), we can improve our search for self-similarity in the “filter bank” structure of Fig. 3.1. If we restrict ourselves to the band-pass IMFs ($k = 2$ to 6), self-similarity means that

$$\mathcal{S}_{k',H}(f) = \rho_H^{\alpha(k'-k)} \mathcal{S}_{k,H}(\rho_H^{k'-k} f) \quad (3.5)$$

for some α and any $k' > k \geq 2$. Consequently, the power spectra of all IMFs should collapse onto a single curve when properly renormalized. Indeed, setting $\alpha = 2H - 1$ verifies this assumption as the corresponding renormalizations converge to the same template (see Fig. 3.3). Even if some low frequency discrepancies are present (especially when $H < \frac{1}{2}$), these diagrams support our claim that, to a first approximation, EMD acts on fGn as a dyadic filter bank of constant- Q band-pass filters.

*The accuracy of this approximation is slightly dependent on the manner used to compute the IMFs and, in particular, on the choice of the stopping criterion used for the sifting process. Further studies will be necessary to clarify this point.

3.3. A deterministic perspective in the time domain

Our second approach to the characterization of the filter-like structure of EMD is constructed deterministically and in the time domain. Therefore, we seek to obtain an equivalent impulse response of the analysis.

3.3.1. Model and simulations

Finding the impulse response of a system usually amounts to observing its output when excited with a Dirac pulse $\delta(t)$ or, in discrete-time, a function $\delta[n]$ that is zero everywhere except when $n = 0$ with $\delta[0] = 1$. Doing so is not possible here since such an input signal would not consist of enough local extrema to initiate the algorithm. An alternative is to consider an idealized pulse as the limit of a noisy pulse as the signal-to-noise ratio goes to infinity. From this point of view, we model the noisy pulse $\delta_\varepsilon[n] := \delta[n] + \varepsilon x_{1/2}[n]$, and the effective IMFs are defined as

$$d_k[n] = \lim_{\varepsilon \rightarrow 0^+} \mathbb{E} \{d_{k,\varepsilon}[n]\},$$

where $d_{k,\varepsilon}[n]$ denotes the k -th IMF of $\delta_\varepsilon[n]$.

In practice, we used in our simulations zero mean unit variance Gaussian white noise $x_{1/2}[n]$, with $\varepsilon = 0.02$ (corresponding to a signal-to-noise ratio (SNR) of 34 dB, with SNR defined as $10 \log_{10}(1/\text{var}\{\varepsilon x_{1/2}[n]\})$). The data length of each realization has been fixed to $N = 256$, and simulations have been carried out on $J = 5000$ independent realizations.

3.3.2. Equivalent impulse responses

The result of the simulation (average EMD for our slightly noisy pulse) is plotted in Fig. 3.4. Again, this figure shows a striking resemblance to what we would have obtained by using a wavelet analysis, the different averaged IMFs apparently all having the same shape for each order k .

To understand our results, we again seek a self-similar structure which would reduce all of the IMFs to one universal waveform thanks to a well-chosen renormalization. Figure 3.5 shows that this reduction is indeed possible. In constructing our waveform, we first plotted the maximum amplitude of the different IMFs as a function of their index. Doing so allowed us to identify an exponential variation of the form:

$$d_k[0] = 2^{C-pk},$$

where $p \approx 0.85$. In the second step, we found that a dilation factor $\alpha = 2^p \approx 1.80$ gave a superposition of the different waveforms (and their spectra) which we could express as

$$d_k[n] = \frac{1}{\alpha^k} \psi \left(\left\lfloor \frac{n}{\alpha^k} \right\rfloor \right),$$




Figure 3.4: Impulse response. The EMD equivalent impulse response was obtained by averaging each IMF by using a large number of decompositions computed with a slightly noisy pulse. Here, 5 000 independent realizations from 256 data points were simulated with a signal-to-noise ratio of 34 dB. The first frame shows the average pulse and each successive frame corresponds to the ensemble average of the first six IMFs normalized in amplitude mode by mode.

where $\psi(t)$ is a reference waveform and analogous to a mother wavelet in a multiresolution analysis. For our analysis, the interpolation scheme used in the EMD was a cubic spline (see Huang et al. 1998; Rilling et al. 2003; available online at <http://perso.ens-lyon.fr/patrick.flandrin/emd.html>). Note how very similar $\psi(t)$ is to a cubic spline wavelet (third convolution power of the Haar wavelet).

3.4. Selected applications

If we accept that EMD may be characterized in some cases as a “spontaneous” filter bank, then several potentially useful applications are immediately suggested.

3.4.1. EMD-based estimation of scaling exponents

Our first application concerns the estimation of the Hurst exponent H for fGn based on the EMD spectral analysis described in Section 3.2. Given the self-similar relation (3.5) for PSDs for band-pass IMFs (index $k > 1$), we can deduce how the variance should evolve as a function of k . Assuming that (3.5) holds for any $k' > k \geq 2$ and

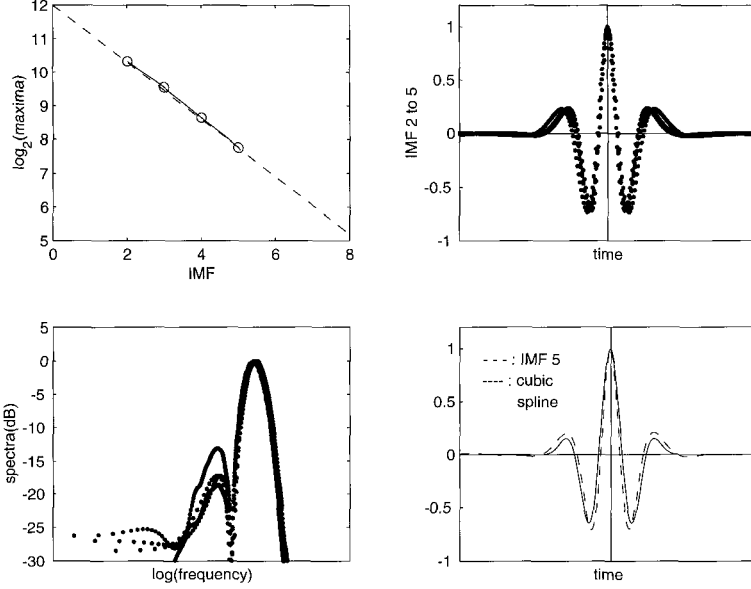


Figure 3.5: Self-similarity. The logarithm of the maximum amplitude of IMFs 2 to 5 in Fig. 3.4 is given by a linear function of the IMF index (top left diagram). Renormalizing these IMFs, either in time (top right) or in frequency (bottom left), yields an unique curve. The “impulse response” is similar to a cubic spline wavelet (bottom right).

$\alpha = 2H - 1$, we have

$$\begin{aligned}
 V_H[k'] &:= \text{var } d_{k',H}[n] = \int_{-1/2}^{1/2} S_{k',H}(f) df \\
 &= \rho_H^{\alpha(k'-k)} \int_{-1/2}^{1/2} S_{k,H}(\rho_H^{k'-k} f) df \\
 &= \rho_H^{(\alpha-1)(k'-k)} V_H[k],
 \end{aligned}$$

which leads to

$$V_H[k] = C \rho_H^{2(H-1)k}. \quad (3.6)$$

The IMF variance should be an exponentially decreasing function of the IMF index with a decay rate which is a linear function of the Hurst exponent H . Experimental evidence for this behavior is given in Fig. 3.6 where a semi-log diagram (in base 2) gives the (energy-based) empirical variance estimate

$$\hat{V}_H[k] := \frac{1}{J} \sum_{j=1}^J \left[\frac{1}{N} \sum_{n=1}^N \left(d_{k,H}^{(j)}[n] \right)^2 \right] \quad (3.7)$$

as a function of the index k . From a logarithmically linearized version of (3.6),




Figure 3.6: Estimated IMF \log_2 -variance in the case of fractional Gaussian noise. The values of the empirical (energy-based) variance estimates are given by dotted lines for different values of the Hurst exponent H . The error bars correspond to the standard deviations associated with the 5 000 realizations run in the study. The mean value of the estimated Hurst exponents is also given, based on a weighted linear fit within the IMF indices range $k = 2$ to 6. For clarity, all curves have been arbitrarily shifted along the vertical axis to avoid overlapping. (Originally published in *Int. J. Wavelets Multiresolut. Inform. Process.*, **2**, 477-496, ©2004 World Scientific.)

straight lines may be fitted to the different curves. The slope κ_H then gives an estimated Hurst exponent \hat{H} via

$$\hat{H} = 1 + \frac{\kappa_H}{2}. \quad (3.8)$$

Figure 3.6 shows that (3.6) holds only for IMF indices $k > 1$. Furthermore, the error increases as H becomes small (typically, the model fits the data reasonably well for $H > \frac{1}{4}$).

To better understand our ability to estimate H from the slope of a “log-energy vs. IMF index” diagram, one must consider not only the evolution of the variance as a function of the modes, but also the possible correlations which may exist within and between modes. To this end, we focus on band-pass IMFs ($k > 1$) and evaluate the two-dimensional correlation function[†]

$$D_H[k', n'] := \mathbb{E} \{d_{k,H}[n] d_{k+k',H}[n + n']\}$$

[†]The definition of this quantity is based on the implicit assumption that IMFs are, jointly, second-order stationary processes.

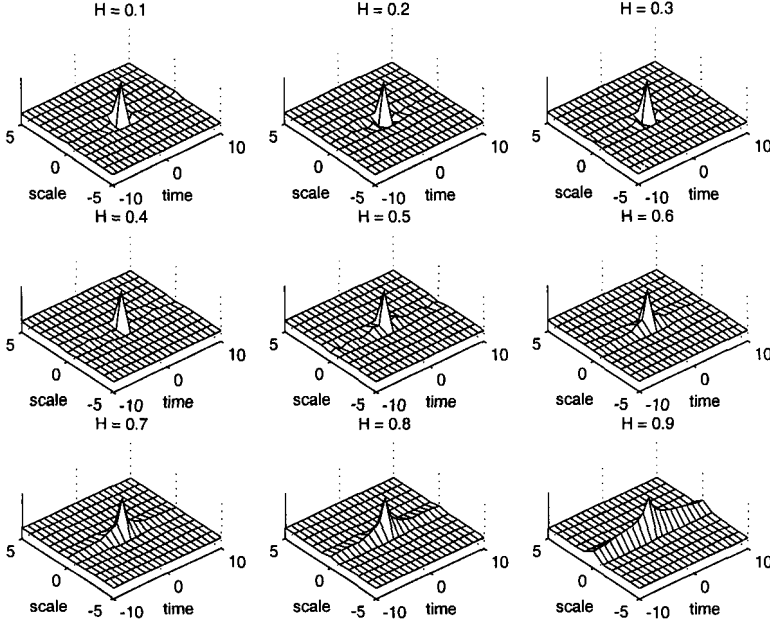


Figure 3.7: Two-dimensional correlation function of the IMF matrix in the case of fractional Gaussian noise. For each Hurst exponent H , the graph displays the quantity $|\hat{D}_H[k', n']|$ given by (3.9) as a function of time and scale (IMF index lag).

by using the averaged empirical estimate

$$\hat{D}_H[k', n'] := \frac{1}{J} \sum_{j=1}^J \left(\frac{1}{NK} \sum_{k=2}^{K-|k'|} \sum_{n=1}^{N-|n'|} d_{k,H}^{(j)}[n] d_{k+|k'|,H}^{(j)}[n+|n'|] \right), \quad (3.9)$$

with $|n'| \leq N-1$ and $|k'| \leq K-2$. Here, K denotes the largest IMF index minus 1, and we have discarded the residual. This two-dimensional correlation function of the full IMF matrix is plotted in Fig. 3.7 and shows that modes with different indices are essentially uncorrelated. The only significant values of $\hat{D}_H[k', n']$ correspond to $k' = 0$, i.e., to intra-scale correlations, with a correlation decay which becomes slower as H is increased.

The effects of using our estimate of the Hurst exponent H given by (3.8) and the slope κ_H deduced from (3.7) on (3.9) are twofold. First, because of the non-zero intra-scale correlations, the variance estimate $\hat{V}_H[k]$ given by (3.7) would experience large fluctuations, especially when the Hurst exponents H and IMF indices k are large. Second, the negligible inter-scale correlations should allow for an estimate of the slope κ_H from a weighted linear regression from a semi-log diagram of $\log_2 \hat{V}_H[k]$ vs. k . Further results from the effective performance of this EMD-based estimator of H (and comparisons with wavelet-based approaches) can be found in Flandrin and Gonçalves (2003).

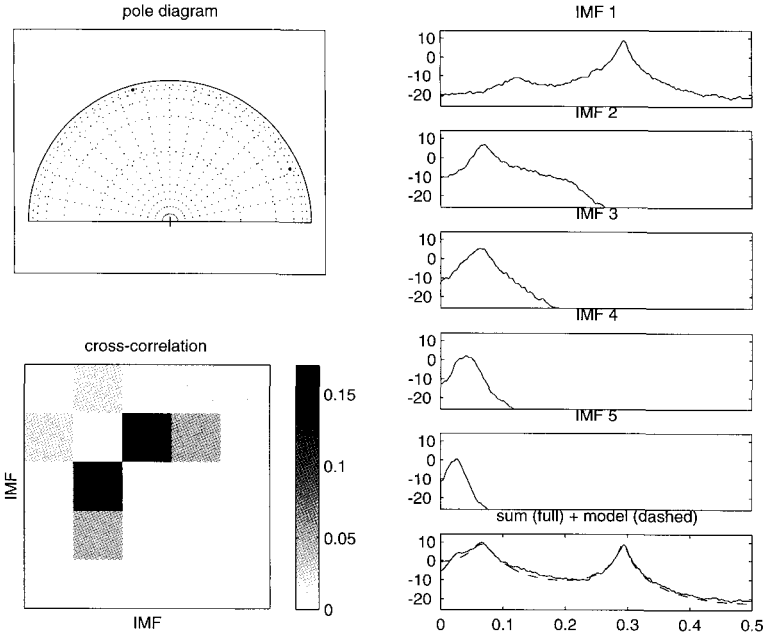


Figure 3.8: EMD as a spectrum analyzer. In the case of an AR(4) process, whose pole constellation in the upper half unit circle is plotted on the top left diagram, the ensemble averaged spectral analysis (over 50 realizations) of the first five IMFs is given in the right column of the figure. The bottom frame displays the cumulative spectrum (solid line) obtained by summing up the five spectra and compares it to the model’s spectrum (dotted line). The bottom left diagram is a schematic of the normalized IMF cross-correlation (3.10) with the (unit) diagonal artificially forced to be zero so as to enhance the gray scale dynamic range (IMF indices grow from left to right and from top to bottom).

3.4.2. EMD as a data-driven spectrum analyzer

If one accepts that EMD behaves as a homogeneous filter bank for processes whose (full) spectrum varies monotonically, one can further investigate how this method decomposes processes with a less regular spectrum. Figure 3.8 shows the preliminary results obtained for an auto-regressive (AR) process of order 4. While EMD does achieve a filter bank-like decomposition in this case, the interpretation of these modes requires some caution. With regard to the first IMF, the selected main frequency band is fully data-driven and automatically adapted to the highest frequency resonance. On the other hand, as noted earlier with regards to the non-negligible capture of low frequencies by IMF 1, some contributions at lower frequencies also occur. These contributions may include resonances at lower frequencies; they can also correspond to artifacts which must be compensated by IMFs of higher orders. These situations may be identified by examining intermodal correlation coefficients: The larger the correlation is, the less significant is the splitting into separate components. A quantitative evaluation of this intermodal correlation is plotted in Fig. 3.8,

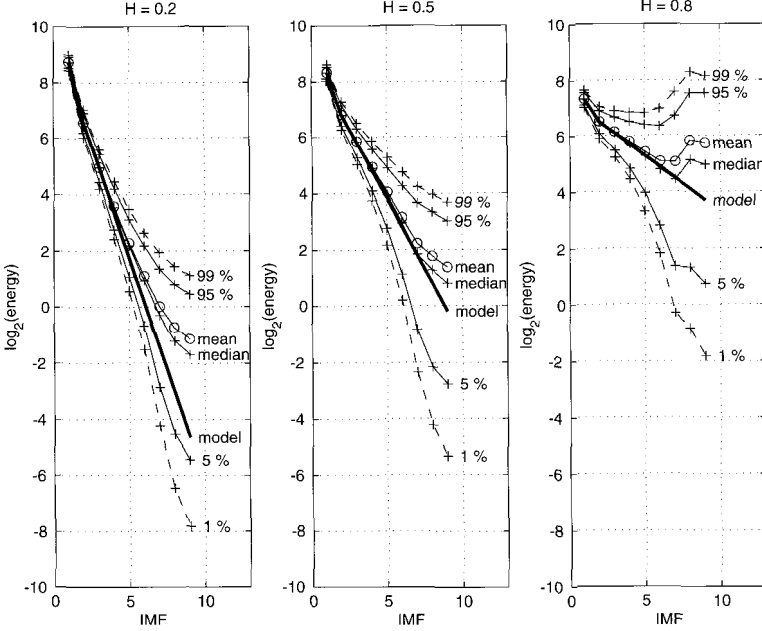


Figure 3.9: Experimental “modegrams” in the case of fractional Gaussian noise. For three different values of the Hurst exponent, the statistical characteristics (mean, median, confidence intervals) of the logarithm of the estimated EMD variance have been plotted as a function of the IMF index. The linear model (3.11) is also plotted. (Originally published in *Int. J. Wavelets Multiresolut. Inform. Process.*, **2**, 477–496, ©2004 World Scientific.)

where

$$\Theta[k, k'] := \left| \frac{1}{J} \sum_{j=1}^J \frac{\Theta^{(j)}[k, k']}{\sqrt{\Theta^{(j)}[k, k] \Theta^{(j)}[k', k']}} \right| \quad (3.10)$$

and

$$\Theta^{(j)}[k, k'] := \frac{1}{N} \sum_{n=1}^N d_{k,H}^{(j)}[n] d_{k',H}^{(j)}[n].$$

By definition, we have $0 \leq \Theta[k, k'] \leq 1$. Figure 3.8 clearly shows that the non-negligible values (as compared to 1) of $\Theta[k, k']$ correspond to index pairs (k, k') for which the IMF DSPs have a large amount of frequency overlap.

3.4.3. Denoising and detrending with EMD

A detailed knowledge of IMF statistics in situations where noise is present can help in gauging the significance of a given mode. This idea, which has been pioneered by Wu and Huang (2004), can be used to separate a signal from noise. Two possible methods, namely *denoising* (by removing those modes identified as noise) and *detrending* (by keeping only them), can be used.




Figure 3.10: Relative confidence intervals. The (base 2) logarithm of the relative confidence intervals given in Fig. 3.9, i.e., $\log_2(\log_2(T_H[k]/W_H[k]))$, behaves essentially linearly as a function of the IMF index k , suggesting (3.12). For each of the three values of H , the crosses (circles) correspond to a confidence interval of 95% (99%), the dotted (dashed) lines refer to the cases where the reference $W[k]$ is chosen by using the mean (median) of the IMF energies over the realizations, and the solid lines indicate the corresponding best linear fit.

With regards to the variability of the variance estimate, Fig. 3.6 gives a rough, second-order indication based on the observed standard deviation. A greater appreciation can be gained from Fig. 3.9, in which the experimental mean, median and various confidence intervals are plotted for $H = 0.2, 0.5$ and 0.8 , as well as

$$\log_2 V_H[k] = \log_2 \hat{V}_H[2] + 2(H - 1)(k - 2) \log_2 \rho_H \quad (3.11)$$

for $k \geq 2$, which was derived from (3.6). This series of simulations (which was carried out on 10 000 realizations of 2048 data points in each case) shows increasingly larger fluctuations for modes as the indices increase.[†] This finding agrees with (and is a generalization of) the findings reported in Wu and Huang (2004) for the case of white noise. Figure 3.10 also suggests that we may parameterize $T_H[k]$ by using the formula:

$$\log_2(\log_2(T_H[k]/W_H[k])) = a_H k + b_H, \quad (3.12)$$

where $W_H[k]$ denotes the H -dependent variation of the IMF energy. As noted earlier,

[†]The skewed (marginal) distribution of these “modegrams” yields better agreement if the linear model (3.12) uses the median rather than the mean of the realizations.




Figure 3.11: Denoising and detrending. An example of an amplitude-modulated, low-frequency oscillation embedded in fractional Gaussian noise with a Hurst exponent $H = 0.3$ is plotted in (b). The estimated energies of the seven IMFs are plotted in (a) as the thick line, together with the “noise only” model (thin line) and the 99% confidence interval (dotted line). The partial reconstruction obtained by adding the EMD residual and IMFs 5 to 7 [only those whose energies exceed the threshold in (a)] is plotted in (c) as a solid line and is superimposed on the actual signal component (dotted line). The partial reconstruction of IMFs 1 to 4 is plotted in (d).

Table 3.1: Confidence Interval Parameters for the Linear Model (3.12).

H	β_H	$a_H(95\%)$	$b_H(95\%)$	$a_H(99\%)$	$b_H(99\%)$
0.2	0.487	0.458	-2.435	0.452	-1.951
0.5	0.719	0.474	-2.449	0.460	-1.919
0.8	1.025	0.497	-2.331	0.495	-1.833

the best linear fit occurs when the median of the IMF’s energy is used to compute $W_H[k]$ over the realizations. The parameters a_H and b_H can be deduced from the simulation results in Fig. 3.9, and their values are reported in Table 3.1. In practice, $W_H[1]$ can be estimated from

$$\hat{W}_H[1] = \sum_{n=1}^N d_{1,H}^2[n], \quad (3.13)$$

and subsequent values of $W_H[k]$ are given by

$$\hat{W}_H[k] = C_H \rho_H^{-2(1-H)k}, \quad k \geq 2, \quad (3.14)$$

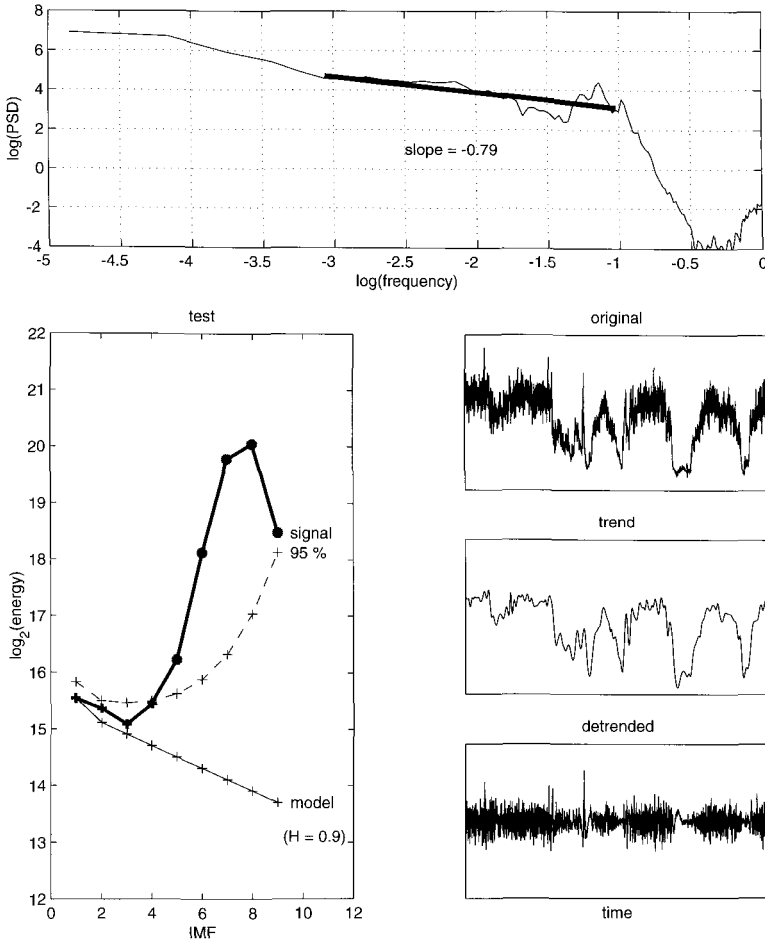


Figure 3.12: Denoising and detrending of a heart-rate signal. Top diagram: Signal spectrum in log-log coordinates (thin line) with a linear fit of slope $p \approx -0.8$ in the mid-frequency range (thick line), supporting a fGn model with Hurst exponent $H = (1 - p)/2 \approx 0.9$. Bottom diagram: Model-based detrending. Left: Estimated energy of nine IMFs, plotted as the thick line, together with the “noise only” model given by $H = 0.9$ (thin line) and the 95% confidence interval (dotted line). Top right: Original signal. Middle right: Estimated trend obtained from the partial reconstruction with IMFs 5 to 9 (only those whose energies exceed the threshold in the left diagram) and the residual. Bottom right: Detrended signal obtained from the partial reconstruction with IMFs 1 to 4.

where $C_H = \hat{W}_H[1]/\beta_H$. The parameter β_H used to compute C_H can, in turn, be estimated from the data displayed in Fig. 3.9, and its values are given in Table 3.1.

Given these results, a possible strategy for denoising a signal corrupted by fGn (with a known H) is as follows:

- (1) Assuming that the first IMF captures most of the noise, estimate the noise level in the noisy signal by computing $\hat{W}_H[1]$ from (3.13).

- (2) Estimate the “noise only” model by using (3.13) and (3.14).
- (3) Estimate the corresponding model for a chosen confidence interval from (3.12) and Table 3.1.
- (4) Compute the EMD of the noisy signal, and compare the IMF energies by using the confidence interval as a threshold.
- (5) Compute a partial reconstruction by keeping only the residual and those IMFs whose energy exceeds the threshold.

An alternative strategy for detrending fGn-type noise process consists of computing the complementary partial reconstruction based on only those IMFs whose energy is below the threshold.

A simple example of the EMD approach to denoising and detrending is given in Fig. 3.11, which presents the case of an oscillatory, low frequency waveform embedded in fractional Gaussian noise. A companion example containing actual data (heart-rate variability) is given in Fig. 3.12.

3.5. Concluding remarks

We have shown that EMD achieves a specific form of hierarchical filtering. This result is in agreement with the intuition associated with the EMD principle. However, because EMD still lacks a sound theoretical foundation, a careful and detailed analysis based on extensive numerical simulations was necessary for asserting and quantifying this behavior. In the cases shown here, we observed the “spontaneous” emergence of an equivalent filter bank structure which has the advantage of being fully data-driven. Furthermore, because it is local in time, this structure can adapt automatically to nonstationary situations with greater flexibility than other approaches using a pre-determined decomposition scheme. Although some possible applications have been outlined, their potential usefulness will require further studies that compare EMD to alternative methods for specific tasks while endeavoring to make the theory more rigorous.

References

- Abry, P., P. Flandrin, M. Taqqu, and D. Veitch, 2000: Wavelets for the analysis, estimation, and synthesis of scaling data. *Self-Similar Network Traffic and Performance Evaluation*, K. Park and W. Willinger, Eds., Wiley, 39–88.
- Coughlin, K. T., and K. K. Tung, 2004: 11-year solar cycle in the stratosphere extracted by the empirical mode decomposition method. *Adv. Space Res.*, **34**, 39–88.
- Embrechts, P., and M. Maejima, 2002: *Selfsimilar Processes*. Princeton Univ. Press, 111 pp.
- Flandrin, P., 1999: *Time-Frequency/Time-Scale Analysis*. Academic Press, 386 pp.
- Flandrin, P., and P. Gonçalves, 2003: Sur la décomposition modale empirique. *Proc. 19ème Coll. GRETSI sur le Traitement du Signal et des Images*, Paris, 149–152.

- Flandrin, P., and P. Gonçalves, 2004: Empirical mode decompositions as data-driven wavelet-like expansions. *Int. J. Wavelets Multiresolut. Inform. Process.*, **2**, 477–496.
- Flandrin, P., G. Rilling, and P. Gonçalves, 2004: Empirical mode decomposition as a filter bank. *IEEE Signal Process. Lett.*, **11**, 112–114.
- Fournier, R., 2002: *Analyse stochastique modale du signal stabilométrique. Application à l'étude de l'équilibre chez l'Homme*. Thèse de Doctorat, Univ. Paris XII Val de Marne, 263 pp.
- Huang, N. E., Z. Shen, S. R. Long, M. L. Wu, H. H. Shih, Q. Zheng, N. C. Yen, C. C. Tung, and H. H. Liu, 1998: The empirical mode decomposition and Hilbert spectrum for nonlinear and non-stationary time series analysis. *Proc. R. Soc. London, Ser. A*, **454**, 903–995.
- Mandelbrot, B. B., and J. W. van Ness, 1968: Fractional Brownian motions, fractional noises and applications. *SIAM Rev.*, **10**, 422–437.
- Mallat, S., 1998: *A Wavelet Tour of Signal Processing*. Academic Press, 577 pp.
- Neto, E. P. S., M. A. Custaud, C. J. Cejka, P. Abry, J. Frutoso, C. Gharib, and P. Flandrin, 2004: Assessment of cardiovascular autonomic control by the empirical mode decomposition. *Method. Inform. Med.*, **43**, 60–65.
- Rilling, G., P. Flandrin, and P. Gonçalves, 2003: On empirical mode decomposition and its algorithms. *IEEE-EURASIP Workshop on Nonlinear Signal Image Process. NSIP-03*, Grado, Italy.
- Wood, A. T., and G. Chan, 1994: Simulation of stationary processes in $[0, 1]^d$. *J. Comp. Graph. Stat.*, **3**, 409–432.
- Wu, Z., and N. E. Huang, 2004: A study of the characteristics of white noise using the empirical mode decomposition method. *Proc. R. Soc. London, Ser. A*, **460**, 1597–1611.
- Wu, Z., E. K. Schneider, Z. Z. Hu, and L. Cao, 2001: The impact of global warming on ENSO variability in climate records. *COLA Technical Report, CTR 110*, 25 pp.

Patrick Flandrin

*Laboratoire de Physique (UMR 5672 CNRS), Ecole Normale Supérieure de Lyon,
46 allée d'Italie, 69364 Lyon Cedex 07 France
flandrin@ens-lyon.fr*

Paulo Gonçalves

*INRIA Rhône-Alpes, On leave at IST-ISR, Av. Rovisco Pais, 1049-001 Lisbon,
Portugal
Paulo.Goncalves@inria.fr*

Gabriel Rilling

*Laboratoire de Physique (UMR 5672 CNRS), Ecole Normale Supérieure de Lyon,
46 allée d'Italie, 69364 Lyon Cedex 07 France
grilling@ens-lyon.fr*