

Visualization of Scientific Video Data Using KL Decomposition

K. A. Robbins

Abstract—Fast methods are developed for visualizing and classifying certain types of scientific video data. These techniques, which are based on KL decomposition, find a best coordinate system for a data set. When the data set represents a temporally ordered collection of images, the best coordinate system leads to approximations that are separable in time and space. Practical methods for computing this best coordinate system are discussed, and physically significant visualizations for experimental video data are developed. The visualization techniques are applied to two experimental systems — one from combustion and the other from neurobiology to show how relevant information can be quickly extracted from video data. These techniques can be integrated into the video acquisition process to provide real-time feedback to the experimentalist during the operation of an experiment.

Keywords—Scientific visualization, real-time visualization, video analysis.

I. INTRODUCTION

The development of fast techniques for visualization of scientific video data is motivated by a need for immediate feedback during the operation of laboratory experiments. Scientists use such feedback to select experimental parameters and follow behavior as a function of these parameters. Such *experimental steering* is analogous to *computational steering*. Computational steering is the on-line management of an application while it is executing. (See [26] for an up-to-date bibliography on this subject).

Visualization is an important component of computational steering [3], [12], but the emphasis of much of the current work in this area is on sophisticated visualization techniques that give the user insight into the detailed structure of a particular computational run [25]. In experimental systems that continuously produce results or respond rapidly to parameter changes, the requirements for speed in visualization are quite demanding. Limited state information is available, narrowing the field of candidate visualizations. Furthermore, the emphasis is on classification and comparison rather than on determining detailed structure in a particular case. Ideally such visualization techniques would be integrated directly into the video acquisition system to provide feedback in real time [16].

The management of visual information is an important emerging area of research because of the rapid growth of video archives on the Internet [8]. Techniques such as automatic abstracting [15], [27], [28] will be important for organizing and archiving video data as it is acquired, but these techniques are not fast enough to be used practically in the laboratory for steering. There are currently no techniques that work in real time for general video libraries,

but the restricted domain used in the current work allows some progress to be made in this direction.

This paper emphasizes simple visual metaphors that can be constructed in near real time to help a user classify and understand scientific video data that varies rapidly in time and space. These visualizations can also be used in conjunction with more traditional representations of high-dimensional scientific data [9]. The visualizations are illustrated using two experimental systems — a combustion experiment from the University of Houston [6], [7] and a neurobiological experiment from the University of Texas at San Antonio [19], [20]. A brief description of these experiments is presented in Section II. Section III gives an overview of the theoretical results that underpin the visualization techniques and their efficient implementation. Section IV presents visual metaphors and interprets visualizations based on these metaphors for the two applications. Section V looks at sampling and resolution issues associated with these techniques, and Section VI discusses the limitations of the techniques and their application to numerical data.

II. MOTIVATING EXPERIMENTS

In the combustion experiment at the University of Houston, scientists establish a flat flame above a porous plug burner in a highly controlled combustion chamber and photograph the flame front with a CCD camera. Because of its variety of dynamic states and precise controllability, this experiment is being used as a testbed for the study of dynamic pattern formation [13], [14], [17]. The experiment has four controlling parameters — the chamber pressure, the types of fuel and oxidizer, the ratio of fuel to oxidizer, and the flow rate of the fuel-oxidizer mixture through the burner. Hundreds of different dynamic states have been observed in this experiment as the parameters are varied. Fig. 1 shows sequences of frames from experimental videotape for three of the simpler states: a) a state in which two rings of cells counter-rotate, b) a standing wave and c) a hopping state. In the counter-rotating ring state, each ring appears to act as a unit, while the standing wave alternates between two configurations each containing two cells. The standing wave configurations are arranged at 90° relative to each other. In the hopping state, each of four cells expands and moves in turn to a new rotational position.

The neurobiological experiment from the University of Texas at San Antonio studies the response of the visual cortex of the pond turtle to light stimulus [19], [20]. An intact eye and visual cortex are subjected to visual stimulation, and the response is imaged onto a 24×24 pho-




Fig. 1. Five consecutive frames of video from four different experiments: a) a counter-rotating ring state from the combustion experiment, b) a standing wave state from the combustion experiment, c) a hopping state from the combustion experiment and d) the leading edge of the response to a visual stimulus in the neurobiological experiment.

todiode array at rates up to 1000 frames/second to produce high-speed movies of membrane electrical activity. The bottom row of Fig. 1 shows five frames from an experiment in which the stimulus is a pulse of diffuse light. Senseman [19] has shown that the traces from individual elements in the photodiode array closely agree with measurements obtained from electrode probes inserted at the same positions in the cortex.

Both of these spatiotemporal systems have intrinsic dynamic variables that are related to light intensity. In the flame experiment, the light intensity is proportional to the flame temperature, while in the cortex experiment, the light absorption is proportional to cell membrane potential. Both systems appear to depend on two spatial variables. These features allow us to make an argument that video captures essential dynamics.

The systems also have some important differences. The combustion experiment has circular symmetry, and changes in solution symmetry may reflect experimental transitions. There is no obvious symmetry in the biological data. The neurobiological experiment has responses that are temporally localized, while the combustion experiment appears to represent statistically stationary behavior when run for long times at fixed parameters. The common and contrasting characteristics of these experiments are useful for developing visualization techniques that are generally applicable to scientific video data.

III. KL DECOMPOSITION

The visual metaphors in this paper are based on the Karhunen-Loève (KL) decomposition, which is sometimes referred to as principal component analysis or the Hotelling transformation. KL decomposition was derived as early as 1907 in a famous paper by Schmidt [18]. The connection between the classical derivation and current

techniques is described in [21], [23]. In the classical derivation, the technique was used to find an optimal sum-of-products approximation for measurable functions. That is, given a function $u(x, t)$, find a sum-of-products representation in terms of functions $a_k(t)$ and $\phi_k(x)$ that minimizes the squared error for any value of M using an appropriate norm:

$$\| u(x, t) - \sum_{k=1}^M a_k(t) \phi_k(x) \|^2$$

The approximation of u by this sum of products for finite M is called a KL decomposition. The entire collection of ϕ_k 's form an optimal orthogonal coordinate system in which to express u . The transformation of u to this coordinate system is called a KL transformation. When $u(x, t)$ represents the solution of a partial differential equation, the ϕ_k 's can naturally be interpreted as the normal modes or eigenfunctions.

In this paper, \mathbf{u} is a real random vector of length N representing temporal snapshots of spatial behavior. The best basis can be computed by solving:

$$\mathbf{R}\phi_k = \lambda_k \phi_k \quad 1 \leq k \leq N$$

where \mathbf{R} is the expected value of the autocorrelation matrix, $E[\mathbf{u}\mathbf{u}^T]$. If the random vectors have zero mean, \mathbf{R} is the correlation matrix. The ϕ_k 's are the eigenvectors of the $N \times N$ matrix \mathbf{R} , and λ_k are the corresponding eigenvalues. If Φ is the matrix whose columns are the eigenvectors ϕ_k , the KL transformation can be expressed as:

$$\mathbf{u} = \Phi \mathbf{v} = \sum_{k=1}^N v_k \phi_k$$

where $v_k = \mathbf{u} \cdot \phi_k$. The ϕ_k 's are sometimes called *empirical eigenfunctions*, KL eigenfunctions or principal components. In the current application the ϕ_k 's are images, so this paper also uses the term *KL eigenimages* to emphasize this point.

A. Energy Spectrum

Since \mathbf{R} is a real symmetric matrix, it has non-negative real eigenvalues, and the ϕ_k 's can be ordered in decreasing order of the corresponding eigenvalues $\lambda_1 \geq \lambda_2 \geq \dots \geq 0$. The relative size of λ_k reflects the amount of data set energy along the direction defined by ϕ_k . That is:

$$E_k = \frac{\lambda_k}{\sum_{j=1}^N \lambda_j}$$

is the fraction of the data set energy that is contained in the direction of the k -th eigenimage. E_k is called the *spectrum* of the KL decomposition.

The energy spectra for the four experiments of Fig. 1 are displayed as bar charts in Fig. 2. Superposed on each chart is a smoothed graph of the cumulative energy, $\sum E_k$. The figure has a horizontal line drawn at .75 to guide the

eye. The spectrum is useful for determining how many dimensions (eigenimages) are needed to describe the data set. The data sets corresponding to the counter-rotating ring state, the standing wave and the cortex are of low dimension. The hopping state, which is considerably more complicated, needs 11 eigenimages to capture 75% of the data set energy and more than 30 to capture 99%. While Sirovich [22] uses a cutoff of .99 in determining how many KL eigenfunctions to keep, Palacios et al. [13] found that a cutoff of .75 is often sufficient for defining the general behavior. Some of the visualization techniques discussed in Section IV address the issue of how many eigenimages are needed to reproduce observed visual features of the experimental video.

The energy spectra of Fig. 2 reveal other characteristics of their respective systems. The eigenimages for the counter-rotating ring and the hopping states appear to come in pairs having roughly equal energy, while the standing wave and the cortex data have a single dominant mode.

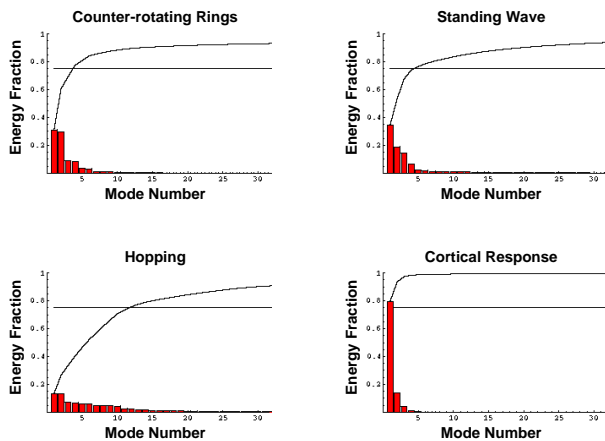


Fig. 2. The distribution of energy contributed by the KL eigenimages to the data sets of Fig. 1.

B. The Eigenimages

When applying KL decomposition to a scientific data set of images representing dynamic behavior, one would not generally expect the resulting eigenimages to resemble objects in the data set but rather to resemble normal modes. The relationship of the normal modes to the observed behavior is not always clear. In the simple 1-dimensional example of a vibrating string, the amplitude of vibration is the observed behavior, while the normal modes are sines and cosines. By understanding the relative importance of the contributions of the modes, one can understand the structure of the dynamic behavior.

In systems with symmetry, such as the combustion experiment, the KL modes give additional information about system structure. The KL modes have at least as much symmetry as the data itself [4], and an observed change in symmetry of the modes reflects a transition or bifurcation in the data set that may not be visible in the original clip. These observations make the calculation of KL

eigenimages useful for detecting changes in systems with symmetry even if no dynamical behavior is considered.

Fig. 3 shows the eigenimages for the four clips of Fig. 1. The symmetries in the first three sets of eigenimages are apparent. The first pair of eigenimages for the counter-rotating ring state, Mode 1 and Mode 2 in Fig. 3a, are approximately rotated versions of each other. Such pairs produce patterns that rotate in circular geometry in much the same way as a superposition of sines and cosines produces traveling waves in rectangular geometry. Notice that Mode 5 and Mode 6 in Fig. 3a, have twice the spatial frequency of the first mode pair, indicating that the pair spatially refines the basic pattern in the outer ring of cells captured by the first pair. Pairing is also observed in the hopping state of Fig. 3c. Mode 1 and Mode 2, Mode 3 and Mode 4, and Mode 5 and Mode 6 each form phase-shifted pairs that contain roughly the same amount of energy.

The leftmost image in each row of Fig. 3 corresponds to the coordinate origin. This image was subtracted from the entire data set prior to the KL decomposition. In the first three examples, the average of the data set was selected as the origin. This choice effectively centers the coordinate system in the data. One would generally expect that a basis calculated relative to a centered coordinate system would better reflect the structure of the underlying data set than a basis calculated relative to a distant coordinate origin. The basis images obtained when the average is removed from the data set are uncorrelated.

In some circumstances, removal of the average is not appropriate. The biological data represents the initiation and subsequent decay of a response to a stimulus. The mean of this data set can be changed arbitrarily by including more images prior to the stimulus onset or by extending the sample after the response decay. The KL decomposition for this data set was performed without moving the coordinate origin.

Sometimes there is a physically meaningful coordinate origin that is not the data set mean. For example, when studying the behavior near a particular type of solution, one might move the origin to this point by subtracting an image corresponding to the solution prior to performing basis calculations. The resulting KL decomposition provides an effective fisheye view that emphasizes behavior near that solution.

Another case where one would not subtract the mean is in the comparison of multiple data sets such as those obtained by varying a parameter to study a transition. Subtraction of the mean for each data set prior to the KL decomposition results in a comparison of the data sets relative to different origins. On the other hand treating the collection of data sets as a whole and removing the mean of the entire collection may not result in a physically meaningful point of origin. The latter method has the added disadvantage that the mean may change when more data sets are added to the collection, making direct comparison with previous results more difficult.

The connection of the KL decomposition to normal modes is important for the current applications. Each of




Fig. 3. The eigenimages for the experiments represented by the video clips of Fig. 1. In each row, the leftmost image represents the origin of the coordinate system used in determining the KL basis. The images are ordered from left to right by decreasing magnitude of their energy contribution to the data set.

these systems has complex spatial and temporal behavior that is most likely described by partial differential equations depending on two spatial variables. A physical argument can be made that the video data captures, to sufficient resolution, the spatial and temporal dynamics of these systems. The KL decomposition then produces a separable approximation (sums of products of functions that depend only on time or only on spatial variables) to the system. When the underlying system in fact has a separable solution, the resulting KL eigenimages are physically significant. These observations motivate the visualizations that are developed in this paper.

C. Direct Computation versus the Method of Snapshots

In video applications, the data sets are digitized video clips of M frames. Each frame is a greyscale image of $n \times m$ pixels stored in row major form to produce a vector of size $n \times m = N$. In the straightforward implementation of KL decomposition, images of dimensions 128×128 produce vectors of length 16,384. The corresponding autocorrelation matrix has dimensions $16,384 \times 16,384$. While the solution of the eigenvalue problem for such a large matrix is doable, it is not attractive for real-time visualization.

When the number of images in the data set (M) is much less than the size of the images (N), the computation can be simplified using the *method of snapshots* due to Sirovich [21]. The method of snapshots is based on the observation that the space spanned by M data vectors must be of dimension $\leq M$ regardless of the length of the vectors. Also, the space spanned by the data and by the

KL eigenimages is the same, so the eigenimages can be written as linear combinations of the original images:

$$\phi_k = \sum_{i=1}^M v_i^k \mathbf{u}_i \quad (1)$$

Sirovich uses these observations to show that $\mathbf{v}^k = (v_1^k, v_2^k, \dots, v_M^k)$ is a solution of $\mathbf{C}\mathbf{v}^k = \lambda_k \mathbf{v}^k$ where \mathbf{C} is an $M \times M$ real symmetric matrix whose (i, j) -th entry is computed as $\frac{1}{M} \mathbf{u}_i \cdot \mathbf{u}_j$. In the method of snapshots, \mathbf{C} is computed directly, and its eigenvectors (\mathbf{v}^k) are then determined. The original data set and these \mathbf{v}^k 's are substituted into Equation 1 to determine the ϕ_k 's. When $M \ll N$, the method of snapshots produces the ϕ_k 's considerably faster than a direct computation.

In the direct method, each of the $N(N-1)/2$ elements in the upper triangular portion of \mathbf{R} is computed by averaging M products. It takes a total of $N(N-1)M$ arithmetic operations to compute \mathbf{R} . A standard technique for finding eigenvectors might take $O(N^3)$ operations. Thus the order of the direct computation is $O(N^3 + N^2M)$.

In the method of snapshots, each of the $M(M-1)/2$ elements in the upper triangular portion of \mathbf{C} is evaluated by computing a vector dot product of length N and then dividing by M . Thus it takes $M(M-1)N$ operations to compute \mathbf{C} . Once the eigenvalue problem for \mathbf{C} has been solved it takes M^2N operations to convert the v_i^k 's into the ϕ_k 's. If it takes $O(M^3)$ arithmetic operations to find the eigenvectors of \mathbf{C} , the resulting computation is $O(M^3 + M^2N)$. When M is 100 and N is 16,384,

the ratio of number of operations for the direct method to the number of operations for the snapshot method is about $4,000\alpha$, where $\alpha > 1$ is the coefficient in the leading order term of the count of operations for solving the eigenvalue problem. The method of snapshots can be applied to 100 images of dimension 128×128 in well under a minute, while the direct method takes several hours on the same machine.

IV. METAPHORS FOR FAST VISUALIZATION

There is an extensive literature on the application of KL decomposition to images [10], [24]. In one well-known study, Sirovich [21] applied this technique to a collection of images containing male faces and obtained a low-dimensional basis. The resulting eigenimages were distinguishable as faces, but they did not correspond to particular faces in the original set. Aubry et al. [1] have applied KL decomposition to find coherent structures in boundary layers. KL decomposition has been used more recently to understand the structure of rotating and modulated rotating patterns in combustion [13], [14].

This paper uses KL decomposition as a basis for fast visualization techniques. The visualizations help the experimentalist identify the part of the visual observations that is contributed by each of the eigenimages. Most of the visualizations involve some type of animation which is difficult to convey on paper. Movies of selected animations can be viewed at [29].

A. Reconstruction

Reconstruction approximates the original data set by using combinations of KL modes. If $\phi_1, \phi_2, \phi_3, \dots, \phi_M$ are the eigenimages for a data set consisting of M video frames, $\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3, \dots, \mathbf{u}_M$, the m -th order successive reconstruction for \mathbf{u}_i is $a_1(i)\phi_1 + a_2(i)\phi_2 + \dots + a_m(i)\phi_m$ where $a_k(i)$ is the dot product of \mathbf{u}_i with ϕ_k . That is, $a_k(i)$ is the projection of \mathbf{u}_i on the direction represented by ϕ_k .

Fig. 4 shows successive reconstructions of a single frame of the example clips of Fig. 1. The first image in each row consists of the projection on Mode 1, and the second image is the sum of the projections on Mode 1 and Mode 2. Similarly the four mode reconstructions use the contributions of Mode 1 through Mode 4, and the six mode reconstructions use the contributions of Mode 1 through Mode 6. The rightmost image is the original video frame. By comparing the reconstructions side-by-side, the user can visually separate the contributions of the individual eigenimages. These contributions are more apparent when the side-by-side displays are made into a movie whose speed can be controlled by the viewer. As the movies play, the motion allows the eye to pick out the contribution of the individual modes to the overall visual appearance.

The reconstructions of Fig. 4 display five images side-by-side. Five images across is probably the upper limit for an understandable animation display. In the counter-rotating ring state reconstructions of Fig. 4a, the eigenim-




Fig. 4. Successive reconstructions of a frame from a) the counter-rotating ring state, b) the standing wave state, c) the hopping state and d) cortical response. The rightmost image in each row is the original frame from the data set. The origin images shown at the left in Fig. 3 are added to each reconstruction.

ages can be grouped in phase-shifted pairs. The hopping state reconstructions in Fig. 4c are not of high enough order to show the actual hopping. The energy spectrum in Fig. 2 shows a natural grouping of eigenimage energy indicating that reconstructions 2, 6, 10 and 15 should be animated for the hopping.

Sometimes *selective reconstruction* provides more information about how the components interact to produce the movement. In this type of visualization, the user selects a set of not necessarily consecutive eigenimages to project the video data onto and reconstructs a movie to play next to the original video. In the cortical data, for example, the contribution of a dominant eigenimage visually swamps the effects due to other eigenimages. By selecting a reconstruction based on eigenimages 2, 3 and 4, the user can better understand the mechanism by which the components captured by these images contribute to the response.

Another method of visualizing reconstructions is the animated bar chart shown in Fig. 5. The figure is a single snapshot of an animation frame of the first, second and sixth order successive reconstructions of the standing wave example of Fig. 1. The bar chart shows the instantaneous values of the coefficients $a_k(i)$ used in the reconstruction. The image on the right, which shows the corresponding reconstruction, provides cues about the relative contributions of the eigenimages to the visual appearance of the video. By running the animation of this standing wave example in slow motion, one can clearly see that the sign of a_2 determines the orientation of the two cells. The motion of the bars going up and down allows the viewer to see how the modes move relative to each other to produce a visual effect.

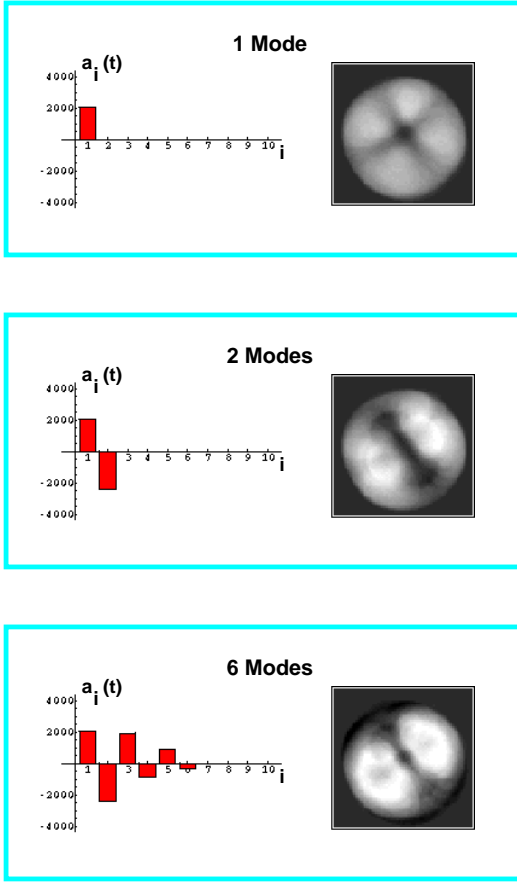


Fig. 5. This bar chart animation shows the instantaneous contributions of eigenimages to the motion in the standing wave state. The images on the right are the corresponding reconstructions. The chart uses successive reconstructions based on 1, 2 and 6 modes, respectively.

B. Phase Plane Visualization

Fig. 6 shows an example of a phase plane visualization for the counter-rotating ring state. In each plot a pair of projections $(a_{k_1}(i), a_{k_2}(i))$ is plotted with the sample number i as a parameter. Since i corresponds to time, these plots provide a dynamic representation of the time behavior of the data set. Each phase plane plot is animated by displaying the point corresponding to the current frame as a large dot with the entire plot as background. As the animation progresses, the time behavior becomes apparent.

The plots of Fig. 6 are called *phase plane plots* because when the data set represents a separable solution to a partial differential equation, the time dependence is described by functions $a_k(t)$ that satisfy a system of nonlinear ordinary differential equations. The behavior of such systems can be represented by plotting pairs of the dependent variables with time as a parameter. Steady solutions appear as single points and periodic behavior as closed curves.

Fig. 6 shows that the outer ring (Mode 1 and Mode 2) and the inner ring (Mode 3 and Mode 4) each rotate periodically. The points in the plot of a_1 versus a_3 cover a rectangular region even though a_1 and a_3 are individually

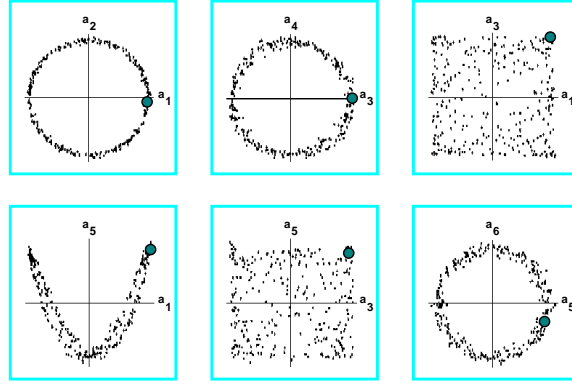


Fig. 6. Phase plane visualizations for the counter-rotating ring state using 300 images of dimension 64×64 . Each small dot represents the projection of a frame of video on the indicated eigenimages. The large dot represents the projection of the current frame.

periodic. This graph implies that the rotational frequencies of the inner and outer rings are incommensurate and that this state can be identified with motion on a torus.

The torus is more clearly visible when a_1 , a_2 and a_3 are plotted against each other in a single plot such as the one shown in Fig. 7. The large dot indicates the instantaneous values of the projections. The upper image on the right shows the corresponding three-mode reconstruction and the lower image is the corresponding frame from the original movie. As the animation plays, the large dot moves to mark the current projection values allowing the user to follow the relationship between the position of the projection on the torus and the behavior of the original data.




Fig. 7. A three-dimensional plot of projections for the counter-rotating ring state. The upper right image is reconstructed from the first three modes corresponding to the projection values marked by the large dot. The lower right image is the corresponding frame in the original video.

C. Visualization of the KL Spectrum

The energy spectrum E_k gives the distribution of energy in the entire data set. This global distribution may not provide a true picture of eigenimage contributions in systems with irregular temporal behavior. For example, if an eigenimage contributes significantly to one or two frames and not to any others, its E_k may be very low. Such isolated contributions can be visualized with a *spectrogram*, which provides a Gantt chart of the energy dis-

tribution in the data set. A spectrograph is an array, S , in which $S[k, i]$ is the energy contributed to the i -th image in the data set by the k -th eigenimage. The last row of S contains the *residual*, which is the amount of energy not accounted for by the selected basis. The spectrograph can be displayed using a pseudocolored or greyscale image to provide a visual picture of how the eigenimages contribute to each frame.

Fig. 8 shows two versions of the spectrograph for the hopping state of Fig. 1. In both graphs black indicates the maximum energy, while light indicates a low energy. Energies that are below a user-specified level are mapped to white to help clarify the pattern of energy distribution. The spectrograph on the top uses a cutoff of 0.01 of the maximum energy, while the spectrograph on the bottom uses a cutoff of 0.1. These spectrographs clearly show the irregularity of the hopping state. The frames in which eigenimages 7 through 10 contribute a maximum amount of energy clearly indicate an irregularity in the motion that is not visible from direct viewing.

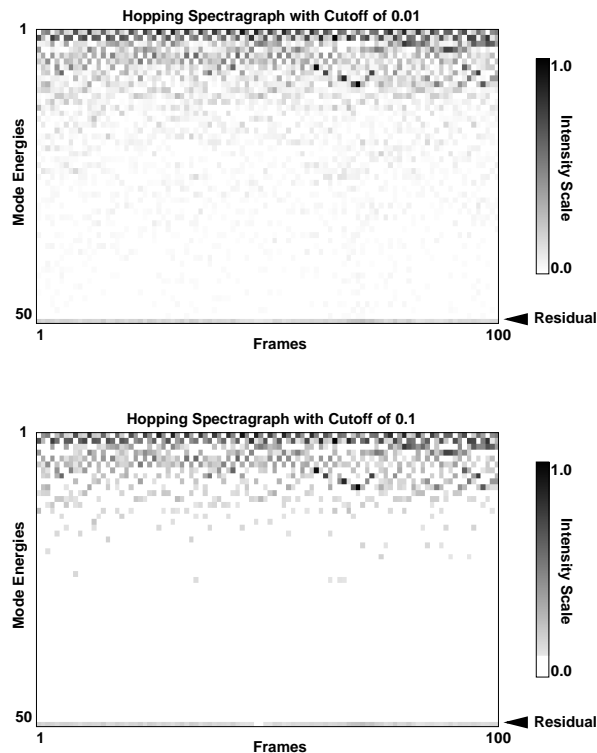


Fig. 8. Spectrographs for the hopping state of Fig. 1 showing the distribution for the 50 most important eigenimages for 100 frames. The top spectrograph uses a cutoff of 0.01 of the maximum energy, while the bottom spectrograph uses a cutoff of 0.1.

The spectrographs of Fig. 8 are scaled by the maximum energy contained by any eigenimage for any frame. This type of scaling works well when the energy distribution across frames is fairly uniform. For data sets in which a few frames contain much greater energy than the rest, the scaled spectrograph is more visually effective. Each column in the spectrograph contains the energy distribution for a particular frame, the column entries in a scaled spec-

trograph are divided by the maximum value in that column to emphasize the distribution in the frame. Fig. 9 shows a comparison of the spectrograph to the scaled spectrograph for the biological data shown of Fig. 1d. The video record contains 125 frames, and the contributions of 15 eigenimages are shown. The initial 20 frames correspond to the resting level prior to the response to visual stimulation and contain very low energy. The resting response is followed by a sharp rise in response that dominates the energy of the data set, so very little structure is visible in the spectrograph. The scaled spectrograph emphasizes the structure in the initial sharp rise and in a later secondary peak.

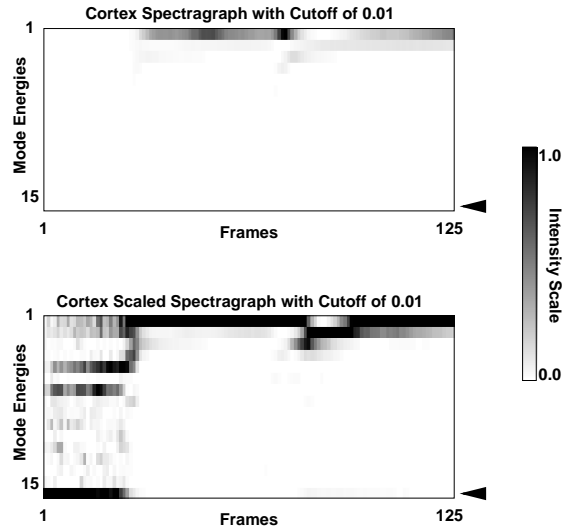


Fig. 9. A comparison of the spectrograph and the scaled spectrograph for the biological data in Fig. 1d. The spectrographs show 125 frames with 15 eigenimages. An energy cutoff of 0.01 was used in both cases. The dark arrows indicate the frame residuals.

The spectrographs can be used in animations by superposing a vertical tuning line that moves horizontally as time progresses. The tuning line marks the position of the current frame on the spectrograph. The spectrograph can be displayed next to other animations to provide a road map of the animation as it progresses.

D. On-the-Fly Visualization

The techniques discussed in this paper require a set of basis vectors to define the coordinate system in which to project the data. In an experimental setting, a scientist would capture a small set of frames and perform a KL decomposition using the method of snapshots. The resulting set of KL modes can be used for visualization as long as the basis captures sufficient data set energy as indicated by the spectrographs. Additional basis vectors can be added later, or a visualization basis can be recomputed if the residuals become large.

Once a KL basis has been selected, side-by-side reconstructions, bar charts and phase plane plots are low overhead methods of display of the projection information. The phase plane plots accumulate historical projection data. The display can use different colors for the




Fig. 10. The eigenimages calculated for the hopping mode using data sets with different spatial resolutions: 128×128 images (top row), 64×64 images (second row), 32×32 images (third row), 16×16 images (fourth row) and 8×8 images (last row).

background dots to indicate different time periods (early or late in the experimental run). The spectrograms are naturally displayed as strip charts, with new values panning to the right in long runs as the video is generated. The scientist can also use spectrograms based on projection energy values averaged over a certain number of frames as a check that the experiment remains statistically stationary during the experimental run.

V. SAMPLING ISSUES

One of the difficulties with visualization in the laboratory is that the experimental setup usually determines the degree of spatial and temporal resolution that is available for processing. This section discusses some general guidelines for adjusting resolution, when such a luxury is available.

The well-known Nyquist criterion in signal processing specifies that time signals should be sampled at least twice the frequency of the highest significant frequency component (the Nyquist frequency). The Nyquist criterion is applicable for both time and space, and one can compute power spectra of the images to get rough estimates to use as a guide in selection.

For fast visualization, one would like to compute at the lowest resolution possible. Fig. 10 shows the results of applying KL decomposition to data sets with different im-

age resolutions for the hopping state. There is little visible change in the results for resolutions greater than 32×32 . A rule of thumb is to start with the lowest resolution that contains at least four pixels across significant structures. If one has the luxury of an adjustable resolution, one can then do a trial with double the spatial resolution and see if there are significant differences. Even at extremely low spatial resolutions such as the 8×8 images shown in the last row of Fig. 10, some structure is visible. When the 16×16 eigenimages are block-expanded to 64×64 images and smoothed, they appear very similar to the eigenimages obtained directly from the 64×64 image data set.

Another issue is how long a clip is needed to capture the significant behavior. Dimensional arguments say that one must take a clip at least as long as the number of significant eigenimages. We compared the eigenimages for the hopping state for clips of 50 images, 100 images, 250 images and 500 images. Very little difference was apparent in the eigenimages calculated from these different clips. The clip must capture the behavior that is to be visualized. A rule of thumb is to take a clip that is at least twice as long as the number of significant eigenimages with a sampling frequency based on the Nyquist criterion.

The visualizations discussed in this paper rely on temporal ordering of the data set in order to have meaningful animations. The calculation to determine the eigenim-

ages need not use an ordered data set. Thus, one can use random sampling to select 100 frames from a long clip and use the resulting eigenimages as a basis for incoming video which is sampled at a rate higher than the Nyquist frequency. Random sampling is particularly useful when the underlying process shows intermittent behavior.

A critical sampling issue in the application of this technique is the dimensionality of the spatial data. The two examples considered in this paper were essentially two-dimensional, and the video data captured the variation in both spatial dimensions. In these cases, there is a theoretical basis for arguing that the eigenimages are related to normal modes and thus have physical significance. However, if the data is a lower-dimensional projection of spatial information (e.g. two-dimensional video of a phenomenon that varies in three spatial dimensions), the method has no such physical interpretation, and the results may not reflect the actual system behavior. This point is illustrated by the phase plane plots of Fig. 11. The top row of phase plots were calculated from counter-rotating ring data sets that had an image resolution of 8×8 . The plots have the same general behavior as that calculated from the 64×64 images shown in Fig. 6. The second row of phase plane plots were calculated from a one-dimensional data set consisting of the middle row extracted from each 64×64 image in the original data set. Although this data set contains the same number of pixels as the 8×8 case, it is one-dimensional rather than two-dimensional, and the phase plane plots do not reflect the observed behavior.

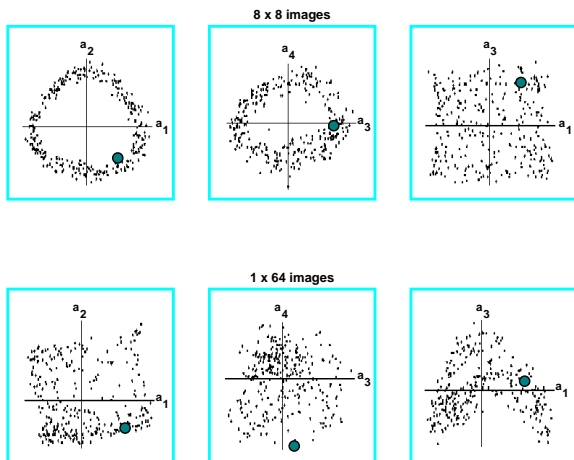


Fig. 11. The phase plane projections for the counter-rotating state at two different spatial resolutions. The row of phase plane plots were calculated from 300 images of dimension 8×8 ; the bottom row of phase plane plots were calculated from the same 300 images using a one-dimensional cut at row 32 in 64×64 images.

VI. DISCUSSION

The visualization techniques described in this paper can be performed quickly to provide feedback about the structure of experimental video. This feedback can be used to classify observations during an experimental survey. The visualizations are most effective when the user can select

animations to view side-by-side and can precisely control viewing speed. All of the visualization displays can be generated quickly without special purpose hardware and are designed for on-the-fly generation.

The techniques rely on the assumption that the samples, in this case frames of video, truly capture the spatial behavior. In the two experimental systems illustrated here, the light intensity is a dynamic variable and the systems depend on only two spatial variables. The same visualizations can also be applied to three-dimensional data, but the data samples must capture the spatial dependence on all three variables. This is unlikely for experimental video data from a single camera, but it is appropriate for the output of numerical simulations of partial differential equations.

These visualization techniques can also be applied to numerical data in any number of spatial dimensions, provided that appropriate spatial information is available. For a numerical simulation that depends on three spatial dimensions, the frames should consist of three-dimensional snapshots. Reconstructions of three-dimensional data may not be possible in real time, but the phase plane representations and the spectrograms are independent of the spatial dimension. Reconstructions in three-dimensions can be done as two-dimensional projections so that a comparison can be made between the KL approximations and the original data.

A KL decomposition which is low-dimensional provides a basis for reducing the amount of data that needs to be retained for other types of visualizations. KL decomposition generally gives a better lower-dimensional representation of the data set than would be obtained from dataset independent transformations such as Fourier series or wavelets. It has not been established whether these techniques are applicable to more general types of video data sets or under more general types of coordinate transformations such as those that arise from projection pursuit [5] or independent component analysis [11].

One of the advantages of KL decomposition is its theoretical relationship to separability. When the KL spectrum is high-dimensional or the projections no longer fall on well-defined curves, a break-down in the separability assumption may be indicated. Broomhead et al. [2] have looked at this issue for certain classes of numerical models. Their work indicates that the breakdown may be an indication of an inertial manifold with high curvature. They propose an adaptive basis technique which may provide the underpinnings for visualizing such systems.

Acknowledgments

The author is grateful to Michael Gorman and David Senseman, the experimentalists who provided the data visualized in this paper. They have been invaluable sounding boards for the visualization work. The author would also like to thank Antonio Palacios, Gemunu Gunaratne and Marty Golubitsky for helpful discussions and feedback. Danny Brown, Deepak Kumar and Jennifer Olveda assisted in the data handling and programming support.

Some of the animations produced for this paper were performed in Mathematica 3.0. This work was partially supported by Office of Naval Research grant N00014-97-0029, NSF grant ACI-9721348 and a UTSA Research Development Award.

REFERENCES

- [1] N. Aubry, P. Holmes, J. L. Lumley and E. Stone, "The dynamics of coherent structures in the wall region of a turbulent boundary layer," *Journal of Fluid Mechanics*, **192** 113-172, 1988.
- [2] D. S. Broomhead, R. Indik, A. C. Newell and D. A. Rand, "Local adaptive Galerkin bases for large-dimensional dynamical systems," *Nonlinearity*, **4** 159-197, 1991.
- [3] J. X. Chen, D. Rine and H. D. Simon, "Advancing interactive visualization and computational steering," *IEEE Computational Science and Engineering*, **3**(4) 13-17, 1996.
- [4] M. Dellnitz, M. Golubitsky and N. Nicol, "Symmetry of attractors and the Karhunen-Loève decomposition," in *Trends and Perspectives in Applied Mathematics* Springer-Verlag Applied Mathematical Sciences Series **100** L. Sirovich, Ed. 73-108, 1994.
- [5] J. Friedman and J. Tukey, "Exploratory projection pursuit," *J. Amer. Statistical Assoc.*, **82** 249-266, 1987.
- [6] M. Gorman, M. el-Hamdi and K. A. Robbins, "Rotating and modulated rotating states of cellular flames," *Combustion Science and Technology*, **98** 25-35, 1994.
- [7] M. Gorman, M. el-Hamdi and K. A. Robbins, "Hopping motion in ordered states of cellular flames," *Combustion Science and Technology*, **98** 47-56, 1994.
- [8] A. Gupta, S. Santini and R. Jain, "In search of information in visual media," *Communications of the ACM*, **40**(12) 35-42, 1997.
- [9] P. Keller and M. Keller, *Visual Cues*, IEEE Press, 1993.
- [10] A. K. Jain, *Fundamentals of Digital Image Processing*, Prentice Hall, 1989.
- [11] S. Makeig, T-P. Jung, A. J. Bell, D. Ghahremani and T. J. Sejnowski, "Blind separation of event-related brain responses into independent components," *Proc. Natl. Acad. Science, USA*, **90**(20) 10979-10984, 1997.
- [12] G. Nielson, H. Hagen and H. Mueller (eds), *Scientific Visualization: Surveys, Methodologies and Techniques*, IEEE Computer Society Press, 1996.
- [13] A. Palacios, G. Gunaratne, M. Gorman and K. A. Robbins, "Cellular pattern formation in circular domains," *Chaos* **7**(3) 463-475, 1997.
- [14] A. Palacios, G. Gunaratne, M. Gorman and K. A. Robbins, "Karhunen-Loève analysis of spatiotemporal flame patterns," *Physical Review E* **57**(5) 1-14.
- [15] S. Pfeiffer, R. Lienhart, S. Fischer and W. Effelsberg, "Abstracting digital movies automatically," *J. Visual Communication and Image Representation* **7**(4) 345-353, 1996.
- [16] K. A. Robbins, "Exact-time and near real-time visualization for the experimentalist," *Proc. 1997 ACM Workshop on New Paradigms in Information Visualization and Manipulation*, 66-68, 1997.
- [17] K. A. Robbins and M. Gorman, "Using image processing and computer animation to analyze spatial and temporal dynamics in a pattern-forming system," *Computers in Physics*, **12**(1) 73-81, 1998.
- [18] E. Schmidt, "Zur Theorie der linearen und nichtlinearen integralgleichungen. I Teil: Entwicklung willkürlicher funktion nach systemen vorgeschriebener," *Mathematische Annalen*, **63** 433-476, 1907.
- [19] D. M. Senseman, "Correspondence between visually evoked voltage-sensitive dye signals and synaptic activity recorded in cortical pyramidal cells with intracellular microelectrodes," *Visual Neuroscience*, **13** 963-977, 1996.
- [20] D. M. Senseman, "Diffuse retinal stimulation triggers a wave of depolarization in spiny pyramidal cells that propagates across the surface of the turtle cortex via recurrent feedback," *Visual Neuroscience*, in press.
- [21] L. Sirovich, M. Maxey and H. Tarman, "Low-dimensional procedure for the characterization of human faces," *J. Optical Society*, **4A** 519-524, 1987.
- [22] L. Sirovich, "Turbulence and the dynamics of coherent structures," *Quart. Applied Mathematics*, **XLV** 561-590, 1987.
- [23] L. Sirovich and R. Everson, "Management and analysis of large scientific datasets," *International Journal of Supercomputer Applications*, **6**(1) 50-68, 1992.
- [24] L. Sirovich, R. Everson, E. Kaplan, B. W. Knight, E. O'Brien and D. Orbach, "Modeling the functional organization of the visual cortex," *Physica D*, **96** 355-366, 1996.
- [25] V. E. Taylor, J. Chen, T. L. Disz, M. E. Papka and R. Stevens, "Interactive virtual reality in simulations: Exploring lag time," *IEEE Computational Science and Engineering*, **3**(4), 46-54, 1996.
- [26] J. S. Vetter, "Computational steering annotated bibliography," *ACM SIGPLAN Notices*, **32**(6), 40-44, 1997.
- [27] B.-L. Yeo and B. Liu, "Rapid scene analysis on compressed video," *IEEE Trans. Circuits, Systems and Video Technology*, **5**(6) 533-544, 1995.
- [28] M. Yeung and B.-L. Yeo, "Video visualization for compact presentation and fast browsing of pictorial content," *IEEE Trans. Circuits, Systems and Video Technology*, **7**(5) 771-785, 1997.
- [29] <http://vip.cs.utsa.edu/KLvisual/KLtrans.html>



Kay Robbins received BS and PhD degrees from the Massachusetts Institute of Technology. She holds the position of Professor of Computer Science at the University of Texas at San Antonio. Her current research interests include real-time processing, scientific visualization and multimedia. She is a member of IEEE, ACM, AGU, APS and the Society for Neuroscience. She has participated in long-term collaborations with scientists from several different disciplines.