

Visualizing Differences in Movies of Cortical Activity

Kay A. Robbins* and David M. Senseman[†]

Abstract

This paper discusses techniques for visualizing structure in video data and other data sets that represent time snapshots of physical phenomena. Individual frames of a movie are treated as vectors and projected onto a low-dimensional subspace spanned by principal components. Movies can be compared and their differences visualized by analyzing the nature of the subspace and the projections of multiple movies onto the same subspace. The approach is demonstrated on an application in neurobiology in which the electrical response of a visual cortex to optical stimulation is imaged onto a high-speed photodiode array to produce a cortical movie. Techniques for sampling movies over a single trial and multiple trials are discussed. The approach provides the traditional benefits of principal component analysis (compression, noise reduction and classification) and also allows the visual separation of spatial and temporal behavior.

CR Categories and Subject Descriptors: I.3.3 [Computer Graphics]: Picture/Image Generation – Viewing Algorithms.
Additional Keywords: Scientific visualization, animation, video analysis.

1 INTRODUCTION

Video display has long been an end product of scientific visualization, particularly in studies of time-dependent numerical models [10]. Video visualizations often incorporate graphical overlays or generate isosurfaces to guide the viewer in understanding the phenomena [7]. Virtual reality techniques allow scientists to interact with visualizations [2, 3, 24], but these interactive systems usually require special equipment and extensive setup.

While the display of video generated from experiments or models is the goal of some visualization work, it is the starting point for understanding structure [5, 14]. When viewing a video, the scientist tries visually to extract physical meaning from moving patterns that may be direct representations of physical quantities or may be visual models such as isosurfaces or streamlines. In looking for fundamental relationships, the viewer tries to understand how the patterns evolve and to compare behavior at different parameters. In all but the simplest cases, it is impossible to quantify how similar two movies are even when they are viewed side by side. Because movies can now be easily disseminated via the World Wide Web, the problem of analyzing and comparing movies of scientific phenomena has become an important visualization problem [25, 26].

This paper adapts an established technique for statistical analysis [23], principal component analysis (PCA), to visualize structure and to compare movies. References to the extensive literature on

the application of principal component analysis to data sets of images can be found in [6, 20, 21]. The problem of finding a good basis for time-varying data sets is sometimes called the *blind source separation problem* [1], and more general techniques such as independent component analysis (ICA) [8, 11] can also be used as a starting point for the visualizations. A nice example of using ICA to produce a basis from signals generated by EEG responses of human subjects to auditory events can be found in [9].

This work emphasizes the visualization of the time behavior of projections of data set images on an appropriate basis. In cases where the approach is applicable, the results allow researchers to organize and compare data across multiple movies. The next section introduces a motivating application in neurobiology and raises some of the questions that might be answered by the comparison technique. Section 3 describes the procedure used in this paper for finding low-dimensional subspaces and illustrates how projections on these subspaces can be used in visualization. Section 4 extends the basic method to compare movies. Section 5 discusses how these techniques can be applied more generally and addresses some limitations. Section 6 provides a summary of the results.

2 MOTIVATION

The movie analysis methods presented in this paper are illustrated using data from experiments to determine the cortical response to visual stimuli in pond turtles. This animal model was chosen because its relatively flat, unstructured visual cortex is well-suited to high-speed imaging techniques. In addition, the unique turtle diving physiology allows an intact visual system to be used without requiring oxygenation. The visual cortex is stained with voltage sensitive dye [18], and the light transmitted through the cortex is imaged onto a 24×24 element photodiode array at a rate of up to 1000 frames/second. The transluence of the dyed tissue is directly proportional to electrical activity in the tissue, so the resulting high-speed, low-resolution movie is a direct measurement of cortical activity [19]. Each trial data set used in this paper consisted of 576 frames acquired at a rate of 353 Hz.

Figure 1 shows the cortical response recorded during two experimental trials in the same animal using the same diffuse flash stimulus. Selected frames at the same times relative to the stimulus onset (point A) are shown for Trial 3202 (top row) and Trial 3207 (bottom row). The center graphic shows traces from a single photodiode for the trials at the point marked by a white arrow in image E in the bottom row. Notice that while the initial responses appear similar for both trials, Trial 3207 has a sharp secondary spike similar to those that appear during seizure activity.

* Division of Computer Science, University of Texas at San Antonio, 6900 North Loop 1604 West, San Antonio, Texas 78249-0664, krobbins@utsa.edu

[†] Division of Life Sciences, University of Texas at San Antonio, 6900 North Loop 1604 West, San Antonio, Texas 78249-0662, senseman@utsa.edu

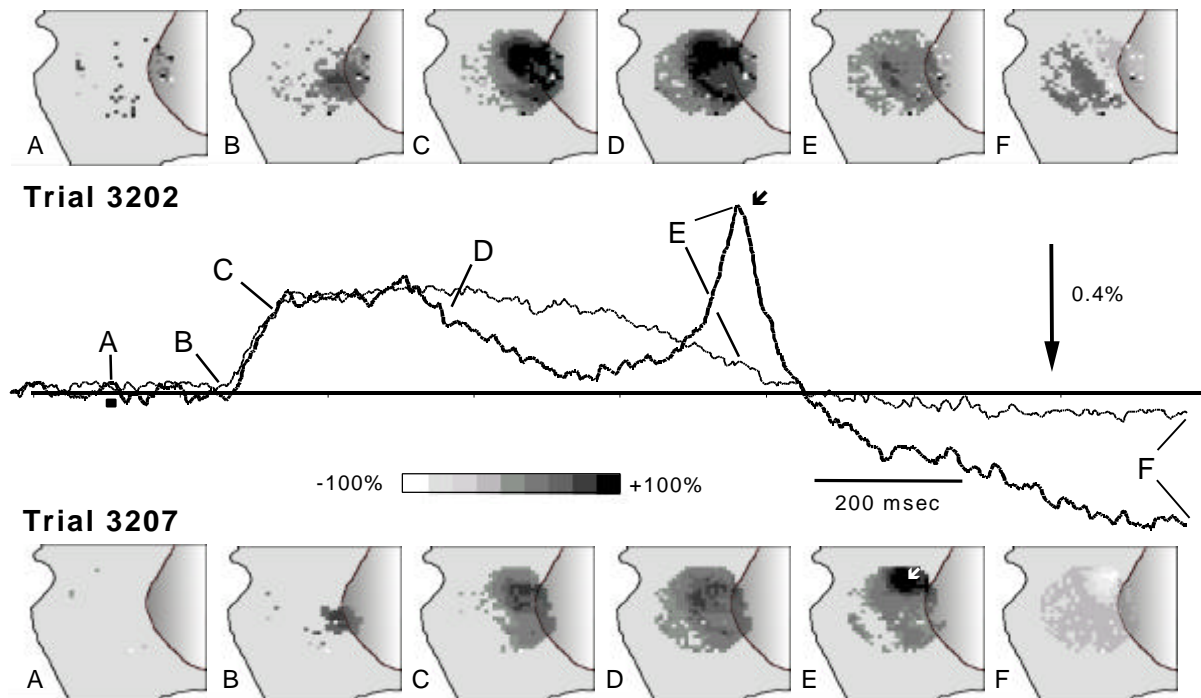


Figure 1 A comparison of cortical response from two trials using a diffuse light flash stimulus. The upper row contains selected images from Trial 3202, while the lower row has the corresponding frames from Trial 3207. The dark center trace is the output of the diode marked by a white arrow in frame E of Trial 3207. The light trace is the corresponding output for Trial 3202. The small square marks stimulus onset, and the vertical arrow shows the scale of the signal relative to the DC component. (Also see the color plate at the end of these proceedings.)

This series of frames illustrates some of the problems that arise in classifying and comparing videos. How similar are these trials? Does the wave of initial response have the same spatial structure across trials? If a similar spatial structure can be established, is it similar to that seen in other animals or in the same animal for different stimuli? At what point in the response can a difference between the two videos be detected? How can this difference be quantified?

This paper addresses some of these issues using methods based on finding the right coordinate system in which to view the movies. By right coordinate system, we do not mean the traditional rotation of a three-dimensional view in physical space. Rather, an abstract view is taken: the movie consists of a series of $n \times m$ images or frames, where each image represents a vector in a high-dimensional space (of dimension $n \times m$). As the movie plays, the images sweep out a curve in this high-dimensional space giving a parametric representation of the movie. We seek a rotation in this high-dimensional space so that the curves defining the movies are confined to a low-dimensional subspace. If such a subspace can be found, a movie can be analyzed by looking at its coordinates in this subspace,

and the structure of the subspace itself provides a method of classification. If two movies project to the same low-dimensional subspace, side-by-side comparisons can be made using simple visualization techniques.

3 OVERVIEW OF THE TECHNIQUE

The first step in the visualization process is to find an appropriate coordinate system for viewing the movie. This paper uses principal component analysis, a well-known technique for finding a basis or coordinate system that minimizes the correlation of the data in the transformed coordinate system [21]. In the movie visualization technique, however, the requirement of optimality is relaxed. Principal component analysis is applied to *selected* frames in hopes of obtaining a *pretty good basis*, that is, one in which a few basis vectors capture the significant behavior. Because the calculation may not use the entire data set, there is no guarantee that the correlation will be minimized over the entire data set. Furthermore, comparison techniques may use a basis derived from combining movies or by selecting a representative movie.

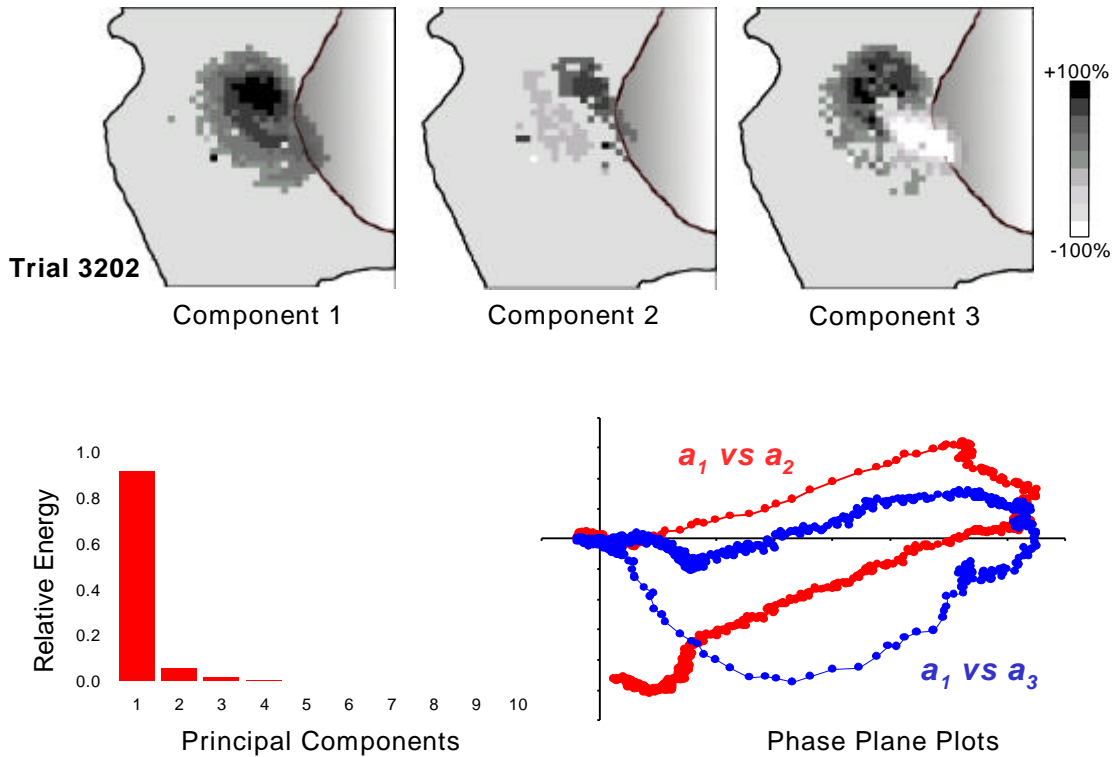


Figure 2 When principal component analysis is applied to Trial 3202, 98% of the energy is contained in the first three components as shown by the bar chart in the lower left. The three dominant components are shown across the top. The curves swept out by the projections on the three dominant components as a function of frame number are shown in the lower right. (Also see the color plate at the end of these proceedings.)

Each point in the data set sample is an image extracted from a frame of video. Strategies for selecting a sample are discussed in the next section. The images in the sample are converted to vectors of length $N = n \times m$ by writing them in row major form. The new coordinate basis is calculated by solving the eigenvalue problem:

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

where \mathbf{R} is the expected value of $\mathbf{u} \mathbf{u}^T$ for the collection of data set vectors \mathbf{u} . In other words, if the i -th image is \mathbf{u}_i , then $\mathbf{u}_i \mathbf{u}_i^T$ is an $N \times N$ matrix. The average of these matrices over the collection is \mathbf{R} . The eigenvectors of \mathbf{R} , also referred to as eigenimages or principal components, are the basis vectors for the new coordinate system. The eigenvalues of \mathbf{R} , λ_k , measure how much of the data set energy is captured by the corresponding eigenvectors, ϕ_k [13].

Once a suitable coordinate system has been determined, the entire movie can be reconstructed [15]. *Reconstruction* refers to using approximations to the original data set consisting of linear combinations of the components $\phi_1, \phi_2 \dots$. The j -th order successive reconstruction for the i -th frame of video, \mathbf{u}_i , is $a_1(i) \phi_1 + a_2(i) \phi_2 + \dots + a_j(i) \phi_j$ 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 . The eigenvectors ϕ_k are indexed in decreasing order of contribution to overall data set energy.

Figure 2 shows the result of applying principal component analysis to Trial 3202. The images across the top of the figure are the three most important components overlaid on an image of the experimental preparation. The original images, which had $24 \times 24 = 576$ pixels, were vectors in a space of dimension 576. Ninety eight percent of the energy is contained in a 3-

dimensional subspace spanned by ϕ_1 , ϕ_2 and ϕ_3 . The curves swept out by the projections on these components (a_1 , a_2 , and a_3) are plotted versus each other in the lower right of Figure 2.

A reconstructed movie for Trial 3202 using the ten most important components is visually indistinguishable from the original movie. Although component 1 contains 98% of the data set energy, the reconstruction using this single component, $a_1(i)\phi_1$, does not reproduce the observed behavior. Component 1 has a single maximum. The projection on this component, a_1 , is always positive, rising from zero and then falling back to zero as the response subsides. The single-component reconstruction resembles a hill that rises and falls, while the original movie gives the visual impression of a wave of activity sweeping from the lower right to upper left. Component 3 appears as a hill and valley combination. The projection on this component, a_3 , is negative on the initial rise, placing the hill on the lower right. As time progresses, a_3 reverses sign before returning to zero. As a result, the hill and valley of component 3 are reversed, indicating how the visual appearance of propagation might be achieved. Component 2 captures the negative intrinsic signal at the end of the response. This intrinsic signal is due to the saturation effects of the dye [18].

3.1 Removal of the Mean

In principal component analysis the data set is often normalized to have zero mean so that the coordinate system origin is centered in the data. Simple geometric arguments can be made that centering usually gives better resolution of data set structure. The vectors representing uncentered data with a large mean are highly correlated, and differences in the individual data points are small compared with the data set mean. When the mean is removed from the data set, however, the vectors representing the individual points are clearly distinguishable.

For visualization applications, there are sometimes physical arguments why the origin should not be centered. In the biological data, the movie is recorded starting at a fixed time prior to stimulus onset and runs for a fixed duration. The leader and trailer lengths for this movie, corresponding to periods of inactivity, are chosen arbitrarily. Increasing the length of these inactive periods shifts the center of the data toward the origin. To eliminate the influence of trailer length on the results, the average was not removed. Instead, the image corresponding to resting activity at the point of stimulus onset was subtracted from the entire data set.

When several data sets are combined or a data set is sampled, the selection of an appropriate data set origin is more complex. When data sets do not have zero mean, principle component analysis does not guarantee that the resulting basis minimizes the correlation. Effective visualizations may not require a basis that minimizes the correlation, and in general, the placement of the origin should be based on physical considerations.

3.2 Implementation for Larger Images

In video applications, the data sets are digitized video clips of M frames. Each frame is a gray scale image of $n \times m$ pixels stored in row major form to produce a vector of size $N = n \times m$. For the biological application discussed in this paper, $N = 576$ and the principal components can be computed quite quickly. In more typical video applications, the images are larger, and a direct computation of the correlation matrix requires a substantial effort. The direct implementation for a data set of 128×128

images produces vectors of length 16,384, and the corresponding autocorrelation matrix has dimensions $16,384 \times 16,384$. The solution of the eigenvalue problem for such a large matrix is not attractive for 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. One can convert the problem to a dual problem that involves the solution of an eigenvalue problem of size $M \times M$. When $M \ll N$, the method of snapshots produces the ϕ_k 's considerably faster than a direct computation. A standard technique for finding eigenvectors might take $O(N^3)$ operations. The order of the direct computation is $O(N^3 + N^2M)$, while the order for the method of snapshots 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 operation count for solving the eigenvalue problem. The method of snapshots can be applied to 100 images of dimension 128×128 in well under a minute on an ordinary workstation, while the direct method takes many hours if it can be done at all.

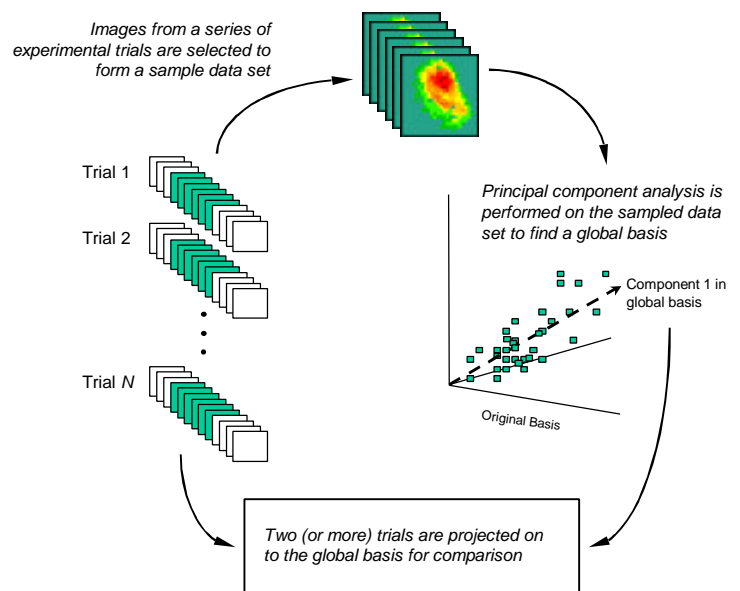


Figure 3 A global basis is computed from a data set consisting of samples from multiple trials. The complete trials are then projected on the basis.

4 MOVIE COMPARISONS

While the principal component analysis of a single movie is useful, it is difficult to use this analysis to assess differences among movies because each trial is projected onto its own coordinate system. This section discusses methods for combining data sets to find a common coordinate system or *global basis* for comparing movies. The selection of a sampling method for combining data sets is an empirical process that should be guided by physical considerations. The straightforward method of simply combining all data sets into a single large data set and finding the components may work in many circumstances. However, if one is particularly interested in the details of a certain type of behavior, sampling techniques are more appropriate.

The general process is described schematically in Figure 3. Samples selected from multiple data sets are combined into a single large data set, and principal component analysis is performed. Once a global basis is found, the most important vectors are selected and the original data sets are projected on this collection. The resulting projections are compared in order to quantify the differences between the movies. There is no guarantee that the basis derived from sampling multiple data sets will capture the behavior of the entire data sets. The *residual* for an image with respect to a collection of vectors is the projection on the complement of the space spanned by the selected vectors. In other words, the residual is the portion of the image that is not accounted for in the new coordinate system. One can improve the global basis by incorporating additional vectors that span the space of significant residuals.

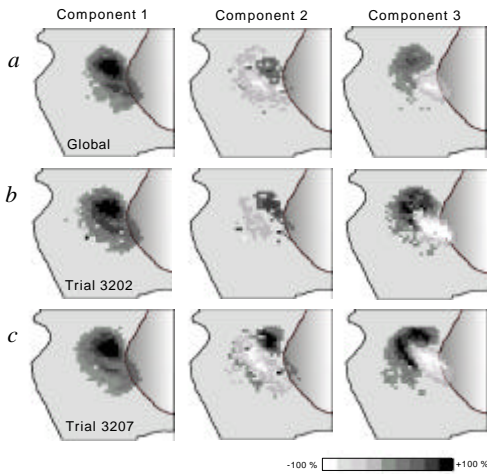


Figure 4 A global basis was computed from a sampled data set consisting of 150 frames from 8 trials. The first three components of the resulting global basis are shown in the top row. The corresponding components from the bases computed from all frames of Trials 3202 and 3207 are shown in the second and third rows respectively.

As an illustration of the technique in the biological application, we created a combined data set of frames selected from eight different trials using a diffuse light stimulus. Each trial consisted of 576 frames, and stimulus onset occurred at frame 50. Frames 50 to 199 were selected from each trial and combined into a single data set. This selection was chosen to capture the initial rise in response for all data sets (corresponding roughly to the interval between frames A and C in Figure 1). Principal component analysis was applied to obtain a global basis. Figure 4 shows a comparison of first three components obtained using this technique with the corresponding components obtained by using all of the frames in each individual trial.

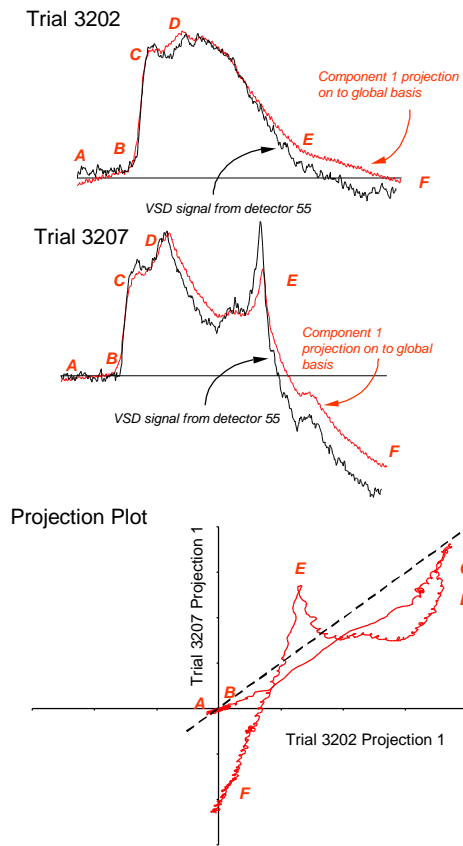


Figure 5 The projection of Trial 3202 on component 1 (a_1) is compared to its trace at detector 55 in the top graph. A similar comparison is made for Trial 3207 in the middle graph. The lower graph compares a_1 for the two trials.

Figure 5 compares the behavior of component 1 for Trial 3202 and Trial 3207. Component 1 clearly captures a significant portion of the response. The second peak in Trial 3207 was not included in the sample used to calculate the global basis, yet it is also captured by the most important component. These results provide evidence that the first component has biological

significance. The projection plot in the bottom graph of Figure 5 shows how the amplitudes of the projections on component 1 vary between the trials. When the curve follows along the dotted 45° line (as in the initial rise between A and C), the trials behave similarly. The departure of the curves at point D marks the time at which the movies begin to differ significantly in behavior.

It was interesting to note that the global basis better captured the energy of each entire individual trial than the individual basis for

that trial even though the global calculation was based only on the initial rise. In the individual basis calculations, components 2 through 6 clearly contained some experimental noise. The combination of multiple trials (even though the stimulus varied in intensity) reduced the weighting of individual noise components in the correlation.




Figure 6 Spectragraphs for Trials 3202 and 3207 using the global basis computed from the rise times on 9 trials. The values are normalized by the maximum energy in the movie. Projections containing less than 1% of the maximum energy are not displayed. The center graphic in each spectragraph shows the projection of the data on the first component.

4.1 Spectragraphs

The energy spectrum, $E_k = \lambda_k / \sum \lambda_k$, is the fraction of the data set energy contained along the direction of the corresponding component. Traditionally, the energy spectrum is used to characterize the relative importance of the individual basis vectors. Although the energy spectrum gives the overall distribution of energy, it may not provide a true picture for systems with irregular temporal behavior. If a few frames contain a significant amount of component ϕ_k but most frames do not, E_k will be low, but component ϕ_k may be physically significant.

Contributions of individual components can be visualized using a *spectragraph* to plot the square of each projection versus frame number for the data set. The spectragraph can be displayed using a pseudocolored or gray scale image to provide a visual picture of how the eigenvectors contribute to each frame as shown in Figure 6. The graph plots frame number on the horizontal axis versus energy contained along a specified component direction on the vertical axis. The last row contains the *residual*, the amount of energy not accounted for by the basis.

The two spectrograms in Figure 6 were calculated by projecting Trials 3202 and 3207 on the global basis as previously described. Each entry in the spectrogram is divided by the maximum energy in the graph. Figure 6 uses a continuous gray scale mapping with black indicating the maximum energy and white indicating less than 1% of the maximum. Recall that only frames 50 to 199 were used to determine the basis. There is surprisingly little residual in frames 200 to 576, even though they were not used in the basis calculation. The most interesting result to come out of the analysis was the observation that the second spike in Trial 3207 was dominated by the first component and thus had the same spatial structure as the original rise.

5 DISCUSSION

Visualization of video using projections on principal components is a reasonably inexpensive technique for analyzing structure and comparing multiple videos. An explanation of why the method appears to work well for many systems dates back to its origin. Principal component analysis was originally introduced as a technique for finding an optimal sum-of-products approximation $\sum_k a_k(t)\phi_k(x)$ for measurable functions [17]. When principal component analysis is applied to data sets whose underlying behavior is separable, the $a_k(i)$'s give the time evolution of fixed spatial components. Furthermore the resulting components reflect underlying system symmetries and can be used to reason about allowed bifurcations of system behavior as parameters are varied [4].

Principal component techniques have been applied to medical data for many years with mixed results [22]. The visualization techniques described in this paper are not directly applicable to data sets that have a large number of significant components or that have the dominant component(s) consisting of noise rather than important structure.

The computation of global bases for comparison of movies is only applicable when individual images within the movies are spatially registered. In the neurobiological application, all movies acquired for the same preparation were spatially registered. Registration across preparations is difficult because of different brain sizes and variance in topology across animals. Registration is a difficult issue in many medical applications, but it is not a problem in numerical applications or in experimental applications with well-defined geometry. An example of the latter is pattern formation on a circular burner in a combustion experiment [12, 13]. These techniques have been successfully applied to the flat flame front, because the underlying system depends on two spatial variables and the light intensity is directly proportional to the temperature, a dynamic variable of the system. Long movies of 64×64 and 128×128 pixel images were analyzed using the method of snapshots. Movies of the original data and the derived visualizations can be viewed at <http://vip.cs.utsa.edu/flames/klvisual/klvisual.html>.

Another caveat in applying this technique relates to the sampling. Robbins [16] has shown that the data (in this case images) must reflect the actual spatial dimension of the system in order for the eigenvectors to be identifiable with normal modes. When principal component analysis was applied to one-dimensional cross-sections of a movie of a two-dimensional, quasi-periodic system, the resulting time behavior was not quasi-periodic and the one-dimensional components were not cross-sections of the two-dimensional components. Since video

data is intrinsically two-dimensional, this technique can only accurately resolve systems with two spatial dimensions from video data unless stereoscopic information or slice information is used to reconstruct a third spatial dimension. Makeig et al. [9] showed that independent components computed from scalp EEG measurements agreed with results from signals recorded directly from the cortex in human patients provided the signals were statistically independent in time, but the computations did not agree otherwise. Independence in time is not an issue for signals from the turtle cortex or combustion experiments because of the intrinsic two-dimensional nature of these systems. While spatial dimensionality problems are often an issue for experimental measurements, data sets generated from numerical simulations usually have the correct spatial resolution and sufficient temporal resolution, and these techniques can be applied directly to time snapshots of numerical output.

6 CONCLUSION

The techniques described in this paper are based on projections of movies onto low-dimensional subspaces. These techniques work particularly well for low-dimensional systems and high-dimensional systems close to onset, where the number of significant components is not too large. Principal component analysis is applied to a sample of images from a movie or collection of movies. The projections onto the resulting components are then visualized using reconstructions, phase plane plots and spectrograms. If the low-dimensional projections capture a significant proportion of data set energy, the similarity of a collection of movies can be ascertained. A key to successful application is the selection of the samples based on appropriate physical considerations. In the neurobiological application discussed in this paper, it was found that a low-dimensional basis formed by combining multiple trials gave a better representation of each data set than the basis computed directly. A comparison of projections on this global basis across trials enabled researchers to quantify differences in response. The technique also allows researchers to follow transitions in response as system parameters are varied. The surprising success of the technique can be attributed to the observation that a time varying superposition of a few fixed independent spatial patterns can result in a visual appearance that seems unstructured and chaotic. While the human eye cannot pick out the underlying correlation, principal component analysis can.

Acknowledgments

The authors would like to thank Michael Gorman, Antonio Palacios, Gemunu Gunaratne and Marty Golubitsky for helpful discussions. This work was partially supported by the Office of Naval Research (N00014-97-0029), the National Science Foundation (BNS-8507594; ACL-9721348), the National Institutes of Health (GM 08194, 07717), the Texas Advanced Research Program (ARP 2222) and the UTSA Faculty Research Award Program.

References

1. A. J. Bell and T. J. Sejnowski. An information-maximization approach to blind separation and blind deconvolution. *Neural Computation*, 1129-1159, 1995.
2. S. Bryson, S. Johan and L. Schlect. An extensible interactive visualization framework for the virtual wind tunnel. *Proc. 1997 Intl. Symposium on Virtual Reality*, 1997.
3. J. X. Chen, D. Rine and H. D. Simon, eds. Special Issue, Advancing interactive visualization and computational steering. *IEEE Computational Science and Engineering*, 3(4), 1996.
4. M. Dellnitz, M. Golubitsky and N. Nicol. Symmetry of attractors and the Karhunen-Loeve decomposition. in *Trends and Perspectives in Applied Mathematics*, Springer-Verlag Applied Mathematical Sciences Series 100 L. Sirovich, Ed. 73-108, 1994.
5. A. Gupta, S. Santini and R. Jain. In search of information in visual media. *Communications of the ACM*, 40(12) 35-42, 1997.
6. K. Jain. *Fundamentals of Digital Image Processing*. Prentice Hall, 1989.
7. B. Jobard and W. Lefer. The motion map: Efficient computation of steady flow animations. *Proc. Visualization '97*, 323-328, 1997.
8. J. Karhunen, L. Wang and R. Vigarito. Nonlinear PCA type approaches for source separation and independent component analysis. *Proc. 1995 IEEE Intl. Conf. On Neural Networks - ICNN '95*, 995-1000, 1995.
9. S. Makeig, T.-P. Jung, A. J. Bell, D. Ghahremani and T. J. Sejnowski. Blind separation of auditory event-related brain responses into independent components. *Proc. National Academy of Sciences*, 94(20) 10979-10984, 1997.
10. G. Nielson, H. Hagen and H. Mueller (eds). *Scientific Visualization: Surveys, Methodologies and Techniques*. IEEE Computer Society Press, 1996.
11. E. Oja. PCA, ICA and nonlinear Hebbian Learning. *Proc. Intl. Conf. on Artificial Neural Networks, ICANN-95*, 89-94, 1995.
12. A. Palacios, G. Gunaratne, M. Gorman and K. A. Robbins. Cellular pattern formation in circular domains. *Chaos*, 7(3) 463-475, 1997.
13. A. Palacios, G. Gunaratne, M. Gorman and K. A. Robbins. A Karhunen-Loeve analysis of spatiotemporal flame patterns. *Physical Review E*, 57 5958-5971, 1998.
14. S. Pfeiffer, R. Lienhart, S. Fischer and W. Effelsberg. Abstracting digital movies automatically. *J. Visual Communication and Image Representation*, 7(4) 345-353, 1996.
15. K. A. Robbins, A. Palacios, and M. Gorman. Visualization, animation and K-L decomposition of spatiotemporal dynamics in a pattern-forming system. *Proc. 4th Experimental Chaos Conference*, 1997.
16. K. A. Robbins. Analysis of Scientific Video Data Using KL Decomposition. UTSA Technical Report, 1997.
17. E. Schmidt. Zur Theorie der linearen und nichtlinearen Integralgleichungen. I Teil: Entwicklung willkürlicher funktion nach systemen vorgeschriebener. *Mathematische Annalen*, 63 433-476, 1907.
18. 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.
19. 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.
20. L. Sirovich, M. Maxey and H. Tarman. Low-dimensional procedure for the characterization of human faces. *J. Optical Society*, 4A 519-524, 1987.
21. L. Sirovich and R. Everson. Management and analysis of large scientific datasets. *International Journal of Supercomputer Applications*, 6(1) 50-68, 1992.
22. 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.
23. G. W. Stewart. On the early history of the singular value decomposition. *SIAM Review*, 35(4) 551-566, 1993.
24. 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.
25. B.-L. Yeo and B. Liu. Rapid scene analysis on compressed video. *IEEE Trans. Circuits, Systems and Video Technology*, 5(6) 533-544, 1995.
26. 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.