SPECTRAL ANALYSIS OF ALIGNMENT IN MANIFOLD LEARNING

Hongyuan Zha[*]

Dept. of Computer Science and Engineering
The Pennsylvania State University
University Park, PA 16802,

Zhenyue Zhang[†]

Dept. of Mathematics
Zhejiang University, Yuquan Campus
Hangzhou, 310027, P. R. China

ABSTRACT

Local manifold learning methods produce a collection of overlapping local coordinate systems from a given set of sample points. Alignment is the process to stitch those local systems together to produce a global coordinate system and is done through the computation of the eigen-subspace of a so-called alignment matrix. In this paper, we present an analysis of the eigen-structure of the alignment matrix giving both necessary and sufficient conditions under which the null space of the alignment matrix recovers the global coordinate system. We also show by analyzing examples that the gap in the spectrum of the alignment matrix is proportional to the size of the overlap of the local coordinate systems. Our results pave the way for gaining better understanding of the performance of local manifold learning methods.

1. INTRODUCTION

For a *parameterized manifold* of dimension d defined by a mapping $f : \Omega \rightarrow \mathcal{R}^m$, where $d < m$, and Ω open in $\mathcal{R}^d$, suppose we have a set of points $x_1, \cdots, x_N$ sampled possibly with noise from the manifold, i.e.,

$$x_i = f(\tau_i) + \epsilon_i, \quad i = 1, \ldots, N, \tag{1.1}$$

where ϵ_i's represent noise. We are interested in recovering the τ_i's and/or the mapping $f(\cdot)$ from the x_i's. This process is generally known as manifold learning or nonlinear dimension reduction [3, 5]. A class of so-called local manifold learning methods start with estimating a collection of overlapping local coordinate systems around each sample points[1] and align (either implicitly or explicitly) those local coordinate systems to obtain a global one and thus recovering the τ_i's. Examples of those local methods include LLE (Local Linear Embedding) [3], manifold charting [1], Hessian LLE [2] and LTSA (Local Tangent Space Alignment) [6]. The alignment idea in particular was also discussed in [4]. The actual alignment process is done by computing the eigenvectors of a so-called alignment matrix corresponding to the smallest eigenvalues that have a gap to the other eigenvalues. In practical applications, however, this gap is not clear, and as we will see later that unwanted eigenmodes can be mixed in giving rise to inaccurate recovery of the τ_i's.

The focus of this paper is to gain a better understanding of the alignment process by analyzing the eigen-structure of the alignment matrix. In particular, we derive both necessary and sufficient conditions under which the null space of the alignment matrix recovers the τ_i's given that each local coordinate system is computed exactly. By way of analyzing a simple case of two overlapping local pieces, we also show that the gap in the spectrum of the alignment matrix is proportional to the *square* of the size of the overlap of the local coordinate systems. Proofs of the results will be omitted due to space constraints

NOTATION. We use e to denote a column vector of all ones the dimension of which should be clear from the context. $\mathcal{N}(\cdot)$ and $\mathcal{R}(\cdot)$ denote the null space and range space of a matrix, respectively. For an index set $I = [i_1, \ldots, i_k]$, $A(:, I)$ denotes the submatrix of A consisting of columns of A with indices in I. A similar definition is used for the rows of a matrix.

2. AN ILLUSTRATIVE EXAMPLE

In this section, we present a simple example to illustrate some important issues in the alignment process. We generate N two-dimensional points $X = [x_1, \ldots, x_N]$, where $x_i = [t_i \cos(t_i), t_i \sin(t_i)]^T$ with $t_1, \ldots, t_N$, $N = 100$, equally spaced in $[\pi/5, 2\pi]$. The exact arc length coordinates for each of the points can be computed as

$$\tau_i = \int_{t_0}^{t_i} \sqrt{1 + t^2} dt,$$

where t_0 is a constant chosen such that the mean of $\tau_1, \ldots, \tau_N$ is zero. First we choose 19 sections (subsets of points in X)

[*]The work of this author was supported in part by NSF grants DMS-0311800 and DMS-0405681.

[†]The work of this author was supported in part by NSFC (project 60372033), the Special Funds for Major State Basic Research Projects (project G19990328), and NSF grants CCF-0305879 and DMS-0405681.

[1]Actually, each local coordinate system can cover a much larger piece of the manifold, and we will call it a local patch.

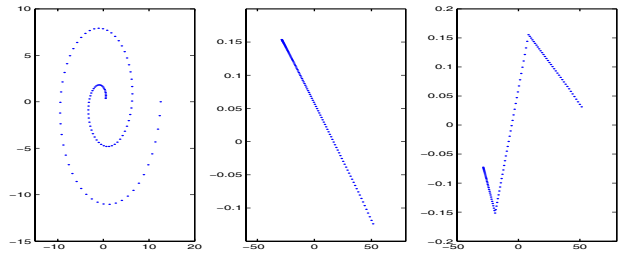

Fig. 1. Plots of the spiral data points and the basis vectors of Φ.

$X_i = X(:, I_i)$, $i = 1, \ldots, 19$, where the index sets I_i are defined by

$$I_i = (5(i-1)+1) : (5i+2), \quad i = 1, \ldots, 18, \quad I_{19} = 91 : 100.$$

It is easy to see that each pair of two consecutive sections share two points. We construct the alignment matrix Φ by computing the orthogonal projections Φ_i such that $\Phi_i[e, \Theta_i^T] = 0$, where Θ_i is the orthogonal projection of X_i on to the tangent space $\mathrm{span}(X_i - \bar{x}_i e^T)$, the column space of $X_i - \bar{x}_i e^T$, and $\bar{x}_i$ denotes the column mean of X_i (the details of this process for Φ is discussed in the next section). An orthonormal basis for the null space of Φ is given by $[e, \tau]$ with $\tau = (\tau_1, \ldots, \tau_N)^T$. In the left panel of Figure 1 we plot the given set of two-dimensional data points. In the middle panel, we plot the components of the computed τ_i's against the original τ_i's. The plotted points are approximately on a straight line, indicating an accurate recovery of the original τ_i's.

If the minimal number of the shared points is reduced to be one, the null space of Φ will have dimension more than two. The null space not only contains e and τ but also other unwanted vectors. For example, if we delete the last columns in the two sections X_6 and X_{13}, respectively, Φ then has three linearly independent eigenvectors corresponding to the zero eigenvalue which include e, τ and another vector. It is not clear how to separate out τ from the null space of Φ if other unwanted eigenvectors are mixed. In the right panel of Figure 1, we plot such an eigenvector against τ, showing that it is not proportional to τ anymore.

The above example illustrates the importance of the null space structure of the alignment matrix for manifold learning. The goal of this paper is to gain a better understanding of this structure. We now proceed more formally to our discussion.

3. ALIGNMENT MATRIX AND ITS NULL SPACE

Let $T = [\tau_1, \ldots, \tau_N] \in \mathcal{R}^{d \times N}$ correspond to the N d-dimensional parameter vectors given in (1.1). A submatrix of T consisting of a subset of the column vectors in T is called a section of T, each of them corresponding to the parameter vectors in a local patch of the underlying manifold. For a collection of sections $\{T_1, \ldots, T_s\}$, we denote $T_i = [\tau_{i_1}, \ldots, \tau_{i_{k_i}}] \in \mathcal{R}^{d \times k_i}$. Before we proceed, we need the following definition.

Definition 3.1. *Let $S_x = \{x_i\}_{i=1}^{\alpha}$ and $S_y = \{y_i\}_{i=1}^{\beta}$ be two sets of column vectors of the same dimension. Denote by $S_z = S_x \cap S_y = \{z_1, \ldots, z_\gamma\}$ the set of column vectors that are in the intersection of S_x and S_y. We say the two sets S_x and S_y are* fully overlapped *if $[z_1, \ldots, z_\gamma] - \bar{z}e^T$ is of full row-rank, where $\bar{z}$ is the mean of the vectors in S_z, i.e, the dimension of the affine space spanned by $[z_1, \ldots, z_\gamma]$ is the same as the dimension of the columns vectors.*

It is easy to verify that $[z_1, \ldots, z_\gamma] - \bar{z}e^T$ is of full row-rank if and only if $[e, [z_1, \ldots, z_\gamma]^T]$ is of full column-rank.

We now consider a collection of sections $\{T_1, \ldots, T_s\}$ of T such that

$$\cup_{i=1}^{s} \{\tau_{i_1}, \ldots, \tau_{i_{k_i}}\} = \{\tau_1, \ldots, \tau_N\},$$

and we are interested in a particular kind of collections of sections which is defined below.

Definition 3.2. *We say $\{T_1, \ldots, T_s\}$ is* connectedly overlapped*, if there is a permutation $\ell_1, \cdots, \ell_s$ of $1, \cdots, s$ such that $\{\cup_{i=1}^{k} T_{\ell_i}, T_{\ell_{k+1}}\}$ is fully overlapped for $k = 1, \cdots, s-1$. Specially, we say T_i itself is connectedly overlapped, if $[e, T_i^T]$ is of full column-rank.*

A subset $\{T_{i_1}, \ldots, T_{i_k}\}$ is said to be a maximally connectedly overlapped subset of $\{T_1, \ldots, T_s\}$, if $\{T_{i_1}, \ldots, T_{i_k}\}$ itself is connectedly overlapped and no other section can be added so that the resulting set is still connectedly overlapped.

For a section T_i, let P_i be the orthogonal projection onto the orthogonal complement $\mathrm{span}^{\perp}([e, T_i^T])$ of $\mathrm{span}([e, T_i^T])$, i.e., $P_i[e, T_i^T] = 0$. Embed P_i in an N-by-N matrix Φ_i such that the (i_p, i_q) element of Φ_i is the (p, q) element of P_i, and zero entries elsewhere. The i-th orthogonal projection Φ_i in the higher dimensional space $\mathcal{R}^N$ can be represented as

$$\Phi_i = S_i P_i S_i^T, \tag{3.2}$$

where $S_i \in \mathcal{R}^{N \times k_i}$ is the selection matrix corresponding to the index set $I_i = \{i_1, \ldots, i_{k_i}\}$ of T_i such that $T S_i = T_i$. For two index sets, I_i and I_j, we denote by $I_{ij} = I_i \cap I_j$ the intersection of I_i and I_j, the corresponding intersection of the two sections is denoted by $T_{ij} = T(:, I_{ij})$.

The so-called *alignment matrix* is the sum of the s orthogonal projection matrices $\{\Phi_i\}$, assuming there are s local patches. The alignment matrix Φ is defined as

$$\Phi = \sum_{i=1}^{s} \Phi_i = [S_1, \ldots, S_s] \, \mathrm{diag}(P_1, \ldots, P_s) [S_1, \ldots, S_s]^T. \tag{3.3}$$

We first characterize the null space of Φ when there are only two overlapping local patches, i.e., two sections. This will pave the way for analyzing the more general case.

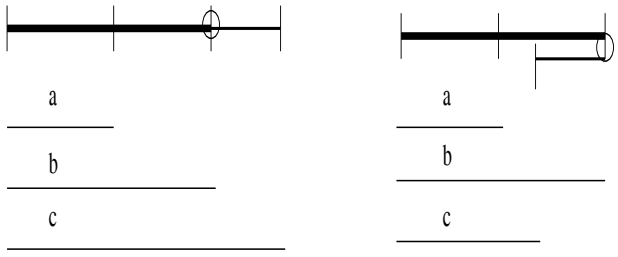

Fig. 2. Two possible layouts for the global coordinates.

Theorem 3.1. *Assume* $s = 2$, *and* $[e, T_1^T]$ *is of full column-rank. Then we have*

1) If $\mathcal{N}(\Phi) = \text{span}([e, T^T])$, *then*

$$\text{rank}([e, T_{12}^T]) = \text{rank}([e, T_2^T]).$$

2) If $[e, T_2^T]$ *is of full column-rank, then* $\{T_1, T_2\}$ *is fully overlapped if and only if* $\mathcal{N}(\Phi) = \text{span}([e, T^T])$.

EXAMPLE 1. We use a simple example with $d = 1$ to illustrate the case when the two sections are not fully overlapped. Consider the situation depicted in Figure 2. The first section consists of three points (denoted by short vertical bars) connected by a thick line, and the second one two points connected by a thin line. The two sections share a single point denoted by a circle, i.e., the thick line segment and the thin one corresponding to the two sections are connected at a single point. Within the first line segment, let the distance between the first and the second point be θ_1 and that between the second point and the third point be θ_2, and within the second line segment, let the distance between the first and the second point be θ_3. In this case, it can be shown $\dim(\mathcal{N}(\Phi)) = 3$ and the three basis vectors can be chosen as $[1, 1, 1, 1]^T$, $[\eta, \eta + a, \eta + b, \eta + c]^T$ with $a = \theta_1, b = \theta_1 + \theta_2, c = \theta_1 + \theta_2 + \theta_3$ and η an arbitrary real number (this corresponds to the left panel of Figure 2), $[\eta, \eta + a, \eta + b, \eta + c]^T$ with $a = \theta_1, b = \theta_1 + \theta_2, c = \theta_1 + \theta_2 - \theta_3$ and η an arbitrary real number (this corresponds to the right panel of Figure 2 where the second line segment folds back to the first one).

Now we present results for the general $s > 2$ case. We first present some necessary conditions.

Theorem 3.2. *Let* $[e, T_i^T]$ *be of full column-rank for* $i = 1, \ldots, s$. *If*

$$\mathcal{N}\{\Phi\} = \text{span}\{[e, T^T]\},$$

then

1) $\{T_1, \ldots, T_s\}$ *is connectedly overlapped, or*

2) $\{T_1, \ldots, T_s\}$ *is not connectedly overlapped, and for any maximally connectedly overlapped subset* $\{T_{i_1}, \ldots, T_{i_k}\}$, $\{[T_{i_1}, \ldots, T_{i_k}], [T_{i_{k+1}}, \ldots, T_{i_s}]\}$ *is fully overlapped.*

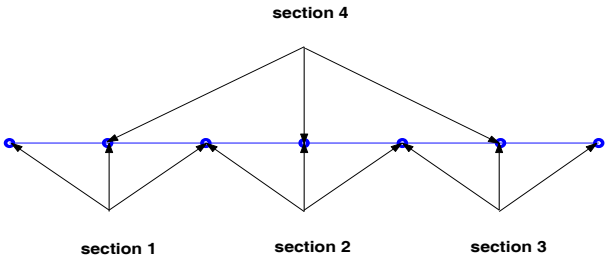

Fig. 3. Overlapping patterns of four sections.

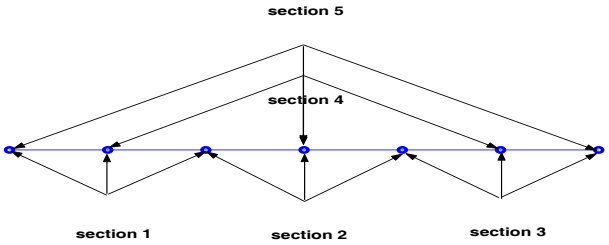

Fig. 4. Adding one more section.

EXAMPLE 2. We show that the condition 2) in Theorem 3.2 is not sufficient. In Figure 3 we plot an example of four sections with no pair of two sections fully overlapped. Obviously, each section T_i itself is a maximal connectedly overlapped subset and the union of its complement sections is fully overlapped with T_i. So the condition 2) in Theorem 3.2 is satisfied. However, $\mathcal{N}(\Phi) \neq \text{span}([e, T^T])$.

To give sufficient conditions, it is required to impose more constraints on a maximal connectedly overlapped subset and/or it complement subset. Below we give such sufficient conditions imposing a constraint of connected overlap on the complement subset.

Theorem 3.3. *Let* Φ *be defined as before. If*

1) $\{T_1, \ldots, T_s\}$ *is connectedly overlapped, or*

2) $\{T_1, \ldots, T_s\}$ *is not connectedly overlapped, but there is a maximally connectedly overlapped subset* $\{T_{i_1}, \ldots, T_{i_k}\}$ *such that* $\{T_{i_{k+1}}, \ldots, T_{i_s}\}$ *is also a connectedly overlapped subset and* $\{[T_{i_1}, \ldots, T_{i_k}], [T_{i_{k+1}}, \ldots, T_{i_s}]\}$ *is fully overlapped, then*

$$\mathcal{N}\{\Phi\} = \text{span}\{[e, T^T]\}.$$

EXAMPLE 3. If we add section 5 in Example 2, see Figure 4, then for the resulting collection of the five sections, condition 2) in Theorem 3.2 is also true but the conditions of Theorem 3.3 are not satisfied. However, we still have the null space of Φ that is spanned by $[e, T^T]$.

4. THE PRACTICAL CASE

The construction of the alignment matrix in the Section 3 assumes that we know the global coordinates τ_i's for each section $\hat{T}_i$. In practice, this information is not available. Interestingly, the above theorem is still true if for each T_i we know $\theta_1^{(i)}, \ldots, \theta_k^{(i)}$ which differs from $\Theta_i = [\tau_1^{(i)}, \ldots, \tau_k^{(i)}]^T$ by an affine transformation i.e., there are vector c_i and non-singular matrix G such that $T_i = c_i e^T + G\Theta_i$. We have

$$T_i^T = [e, \Theta_i^T] \begin{bmatrix} c_i^T \\ G^T \end{bmatrix}, \quad \Theta_i^T = [e, T_i^T] \begin{bmatrix} -(G^{-1}c_i)^T \\ G^{-T} \end{bmatrix}.$$

It follows that

$$\mathrm{span}([e,\, T_i^T]) = \mathrm{span}([e,\, \Theta_i^T])$$

Therefore the following theorem is true.

Theorem 4.1. *Let $T_i, i = 1, \ldots, s$, be s sections of T, $\Theta_i = [\theta_1^i, \ldots, \theta_{k_i}^i]$ be an coordinates that are affine to T_i, and P_i be the orthogonal projection onto the orthogonal complement of $[e, \Theta_i^T]$. Then $\mathrm{span}\{[e, T^T]\} \subset \mathcal{N}\{\Phi\}$ with Φ defined as before by P_i here. Furthermore, if $\{\Theta_i\}$ is connectedly overlapped, then $\mathcal{N}\{\Phi\} = \mathrm{span}\{[e, T^T]\}$.*

5. THE SMALLEST NONZERO EIGENVALUE OF THE ALIGNMENT MATRIX

How well the null space of Φ is determined depends on the size of its smallest nonzero eigenvalue. This has significant ramifications when we need to use the small eigenvectors of Φ to recover the τ_i's. Most of the results in this section will be for the special case of $s = 2$ although generalization to the more general case can be done as well. We first give a characterization of the smallest nonzero eigenvalue of Φ.

Theorem 5.1. *Let $P_i = Q_i Q_i^T$ be the orthogonal projections onto the orthogonal complements of $[e, T_i^T]$, and $\tilde{Q}_i$ the embedded one of Q_i. Denote by $G = (G_{i,j})$ with $G_{ii} = 0$ and $G_{ij} = \tilde{Q}_i(I_{ij}, :)^T \tilde{Q}_j(I_{ij}, :)$, $i \neq j$. Then the smallest nonzero eigenvalue of Φ is given by*

$$\lambda = 1 - \max\{\, \lambda :\ \lambda \in \lambda(G),\ \lambda < 1 \,\}.$$

Specially, if $s = 2$, then

$$\lambda = 1 - \max\{\, \sigma :\ \sigma \in \sigma(G_{12}),\ \sigma < 1 \,\}.$$

With the help of the above theorem, we now presents some quantitative estimates of the smallest nonzero eigenvalue of Φ for the case of $s = 2$ assuming $T = [U_1, V, U_2]$ and $T_1 = [U_1, V]$ and $T_2 = [V, U_2]$, in particular we will show that the smallest nonzero eigenvalue of Φ is proportional to $(\sigma_d(V - \bar{v}e^T))^2$ where $\sigma_d(\cdot)$ is the d-th largest singular value of a matrix. In a sense, $\sigma_d(V - \bar{v}e^T)$ measures the *strength* of the overlap between T_1 and T_2. We will consider the case of $d = 1$.

Theorem 5.2. *Let $T = [u_1^T, v^T, u_2^T]$ with $u_i \in \mathcal{R}^{k_i}$, $i = 1, 2$, $v \in \mathcal{R}^\ell$, u_1 and u_2 do not share any components. Assume $T_1 = [u_1^T, v^T]$ and $T_2 = [v^T, u_2^T]$, ie., sections T_1 and T_2 share ℓ points which are the components of v. Denote by $\bar{v}$ the mean of v. Then the smallest nonzero eigenvalue of Φ is $O(\|v - \bar{v}e\|_2^2)$.*

For Theorem 5.2, we conjecture that similar results should hold for $d > 1$ case by replacing $\|v - \bar{v}e\|_2$ with $\sigma_d(V - \bar{v}e^T)$. However, for $s > 2$ case, it is still not clear how to formulate the concept of the strength of overlaps in this general case.

6. CONCLUDING REMARKS

The spectral properties of the alignment matrix plays an essential role in using local methods for manifold learning. The results proved in this paper represent the first step towards a better understanding of those spectral properties and their interplay with the geometric properties of the underlying manifold. There are still several issues which deserve further investigation including 1) deriving a set of conditions which are both necessary and sufficient for the null space of the alignment matrix to recover the global coordinates; 2) improving the quantitative results under more general conditions; and 3) incorporating the effects of noise.

7. REFERENCES

[1] M. Brand. Charting a manifold. *Advances in Neural Information Processing Systems* 15, MIT Press, 2003.

[2] D. Donoho and C. Grimes. Hessian Eigenmaps: new tools for nonlinear dimensionality reduction. *Proceedings of National Academy of Science*, 5591-5596, 2003.

[3] S. Roweis and L. Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290: 2323–2326, 2000.

[4] Y. Teh and S. Roweis Automatic Alignment of Hidden Representations. *Advances in Neural Information Processing Systems*, 15, MIT Press, 2003.

[5] J. Tenenbaum, V De Silva and J Langford. A global geometric framework for nonlinear dimension reduction. *Science*, 290:2319–2323, 2000

[6] Z. Zhang and H. Zha. Principal Manifolds and Nonlinear Dimensionality Reduction via Tangent Space Alignment. *SIAM J. Scientific Computing,* 26:313–338, 2004.