Non-Linear Embeddings and the Underlying Dimensionality of Reflectance Spectra and Chromaticity Histograms

Brian Funt and Dejan Kulpinski
School of Computing Science
Simon Fraser University, Vancouver, Canada
and
Vlad Cardei
Polaroid Corporation
Wayland, Massachusetts

Abstract

We have used the locally linear embedding[2] and Isomap[1] techniques to process high-dimensional color and spectral data. These techniques allowed us to create low-dimensional embeddings of the original data. In particular, the dimensionality is significantly lower than that obtained by principal components analysis. The data that we processed included spectral reflectances, illuminant spectra (both 101 dimensional) and chromaticity histograms (251 dimensional).

Isomap was useful in determining the dimensionality of the data in question. Using the Isomap technique, we found that the dimensionality for the spectral reflectances is between 3 and 4. For the chromaticity histograms, we found that the 251 original histogram dimensions were transformed into 5 to 6 dimensional space. In addition to providing an estimate of the underlying dimensionality of the data, both Isomap and LLE technique were used in producing the low-dimensional embeddings of the high dimensional data.

Introduction

There has been a lot of interest recently in non-linear embedding techniques. In this paper we investigate the application of these techniques to the problem of dimensionality reduction in the context of color. The non-linear embedding techniques are intended to uncover manifolds embedded within higher dimensional spaces. In other application areas, the methods have been shown to be quite effective. Intuitively speaking, in contrast to linear methods which linearly project the data onto a basis, the non-linear methods cling to the underlying manifold by tracing from one data point to its neighboring data points in the high dimensional space. For example, the underlying two-dimensional structure of a coiled spring can be discovered and unraveled into the plane (see the 'Swiss roll' example in Isomap[1] and LLE[2]).

In the color field, linear models of surface reflectance and illumination based on PCA (principal components analysis) have been widely used.[4,6,7] In general, the lower the dimension obtained by the analysis the better. Some methods such as the Maloney-Wandell[7] color constancy algorithm, require the dimensionality of reflectance spectra to be strictly less than the number of sensor classes. Previous approaches to dimensionality reduction include considering linear models of logarithms of reflectance spectra[16] and analyzing the structure of the linear model coefficients[17] of illumination spectra. Although non-linear methods cannot simply be substituted for the linear ones in most of these applications, it is still interesting to establish whether the data might have a lower dimensional description than previously realized.

Spectral Reflectance and Illuminant Spectra

We analyzed a set of 1996 reflectance spectra which include the Krinov[13] Kodak[14] and Munsell data sets. The results of PCA and Isomap are compared in Figures 1 and 2. As can be seen, the residual variance drops much more quickly in the Isomap case than the PCA case.

Generally speaking, the Isomap residual variance is similar to the PCA residual variance but a full dimension lower. For example, the non-linear embedding requires only 3 dimensions to reduce the residual variance to 0.03 while the PCA requires 4. Tenenbaum et. al.[1] suggest looking for the "elbow" in the each curve as a method of estimating the underlying dimensionality. Using this technique, we conclude the dimensionality of the structure Isomap finds is 3 or less.

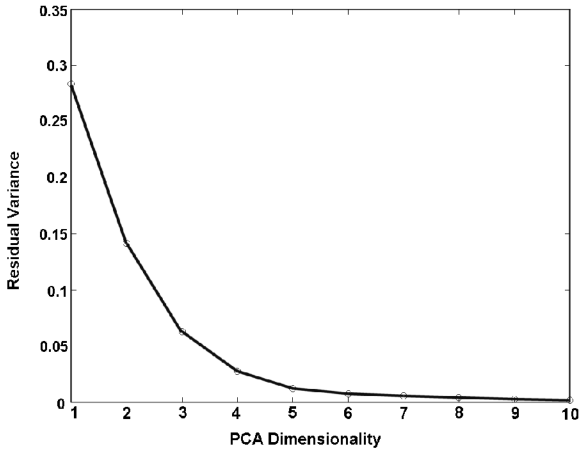

Figure 1.

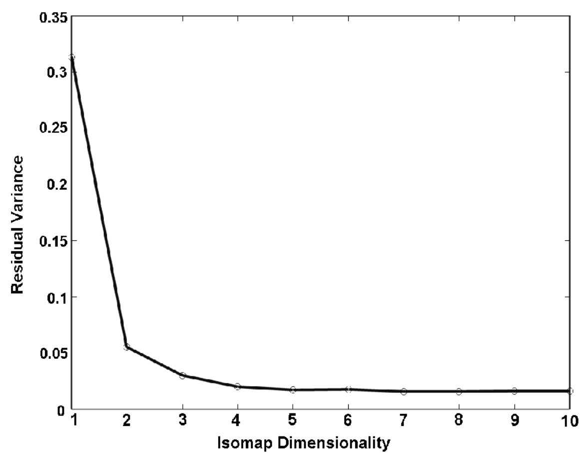

Figure 2.

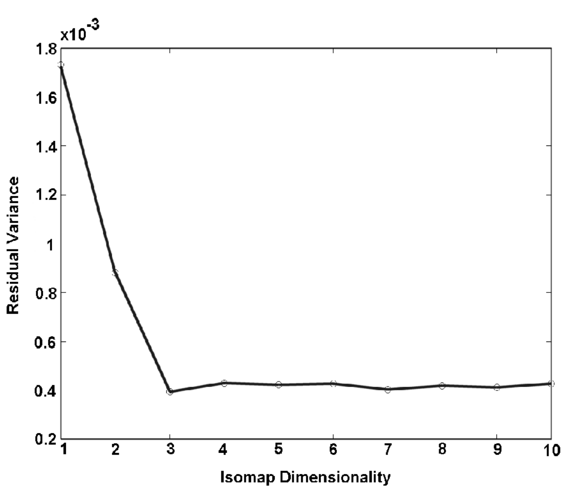

Figure 3.

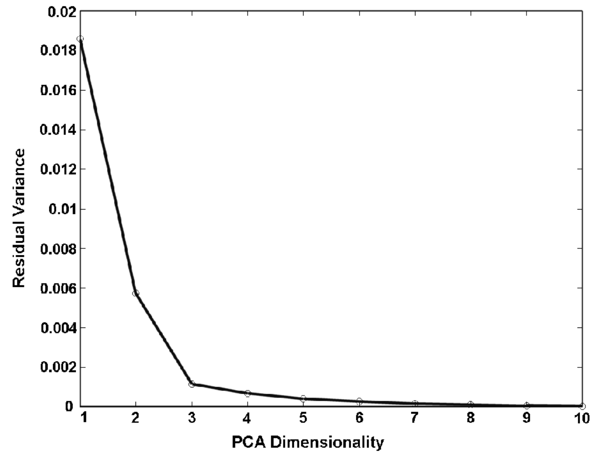

Figure 4.

We also did an Isomap analysis of a set of spectra of real 102 illuminant sources[3] that are available on the web.[9] Plots of the residual variance as a function of dimensionality are shown in Figures 3 and 4. The residual variance is very low even for just 1 dimension, although it is only at 3 dimensions that the residual variance levels off.

Of course, correlated color temperature is a standard one-dimensional parameterization for illuminants, so it is not unexpected to find a low-dimensional representation; however, it is reassuring to see how closely the data can be modeled by a non-linear 1-dimensional model. One application of 1-dimensional approximations for illuminants has been in some illumination estimation algorithms.[12,15]

Chromaticity Histograms

Color histograms and chromaticity histograms have been used widely in object recognition[8,10] and illumination estimation.[5,11] Here we use Isomap to analyze the structure of the space of binarized chromaticity histograms.

A chromaticity histogram of an image is a count of the number of pixels occurring in the image as a function of the pixels' chromaticities. A binarized chromaticity histogram is simply a chromaticity histogram in which the histogram bin counts have been set to zero or one depending on whether the original count was above or below a small threshold. A '1' in a binarized chromaticity histogram indicates the presence of that chromaticity in the image. Using a threshold avoids the problem of noise causing spurious chromaticities to appear in the histograms. For this analysis, we tried both *rg* and *rb* chromaticity spaces, where

$$r=R/(R+G+B);\ g=R/(R+G+B);\ b=R/(R+G+B)$$

R,G, and B are raw digital camera output values.

We used a library of 1050 images taken using different digital cameras, under different lighting conditions, both indoors and out, and including, for example, portraits of people, scenes from nature, Macbeth charts, buildings, cars.

Each image contains a standard grey card somewhere within it for use in evaluating its color cast. The images were also taken using a variety of settings on the cameras in terms of flash or no flash, automatic white balance on or off, manual white balance set to tungsten or daylight, and so forth. The cameras did not have a gamma option, so we presume some gamma-like tone correction was applied to all images.

Before generating an image's chromaticity histogram, the image is first pre-processed to remove very bright or clipped pixels (i.e., R or G or B near to the camera maximum) and very dark pixels. Both clipped and dark pixels lead to unreliable chromaticity data. The chromaticity space (a triangle with vertices at the origin, (0,1) and (1,0)) is coarsely quantized into 450 distinct chromaticity values. The resulting histograms therefore also contain 450 bins. We can think of the histograms as points in a 450-dimensional space. Although in principle all 450 bins could get used, in practice only 251 of them actually do get used by at least one of the 1050 images. Hence, the effective dimensionality of the space of chromaticity histograms of the entire set of images is only 251.

We hypothesized that the chromaticity histogram data might have a much lower dimensionality then initial 251 dimensions of the input space. In, addition we hypothesized that the main dimensions might correspond to variations in image color cast caused by scene illumination conditions and camera white balance settings.

We processed the chromaticity histogram data using both the LLE and Isomap algorithms. Incidentally, the computation times for this amount of data can be up to an hour. Figure 5 shows the residual variance versus the sub-space dimension. Once again using the Isomap "elbow" technique to establish the dimensionality of the data, it appears that the histograms fall in a subspace of 5 to 6 dimensions.

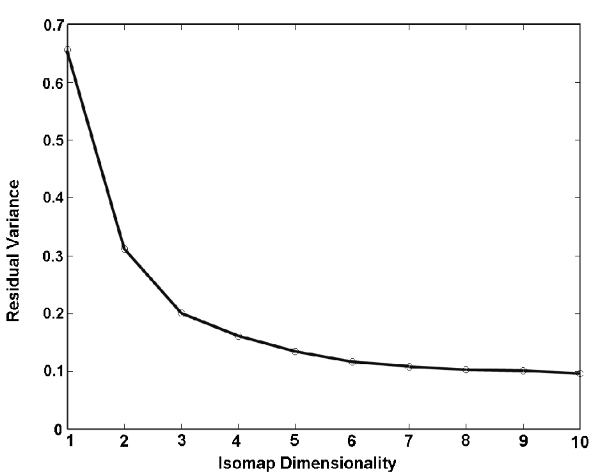

Figure 5.

To visualize the nature of the embedding we have plotted thumbnails of the images at the embedding location of their histograms. The columns of Figure 6 show slices orthogonal to the first 2 dimensions of the embedding of the histograms. We have also plotted, but do not show due to space limitations, similar plots with the chromaticity of the each image's grey card at the histogram's embedding location. In either representation (especially with larger and more numerous slices than we can show here), there appears to be a clear variation in color cast along dimensions 2 and 3. Roughly speaking, dimension 2 reflects an orange-to-cyan variation. This corresponds to a variation in red versus blue for a roughly constant amount of green. Similarly, dimension 3 reflects a change from cyan to purple. This corresponds to a variation in red versus yellow. We were somewhat surprised that color casts do not appear to be the dominant aspect of dimension 1. Our impression is that dimension 1 reflects the number of non-zero histogram bins. In other words, the primary variation along dimension 1 appears to be the number of distinct colors in the image.

Conclusion

Reflectance spectra lie in a non-linear subspace of dimension possibly as low as 3. This is a significant reduction in dimension for the equivalent dimensionality of a linear subspace, which for our reflectance data is at least 4. Chromaticity histograms also lie in a subspace of quite low dimension, with some of the primary dimensions reflecting the overall color cast of the corresponding image.

References

1. Joshua B. Tenenbaum, Vin de Silva, and John C. Langford, A Global Geometric Framework for Nonlinear Dimensionality Reduction, *Science*, Dec 22 2000: 2319-2323.
2. Sam T. Roweis and Lawrence K. Saul, Nonlinear Dimensionality Reduction by Locally Linear Embedding, *Science,* Dec 22 2000: 2323-2326
3. Kobus Barnard, Lindsay Martin, Brian Funt, and Adam Coath, A Data Set for Color Research, *Color Research and Application*, (to appear).
4. G. Finlayson, M. Drew and B. Funt, Color Constancy: Generalized Diagonal Transforms Suffice , *J. Opt. Soc. Am. A,* 11(11), 3011-3020, 1994a.
5. B. Funt, V. Cardei and K. Barnard, Method of Estimating Chromaticity of Illumination Using Neural Networks, *U.S. Patent 5,907,629.*
6. D.B. Judd, D.L. MacAdam and G.W. Wyszecki, Spectral distribution of typical daylight as a function of correlated color temperature, *J. Opt. Soc. of Am. A*, Vol. 54, 1031, 1964.
7. L. Maloney and B. Wandell, Color constancy: a method for recovering surface spectral reflectance, *J. Opt. Soc. Am. A*, 3:29-33, 1986.
8. M.J Swain and Ballard, Color Indexing, *Intl. Journal of Computer Vision*, 7:1, pp11-32, 1991.
9. Computational Vision Lab, Simon Fraser University, Data for Computer Vision and Computational Colour Science (http://www.cs.sfu.ca/~colour/data/)

Figure 6 Each of the 4 panes in the left column shows a slice orthogonal to dimension 1 of the subspace embedding with subspace dimension 2 varying vertically and 3 horizontally. The slices are ordered from top to bottom according to the intersection point along the dimension-1 axis from negative to positive. Each pane shows a collage of thumbnail images with each thumbnail plotted at the embedding location of the corresponding image's chromaticity histogram. The right-hand column is similar except each of the 4 panes shows a slice orthogonal to dimension 2 of the subspace embedding with subspace dimension 1 varying vertically and 3 horizontally.

10. Funt, B.V., and Finlayson, G., Color Constant Color Indexing, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17 (5), 522-529, May 1995
11. G. Finlayson, P. Hubel and S. Hordley, Color by Correlation, *Proc Fifth Color Imaging Conf.*, pp. 6-9, Scottsdale, November 1997.
12. S. Tominaga, S. Ebisui and B. Wandell, Color temperature estimation of scene illumination, *Proc. Seventh Color Imaging Conference*, pp. 42-47, Scottsdale, Nov. 1999.
13. E. L. Krinov, Spectral Reflectance Properties of Natural Formations, *National Research Council of Canada* (1947).
14. M.J. Vrhel, R. Gershon and L.S. Iwan, Measurement and analysis of object reflectance spectra, *Color Research and Application*, vol. 19, pp 4-9, 1994.
15. G. Finlayson, Color Constancy at a Single Pixel, CIC7
16. B. Funt, Modeling Reflectance by Logarithmic Basis Functions, *Proc. First Color Imaging Conference*, Scottsdale, Nov. 1993.
17. J. DiCarlo and B. Wandell, Illuminant Estimation: Beyond the Bases, *Eighth Color Imaging Conference: Color Science*, Systems and Applications, p.p. 91-96, 2000.

Acknowledgments

Financial support provided by the Natural Sciences and Engineering Research Council of Canada.