EXTENSION OF MUTUAL SUBSPACE METHOD FOR LOW DIMENSIONAL FEATURE PROJECTION

Dragana Veljkovic[1], Kay A. Robbins[1], Doug Rubino[2] & Nicholas G. Hatsopoulos[3]

[1]University of Texas at San Antonio, Department of Computer Science
[2]University of California at San Diego, Department of Neuroscience
[3]University of Chicago, Department of Organismal Biology & Anatomy

ABSTRACT

Face recognition algorithms based on mutual subspace methods (MSM) map segmented faces to single points on a feature manifold and then apply manifold learning techniques to classify the results. This paper proposes a generic extension to MSM for analysis of features in high-throughput recordings. We apply this method to analyze short duration overlapping waves in synthetic data and multielectrode brain recordings. We compare different feature space topologies and projection techniques, including MDS, ISOMAP and Laplacian eigenmaps. Overall we find that ISOMAP shows the least sensitivity to noise and provides the best association between distance in feature space and Euclidean distance in projection space. For non-noisy data, Laplacian eigenmaps show the least sensitivity to feature space topology.

Index Terms—*Feature extraction*, *distance measurement, multidimensional systems, visualization*

1. INTRODUCTION

Mutual subspace methods (MSM) [13] for video face recognition transform faces segmented from sequences of video images to points on a feature manifold. Each point on the feature manifold represents a subspace that captures the distribution of segmented video corresponding to one person's face. Clustering or other distance-based techniques are then applied on the feature manifold to classify faces in the original space [1, 3, 5]. In other words, the face-classification problem is lifted to an abstract space in which manifold learning algorithms can be applied. Lifting works well for face recognition because the segmented faces are highly correlated and lie in low-dimensional subspaces [3].

Recordings of brain signals from implanted multielectrode arrays have very different features than segmented faces. However, the structure in short time windows appears to be low-dimensional, suggesting that recent advances in manifold learning might be useful in understanding and organizing this data. In brain recordings, the data are characterized by short-duration waves traveling in many directions, making the datasets difficult to analyze using standard Fourier techniques [8]. We extend the ideas of MSM to these datasets by representing short segments of video by single points in a lifted space. However, the segments no longer contain face images but traveling waves. Because waves are highly correlated signals, we can find linear, low-dimensional subspaces that capture a significant amount of energy. Unlike previously mentioned extensions of MSM, our goal is not classification but rather finding planar projections of the feature manifold that group segments containing waves with similar propagation directions. We demonstrate that the distance between the low-dimensional subspaces is a measure of similarity for the direction of the waves in corresponding segments.

In general, when it is possible to represent local segments of data by low-dimensional subspaces that capture a feature of interest, lifting can be applied to gain insight into the distribution of that feature in the original space. Lifting not only reduces the dimensionality of the data, but also reduces the dataset volume, since an entire segment is mapped to a single point on the feature manifold.

The rest of the paper is organized as follows. Section 2 describes the proposed technique, and section 3 applies the method to real and simulated data. Section 4 offers concluding remarks.

2. PROJECTIONS OF FEATURE MANIFOLDS

The algorithm consists of three phases: mapping to a feature manifold, applying a topology to this manifold and performing a low-dimension projection. Each of these steps needs to be tailored to the particular problem being solved.

2. 1. Mapping to the feature manifold

The first step is to lift the data from the original space to a feature manifold by applying a feature detection test. If the feature detection test on a sequence of video is positive, we compute the coordinates of the corresponding point on the feature manifold.

To illustrate this technique, we consider video recordings containing short duration waves propagating in different directions. We detect segments of video that contain waves by applying the wave subspace test [7]. The

video is segmented using fixed length sliding time windows. Each time frame within the window is converted to a vector after moving the origin to the frame containing the pixel with the smallest amplitude in the window. PCA is performed on the resulting set of vectors. If the second PCA mode is a phase shift of the first PCA mode to within a specified threshold, the time window is considered to have a wave feature. We have found that the first two PCA modes capture sufficient energy to adequately represent the feature. Thus for the reminder of this paper we will represent a feature by the two-dimensional subspace spanned by the two most energetic PCA modes calculated during the wave subspace test.

2. 2. Defining the topology of the feature space

Finding a distance that reflects the intrinsic structure of the target problem is crucial, because different distance metrics will result in different feature space topologies.

In this paper each segment of video that contains a wave is represented by a two-dimensional subspace. Much work has been done on different ways to compute subspace distances [3, 6, 10, 12]. We compare four subspace distances, all of which can be used with subspaces of different dimensions.

Let $\mathcal{A}_i$ be the i^{th} subspace and $\mathbf{A}_i$ be a matrix whose columns form an orthonormal basis for $\mathcal{A}_i$. The projection matrix for $\mathcal{A}_i$ is given by $\mathbf{P}_i = \mathbf{A}_i \mathbf{A}_i^{\mathrm{T}}$.

Subspace distance d_1 is based on principal angles:

$$d_1(\mathcal{A}_1, \mathcal{A}_2) = \sqrt{1 - s^2_{\min}} = \sqrt{1 - \cos^2 \theta_{\max}} \ .$$

Here $s_{\min}$ is the smallest singular value of the matrix $\mathbf{A}_1^{\mathrm{T}}\mathbf{A}_2$. The singular values of $\mathbf{A}_1^{\mathrm{T}}\mathbf{A}_2$ correspond to the cosines of the principal angles, so this distance compares two subspaces using only the largest principal angle [6].

Subspace distance d_2 uses all principal angles [3]:

$$d_2(\mathcal{A}_1, \mathcal{A}_2) = \sqrt{\sum_i \sin^2 \theta_i} = \frac{1}{\sqrt{2}} \left\| \mathbf{P}_1 - \mathbf{P}_2 \right\|_F .$$

Distances can also be induced by mapping either the orthonormal basis matrices $\mathbf{A}_i$ or the corresponding projection matrices $\mathbf{P}_i$ to vectors and applying any vector norm. Since projection matrices are independent of the choice of bases, they are better candidates. Distance d_3 is the L1 norm applied to the difference between the projection matrices, after converting these matrices to vectors [10]. Distance d_4 uses the L1 matrix norm, which corresponds to taking the largest column sum.

2. 3. Low dimensional projection

The previous steps follow closely the ideas of MSM. However, the goal here is not to distinguish between faces of different people but to find a low-dimensional mapping that represents the distribution of features. Therefore, the final step of the proposed technique is finding a low-dimensional projection that preserves distances in feature space. In this paper we compare 4 projection strategies.

Multidimensional scaling (MDS) is a linear projection technique that tries to preserve the pairwise distances between points in the original space [4]. MDS has been widely used in the social sciences to visualize complex relationships based on dissimilarity.

ISOMAP is a manifold learning algorithm that constructs a distance matrix using geodesic distances on the feature manifold and then applies MDS to this matrix to compute the low-dimensional projection [11].

Laplacian eigenmaps tries to preserve the proximity of nearby points during projection [2]. One must specify a weight function W_{ij} that gives the importance of point j in determining proximity to point i. A simple choice is $W_{ij} = 1/k$ for the k nearest neighbors of node i, with the other weights set to zero. We denote this version of the algorithm as LE. An alternative is to use Gaussian falloff for the neighborhood weight. We denote this alternative by LEG.

3. APPLICATION TO DATA WITH WAVES

We apply the proposed technique to two datasets, a synthetic dataset containing simulated waves and an experimental dataset recorded from the motor cortex of a monkey during the execution of a visuomotor task. Both datasets contain 10×10 images recorded at 1000 frames per second.

3. 1. Simulated waves

We generated overlapping sinusoidal waves by superposing the basic wave form:

$$W(x, y, t) = A(t) \cos(k_x x + k_y y + \omega t + \varphi)$$

A(t) is a Gaussian envelope amplitude function with mean 50ms and variance 30ms. This choice of parameters ensures that middle 100ms interval has a dominant wave direction and that no more than 30% of the wave energy is carried to another window. Adjacent waves overlap by approximately 60ms. The direction of wave propagation, θ, is randomly selected from the interval $[0, 2\pi]$. The parameters k_x and k_y are computed as $\cos(\theta)$ and $\sin(\theta)$, respectively. The parameters ω and φ control the speed of wave propagation and the phase offset, respectively.

To test the sensitivity of the approach to the details of feature selection, we applied the technique to the simulated dataset using four projection techniques with three different thresholds in the wave subspace test. A low threshold chooses only signals with nearly perfect waves, while higher thresholds include features that are less wave-like. Distance d_1 is used in all cases. ISOMAP is computed using 25 nearest neighbors, while LE and LEG use 15. The Gaussian kernel parameter in LEG is 2.

Fig. 1 displays the two-dimensional projections for the test. The mappings in the top row are computed using a rigorous wave detection test. The results suggest well-defined shapes, but some areas of the manifold do not have enough sample points. Applying the wave detection test with a very high threshold, as shown in the bottom row,

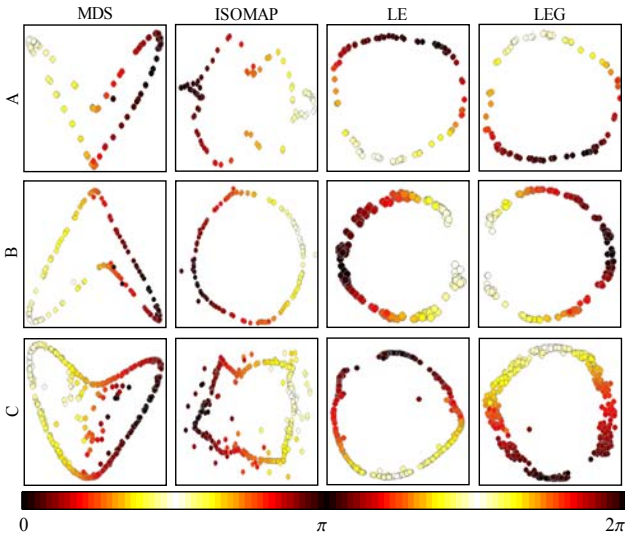

Figure 1. Projections of simulated waves using threshold 0.003 (A), 0.01 (B) and 0.1 (C). Window size is 70ms, overlap 20ms and topology is defined by d_1. ISOMAP uses 25 neighbors. LE and LEG use 15 neighbors. Points are colored based on computed wave direction in [0, 2π], as shown by the color bar.

identifies segments with lower correlations as waves. In this case, the feature manifold has too many noisy points and the shapes of the projections are quite distorted. The middle row illustrates a good selection of threshold.

We also computed the wave direction of each feature independently using a Hilbert transform-based method [8].

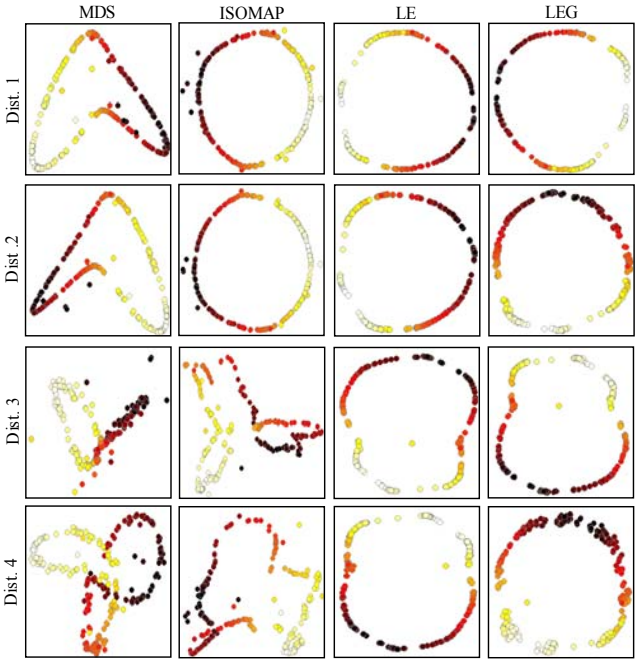

Figure 2. Projections of simulated waves using different topologies for the feature space. Window size is 70ms with overlap 20ms and threshold is 0.02. Colormap and other parameters as in Fig. 1.

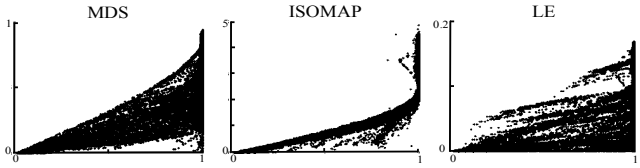

Figure 3. Distance d_1 in feature space vs. Euclidean distance in projection space for the synthetic dataset. Parameters as in Fig. 2.

The computed wave direction is used to color the projection points in all of the figures. The color bar displayed at the bottom of Fig. 1 shows the correspondence between wave direction and color. The smooth variation of colors in the resulting projections confirms that all of the projection techniques preserve wave directions.

Fig. 2 compares the effect of feature space topology on the different projection strategies. As expected, distances d_1 and d_2 give similar projection shapes independent of the projection technique. Laplacian eigenmaps appear to be the least sensitive to topology. MDS projections with distances d_3 and d_4 have smooth color variation, but the crossed curves indicate that projecting into three dimensions rather than two might result in a better projection curve.

Fig. 3 compares the distance in feature space with the distance between the corresponding projections. We analyze three of the projection techniques using d_1 in the feature space and Euclidian distance in the projected space.

The projected distance with MDS is always less than the original distance, and small distances are closely preserved. However, the capacity of the projected space is significantly smaller than that of the original feature space, and a significant number of distant points are projected in relatively close proximity. Thus, geometric interpretation of the MDS projections is somewhat problematic.

ISOMAP, which uses geodesic distances, displays a much better-defined relationship between distances in the two spaces. Points mapped to nearby locations in projection space are most likely to be neighbors in feature space. Notice, however, that except for very close points, the projected distance is larger than the original distance, suggesting a non-planar structure of the feature manifold.

Laplacian eigenmaps attempt to preserve neighborhood relationships but impose no restriction on non-neighboring points. Points that are close in the original space remain close in the projection, but the reverse is not the case. In fact, some points with very large feature space distances are mapped to almost identical locations.

3. 2. Multielectrode recordings of macaque motor cortex

We also applied the technique to the beta frequency band (15−45Hz) of local field potential signals recorded from a multielectrode array implanted in the motor cortex of a macaque monkey during an instructed target task [8]. Low signal to noise ratio and significant volume make this dataset difficult to analyze. The dataset is characterized by very noisy waves and a single preferred direction, although waves from all directions and amplitudes are present.

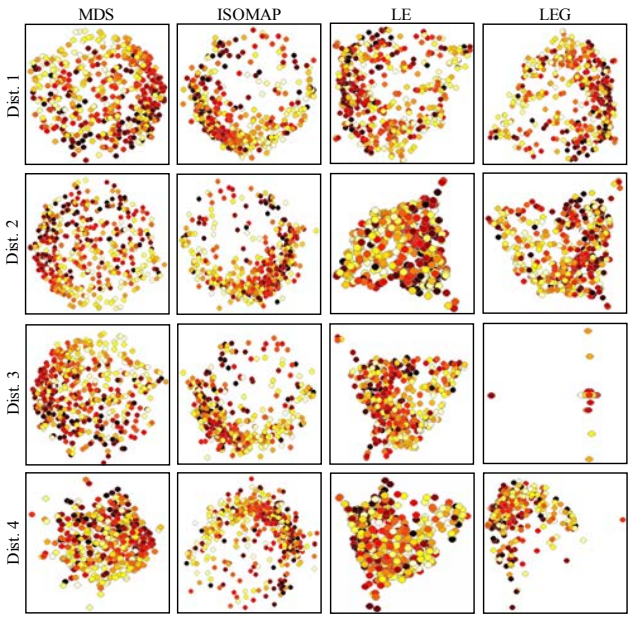

Figure 4. Projections of motor cortex recordings of monkey Rx using threshold 0.01, window size 15ms, and overlap 5ms. LE and LEG use 5 neighbors. Colormap and parameters as in Fig. 1.

Fig. 4 compares the results of the four projection techniques with different feature space topologies. ISOMAP provides the cleanest results, independent of distance measure. Distances d_3 and d_4 are especially sensitive to noise, and LE proved to be unstable using these distances.

Fig. 5 compares the d_1 distance in feature space with Euclidian distance in projection space. Observe that very few points are close in the feature space, indicating high dimensionality and very noisy data. All three projection techniques map nearby points in the feature space to nearby points in the projection space. Unfortunately, a significant number of points that are distant in the feature space are mapped to nearby points in projection space. This indicates that a 2D projection space has insufficient capacity to represent the dataset.

4. CONCLUSION AND FUTURE WORK

This paper applies an MSM strategy to analyze waves. We examine the effect of topology and dimension reduction algorithm by comparing the results using four feature space distance measures and four feature space projection methods. All combinations produce good low-dimensional mappings of wave direction for synthetic data, but ISOMAP appears to do the best job for noisy experimental data.

The feature space topology can also be changed by using a different feature representation. Defining an inner product function in the feature space would allow projections based on Gram and covariance matrices [9].

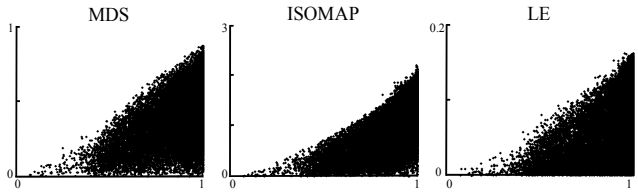

Figure 5. Distances in feature space with d_1 topology vs. distances in projection space for monkey Rx dataset. Parameter as in Fig. 4.

Acknowledgments

We acknowledge support from the Whitehall Foundation (N.G.H.), NIH R01 NS45853-01 (N.G.H.) and NIH Research Centers in Minority Institutions 2G12RR1364-06A1 (K. A. R.). Computational support was provided by the SA Computational Biology Initiative.

6. REFERENCES

[1] Arandjelovic, O. and R. Cipolla (2006). "Automatic cast listing in feature-length films with anisotropic manifold space." Computer Vision and Pattern Recognition 2: 1513-1520.

[2] Belkin, M. and P. Niyogi (2003). "Laplacian eigenmaps for dimensionality reduction and data representation." Neural Computation 15(6): 1373-1396.

[3] Chin, T. J., J. U. K. Schindler, et al. (2005). "Face recognition from video by matching image sets." Digital Image Computing: Techniques and Applications: 188-194.

[4] Cox, T. F. and M. A. A. Cox (2001). Multidimensional Scaling, Chapman and Hall.

[5] Fukui, K. and O. Yamaguchi (2003). "Face recognition using multi-viewpoint patterns for robot vision." International Symposium of Robotics Research: 192-201.

[6] Huynh, D. Q. and A. Heyden (2002). "Robust factorization for the affine camera: Analysis and comparison." International Conference on Control, Automation, Robotics and Vision: 126-131.

[7] Robbins, K. A. and D. M. Senseman (2004). "Extracting wave structure from biological data with application to responses in turtle visual cortex." Journal of Computational Neuroscience 16(3): 267-98.

[8] Rubino, D., K. A. Robbins, et al. (2006). "Propagating waves mediate information transfer in the motor cortex." Nature Neuroscience 9: 1549-1557.

[9] Saul, L. K., K. Q. Weinberger, et al. (2006). Spectral methods for dimensionality reduction, MIT Press: Cambridge, MA.

[10] Sun, X. and Q. Cheng (2006). "On subspace distance." Image Analysis and Recognition 2:81-89.

[11] Tenenbaum, J. B., V. de Silva, et al. (2000). "A global geometric framework for nonlinear dimensionality reduction." Science 290: 2319-2323.

[12] Wang, L., X. Wang, et al. (2006). "Subspace distance analysis with application to adaptive Bayesian algorithm for face recognition." Pattern recognition 39(3): 456-464

[13] Yamaguchi, O., K. Fukui, et al. (1998). "Face recognition using temporal image sequence." Automatic Face and Gesture Recognition: 318-323.