Monthly Notices
of the
ROYAL ASTRONOMICAL SOCIETY

MNRAS **461,** 2250–2256 (2016)
Advance Access publication 2016 June 24

doi:10.1093/mnras/stw1510

Locally linear embedding: dimension reduction of massive protostellar spectra

J. L. Ward[1,2]★ and S. L. Lumsden[2]

[1]*Astrophysics group, Lennard-Jones Building, Keele University, Keele, Staffordshire ST5 5BG, UK*
[2]*School of Physics and Astronomy, E. C. Stoner Building, University of Leeds, Leeds LS2 9JT, UK*

Accepted 2016 June 21. Received 2016 June 20; in original form 2016 January 26

ABSTRACT

We present the results of the application of locally linear embedding (LLE) to reduce the dimensionality of dereddened and continuum subtracted near-infrared spectra using a combination of models and real spectra of massive protostars selected from the Red MSX Source survey data base. A brief comparison is also made with two other dimension reduction techniques; principal component analysis (PCA) and Isomap using the same set of spectra as well as a more advanced form of LLE, Hessian locally linear embedding. We find that whilst LLE certainly has its limitations, it significantly outperforms both PCA and Isomap in classification of spectra based on the presence/absence of emission lines and provides a valuable tool for classification and analysis of large spectral data sets.

Key words: methods: data analysis – stars: protostars – infrared: stars.

1 INTRODUCTION

Advancements in instrumentation and observing facilities have led to an ever increasing rate of production of spectroscopic data. It is therefore an increasingly complex task to extract relevant data from large data sets. Dimension reduction algorithms offer the ability to reduce complex data sets into the lowest number of free parameters necessary to address a specified degree of variation. The most commonly used dimension reduction algorithm applied to astronomical data is principal component analysis (PCA; Deeming 1964; Jolliffe 1986); however, it is most sensitive to globally linear variations and has proven inefficient in classification based on varying emission line strengths (Yip et al. 2004).

Locally linear embedding (LLE) is a manifold mapping dimension reduction algorithm introduced by Roweis & Saul (2000). It has already been successfully applied to the classification of galaxies and QSOs (Vanderplas & Connolly 2009) and the separation of stars, galaxies and quasars (Daniel et al. 2011), both using data from the Sloan Digital Sky Survey (York et al. 2000). It has more recently been applied by Bu et al. (2013) to the classification of stellar subclasses for M-type stars. In these previous studies, it was shown to provide useful distinction between classes of object based on features in the spectra which relate to real physical properties. In all previous applications LLE has proven to reduce the dimensionality of data more efficiently than PCA and is better at separating spectral classes.

LLE attempts to compute low-dimensional embeddings of high-dimensional input data whilst preserving the local neighbourhood of each input vector. The basis of the technique is to map an underlying lower-dimensional manifold upon which the high-dimensional data lies. When applying this technique to spectra, each spectrum is treated as a vector, $\boldsymbol{X}_i$, with a number of dimensions equal to the number of wavelength bins of the spectrum. For example a spectrum with a wavelength range of 1.5–2.5 μm where each element covers 13.351 Å is treated as a 750 dimensional vector. We summarize the LLE algorithm in the following steps.

(i) Taking a set of n vectors $\boldsymbol{X}_i$ with dimensionality D, the k nearest neighbours for each vector are calculated based on Euclidean distance. These k nearest neighbours form the local neighbourhood of $\boldsymbol{X}_i$.

(ii) Linear weights to each of the nearest neighbours are calculated. This is done by minimizing the cost function:

$$\epsilon(W) = \sum_i \left| \boldsymbol{X}_i - \sum_{j=1}^{k} W_{ij} \boldsymbol{X}_j \right|^2 . \quad (1)$$

(iii) The inner products between each vector and each of its nearest neighbours are computed to produce a neighbourhood correlation matrix, $\mathbf{C}_{jk} = \boldsymbol{X}_j \cdot \boldsymbol{X}_k$.

(iv) The reconstruction weights are then computed using

$$W_{ij} = \sum_k \mathbf{C} - jk^{-1}(\boldsymbol{X}_i \cdot \boldsymbol{X}_k + \lambda). \quad (2)$$

To ensure that the embedding is only based on a vectors local neighbourhood, W_{ij} must be equal to 0 if X_j is not one of the nearest neighbours of X_i. The sum of all of the weights to the neighbours of a single vector, λ, is equal to 1 (i.e. $\sum_j W_{ij} = 1$). This ensures that every neighbourhood used in the embedding is equally valid and therefore no single data point can distort the final

★E-mail: j.l.ward@keele.ac.uk

embedding. This is done using the Lagrange multiplier, $\lambda = \alpha/\beta$, where $\alpha = 1 - \sum_{jk} \mathbf{C}_{jk}^{-1}(\boldsymbol{X}_i \cdot \boldsymbol{X}_k)$ and $\beta = \sum_{jk} \mathbf{C}_{jk}^{-1}$.

(v) Finally, the embedded vectors $\boldsymbol{Y}_i$ which will make up the final output data are calculated by minimizing a second cost function:

$$\Phi(Y) = \sum_i |\boldsymbol{Y}_i - \sum_j W_{ij}\boldsymbol{Y}_j|^2, \quad (3)$$

where the weights W_{ij} are known and the set of vectors $\boldsymbol{Y}_i$ need to be calculated. The constraint that $\sum_i \boldsymbol{Y}_i = \mathbf{0}$ is imposed in order to centre the projection on the origin and a second constraint is imposed in order to avoid degenerate solutions;

$$\frac{1}{n}\sum_i \boldsymbol{Y}_i \otimes \boldsymbol{Y}_i = \mathbf{I} \quad (4)$$

where $\mathbf{I}$ is the $d \times d$ identity matrix (d being the dimensionality of the projected vectors).

(vi) This cost function now defines a quadratic form containing the $N \times N$ symmetric matrix $\mathbf{M}_{ij}$:

$$\Phi(\mathbf{I}) = \sum_{ij} \mathbf{M}_{ij}(\boldsymbol{Y}_i \cdot \boldsymbol{Y}_i), \quad (5)$$

where $\mathbf{M}_{ij} = \delta_{ij} - W_{ij} - W_{ij} + \sum_k W_{ki}W_{kj}$ and $\delta_{ij} = 1$ if $i = j$ and $\delta_{ij} = 0$ otherwise.

(vii) The lowest $d+1$ eigenvectors of the matrix $\mathbf{M}_{ij}$ are computed with the exception of the bottom eigenvector (which represents a free translation mode of eigenvalue 0) in order to find the optimum embedding. The embedded output is a representation of the variance between input data and is completely independent of the original form of the data.

The Red MSX[1] Source (RMS) survey (Lumsden et al. 2013) is the largest catalogue of massive young stellar objects (MYSOs) and ultracompact H II regions to date. This survey has led to a series of follow-up observations including near-infrared spectroscopy of the largest collection of MYSO candidate objects to date using the UK InfraRed Telescope (Cooper et al. 2013) and the New Technology Telescope (Ward et al. in preparation). It is from this set that the spectra used in this work have been selected.

In this paper, we apply the LLE algorithm to available *HK* spectra of Galactic MYSOs in order to assess the effectiveness of dimension reduction algorithms as a method of automated classification. In Section 2 the input data will be described and the methodology used will be outlined. The results will be presented in Section 3, while Section 4 will provide a comparison with other dimension reduction algorithms for the same data set.

2 INPUT DATA AND METHODOLOGY

The majority of *HK* spectra of MYSOs can be separated into four distinct types based on the different emission line features that are present and it is believed that these types likely represent an evolutionary sequence. The types and the rational behind the types, are explained in detail in Lumsden et al. (in preparation) but the main emission line classification criteria of the types are outlined in Table 1.

2.1 Real spectra

The real reduced spectra used in this work are presented, along with investigations into the properties of the spectra, in two other papers;

[1] Midcourse Space Experiment (Egan et al. 2003).

Table 1. Formalization of MYSO type classification criteria.

Type	H_2	Br γ	Br10	*Fluor.* Fe II
I	Present	Absent	Absent	Absent
II	Present	Present	Absent	Absent
III	Present	Present	Present	Present
IV	Absent	Present	Present	Present

Cooper et al. (2013) and Ward et al. (in preparation). All spectra were reduced and prepared for input into the LLE algorithm using the FIGARO data reduction package and the DIPSO spectral analysis package. The SNIP function of DIPSO was used to remove data in the range $17800 < \lambda < 19700$ Å as this region of the spectrum has the poorest atmospheric transmission. The spectra were then re-binned into 750 elements in the range 15000–25000 Å using the SCRUNCH operation in FIGARO.

At this point, the data are split into two sets; the dereddened spectra and the continuum subtracted spectra. For the former, the spectra were dereddened using the DERED function in DIPSO with extinction values from Cooper et al. (2013) and Ward et al. (in preparation). The continuum subtracted data had third-order polynomials fitted to the continua which were then subtracted.

The final step in the preparation of the real spectra was to cross-correlate the spectra to a reference point. The reference spectrum used for all cross-correlation was that of the compact H II region G274, chosen for its prominent and extensive set of Brackett series emission lines. Any spectra for which the absolute wavelength shift was greater than one element were discarded. The remaining spectra were then classified into the types shown in Table 1, discarding any that did not fall into any of these categories or where the type was uncertain. This yielded a remaining 184 spectra (94 from Cooper et al. 2013 and 90 from Ward et al. in preparation), with 28 type I sources, 79 type II, 35 type III and 42 type IV. The average real spectra for each type are shown in the upper four panels of Fig. 1.

2.2 Models

Preliminary tests showed that the low number of spectra used in this study did not provide any significant result when the LLE algorithm is applied. We therefore include an additional large set of model spectra to act as a framework for the real spectra. Two main sets of models have been created to complement the two sets of real spectra; a dereddened set of models and a continuum subtracted set. To generate these models a series of simple PYTHON scripts were used.

The continuum of the dereddened set of models was calculated using a blackbody distribution with a randomly generated temperature in the range $0 < T < 40\,000$ K for each spectrum and the wavelength at each element to calculate the appropriate flux. The values in the range 17 800 Å to 19 700 Å were set to zero to match the section which was removed from the real spectra. The emission lines were then included through the addition of Gaussian profiles using the rest wavelength of each emission line as a mean and an amplitude coefficient. The amplitude coefficients were selected so that the mean peak fluxes of the model lines equal the mean peak fluxes of the real emission lines and are related to the peak flux by

$$f_{\rm peak} = A_{\rm e} \times \frac{1}{\sigma_l\sqrt{2\pi}} \times \frac{1}{2\sigma_l^2} \times 10^{-24}, \quad (6)$$

where $f_{\rm peak}$ is the peak flux, $A_{\rm e}$ is the amplitude coefficient and $\sigma_{\rm l}$ is the line width. Finally, randomly generated noise was added to each point along the spectrum using a Gaussian distribution with a

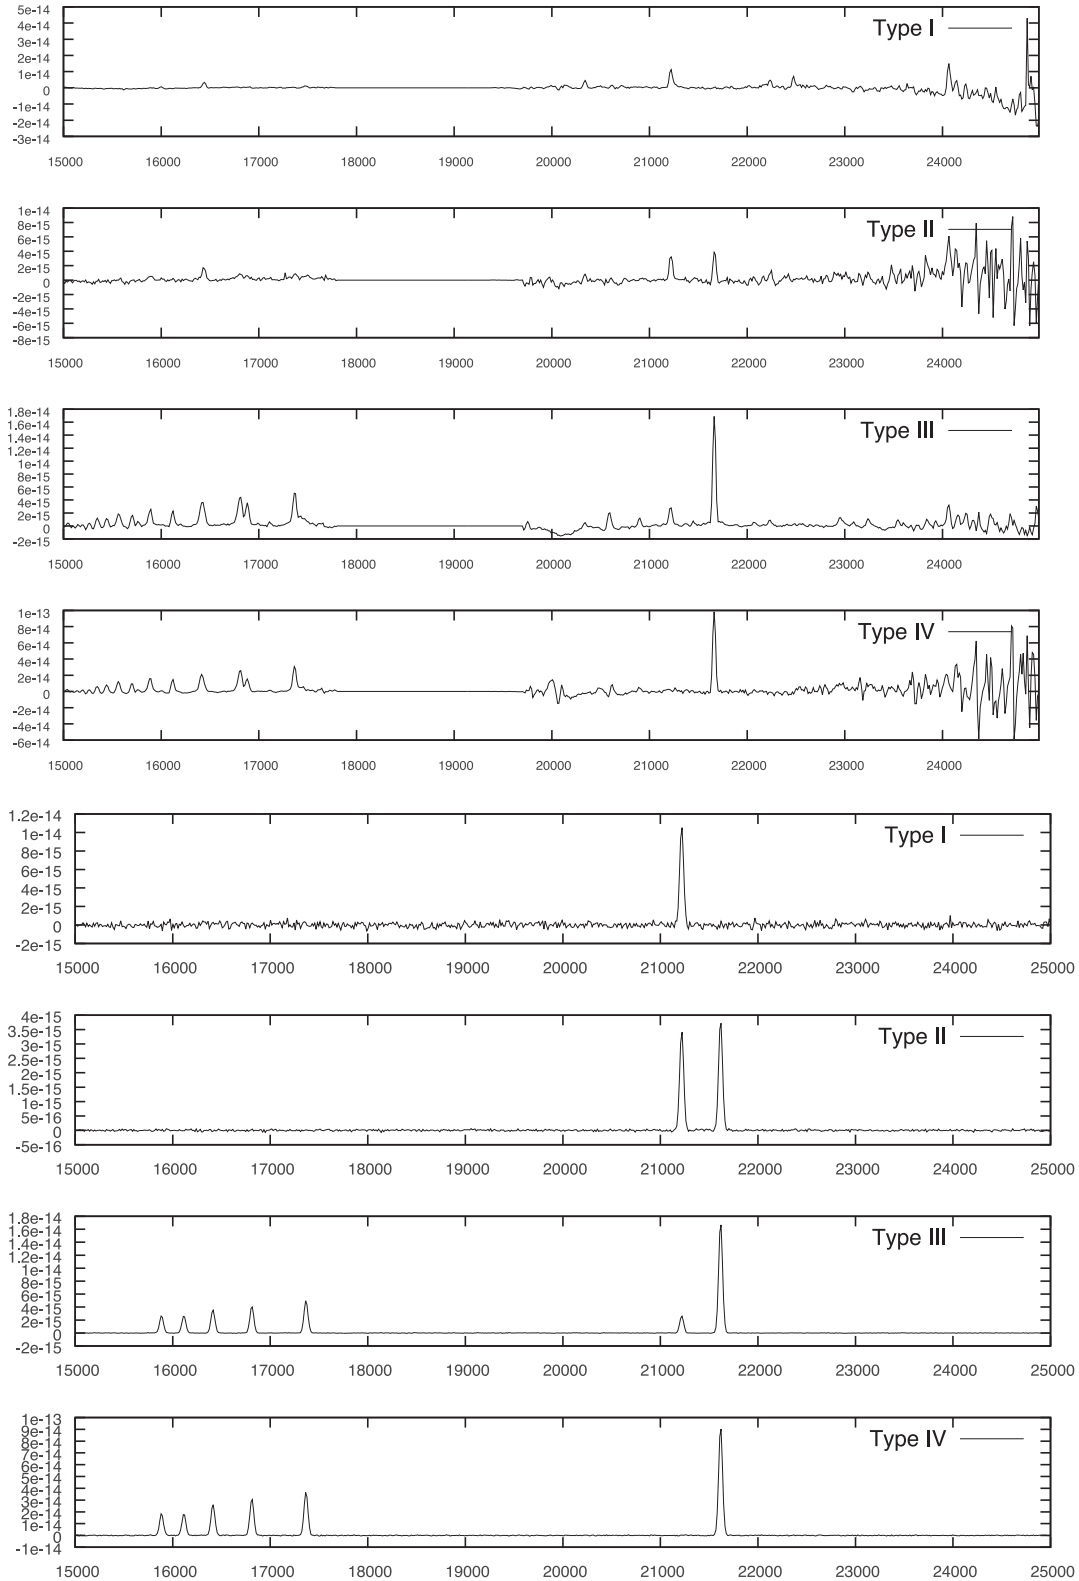

Figure 1. Upper: average spectra for each of the types I-IV calculated from continuum subtracted real spectra. Lower: average spectra for each of the types I-IV calculated from the model spectra.

mean set to zero and where the σ is inversely proportional to λ^4 to simulate the increase in noise from the *K* band to the *H* band which is exaggerated by the dereddening process in real spectra. The high levels of noise in the long wavelength end of the *K* band is neglected but this area of the spectra will not be included in the inputs for the dimension reduction algorithms. A total of 8000 models were generated for both the dereddened and continuum subtracted data sets, equally divided between the four types and with the average

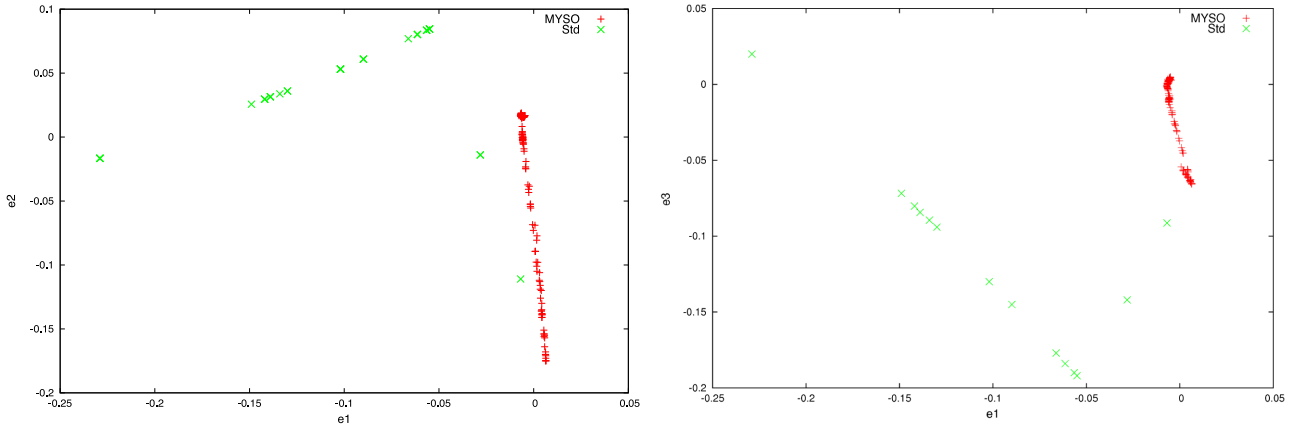

Figure 2. Left: output e2 and e1 values of LLE applied to real MYSO spectra along with the telluric standard star spectra using nearest neighbours, $k = 8$. Right: output e3 versus e1 values of LLE applied to real MYSO spectra with telluric standard spectra.

model spectra for each type shown in the lower four panels of Fig. 1. The average model spectra exhibit significantly lower levels of noise than the average real spectra because of the much larger number of models. All of the models generated for this work were based on making random variations surrounding a set of mean line profiles and a mean continuum (in the case of the dereddened models) which were designed to match those of the average real spectra.

2.3 Methodology

The LLE node (Vanderplas & Connolly 2009) from the Modular Data Processing toolkit (MDP; Zito et al. 2009) was used to perform the LLE. A range of values for the number of nearest neighbours, k were used in order to find the optimum sampling for the data. This was to ensure that the quality of the resulting projections was not limited by the number of nearest neighbours used. The LLE node of MDP determines the output dimensionality required to express a specified percentage of the variance in a similar manner to that suggested by de Ridder & Duin (2002). First a covariance matrix is calculated for each local neighbourhood and an eigenanalysis is performed to find the minimum number of dimensions needed to represent the specified percentage of local variance. The output dimensionality used in the LLE projection is then the mean of the dimensionalities calculated for the local neighbourhoods (Vanderplas & Connolly 2009).

As mentioned earlier, the longer wavelength end of the K band was not used for LLE with the cut-off point at 2.28 μm. This point was chosen in order to avoid the CO bandhead which, as it does not affect the classification of the spectrum, would introduce an unnecessary complication into the data set. It also avoids the high levels at noise at wavelengths greater than 2.35 μm which again is irrelevant to the classification criteria set out in the previous section. All spectra were normalized to a distance of 1 kpc prior to dimension reduction using distances from the RMS survey data base.

3 RESULTS

3.1 Distinguishing MYSOs from standard stars

As a preliminary test, we used reduced and normalized spectra (before dereddening) of both MYSOs and the telluric standard calibration stars from the same observing runs. Prior to the application of LLE to these data, both the standard star spectra and the MYSO spectra were rebinned using SCRUNCH as in the preparation procedure for the main LLE runs. They were then normalized by the mean value of the continuum emission for each spectrum. This ensures that the projection will be independent of total NIR continuum brightness. Fig. 2 shows the three lowest eigenvector output dimensions of the LLE embedding and in both e2 versus e1 space and e3 versus e1 space the MYSO spectra form a line which is clearly distinguishable from the standard spectra. It is therefore apparent that LLE could provide a method for automated classification in larger spectroscopic surveys.

3.2 Classification of MYSOs

Preliminary tests using the 8000 model dereddened spectra found that the spectra are indistinguishable based on type, forming simple polynomial shaped projections in the output eigenspace. This is attributed to the significant continuum emission in the spectra, which represents the dominant source of variation. It is apparent however that for a data set consisting of a large number of less embedded spectra, LLE may provide an efficient method for estimating the temperatures of all of the spectra simultaneously as the temperature is the dominant source of variation between these models (see Fig. 3).

As no distinction can be made between types based on raw or dereddened spectra due to the dominance of continuum emission, the next logical step is to perform LLE on spectra without continuum emission. The process of continuum subtraction based on a simple polynomial fitting is a step which can be easily automated for most spectral data sets and would therefore would be reasonable to include in an initial classification algorithm for a large spectroscopic survey.

Our preliminary tests using only continuum subtracted real spectra of MYSOs yielded no significant results that we were able to interpret. The most likely explanation for this is that there are not enough spectra included in order to produce a reasonable projection of the variation within the data based on any physical properties.

In order to account for 95 per cent of the variation between spectra in our data set only three output dimensions were required when using LLE. The lowest two output dimensions (representing the largest variation) of the LLE projection when applied to our model and real continuum subtracted spectra are shown in Fig 4. When LLE is applied to our model continuum subtracted spectra it is

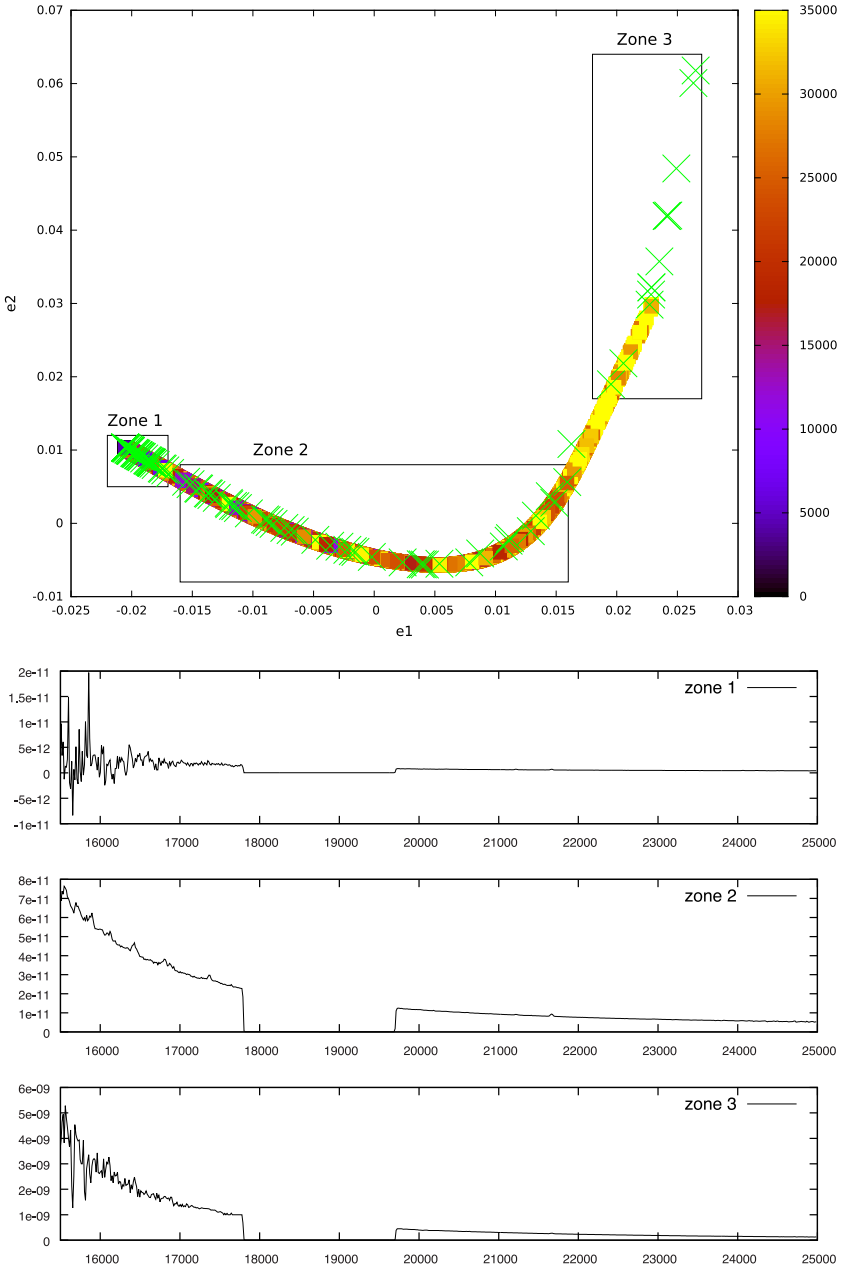

Figure 3. Upper panel: LLE eigenvector outputs e2 versus e1 for model spectra, colour coded by model blackbody temperature and real spectra (green Xs). Lower panel: average spectrum for the real spectra in each of the zones in the left-hand panel.

immediately clear that the separation of the model types is significantly improved compared with the application to dereddened spectra. The projection in e2 versus e1 resembles two main branches, one of type I and II spectra and one of type III and IV spectra with an area of mixing between type II and III spectra close to the point at which the two main branches converge. This point of convergence coincides with the spectra with the lowest signal-to-noise ratios.

The real type I spectra coincide well with their model counterparts, as do the majority of the type II sources. There is considerable mixing between the type III/IV sources and some of the type II sources however. Regardless of any mixing however, a clear trend can be seen from positive e1 and negative e2 values in type I spectra and negative e1 with positive e2 values in types III and IV. With more robust models which accurately represent the Fe II emission and the complete Brackett series it is likely that this distinction will be clarified.

A simple classification script has been written incorporating LLE which automatically assigns a type based on the location of the object in e2/e1 space using the angle and distance of the sources from the point at which the I/II and III/IV branches converge. We find that even with this simplistic approach 63 per cent of the real spectra were successfully determined as belonging to either types I/II or III/IV and for type I/II spectra, this figure rises to 79 per cent. Whilst this leaves many spectra without classification, less than 3 per cent of the spectra were actually mistakenly assigned incorrect types.

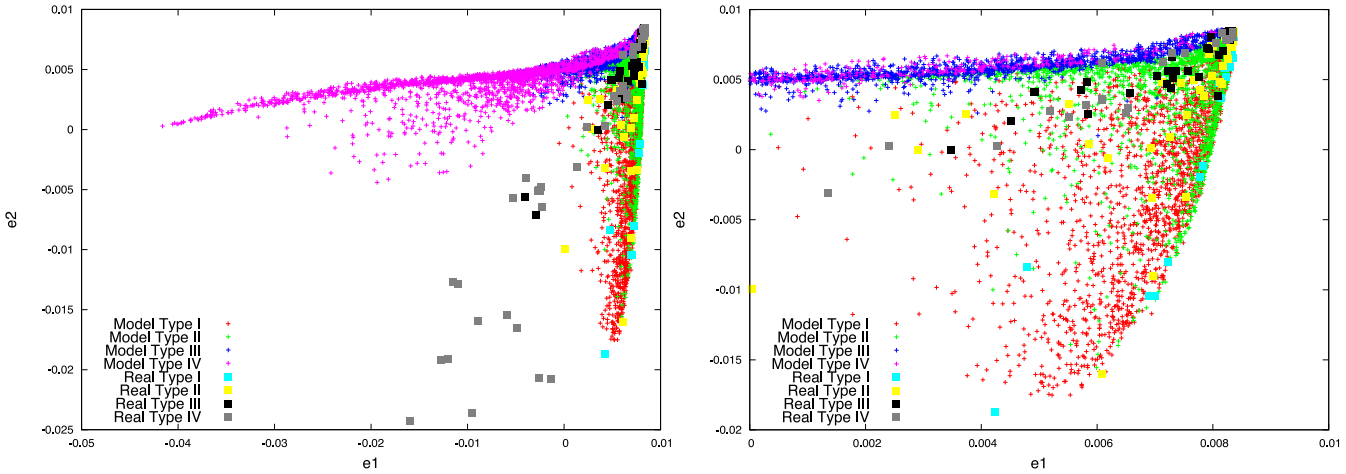

Figure 4. Left: LLE e2 versus e1 output for models plus real spectra with nearest neighbours k=100. Right: same as left but showing only $0 < \mathrm{e1} < 0.01$.

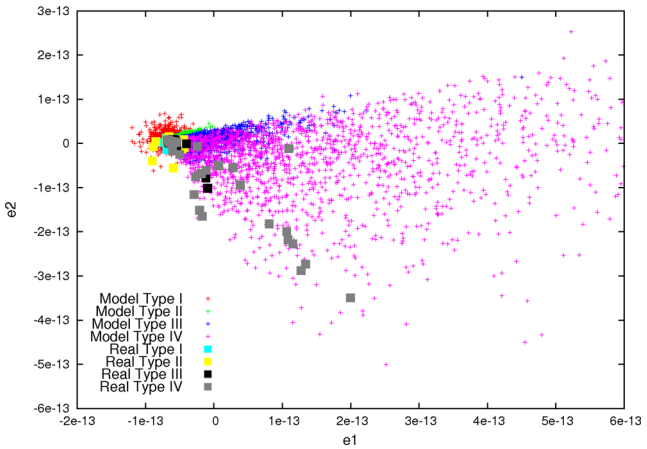

Figure 5. PCA output with dimensions equivalent to those output from LLE. The types of the data points correspond directly to those of Fig. 4.

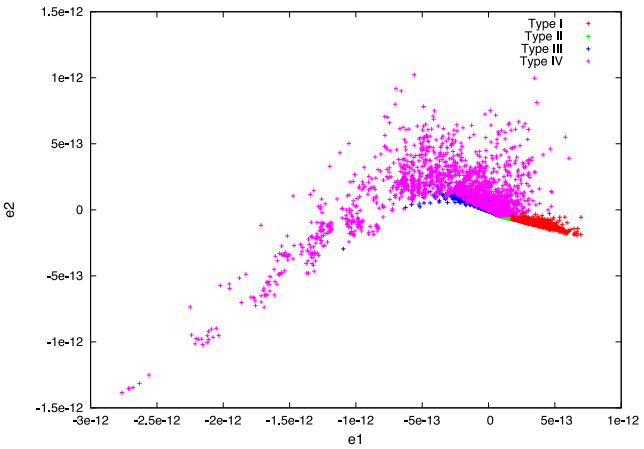

Figure 6. Isomap output with dimensions equivalent to those output from LLE.

4 COMPARISON WITH OTHER DIMENSION REDUCTION ALGORITHMS

We will now compare our outputs of the LLE dimension reduction algorithm with those of the similar, widely used dimension reduction algorithms PCA and Isomap. Finally, we will compare the output of a more advanced version of LLE; Hessian locally linear embedding (HLLE).

4.1 PCA and Isomap

PCA requires 431 output dimensions in order to account for 95 per cent of the variance between our (model and real) spectra. As a dimension reduction algorithm therefore, it has proved to be inefficient compared with LLE. Fig. 5 shows the resulting second and first eigenvector outputs of the application of PCA to the same data set as LLE is applied to in Fig. 4. It is immediately apparent that all of the spectra are clumped together. We find similar effects in all output dimensions and conclude that for such similar spectra, PCA cannot be used to accurately distinguish between our different types in this way. Performing Isomap (Tenenbaum, de Silva & Langford 2000) on our set of model spectra (the output is shown in Fig. 6) we find that a type I and type IV spectra are more adequately separated than when PCA is applied. The Isomap output does however lack much of the finer structure of the LLE output as a result of losing much of the local neighbourhood information which is preserved in the LLE algorithm.

4.2 Hessian locally linear embedding

Finally, we apply the HLLE algorithm (Donoho & Grimes 2003) to the same data to see whether our results with LLE can be improved upon. HLLE differs from the standard LLE algorithm through the use of Hessian estimators in place of the weights in the original LLE procedure. The nearest neighbours are found by calculating Euclidean distances and the matrix $\mathbf{M}$ is produced for each neighbourhood as in LLE. Then a singular value decomposition of $\mathbf{M}$ is performed to obtain tangent coordinates and the Hessian estimator is developed;

$$\mathbf{H}_{i,j} = \sum_{l} \sum_{r} \left(\mathbf{H}^{l}_{r,i} \mathbf{H}^{l}_{r,j} \right), \tag{7}$$

where $\mathbf{H}^{l}$ is a $d(d+1)/2 \times k$ matrix associated with estimated the Hessian over the neighbourhood N_i. The eigenvectors corresponding to the lowest eigenvalues are computed from the matrix $\mathbf{H}$ similar to the LLE analysis of the matrix $\mathbf{M}$. Embedding coordinates are obtained from the matrix $\mathbf{W} = V^T R^{-1/2}$, where $R = VV^T$.

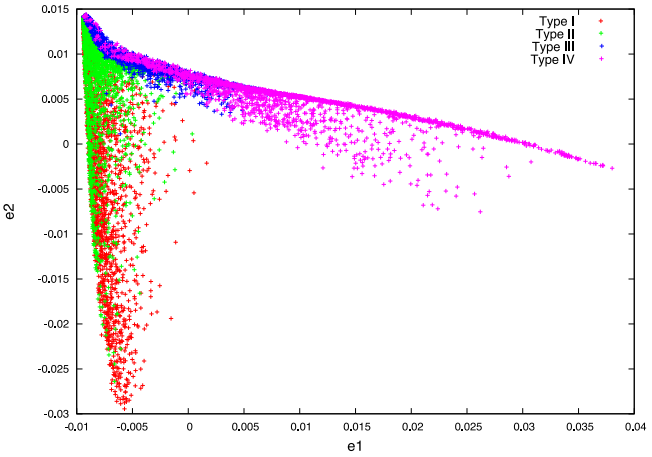

Figure 7. HLLE output with dimensions equivalent to those output from LLE.

The output embedding e2 versus e1 plot for the HLLE algorithm is shown in Fig. 7. When compared to the equivalent output from the standard LLE algorithm in Fig. 4 we find that the application of HLLE gives no marked improvement over that of LLE for this particular data set. This is most likely due to the continuous nature of the model production which provides a data set that is inherently complete without any gaps. It is likely therefore that the HLLE algorithm would provide a greater advantage where no model spectra are available.

5 CONCLUSIONS

We have successfully applied LLE to a relatively small set of real spectra using a framework of 8000 model spectra. The conclusions of this application of LLE to massive protostellar spectra are summarized below.

(i) LLE has successfully been utilized to separate MYSO spectra from telluric standard stars with no training sets or modification to the spectra and it is therefore likely that LLE would be particularly useful in the early analysis of large spectroscopic surveys to classify sources.

(ii) It was not possible to separate the dereddened spectra based on their emission lines because they are dominated by continuum emission. It may be possible in some cases however to determine properties of the continua of spectra and for objects without significant continuum emission (relative to the line emission/absorption) this would not be an issue.

(iii) Continuum subtracted spectra were separated into the expected types using LLE and through a simple automation of the interpretation of the output data almost two-thirds of the spectra were successfully assigned a type whilst only two sources were misclassified.

(iv) LLE has shown itself to be superior to PCA and Isomap in terms of the efficiency of dimension reduction and separation of MYSO spectra based on the presence and absence of relevant emission lines. Little improvement was found when replacing LLE with HLLE but this would likely be more effective with a less continuous set of data than our models.

Whilst computationally more expensive than many of the available alternative dimension reduction algorithms, LLE and its derivatives HLLE and robust LLE present a powerful tool for classifying large spectroscopic data sets, outperforming the more commonly utilized PCA and Isomap algorithms. With the advent of the next generation of observing facilities and ever more advanced multi-object spectrographs, the ability to quickly analyse and classify large samples of spectra will play a fundamental role in future large-scale surveys.

ACKNOWLEDGEMENTS

The authors thank the anonymous referee for his/her useful comments. JLW acknowledges financial support from the Science and Technology Facilities Council of the UK (STFC) via the PhD studentship programme. This paper made use of information from the Red MSX Source survey data base at http://rms.leeds.ac.uk/cgi-bin/public/RMS_DATABASE.cgi which was constructed with support from STFC. This research has made use of the SIMBAD data base, operated at CDS, Strasbourg, France.

REFERENCES

Bu Y., Pan J., Jiang B., Wei P., 2013, PASJ, 65, 81
Cooper H. D. B. et al., 2013, MNRAS, 430, 1125
Daniel S. F., Connolly A., Schneider J., Vanderplas J., Xiong L., 2011, AJ, 142, 203
de Ridder D., Duin R., 2002, Technical Report PH-2002-01, Locally Linear Embedding for Classification. Delft Univ. Technology, Delft, the Netherlands
Deeming T. J., 1964, MNRAS, 127, 493
Donoho D. L., Grimes C., 2003, Proc. Natl. Acad. Sci., 100, 5591
Egan M. P. et al., 2003, The Midcourse Space Experiment (MSX) Point Source Catalog Explanatory Guide. v2.3, Air Force Research Laboratory Technical Report No. AFRL-VS-TR-2003-1589, Air Force Research Laboratory, Hamscom AFB, MA, USA
Jolliffe I. T., 1986, Principal Component Analysis. Springer-Verlag, Berlin
Lumsden S. L., Hoare M. G., Urquhart J. S., Oudmaijer R. D., Davies B., Mottram J. C., Cooper H. D. B., Moore T. J. T., 2013, ApJS, 208, 11
Roweis S. T., Saul L. K., 2000, Science, 290, 2323
Tenenbaum J. B., de Silva V., Langford J. C., 2000, Science, 290, 2319
Vanderplas J., Connolly A., 2009, AJ, 138, 1365
Yip C. W. et al., 2004, AJ, 128, 585
York D. G. et al., 2000, AJ, 120, 1579
Zito T., Wilbert N., Wiskott L., Berkes P., 2009, Frontiers Neuroinformatics, 2, 8

This paper has been typeset from a TEX/LATEX file prepared by the author.