Watching a Travelling Wave

Peter Bloomfield
John F. Monahan

Institute of Statistics Mimeo Series No. 1936

December 1988

Watching a Travelling Wave

Peter Bloomfield

John F. Monahan

December 1988

Department of Statistics

North Carolina State University

Abstract

Travelling waves are a feature of many spatio-temporal series. The use of single side band filtering and the complex singular value decomposition to isolate such waves is described. As with other applications of the singular value decomposition, such as principal components, the results can be difficult to interpret. In the case of travelling waves, however, a dynamic graphical display using color can give much insight.

1. INTRODUCTION

Estis, Davis, Bloomfield and Monahan (1989) describe the analysis of a large set of U.S. sea-level barometric pressure data from eight (1973–76, 79–82) recent winters. The study focused on behavior in the 2–10 day "synoptic" time scale. Since this behavior is largely that of travelling waves, the data analysis procedures were oriented towards detecting such features.

Travelling waves are discussed briefly in Section 2, first in a very simple form and then more realistically. Section 3 contains a description of a two-step analysis for isolating waves of the more general form. The two steps are single side band filtering and the complex singular value decomposition. In Section 4 we describe a graphical procedure for dynamically displaying the results of such analyses.

2. TRAVELLING WAVES

Consider a physical quantity p, such as sea-level barometric pressure, that is defined at each point (x,y) in a plane region and at each time in an interval. Suppose that for some constants R, θ_1, θ_2 and f,

$$p(x,y,t) = R \cos\{2\pi(ft - \theta_1 x - \theta_2 y)\} \ . \qquad (1)$$

Then for $t = t_0$, fixed, the surface $p(x,y,t_0)$ is a cosine wave of amplitude R, with maxima along the lines $\theta_1 x + \theta_2 y = ft_0 + k$ for any integer k. The lines of maxima mark parallel wavefronts equally spaced with wavelength $\ell = (\theta_1^2 + \theta_2^2)^{-1/2}$. As t changes, the wavefronts move with velocity $(\theta_1 f\ell^2, \theta_2 f\ell^2)$. At a fixed point (x_0,y_0), p oscillates with frequency f (cycles per unit time), and with amplitude R. This is a simple example of a *travelling wave* in the plane.

Travelling waves in real phenomena often show three types of behavior that cannot be modeled in such a simple way.

1. The *amplitude* of the wave may vary from place to place and from time to time.

2. The *lines of maxima* may be

   (a) curved

   (b) not parallel

   (c) not equally spaced.

3. The *time behavior* may not be that of a simple periodic oscillation. We can generalize the model so as to allow all three types of behavior without losing tractability. Consider the equation

$$p(x,y,t) = R_1(x,y)R_2(t)\cos[2\pi\{F(t) - \Phi(x,y)\}]. \qquad (2)$$

The amplitude R is now the product of an arbitrary function of the spatial coordinates (x,y) and an arbitrary function of time. Replacing the linear function ft by the arbitrary function F(t) allows waves to move with speeds that change over time, perhaps even reversing their direction. The lines of maxima of the cosine factor at time t_o are the curves $\Phi(x,y) = F(t_o) + k$, and can therefore curve and be unequally spaced. These are not typically the locations of the maxima of the product, but we shall use them to follow progress of wavefront. The velocity of these curves depends both on space and time: at location (x,y), the velocity at time t is

$$\left(\frac{\Phi_x f(t)}{\Phi_x^2 + \Phi_y^2}, \frac{\Phi_y f(t)}{\Phi_x^2 + \Phi_y^2} \right)$$

where

$$\Phi_x = \frac{\partial\Phi}{\partial x}, \quad \Phi_y = \frac{\partial\Phi}{\partial y}, \quad f(t) = \frac{dF}{dt}.$$

The equation (2) is still not completely general, in that the spatial and temporal influences still *separate* in a sense to be made more specific in the next section. However, it is both tractable and sufficiently general to provide a good approximation to many phenomena.

3. EXTRACTING TRAVELLING WAVES

Suppose that we wish to analyze a set of data

$$P_{ij} = p(x_i, y_i, t_j): \quad i = 1,\ldots,m; \; j = 1,\ldots,n \; ,$$

by identifying any travelling waves of the forms (1) or (2).

A common method of analysis is principal components. Here, a low dimensional approximation of the covariance matrix (in space)

$$C_{ii'} = \frac{1}{n-1} \sum_{j=1}^{N} (P_{ij} - \bar{P}_i)(P_{i'j} - \bar{P}_{i'})$$

is obtained from its expansion in terms of its eigenvalues λ_k^2 and eigenvectors $a^{(k)}$

$$C_{ii'} = \sum_{k=1}^{r} \lambda_k^2 \, a_i^{(k)} a_{i'}^{(k)} \quad \text{where } r = \min\,(m,n) \; .$$

The original data can then be expressed by its *singular value decomposition*

$$P_{ij} = \sum_k \lambda_k \, a_i^{(k)} b_j^{(k)}$$

where $b_j^{(k)} = \Sigma \, P_{ij} a_i^{(k)}/\lambda_k$ are the *component scores*, carrying temporal information and where partial sums ($k = 1$ to $s < r$) are least square fits of such reduced rank approximations. This expansion is also knows as *empirical orthogonal function* analysis (Wallace and Dickinson (1972, Wallace (1972), Barnett (1983)) because of the orthogonality relations

$$\sum_i a_k(x_i, y_i) a_{k'}(x_i, y_i) = \sum_j b_k(t_j) b_{k'}(t_k) = \delta_{kk'} = \begin{cases} 1 & k = k' \\ 0 & k \neq k' \end{cases} \; .$$

If the data P_{ij} were dominated by the simple form of travelling wave (1), the covariance matrix would be approximately

$$C_{ii'} = ½R^2[\cos(2\pi\gamma_i)\cos(2\pi\gamma_i) + \sin(2\pi\gamma_i)\sin(2\pi\gamma_i)]$$

where $\gamma_i = \phi_1 x_i + \phi_2 y_i$. Although clearly a matrix of rank 2, principal components cannot extract information on the waveform γ_i. And, more importantly, all information regarding the frequency f is lost.

If the complex portion of a travelling wave were also observed, instead of just the real part, then for data P_{ij} without error

$$P_{ij} = R \exp\{i(2\pi f t_j - \gamma_i)\}$$

and the complex covariance matrix would then be

$$C_{ii'} = R \exp\{i2\pi\gamma_i\}\exp\{-i2\pi\gamma_{i'}\}$$

which is a matrix with rank one. Additional waves at different frequencies would then sum in C. The complex singular value decomposition of the data then is

$$\Sigma\lambda_k a_i^{(k)} b_j^{(k)*}$$

where asterisk (*) denotes complex conjugate.

Here each term k would represent a single wave, instead of two terms for each wave. The key is to add a complex part to the real data so that wave phenomena are more easily detected.

Since a real wave with frequency f can be written as a sum of two complex waves with frequencies ± f, a complex part of the data P_{ij} can be constructed and added to it through the use of a single sided filter.

If the data P_{ij} are transformed into the frequency space through the use of the discrete Fourier transform (DFT)

$$d_i(f_k) = \frac{1}{n}\sum_j P_{ij}\exp\{-2\pi f_k t_j\} \qquad f_k \varepsilon \{-½, +½\} ,$$

then a single sided filter would zero out the negative frequencies, double the values at the positive frequencies, and transforming back produces the

$$\tilde{P}_{ij} = \sum_{\{f_k>0\}} 2d_i(f_k)\exp\{+2\pi i f_k t_j\}$$

which has the appropriate complex part to complement any wave phenomena present in the real part. The complex singular value decomposition of the filtered data $\tilde{P}_{ij}$ would then reveal travelling waves as distinct single components.

For the problem at hand, the original data consists of 180 consecutive observation (90 days) from each of eight winters (Jan.-March, 1973-76, 1979-82). Moreover, this study focussed only on phenomena with 2 to 10 day cycles. As a result, the covariance matrix for the complex principal components analysis/singular value decomposition is the average of the sample covariance matrices of the single side band filtered data. Since the sampling rate was twice daily, the frequency range of interest is $\{f_0,f_1\}$ where $f_0 = 1/20$, $f_1 = ¼$. The single side band filtering is achieved by summing only over $[f_0, f_1]$ in computing the inverse DFT

$$\tilde{P}^w_{ij} = \sum_{f_k \varepsilon [f_0,f_1]} 2d_i(f_k)\exp\{2\pi i f_k t_j\}$$

for winters w = 1,...,8. The complex covariance matrix $\tilde{C}$ is then formed by averaging over winters,

$$\tilde{C}_{ii'} = \frac{1}{8}\sum_w \sum_j \tilde{P}_{ij}\tilde{P}^*_{ij}$$

which is Hermitian ($\tilde{C}^*_{ii'} = \tilde{C}_{i'i}$). A computational advantage can also be exploited by computing $\tilde{C}$ from the DFT of $\tilde{P}\psi_j$ using the discrete form of the Parseval relation

$$\tilde{C}_{ii'} = \frac{1}{8} \sum_{w} \sum_{f_k \varepsilon [f_0, f_1]} d_i^w(f_k) d_{i'}^w(f_k)^* .$$

The eigenvalues λ_k^2 (real and positive) and eigenvectors $a^{(k)}$ of $\tilde{C}$ are then computed, as well as the component scores for winter w

$$b_{wj}^{(k)} = \lambda_k^{-1} \Sigma \; a_i^{(k)} (P_{ij}^w)^* \; .$$

Again notice that $a_i^{(k)}$ carries spatial information and $b_{wj}^{(k)}$ carries temporal information. Travelling wave phenomena of the simpler form (1), or the more general complex from of (2)

$$\{R_1(x,y)\exp[-2\pi i \Phi_k(x_i,y_i)\}\{R_2(t)\exp[2\pi i F(t)]\}$$

can be viewed by examination of the complex quantities $a_i^{(k)}$ and $b_{wj}^{(k)}$ rewritten in polar coordinate form as

$$R_{1,k}(x_i,y_i)\exp\{-2\pi i \Phi_k(x_i,y_i)$$

and

$$R_{2,k}(t_{wj})\exp\{+2\pi i F_k(t_{wj})\}$$

respectively.

Thus, if the data contain travelling waves with frequencies in the band of interest, single side band filtering followed by the complex singular value decomposition will reveal them as strong components. It should be noted that the components that result from the complex singular value decomposition satisfy the (complex) orthogonality relations

$$\sum_i a_k(x_i,y_i) a_{k'}(x_i,y_i)^* = \sum_j b_k(t_j) b_{k'}(t_j)^* = \delta_{k,k'}$$

or

$$\sum_i R_{1,k}(x_i,y_i)R_{1,k'}(x_i,y_i)\ \exp[-2\pi i\{\Phi_k(x_i,y_i) - \Phi_{k'}(x_i,y_i)\}]$$

$$= \sum_j R_{2,k}(t_j)R_{2,k'}(t_j)\exp\ [2\pi i\{F_k(t_j) - F_{k'}(t_j)\}]$$

$$= \delta_{k,k'}\ .$$

It may well happen that waves of physical interest do not satisfy these conditions. Thus in order to aid interpretation it may be necessary to transform some number of components, say the first K, to the form

$$a_k^R(x_i,y_i) = \sum_{\ell=1}^{K} g_{k,\ell}\ a_\ell(x_i,y_i)$$

$$b_k^R\ (t_j) = \sum_{\ell=1}^{K} h_{k,\ell} b_\ell(t_j)\ ,$$

where the matrices $G = ((g_{k,\ell}))$ and $H = ((h_{k,\ell}))$ are related by $H^*G = \mathrm{diag}(\lambda_k)$. The transformed components then satisfy

$$\sum_{k=1}^{K} a_k^R(x_i,y_i)b_k^R(t_j)^* = \sum_{k=1}^{K} \lambda_k a_k(x_i,y_i)b_k(t_j)^*\ ,$$

and thus are simply another way of expressing the approximation to the complex data provided by the first K complex components. Transformations to aid interpretation might aim for localization in space, time, or temporal frequency. We do not address such issues in this article.

4. DISPLAYING A TRAVELLING WAVE

We have seen how a set of data may be approximated by a sum of travelling waves of the form (2). However, numerical determination of the elements of such a wave does not by itself necessarily give insight as to the phenomena that cause the wave. We need to display these elements in appropriate ways.

The temporal component $R_2(t)\exp\{-2\pi iF(t)\}$ is a complex function of one variable, and hence can be displayed by a pair of graphs: $R_2(t)$ vs. t, and F(t) vs. t. An effective form of the latter graph is to present F(t) and (F(t) + 1)*modulo* 2 vs t.

The spatial component $R_1(x,y)\exp\{2\pi i\Phi(x,y)\}$ presents greater difficulties. It is easy to make displays of R_1 and Φ separately, such as contour plots, and much may be learned from such displays. The contour plot of R_1 shows which regions are strongly influenced by the wave, and the contours of Φ are precisely the curves of maxima at different phases of the wave. However, we have found that a single display combining both elements of the spatial component is much easier to interpret.

We constructed such a display by using color. Regions with zero amplitude were represented by a neutral grey. Regions of high amplitude were represented by an intense color, with the hue indicating the phase. A circular set of hues was constructed to represent the maximum intensity, approximating the sequence of hues in the rainbow. Regions of intermediate amplitude were shown with an appropriately muted hue. This resulted in a single, static display that presented both amplitude and phase information.

However, a representation of a travelling wave should really show it travelling. We achieved this as follows. Spatial components are in fact defined only up to an arbitrary scalar complex multiplier of unit magnitude. Thus amplitudes are well defined, but phases are defined only up to an arbitrary added constant. We can make a *dynamic* display of a component by adding progressively larger constants to all of the phases, thus stepping through all of these equivalent versions of the component. In fact, if the corresponding temporal component were a pure sinusoid, this would represent the temporal evolution of the wave. In our application, all the temporal components were roughly approximated by sinusoidal functions, so the dynamic display gave good insight into the contribution of each wave to the overall sea-level barometric pressure. In particular, the variations in the velocity of the wave fronts across the region were self-evident in the dynamic display, whereas they required conscious measurement from the static contour plots.

We have implemented these color displays in two hardware-software environments. In the first, the displays were created on a Vectrix [1]VX*-384 display attached to a DEC VAX-11/750[2] minicomputer running the Eunice [3] emulator of the Unix[4] operating system under the VMS[2] host operating system. The VX-384 provides a 672 x 480 display with 512 on-screen colors chosen from a palette of 2^{24} possible colors (8 bits, or 256 levels, of each of red, green, and blue). We reserved one color for the borders of the display and one to show the outline of the U.S., Canada and Mexico and their state lines, leaving a maximum of 510 colors for the wave.

The data and singular components were defined on a 20 (E-W) x 14 (N-S) grid, which was interpolated and smoothed onto a finer grid (chosen at run time by the user, but defaulting to the maximum resolution of the display) using B-spline surfaces (Newman and Sproull, 1979; degree chosen at run time, defaulting to quadratic). The amplitude and phase for each pixel were computed and allotted to linearly spaced bins. All pixels allotted to the lowest amplitude bin were displayed as a neutral grey. Other pixels were displayed by mixing the neutral grey with an intense hue, the hue determined by the phase bin and the mixing parameter determined by the amplitude bin. The total number of colors was thus

---

[1]Vectrix is a trademark of the Vectrix Corporation.
[2]DEC, VAX and VMS are trademarks of Digital Equipment Corporation.
[3]Eunice is a trademark of the Wollongong Group.

1 + (number of amplitude bins - 1) x (number of phase bins)

and this was required to be $\leq$ 510-2. We found that 11 amplitude bins and 50 phase bins gave a good display. There was a perceptible change in color intensity between adjacent amplitude bins, and these boundaries created a contour plot of amplitude. Changes in hue from one phase bin to the next were less conspicuous, so that hue appeared to change almost continuously along bands of constant amplitude.

In order to create the dynamic display, the screen had to be changed several times a second. The communications channel between the VAX and the VX-384 was limited to 19.2 Kbaud, and therefore the display updates could not be achieved by recreating the display. Intitial setup required nearly 8 minutes of cpu time and 3½ minutes of communication time, partly overlapped, for a total setup time of around 10 minutes for a 672 x 480 display.

However, the display data merely instructs the VX-384 which color from the current selection of 512 to display at a given pixel, and the colors can therefore be changed rapidly specifying a new selection. A small program in 8088 machine code was downloaded to the VX-384, which permuted the selection of colors appropriately on receipt of an instruction from the VAX. When operated with no delays other than video synchronization, this gave around 20 updates per second, which resulted in a visually smooth progression of the colors across the display, and a complete cycle in around 2½ seconds (when using 50 phase bins). Provision was also made for single-stepping the display.

For the second implementation, the C-language code was ported from the VAX to an Amiga[5] personal computer equipped with the Lattice C-compiler.[6]

---

[4]Unix is a trademark of AT&T.
[5]Amiga is a trademark of Commodore Computers.
[6]Lattice C is a trademark of SAS Institute, Inc.

The Amiga display is capable of 320 or 640 pixels horizontally by 200 or 400 pixels vertically. It has a palette of 2^{12} = 4096 colors, and can display 32 of these in 320 pixel horizontal resolution or 16 in 640 pixel resolution. The 400 pixel vertical resolution is achieved by interlacing the display, which results in unacceptable flicker on the color monitor we were using. Accordingly, we found our best results with 320 x 200 resolution and 32 colors. One color was reserved for the map overlay and the background, and the other 31 were used to display the image. We found that 4 amplitude bins (grey plus 3 levels of color) and 10 phase bins gave a useful display, although, necessarily, one that was hampered by the small number of colors. Updates were controlled at 5 per second to give a 2 second cycle time, with somewhat abrupt transitions (36^{o} per update). Again, provision was made for single-stepping. Initial set up for a 320 x 200 display took around 45 minutes.

The Amiga also has a display mode in which all 4,096 colors can be displayed simultaneously. This would provide greater display flexibility, but does not allow for rapid permutation of the displayed colors, and consequently we have not explored this option.

5. CONCLUSIONS

We have described a technique for isolating travelling waves in spatio-temporal data, and for displaying the temporal and spatial aspects of these waves using minicomputer and personal computer hardware/software configurations. Substantive conclusions obtained by the use of these techniques are presented by Estis, Davis, Bloomfield and Monahan (1989). Computational details of the single side band filtering, complex singular value decomposition, and display algorithms are available from the authors.

Linear transformation of a selected number of leading components was mentioned in Section 3 as possibly aiding interpretation. In the meteorological context, one goal might be localization of the transformed *spatial* components. Time domain or frequency domain localization of the transformed *temporal* components would also be reasonable, although mutually incompatible goals. These issues remain to be explored.

REFERENCES

Barnett, T. P. (1983) "Interaction of the Monsoon and Pacific Trade Wind System at Interannual Time Scales. Part I: The Equatorial Zone," *Monthly Weather Review*, Vol. 111, pp. 756-773.

Estis, F. L., J. M. Davis, P. Bloomfield and J. F. Monahan (1986) "Complex Principal Components Analysis of Sealevel Pressure over the GALE Region". (In preparation).

Horel, J. D. (1984) "Complex Principal Component Analysis: Theory and Examples," *Journal of Climate and Applied Meteorology,* Vol. 23, pp. 1660-1673.

Newman, William M. and Robert F. Sproull (1979). *Principles of Interactive Graphics.* New York: McGraw-Hill.

Wallace, John M. and Robert E. Dickinson (1972) "Empirical Orthogonal Representation of Time Series in the Frequency Domain. Part I: Theoretical Considerations," *Journal of Applied Meteorology,* Vol. II, pp. 887-892.

Wallace, John M. (1972) "Empirical Orthogonal Representation of Time Series in the Frequency Domain. Part II: Application to the Study of Tropical Wave Disturbances," *Journal of Applied Meteorology,* Vol. II, pp. 893-900.