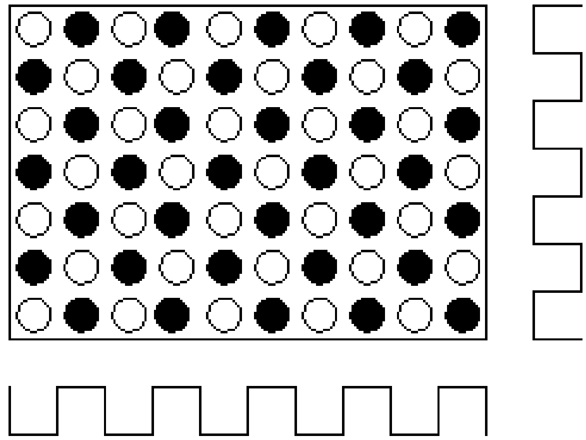

Figure 1: Sample data set with checkerboard pattern, corresponding to data values of +1 and -1, with EOF modes in space and time plotted on sides.

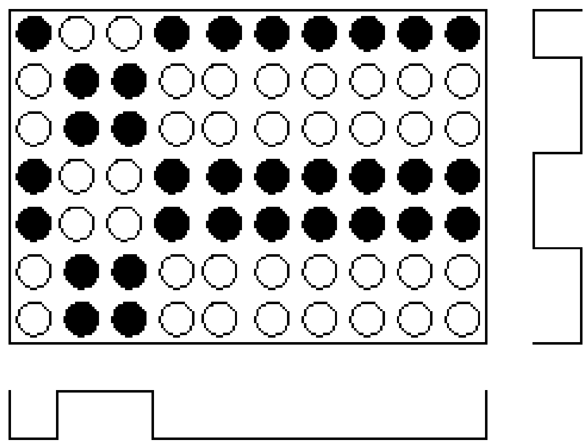

Figure 2: Sample data set with more structured pattern, corresponding to data values of +1 and -1, with EOF modes in space and time plotted on sides.

Empirical Orthogonal Functions

Empirical orthogonal functions represent dominant patterns of variability in a space-time record, using as few spatial patterns as possible. (EOFs go by a variety of names, and you'll also often hear "principal component analysis".) Just to get started, let's look at a few simplified examples.

Figure 1 shows a checkerboard pattern, with black and white data points that can be interpreted as alternating between +1 and -1. You can consider one dimension of the grid to be time and the other space, and for the purposes of this discussion, it isn't important which one is which. If you look at the matrix closely, you will see that all of the rows have the same pattern, and would be identical if you multiply them by either +1 or -1. Similarly, all of the columns have the same structure, though some vary in sign. In other words the variability in this entire domain can be explained by a single mode. The space and time patterns of this mode are plotted on the sides.

Figure 2 shows another idealized data set. In this case, although the pattern is not a checkerboard, it still retains the same characteristic that we saw in Figure 1: all of the rows have the same structure, and all of the columns have the same structure. Again a single mode in space and time is sufficient to reconstruct the entire data matrix. Both of these examples represent cases where EOFs provide efficient means to represent the dominant and potentially complicated patterns of variability within a data set. In these cases, a large matrix of data can be reduced to a much more compact form represented by one vector for spatial variability and one vector for time variability.

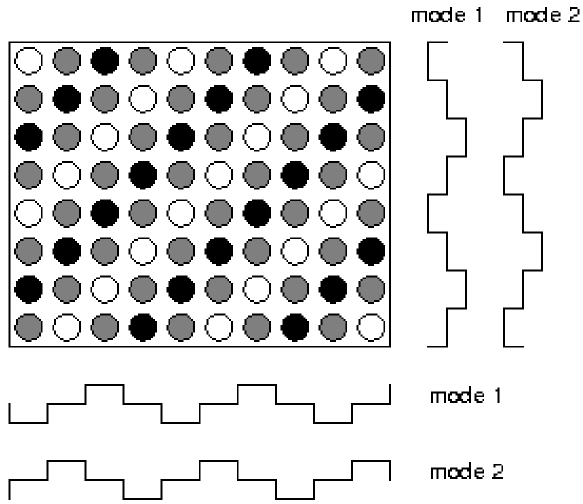

Figure 3: Sample data set with propagating wave pattern, with data values corresponding to +1, 0, or -1. EOF modes 1 and 2 in space and time are plotted along the sides.

What about a more complicated data set? Figure 3 shows a propagating wave pattern. In this case, there are two different distinct rows (e.g. the matrix has rank 2), and we can define two EOF modes. Depending on the exact dimensions of our domain, modes 1 and 2 may represent equal amounts of variance, or either mode 1 or mode 2 could be more energetic. While two EOF modes are required to explain the propagating wave pattern, a two-dimensional Fourier transform would be able to represent this wave as a feature with a single frequency and wavenumber. In this case, EOFs are usable, but they may not represent the most compact way of representing the data. And importantly, although the EOFs will identify two modes, in fact the modes are very much a pair, and it makes sense to analyze their effects jointly.

Just to complete this thought about propagating modes, here is a slightly more complicated example. Figure 4 shows a propagating signal, and the fraction of variance explained by the leading mode EOFs. The first two EOFs capture about 40% of the variability in this system, which isn't bad, but it's not great, considering that all of the variability in this system occurs at a single point in frequency and wavenumber space.

Figure 5 shows the leading order spatial EOFs and the signal recontructed from the first two EOFs. The two EOFs are essentially sine waves that are out of phase—the EOF is using two modes (a sine and a cosine) to reconstruct the original propagating variability. Mathematically, we've approximated the data Y as:

$$\hat{Y} = \sum_{i=1}^{2} U_i S_{i,i} V_i^T = U(:, 1:2) S(1:2, 1:2) V(:, 1:2)^T \quad (1)$$

The resulting signal is a fair representation of the data, but it's a bit awkward compared with the original signal. As noted above, there are really just two problems with this EOF. One is that we've taken a signal that is just a single point in frequency/wavenumber space and represented it with two full EOF modes, so we've lost a lot of compactness. The second problem is that we can't fully distinguish mode 1 and mode 2. They both explain about 18% of the variability, and in this case they're really just two halves of the same oscillatory mode.

Finally, Figure 6 shows an example of a field that has a fair amount of structure. For example, all of the values in the upper right above the diagonal are black. But the rows are not obviously related to each other, and an EOF is unlikely to capture the patterns of variability that define this particular matrix.

Now that we've looked closely at EOF patterns in some test cases, let's quickly review how to compute them efficiently. Here are three methods for finding EOFs for a data matrix Y:

- Find the covariance matrix formally, and then find its eigenvalues and eigenvectors. If the matrix Y has no gaps, then its covariance matrix is just YY^T or Y^TY. In Matlab you can solve the eigenvalue

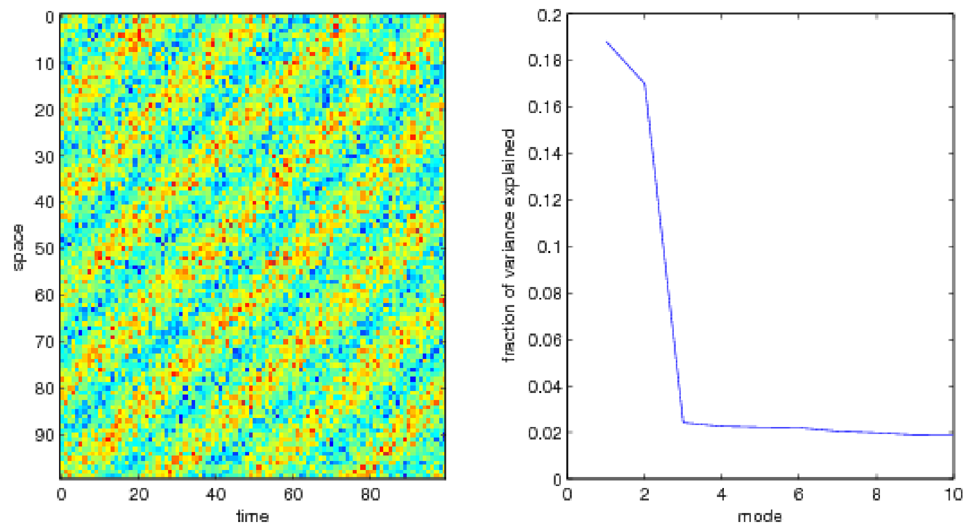

Figure 4: (left) Random data with propagating mode signal. (right) Fraction of variance explained by first few EOFs for this signal.

equation with "[V,lambda]=eig(Y'*Y)". V gives us the empirical modes h in the space domain. To find the empirical modes g in the time domain, we could solve for YY^T. It's probably more efficient to project our data onto the time domain modes and find the appropriate amplitudes, keeping in mind that $g = Yh^T$.

- We can also find eigenvalues and eigenvectors using the singular value decomposition of YY^T or Y^TY. We know that any matrix can be represented with its singular value decomposition USV^T. Since YY^T is symmetric, its singular value decomposition will be symmetric as well: $YY^T = VSV^T$, where S represents the eigenvalues and V the eigenvectors or empirical orthogonal modes. The nice thing about the singular value decomposition is that we've already thought about it for least-squares fitting, and we know that the vectors V are defined to be orthonormal and that the singular values are sorted by size. In this case we can define $h = V$.
- If we're considering the SVD, then why not skip the covariance matrix. We know that any matrix Y can be decomposed into its singular value decomposition. $Y = USV^T$. That means that $YY^T = USV^TVSU^T = US^2U^T$ and $Y^TY = VSU^TUSV^T = VS^2V^2$. Either way U and V for matrix Y are the same singular values that we'd find by computing the SVD of YY^T or Y^TY. The singular values S are a measure of amplitude, and we can interpret them as telling us how much variance is tied up in each mode. S^2 in the SVD of Y are equivalent to S computed above from YY^T. Using the SVD is far and away the easiest strategy because we haven't had to do anything about determining the covariance matrix of our data, and we've managed to represent it with both spatial and temporal modes. (We do have to make sure that we remove a sensible mean.)

Although people often distinguish between temporal and spatial modes by giving one of them an amplitude that is not one, the SVD format suggests that there's really no difference between temporal and spatial modes. Both are automatically normalized to have length 1. The singular values tell us how much variability is associated with each mode, and it often seems more sensible to keep the singular values separate, rather than multiplying them by one or the other of the modes. Thus, instead of saying that $h^T = V_i$ and $g = S_{j,j}U_j$, we can interpret U and V as spatial and temporal modes respectively, and the diagonal of S as corresponding amplitudes.

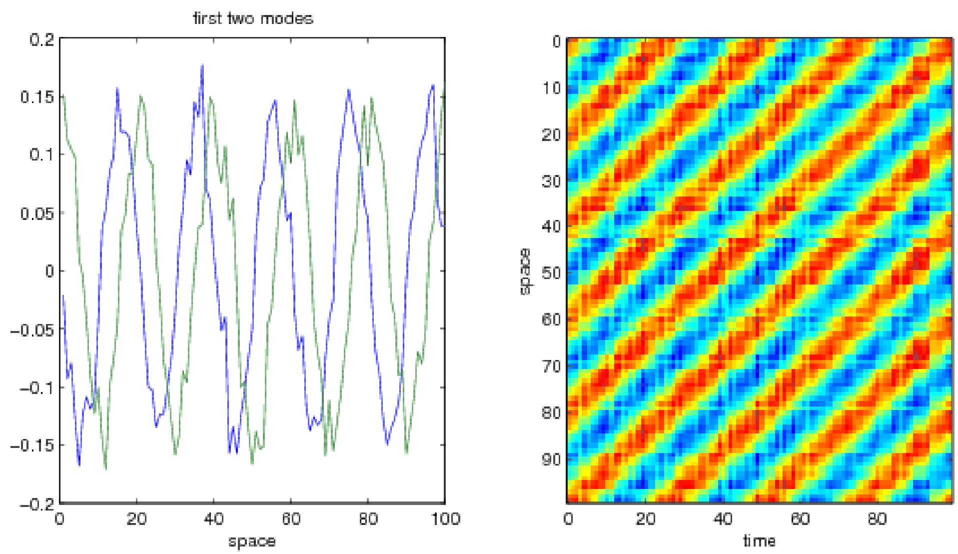

Figure 5: (left) First two spatial modes for random propagating signal in Figure 4. (right) Reconstruction of original signal using first two EOF modes.

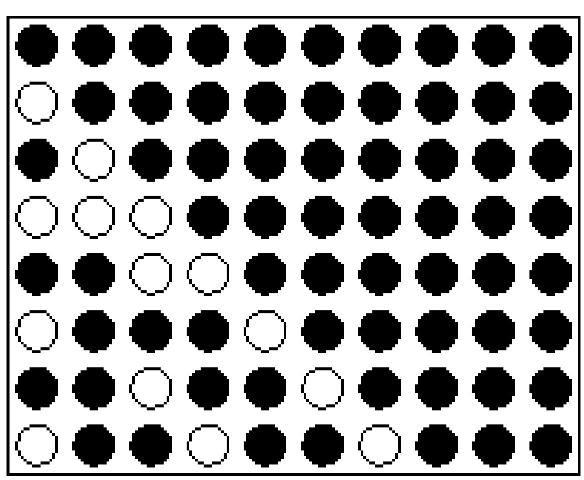

Figure 6: Sample data set with clear patterns but no trivial EOF modes.