Looking for invariance in Locally Linear Embedding

Elodie Maignant
Epione, Inria Sophia-Antipolis, France – Centre Borelli, ENS Paris-Saclay, France
elodie.maignant@inria.fr

Locally Linear Embedding [1] is a method for learning intrinsically low-dimensional data embedded in a high-dimensional space, for example points sampled from an underlying surface. It constructs a neighbourhood-preserving embedding and recovers the global nonlinear structure of the data from locally linear fits.

In the method, neighbourhoods are described using affine barycentric coordinates. This construction makes them invariant to rotations, translations and rescaling and thus should characterise the intrinsic local geometric properties of the data. Therefore, it is expected that the final embedding reflects the intrinsic structure of the data. However, some experiences have shown that these coordinates may not be rigid enough and could allow for unwanted reconstruction patterns [2].

In this work, we try to characterise under which assumptions, especially on the neighbourhood graph, the reconstruction is unique up to affine transformations. At the same time, we look for new invariants of local barycentric coordinates in order to predict the properties of the embeddings constructed by Locally Linear Embedding. Ultimately, we hope to understand better which part of the global structure of the data is recovered by the method and which is not.

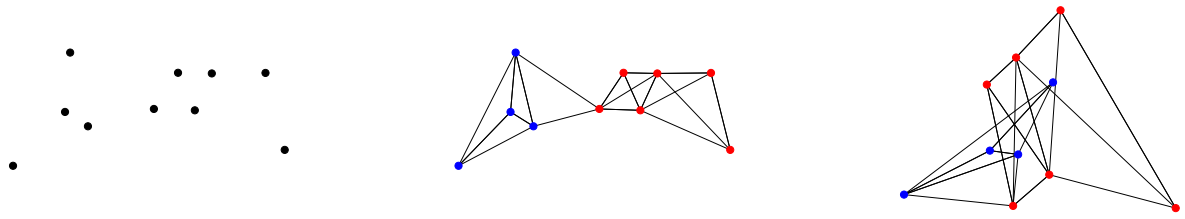

(a) Data points. (b) Neighbourhood graph. (c) Embedded points.

Figure 1: Some effect of the neighbourhood graph's connectivity in Locally Linear Embedding.

Joint work with: Xavier Pennec and Alain Trouvé.

References

[1] Sam T Roweis and Lawrence K Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500):2323–2326, 2000.

[2] Lin, Liren. Avoiding unwanted results in locally linear embedding: A new understanding of regularization. *arXiv preprint*, arXiv:2108.12680, 2021.