Low dimensional representations of MEG/EEG data using Laplacian Eigenmaps

Alexandre Gramfort, Maureen Clerc

HAL Id: inria-00502735
https://inria.hal.science/inria-00502735

Submitted on 15 Jul 2010

LOW DIMENSIONAL REPRESENTATIONS OF MEG/EEG DATA USING LAPLACIAN EIGENMAPS

Alexandre Gramfort, Maureen Clerc

Odyssée Laboratory
ENPC - ENS Ulm - INRIA
France

ABSTRACT

Magneto-encephalography (MEG) and electro-encephalograhy (EEG) experiments provide huge amounts of data and lead to the manipulations of high dimensional objects like time series or topographies. In the past, essentially in the last decade, various methods for extracting the structure in complex data have been developed and successfully exploited for visualization or classification purposes. Here we propose to use one of these methods, the Laplacian eigenmaps, on EEG data and prove that it provides an powerful approach to visualize and understand the underlying structure of evoked potentials or multitrial time series.

Index Terms— Laplacian eigenmaps, dimensionality reduction, EEG, P300, Event-related potentials

1. INTRODUCTION

The problem of dimensionality reduction is inherent to every data analysis procedure but its application can have various motivations, among which computation time, visualization or classification. When dealing with MEG or EEG data, linear dimensionality reduction methods like principal component analysis (PCA) or independent component analysis (ICA) are commonly used to remove artifacts and to select interesting components relatively to the experimental paradigm. What such methods fail to do, because of the linear constraint, is to provide an idea of the global structure of the data.

To circumvent this limitation, the past years have seen a growing interest in non linear dimensionality reduction methods. Among the most recent and popular techniques are the Locally Linear Embedding (LLE) [2], Laplacian eigenmaps [3] and diffusion maps [4]. These techniques construct an adjacency graph between the data points in order to capture the structure of the data set and map them into a lower-dimensional space while preserving the local properties of the graph. Such dimensionality reduction, that aims to map data points which are close - respectively to a certain metric - onto points which are close in the low dimensional space, can be achieved using spectral methods i.e. through the resolution of a generalized eigenvectors problem. Recent work in electromagnetic source
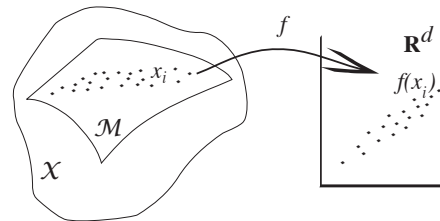

Fig. 1. Non-linear embedding into a low-dimensional Euclidian space

imaging uses a spectral method to constrain the solutions of the inverse problem to belong to a low dimensional space [1].

In this paper, we apply the Laplacian eigenmaps technique on EEG topographies and time series recorded on multiple trials to obtain a spectral representation of our data. We prove that it succeeds in extracting complex structures such as a cycle during the positive deflection of the P300 neural evoked potential or the delays of the neural response in an oddball event-related potentials experimental paradigm.

2. THEORY

The general formulation of dimensionality reduction problems is the following. Let $(x_i)_{i=1,...,N}$ be N objects in a metric space $\mathcal{X}$. Let us assume that these objects are sampled from a low-dimensional manifold $\mathcal{M} \subset \mathcal{X}$, which can be simply reformulated by saying that the (x_i) belong to a low-dimensional structure. Such an hypothesis is not very restrictive since most data satisfy this constraint. In the case of EEG, not every topography reflects a plausible neural activation which shows that valid topographies belong to a subpart of the space $\mathcal{X}$ of lower dimensionality.

Given the metric and the (x_i), the challenge of non-linear dimensionality reduction methods is therefore to recover the structure of $\mathcal{M}$ and to provide a mapping f from $\mathcal{M}$ to a low-dimensional Euclidian space $\mathbb{R}^d$ (cf. Fig. 1). We now present one of the existing algorithms, the Laplacian eigenmaps.

Laplacian Eigenmaps

In order to extract the structure of the manifold $\mathcal{M}$, the first step of the Laplacian eigenmaps method is to build, out of the data, an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ where $\mathcal{V}$ are the vertices (the $(x_i)_{i=1,...,N}$ in our case) and $\mathcal{E}$ are the edges. $\mathcal{G}$ is a weighted graph. For every edge $(i, j) \in \mathcal{E}$, the weight of the edge is noted w_{ij}.

In practice, we chose to use a gaussian kernel with standard variation σ [3]. The w_{ij} are then given using the metric $d_{\mathcal{X}}$ in $\mathcal{X}$ by :

$$w_{ij} = e^{-\frac{d_{\mathcal{X}}(x_i, x_j)^2}{\sigma^2}}$$

Note that to circumvent the problem of choosing the σ, it is possible to set $w_{ij} = 1$ for connected point in the graph and $w_{ij} = 0$ otherwise. However using a gaussian kernel has a strong justification in terms of diffusion processes over the manifold [3, 4]. The Laplacian matrix L of the graph is given by $L = D - W$ where $W = (w_{ij})_{ij}$ and D is a diagonal matrix with $D_{ii} = \sum_j w_{ij}$. D_{ii}, classically called the degree of vertex x_i, can be interpreted as a measure of the empirical density of points around x_i.

The Laplacian eigenmaps method consists in finding the generalized eigenvectors (y_k) solution of :

$$L y_k = \lambda_k D y_k.$$

It can be easily proved that 0 is a trivial eigenvalue but also that L is a symmetric positive matrix which implies that the eigenvalues λ_k are all positive and can be ordered :

$$0 = \lambda_0 \leq \lambda_1 \leq \cdots \leq \lambda_k \leq \lambda_{k+1} \leq \cdots \leq \lambda_{N-1}$$

The mapping f is then given by :

$$f(x_i) = f_i = (y_1(i), y_2(i), ..., y_d(i))$$

L being symmetric, computing only the first eigenvectors can be done efficiently with an iterative method using for example the *eigs* command of Matlab®. For computational efficiency, it is also possible to set the smallest w_{ij} to 0 leading to the diagonalization of a sparse matrix which is even faster.

To give a quick justification of this method, the mapping f is the optimal solution of the following minimization problem (cf. [5]) :

$$f = \arg \min_f \sum_{(i,j) \in \mathcal{N}} w_{ij} \, ||f(x_i) - f(x_j)||^2_{\mathbb{R}^d} \qquad (1)$$

with the constraint :

$$(f_1, \ldots, f_d)^T D (f_1, \ldots, f_d) = I_d$$

If $d = 1$, this constraint avoid the trivial mapping, that consists in mapping all the $(x_i)_i$ to a single point.

Intuitively this method intends to map objects of $\mathcal{X}$ onto points of $\mathbb{R}^d$ under the constraint that, if two objects are close in $\mathcal{X}$ they should be close in the low dimensional space $\mathbb{R}^d$. This representation is the one which optimally preserves the graph structure in a d dimensional space.

For more development on spectral graph theory, we invite the reader to consult [5, 6]. An interpretation as kernel methods of such non linear dimensionality reduction algorithms is presented in [7].

3. MATERIAL AND METHODS

We now present two results of spectral representation of EEG data, and we show what valuable information can be extracted out of two different data sets using Laplacian eigenmaps. In the first experiment, the objects considered are topographies with 32 electrodes, while in the second experiment objects are time series.

3.1. Spectral representation of P300 evoked potential

In this first illustration of the method, we computed a 2 dimensional spectral representation ($d = 2$) of EEG topographies with 32 electrodes ($\mathcal{X} = \mathbb{R}^{32}$). Topographies were extracted from a data set containing a classical P300 neural evoked potential, available as a sample data set in the EEGLAB Matlab®Toolbox [8]. Data were sampled at 128Hz, and 80 trials were averaged to compute the evoked signal. We used the Euclidian distance for $d_{\mathcal{X}}$, but it would be interesting to introduce other distances to see how it affects the comparison of EEG topographies.

Results are presented at three time points indicated as vertical bars in Figure 2-a. The first time point is before stimulation, the second one is at the beginning of the positive deflection of the P300 while the third one is at the end. In Figures 2-b1, b2 and b3, the three corresponding EEG topographies are represented. It appears that topographies before and after the P300 are accumulated in a dense region of the 2D spectral representation presented in Figure 2-c while topographies during the P300 follow a cycle. Since in the first part of the cycle, points are quite far from the points in the last part of the cycle, we can conclude that the topographies at the beginning of the P300, are quite different from the ones in the last part. This implies that the localization of the neural activity evolves during the positive deflection of the P300. The difference between the two topographies, one frontal (Fig. 2-b2) and one more parietal (Fig. 2-b3), is a confirmation. This observation is also present in the analysis in [8].

3.2. Reordering raster plots with a spectral embedding

In this illustration, each x_i is a time series, and $\mathcal{X} = \mathbb{R}^T$ where T is the number of time frames. The distance $d_{\mathcal{X}}$ is again the Euclidian distance. A same experiment has been carried out N times leading to N trials. In practice, a time series could be the evolution of the electric potential on a single
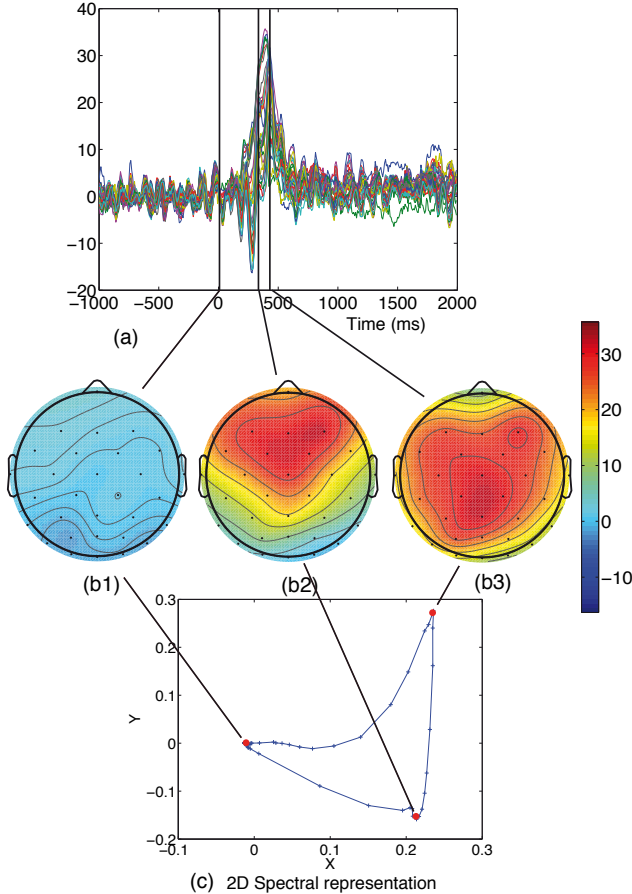
Fig. 2. Spectral embedding of EEG topographies at 3 time points (0 ms, 340 ms and 420 ms)

electrode or alternatively it could represent the activation of a PCA or ICA component during the experiment. As a pre-processing, the time series are centered and normalized. Note that doing this implies that minimizing the Euclidian distance for $d_\mathcal{X}$ is equivalent to maximizing a normalized correlation between time series to measure their similarity in $\mathcal{X}$.

It is reasonable to assume that a similar neural activation occurs in each repetition of the experiment. However, the delay between the response and the stimulation should not be identical across trials. This leads to the intuition that the $(x_i)_{i=1,...,N}$ lie on a noisy one dimensional manifold which can be parameterized by this delay.

This is confirmed by two experiments, presented in Figures 3 and 4. The first experiment is run on synthetic data, consisting in 100 time series ($N = 100$ and $T = 500$) computed from a reference (cf. Fig. 3-e) by adding an autoregressive noise and translating the positive deflection with a random time lag, with a Gaussian probability of standard deviation σ ($\sigma = 50$). The second experiment is run on real acquisitions of oddball event-related potentials, previously used

in [9], where a description of the experimental paradigm can be found. In Figure 3-a and Figure 4-a, which present raster plots of both data sets, the stochastic nature of the time lag is evident.

In both cases, the time series were embedded into a two dimensional space (cf. Figure 3-b and Figure 4-b). It can be noticed that the points in the 2D spectral representation are clustered along an elongated, 1D structure and the first coordinate was used to reorder the time series. By observing the reordered raster plots in Figure 3-d and Figure 4-d, it appears that the first coordinate in the low dimensional space has captured the information about the lag in the neural response. In both subfigures, the trend of the time lags has been highlighted manually by adding a dark line along the maxima of the time series. Note that such a simple "line-draw" highlighting would have been impossible in the unordered trials displayed in the raster plots of Figures 3-a and Figure 4-a, and is made possible after the reordering which was provided by the Laplacian eigenmap embedding.

We extracted manually the time-lags at each trial, and we realigned the time series according to it. An average of all time series was then computed, with and without realignment, leading to the results in Figure 3-e and Figure 4-e. In the synthetic case, we observe a very good match between the average of the realigned time series and the reference that was used to generate the data. As expected, realigning the time series provides bigger deflections by reducing the blurring effect caused by the variable delays of the neural response.

4. ACKNOWLEDGEMENT

5. CONCLUSION AND FUTURE WORK

The above illustrations show that Laplacian eigenmaps are able to reveal, without any prior knowledge on the experimental paradigm, some of the complex structure of EEG data. Such an observation suggests that non linear dimensionality reduction methods can help to provide a better understanding of the underlying neural phenomena. We also pointed out that the formalism behind such methods is very general and perfectly adapted to the manipulation of objects as different as time series or topographies. After these encouraging results, we plan to investigate the role the $d_\mathcal{X}$ distance in order to find more adapted similarity measures for MEG and EEG standard objects. One other goal is to define automatic postprocessing procedures, such as the time-lag detection.
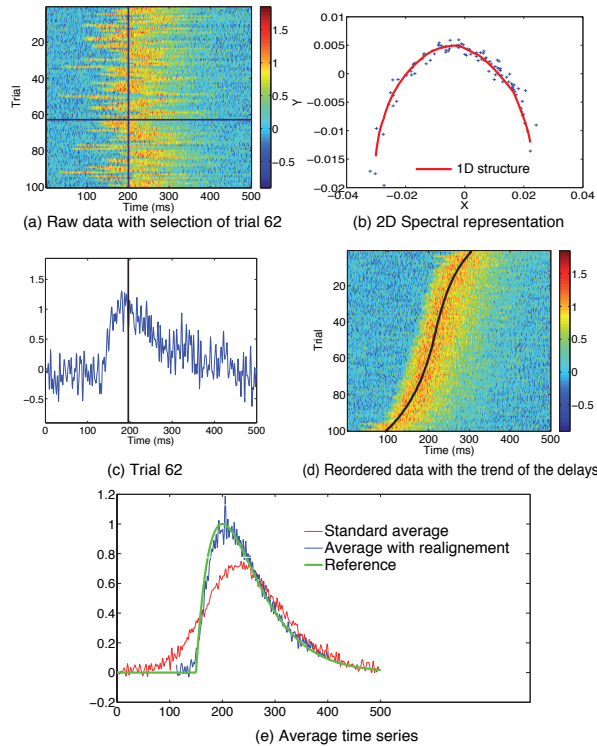
Fig. 3. Illustration on synthetic data

Fig. 4. Illustration with oddball event-related potentials (Recordings on the Cz electrode filtered at 8Hz)

6. REFERENCES

[1] Fijoy Vadakkumpadan and Y. Sun, "Embrioss: Electromagnetic brain imaging by optimization in spectral space," *Proceedings of ISBI*, pp. 1024–1028, Mar 2007.

[2] Sam Roweis and Lawrence Saul, "Nonlinear dimensionality reduction by locally linear embedding," *Science*, vol. 290, pp. 2323–2326, 2000.

[3] Mikhail Belkin and Partha Niyogi, "Laplacian eigenmaps for dimensionality reduction and data representation," *Neural Computation*, vol. 15, no. 6, pp. 1373–1396, jun 2003.

[4] R.R Coifman, S Lafon, A.B Lee, M Maggioni, and Nadler, "Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps," *Proceedings of the National Academy of Sciences*, vol. 102, no. 21, pp. 7426–7431, 2005.

[5] F. Chung, *Spectral graph theory*, CBMS-AMS, 1997.

[6] A. Ng, M. Jordan, and Y. Weiss, "On spectral clustering: Analysis and an algorithm," in *Advances in Neural Information Processing Systems*, 2002, vol. 14.

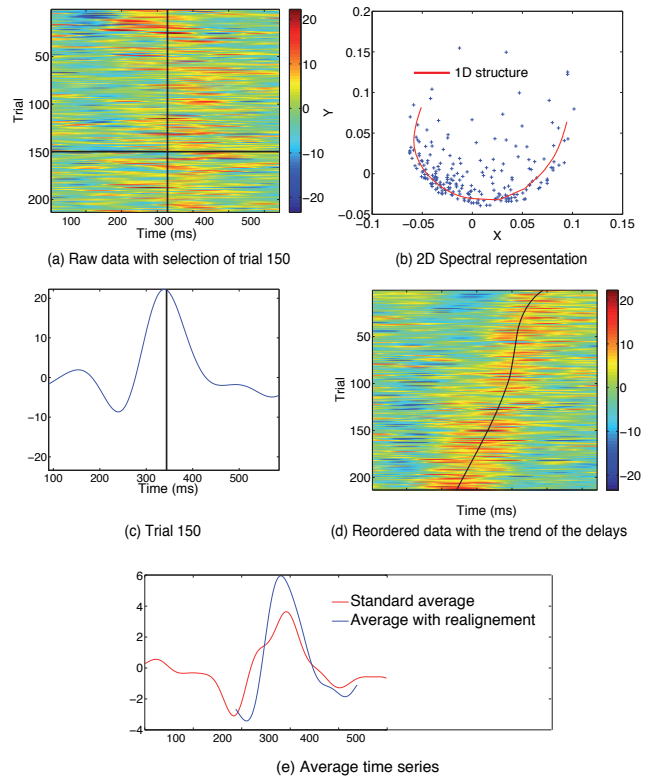[7] Jihun Ham, Daniel D. Lee, Sebastian Mika, and Bernhard Scholkopf, "A kernel view of the dimensionality reduction of manifolds," in *ICML '04: Proceedings of the twenty-first international conference on Machine learning*, New York, NY, USA, 2004, p. 47, ACM Press.

[8] Scott Makeig, Arnaud Delorme, Marissa Westerfield, Tzyy-Ping Jung, Jeanne Townsend, Eric Courchesne, and Terrence J Sejnowski, "Electroencephalographic brain dynamics following manually responded visual targets," *PLoS Biol*, vol. 2, no. 6, pp. e176, Jun 2004.

[9] C.G. Bénar, D. Schön, S. Grimault, B. Nazarian, B. Burle, M. Roth, J.M. Badier, P. Marquis, C. Liegeois-Chauvel, and J.L. Anton, "Single-trial analysis of oddball event-related potentials in simultaneous EEG-fMRI.," *Human Brain Mapping*, vol. in press, 2007.