Issues with the Application of Empirical Mode Decomposition Analysis

[1]Peel, M.C., [1]G.E. Amirthanathan, [2]G.G.S. Pegram, [1]T.A. McMahon and [1]F.H.S. Chiew

[1]Department of Civil and Environmental Engineering, The University of Melbourne, Australia, [2]Civil Engineering, University of KwaZulu-Natal, South Africa, E-Mail: mpeel@unimelb.edu.au

Keywords: *Empirical Mode Decomposition; Hydroclimatology; Interannual variability; Interdecadal variability.*

EXTENDED ABSTRACT

Hydroclimatic variability occurs over time scales ranging from seconds through to millennia. Fluctuations at certain time scales, like monthly, seasonal, annual, interannual and interdecadal, are particularly important for the sustainable management of land and water resources systems. Quantifying the proportion of variation in a hydroclimatic time series due to fluctuations at different time scales is usually done using spectral techniques, like Fourier analysis, which assume a time series to be both linear and stationary. Empirical mode decomposition (EMD), a relatively new form of time series analysis for quantifying the proportion of variation at different time scales, is introduced and key aspects of its application are discussed in this paper.

EMD was originally developed as a form of adaptive time series decomposition, used prior to spectral analysis using the Hilbert transform, for non-linear and non-stationary time series data. In this paper the spectral analysis component (Hilbert transform) is not discussed, only the EMD procedure is addressed. EMD has several advantages over other spectral techniques, in that it is relatively easy to understand and use, the fluctuations within a time series are automatically and adaptively selected from the time series and it is robust in the presence of non-linear and non-stationary data. EMD allows the data to speak for themselves. Robustness in the presence of non-linear and non-stationary data is particularly important for hydroclimatology applications where time series are generally non-stationary (or not long enough to categorically prove stationarity) and they exhibit non-linear characteristics like amplitude and frequency modulation with time.

In this paper the process of applying EMD to a time series is demonstrated using 10 years of monthly precipitation data from Melbourne Regional Office. The Melbourne monthly precipitation time series is decomposed into three intrinsic mode functions (IMFs) and a residual. The conditions defining an IMF are presented. The sifting process used to obtain each IMF and the residual is described. The general features of IMFs are described along with the ability to combine IMFs and the residual to form low frequency or high frequency filters.

Two key decisions in the EMD application process, the rule for deciding when to stop sifting for an IMF and the choice of cubic spline end condition rule are reviewed and discussed in detail. A new cubic spline end condition rule, based on the assumption that the slope of the cubic spline at the end point is equal to zero, is proposed and compared to two other end condition rules from the literature. The comparison is based on three applications of the EMD algorithm, each application with a different end condition rule, to 8135 annual precipitation time series from around the world. The annual precipitation time series have periods of record ranging from 30 up to 299 years and represent a wide range of climatic zones. The proposed end condition rule is found to be the most efficient of the three rules tested, due to the EMD algorithm producing less IMFs when using the proposed rule. The end condition rule proposed in this paper is recommended for future EMD applications.

1. INTRODUCTION

In many fields of science separating signal from noise in a time series is a consistent problem. In physical sciences, like hydroclimatology, short record lengths, non-stationary data and non-linear processes further complicate this problem. Separating process signal from random noise is particularly important for understanding the processes driving an observed variable as well as for forecasting that variable. Techniques for separating signal from noise are therefore of considerable interest to engineers and scientists. In this paper we present and discuss some issues with the application of empirical mode decomposition, a relatively new technique for decomposing non-linear and non-stationary time series data.

2. WHAT IS EMD?

Empirical Mode Decomposition (EMD), developed by Huang *et al* (1998), is a form of adaptive time series decomposition, used prior to spectral analysis (with the Hilbert transform), for non-linear and non-stationary time series data. Traditional forms of spectral analysis, like Fourier, assume that a time series (either linear or non-linear) can be decomposed into a set of linear components. However, as the degree of non-linearity and non-stationarity in a time series increases, the set of linear components describing that time series increases substantially when using Fourier techniques. In the physical sciences, time series are often non-linear and or non-stationary, so Fourier based spectral analysis techniques often produce large sets of physically meaningless harmonics when applied to these problems (Huang *et al* 1999). In contrast the EMD method does not assume a time series is linear or stationary prior to analysis, it lets the data speak for themselves. EMD adaptively decomposes a time series into a set of independent intrinsic mode functions (IMFs) and a residual component. When the IMFs and residual are summed together they form the original time series. The IMFs and residual component may display linear and or non-linear behaviour (amplitude and frequency modulation) depending on the nature of the time series being studied.

3. APPLYING EMD

Application of EMD to a time series is demonstrated using monthly precipitation data for the Melbourne Regional Office (see Figure 1a - Observed) for the period January 1995 to December 2004 (data provided by the Australian Bureau of Meteorology).

Figure 1 shows the four steps required to estimate the first IMF. First, the local extrema, both maxima and minima, are identified (see Figure 1a). Second, cubic splines are fit to the sequences of maxima and minima (see Figure 1b). Third, the mean of the two cubic splines is taken (see Figure 1c) and, fourth, the mean of the cubic splines is subtracted from the original time series and the remainder forms a residual (see Figure 1d). The residual shown in Figure 1d is the first estimate of the first IMF.

An IMF must satisfy two conditions (Huang *et al* 1998):

(i) the number of extrema (sum of maxima and minima) and the number of zero crossings must be equal or differ by one, and

(ii) the mean of the cubic splines (Figure 1c) must be equal to zero at all points.

In Figure 1d the number of extrema (maxima = 39 + minima = 40) = 79, while the number of zero crossings = 75, a difference of 4. The first estimate of the first IMF does not satisfy condition (i) and is therefore not an IMF. The difference between the number of extrema and zero crossings is due to the positive local minima in May 1998 and the negative local maxima in February 2003 (also shown in red circles in Figure 1d). Whenever there are multiple extrema between zero crossings, condition (i) is violated. The mean of the cubic splines, in Figure 1c, is also clearly not zero at any point.

3.1. Sifting to obtain the first IMF

In order to satisfy the above two IMF conditions the first estimate of the first IMF is subjected to a recursive sifting process. The four steps shown in Figure 1, initially applied to the observed time series, are applied to the first estimate of the first IMF, resulting in a new residual (the second estimate of the first IMF). Whether this residual satisfies the two IMF conditions is assessed and, if not, the sifting process continues. The result of the previous sift is used as input for the next sift, until a residual satisfying the two conditions is found.

Once a residual is found that satisfies the two IMF conditions, then this residual is designated the first IMF (see Figure 2b) and it is subtracted from the original observed time series.

3.2. Obtaining subsequent IMFs

The four steps and recursive sifting process are

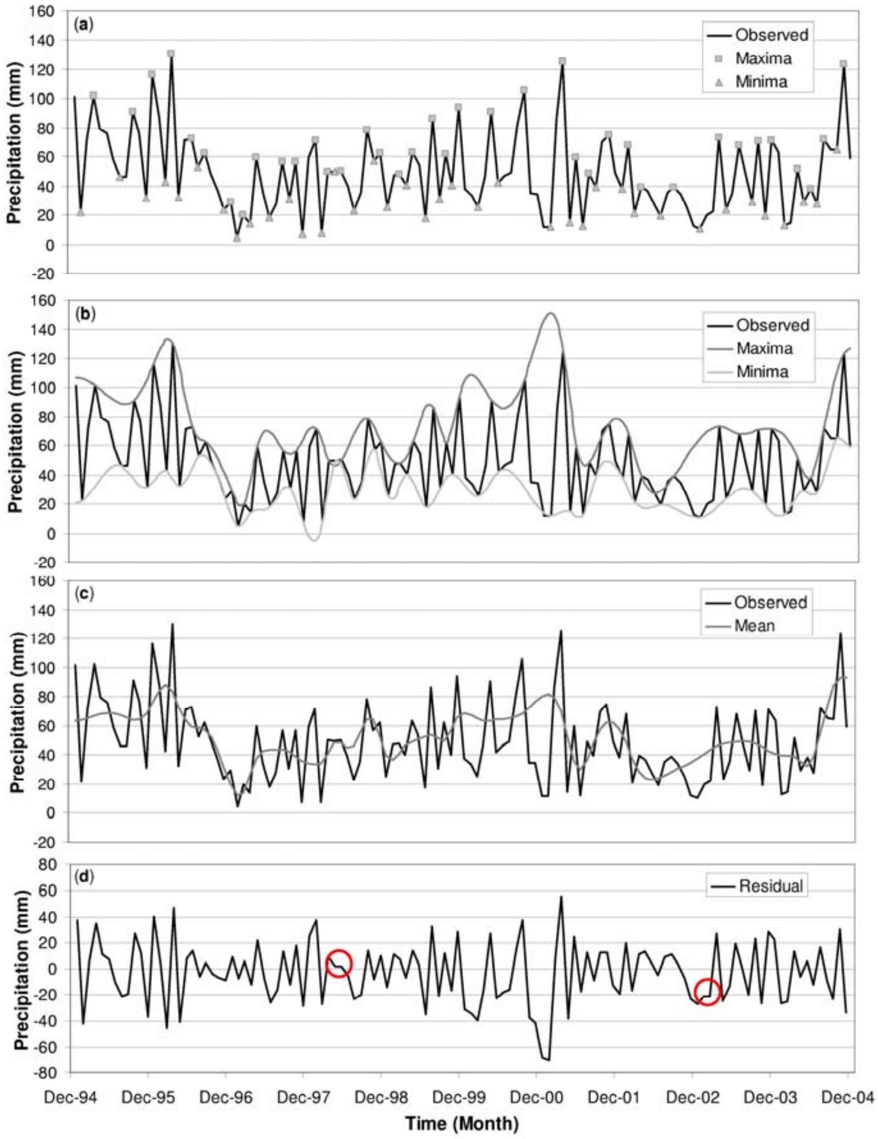

Figure 1. Steps involved in identifying the first IMF (see text for details).

then applied to the residual of the observed time series minus the first IMF to obtain the second IMF (see Figure 2c). As each subsequent IMF is identified it is subtracted from the residual of the observed time series minus the previous IMFs. The process of IMF identification continues until all IMFs are extracted (see Figure 2d for the third and final IMF) and a final residual remains (see Figure 2a - "Residual"). The final residual will be a constant, a monotonic trend or an incomplete (≤ 3 extrema, like that shown in Figure 2a for Melbourne) fluctuation with a cycle longer than the period of record.

3.3. Features of IMFs

The average period of each IMF can be calculated by dividing twice the sample size (2 x N) by the number of zero crossings. The average period and variance of each IMF are shown in Table 1 for the 120 months of Melbourne Regional Office precipitation data ($N = 120$).

The EMD methodology is based on the identification of local extrema, which results in the first IMF containing the highest frequency fluctuations. Once the first IMF is removed from the original time series only lower frequency fluctuations remain, thus subsequent IMFs have progressively lower frequencies (see Figure 2 and Table 1). Wu *et al* (2001) noted that the number of IMFs extracted from a time series is roughly equal to $\log_2(N)$.

Separately each IMF and the residual are significantly correlated (at the 1% level of significance) with the observed time series, but are not significantly cross-correlated for the

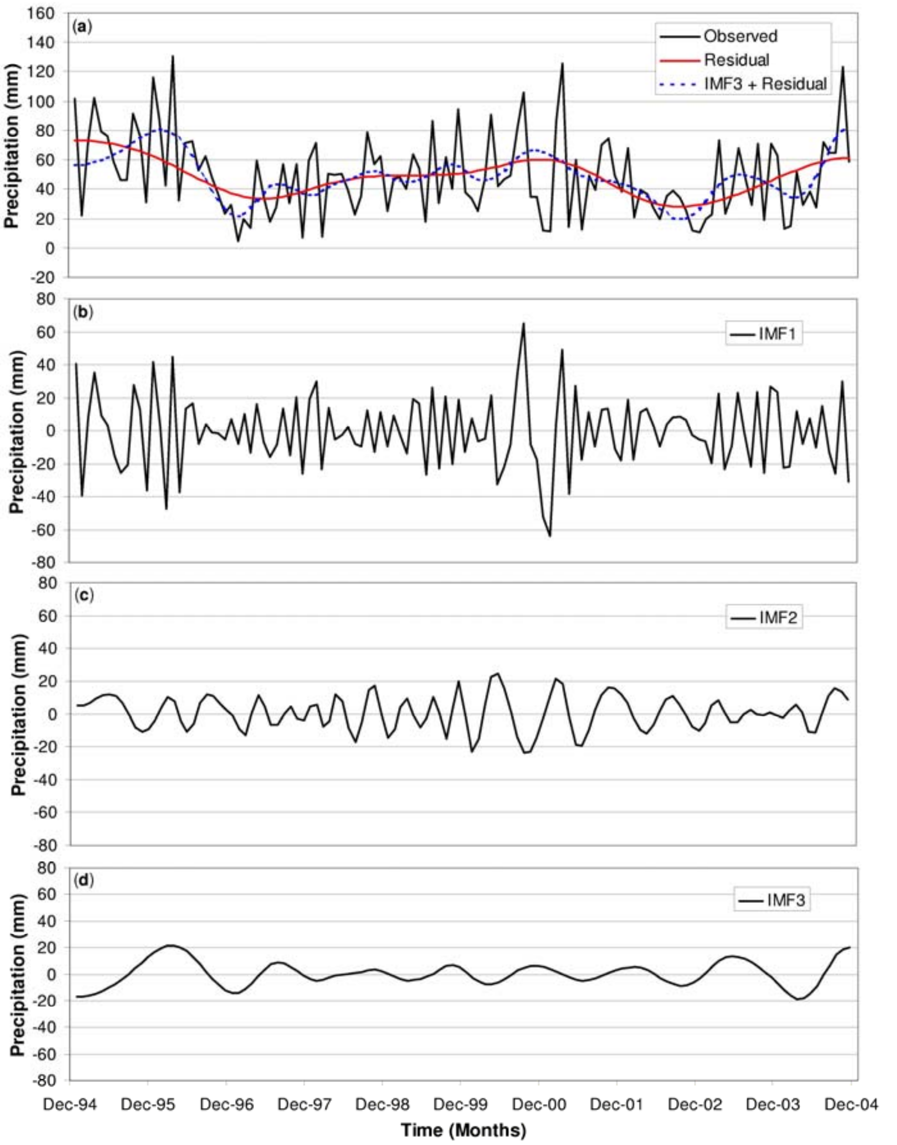

Figure 2. Observed data, IMFs, final residual and low frequency filter (see text for details).

Table 1. Average period and variance of IMFs.

	Average period (months)	Variance (mm^2)	Variance as a % of Observed (ΣIMFs + residual)
Observed		748	
IMF1	3.0	478	64 (59)
IMF2	6.3	106	14 (13)
IMF3	16.0	82	11 (10)
Residual		142	19 (18)
ΣIMFs + Residual		808	108 (100)

Melbourne precipitation data (not shown). The non-significant cross-correlation indicates that the IMFs and residual are independent / orthogonal to each other as expected. This indication is confirmed with an orthogonal index (Huang *et al* 1998) value of –0.04 for the Melbourne data.

Due to the independence of these IMFs, it is possible to sum the variances and determine the percentage of the observed time series variance that is explained by each IMF and the residual (see last column of Table 1). However, the IMF and residual variance do not always sum to the observed variance (in Table 1, there is a +8% difference) due to a combination of rounding error, the non-linearity of the original time series and the introduction of variance by the treatment of the cubic spline end conditions (see next section).

Summation of low frequency IMFs with the residual forms a low pass filter (Gloersen and Huang 2003, Huang *et al* 2003a). This is shown for the Melbourne precipitation data using the summation of IMF3 and the residual in Figure 2a, which represents periods greater than one year.

Similarly high pass filters can be created using the high frequency IMFs. Flandrin *et al* (2004) note that the EMD process is an automatic and adaptive (signal dependent) time-variant filter and that the IMFs should not be considered the result of a predetermined subband (frequency only) filtering process. Frequency filtering is difficult for non-linear / non-stationary data, due to the generation of harmonics (Gloersen and Huang 2003).

4. ISSUES WITH APPLYING EMD

A key feature of the EMD algorithm is its simplicity of application and robustness across a wide range of input time series. However, two issues, the sifting stopping rule and the cubic spline end conditions, require further attention for successful application of EMD.

4.1. Sifting stopping rules

The purpose of sifting is to obtain an IMF that satisfies the two rules outlined in the previous section. Huang *et al* (1998) note that sifting reveals hidden riding waves in the input series (like in Figure 1d, red circles) as well as tending to smooth uneven amplitudes in the resultant IMF. As riding waves are identified the number of extrema change, which is important for satisfying IMF condition (i). Over-sifting each IMF will produce a series of smooth amplitude IMFs where any physically meaningful amplitude variation will be sifted away. So the sifting process involves a trade off between under-sifting (producing incorrectly defined IMFs due to insufficient sifts to reveal all the riding waves) and over-sifting (producing less physically meaningful IMFs).

Initially Huang *et al* (1998) implemented a "Cauchy-like convergence criterion" (Huang *et al* 1999) for automatically deciding when to stop sifting. Based on minimising the difference between residuals in successive sifts to below a predetermined level, this criterion did not explicitly take into account the two IMF conditions, so the predetermined level could be obtained without the two IMF conditions being satisfied (Huang *et al* 2003b).

Rilling *et al* (2003) devised an alternate stopping criterion based on satisfying IMF condition (i) and then trying to satisfy condition (ii) without over-sifting. They used two thresholds, one designed to ensure globally small fluctuations in the mean of the cubic splines from zero, and the second allowing small regions of locally large deviations from zero.

The sifting stopping criterion recommended here is that proposed by Huang *et al* (2003b). Sifting is conducted until condition (i) is satisfied. At this sift, and each subsequent sift, the number of extrema and zero crossings are recorded and compared to those of the previous sift. When the number of extrema and zero crossings remain constant for five successive sifts, the sifting is stopped and the residual is designated as an IMF. Huang *et al* (2003b) found that IMFs produced using this sifting stopping criterion satisfied condition (i), were consistently orthogonal and were not over-sifted.

4.2. Cubic spline end conditions

Huang *et al* (1998) also highlighted the importance of how the end points of the maximum and minimum cubic splines were dealt with as a key practical implementation issue. They were particularly concerned that if the ends were left to oscillate wildly these artefacts of the methodology would propagate inwards and progressively corrupt the subsequent lower frequency IMFs.

The first technique for dealing with the spline end conditions proposed by Huang *et al* (1998) and slightly modified by Coughlin and Tung (2004) is to pad the beginning and the end of the time series with additional "characteristic" or "typical" waves. Huang *et al* (1998) based their additional waves on the two closest maxima and minima, while Coughlin and Tung (2004) based theirs on the closest maximum and minimum.

A simpler methodology proposed and tested by Rilling *et al* (2003) is to "mirrorize" the extrema closest to the edge, rather than pad the time series with extra data. Chiew *et al* (2005) used the average of the two closest maxima (minima) for the maximum (minimum) spline.

Here we present an alternate methodology (SZero) for handling the spline end conditions developed by the third author, based on the assumption that the slope of the cubic spline is equal to zero at the end points. The procedure is shown below.

i) Calculate the maximum (minimum) cubic spline for the series of maxima (minima).

ii) Assuming the slope of the spline equals zero at the end point, project the spline to the end.

iii) Check whether the observed data are less (greater) than the projected maximum (minimum) spline.

   a. If true, then finish.

Table 2. Comparison of number and period of IMFs produced for three cubic spline end condition rules.

	IMF	1	2	3	4	5	6	7	Total
Mirror	No.	8135	8133	7685	4260	959	53		29225
	Mean	3.0	6.4	13.0	24.1	41.2	61.0		
	StD	0.3	1.0	2.9	6.5	11.8	18.7		
Average	No.	8135	8135	7924	5308	1736	175	1	31414
	Mean	3.0	6.3	12.7	23.4	40.6	62.3	149.5	
	StD	0.3	1.0	2.8	6.2	10.9	16.7		
SZero	No.	8135	8132	6759	2650	335	8		26019
	Mean	3.0	6.8	14.8	29.6	49.4	91.8		
	StD	0.3	1.2	3.4	7.7	13.6	38.6		

b. If false, then recalculate the spline assuming that the observed end point is a maximum (minimum).

The last step of the above procedure is required to ensure that the cubic splines contain all the observed points (Huang *et al* 1998). Without this step the projected cubic splines may be exceeded by the observed data. Although not explicitly stated in Rilling *et al* (2003) and Chiew *et al* (2005), this last step is also required in the mirrorizing and average rules.

A comparison of the performance of SZero, along with a mirrorizing rule like Rilling *et al* (2003) and an average rule like Chiew *et al* (2005), is conducted below using a world data set of annual precipitation data (Peterson and Vose 1997). The world precipitation data set includes 8135 time series from a wide range of climatic zones with record lengths of 30 up to 299 years, which provide the EMD algorithm with a realistic range of hydroclimatic conditions to be tested on.

The EMD algorithm was applied to the 8135 annual precipitation records three times, each time with a different cubic spline end condition. Table 2 shows the results classified by the three different end condition rules. In Table 2 the number of times (No.) each IMF is produced by the EMD algorithm for the 8135 stations, the average period (Mean) of each IMF and the standard deviation (StD) of the period for each IMF are shown.

There is no difference, in Table 2, between end condition rules seen in the first IMF, however, for subsequent IMFs a pattern emerges. The average period of the IMFs produced by the SZero end condition rule are greater than the other rules and the standard deviations of the periods are also higher for the SZero rule. This pattern is explained by the fact that the SZero rule requires significantly less IMFs to decompose the annual precipitation data than the other two rules (total IMFs shown in last column of Table 2). In requiring less IMFs to decompose the time series, the SZero rule IMFs must cover a broader range of periods (higher Mean and StD) for a given IMF, it must also be introducing fewer extrema into the process that require additional IMFs to resolve (like the average rule). We conclude that the SZero rule is the most efficient of the three rules, due to the lower number of IMFs required to decompose the observed time series using this method.

The average variances of each IMF and residual as a percentage of the observed precipitation station variance (% of Obs.) for the three end condition rules are compared in Table 3. Ideally for orthogonal IMFs the average total variance of the IMFs and residual should sum to the observed variance (a value of 100 in the last column). The average rule is closest to 100, followed by SZero and the mirror rule. The mirror rule displays significantly higher standard deviations of IMF variance for low frequency IMFs and the residual, indicating that this rule is not consistently handling all the time series. This inconsistency is clearly seen in Figure 3, which summarises the summation of the IMFs and residual variance as a percentage of the observed variance (% of Obs.) at each of the 8135 stations for the three methods. The average rule outperforms the SZero and mirror rules across all time series, with the mirror rule producing very high summed variance compared to the observed variance in the last 5% of stations.

Clearly the mirror rule is inconsistent in enough cases to be discarded. Although SZero decomposes into less IMFs, the average rule consistently produces IMFs where the sum of variances are closer to the observed variance. It is possible that the average rule achieves the consistent summation of variances close to the observed variance by producing significantly more (5395) IMFs than SZero, which may be less physically meaningful. We therefore recommend SZero as the more efficient and physically meaningful spline end condition rule.

Table 3. Comparison of variance of IMFs for three cubic spline end condition rules.

	IMF	1	2	3	4	5	6	7	Residual	Total
Mirror	% of Obs.	60.7	23.2	14.8	10.5	5.3	3.7		14.9	118.9
	StD	0.13	0.15	0.40	0.83	0.11	0.06		0.73	1.62
Average	% of Obs.	59.5	21.4	12.2	7.7	5.3	4.5	1.1	10.6	109.6
	StD	0.13	0.10	0.08	0.07	0.06	0.09		0.11	0.15
SZero	% of Obs.	62.6	24.9	14.4	7.6	3.7	2.4		14.3	116.4
	StD	0.14	0.13	0.09	0.06	0.04	0.03		0.14	0.21

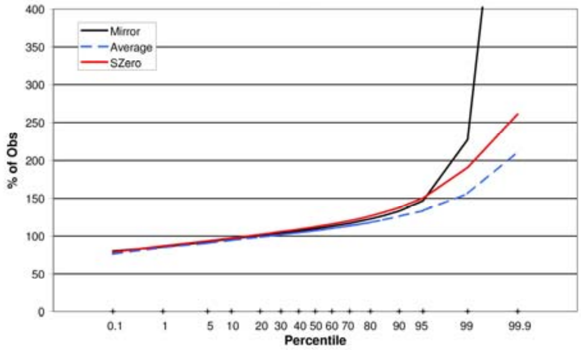

Figure 3. Percentiles of % of Obs. results from 8135 stations for three end condition rules.

5. CONCLUSIONS

In this paper EMD has been introduced and issues with its application have been discussed. In particular the IMF sifting stopping rule and the cubic spline end condition rule have been reviewed, investigated and discussed. A new cubic spline end condition rule has been proposed and its performance compared against two other end condition rules from the literature. The proposed rule, which assumes that the slope of the cubic spline equals zero at the end point, is found to be the most efficient of the three rules. The comparison was based on three separate applications of EMD, each application with a different end condition rule, to 8135 annual precipitation records from around the world.

6. ACKNOWLEDGMENTS

This research was conducted with the financial assistance of the Australian Research Council (grant number DP0449685) and Melbourne Water.

7. REFERENCES

Chiew, F.H.S., M.C. Peel, G.E. Amirthanathan and G.G.S. Pegram (2005), Identification of oscillations in historical global streamflow data using empirical mode decomposition. Seventh IAHS Scientific Assembly at Foz do Iguaçu, Brazil. IAHS Publ. 296, pp: 53-62.

Coughlin, K.T. and K.K. Tung (2004), 11-Year solar cycle in the stratosphere extracted by the empirical mode decomposition method, *Advances in Space Research* 34, 323-329.

Flandrin, P., G. Rilling and P. Gonçalvès (2004), Empirical mode decomposition as a filter bank, *IEEE Signal Processing Letters* 11(2), 112-114.

Gloersen, P. and N. Huang (2003), Comparison of interannual intrinsic modes in hemispheric sea ice covers and other geophysical parameters, *IEEE Transactions on Geoscience and Remote Sensing* 41(5), 1062-1074.

Huang, N.E., Z. Shen, S.R. Long, M.C. Wu, H.H. Shih, Q. Zheng, N.C. Yen, C.C. Tung and H.H. Liu (1998), The empirical mode decomposition and the Hilbert spectrum for nonlinear and non-stationary time series analysis, *Proceedings of the Royal Society London A* 454(1971), 903-995.

Huang, N.E., Z. Shen and S.R. Long (1999), A new view of nonlinear water waves: the Hilbert spectrum, *Annual Review of Fluid Mechanics* 31, 417-457.

Huang, N.E., M.L. Wu, W. Qu, S.R. Long, S.S.P. Shen, and J.E. Zhang (2003a), Application of Hilbert-Huang transform to non-stationary financial time series analysis, *Applied Stochastic Models in Business and Industry* 19, 245-268.

Huang, N.E., M.C. Wu, S.R. Long, S.S.P. Shen, W. Qu, P. Gloersen and K.L. Fan (2003b), A confidence limit for the empirical mode decomposition and Hilbert spectral analysis, *Proceedings of the Royal Society London A* 459(2037), 2317-2345.

Peterson, T.C. and R.S. Vose (1997), An overview of the Global Historical Climatology Network temperature database, *Bulletin of the American Meteorological Society* 78(12), 2837-2849.

Rilling, G., P. Flandrin and P. Gonçalvès (2003), On empirical mode decomposition and its algorithms. IEEE-EURASIP Workshop on Nonlinear Signal and Image Processing NSIP-03, Grado, Italy.

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.