Geometry of Synchronisation Problems

Jacek Brodzki
University of Southampton

INRIA Sophia Antipolis
October 2016

Joint work with Sayan Mukherjee and Tingran Gao, Duke University

Synchronisation, informally

Synchronisation problems deal with questions regarding aligning a collection of objects in a consistent manner.

Given a sequence of objects $(o_1, ..., o_n)$ and a group (or collection of groups) G the objective is to learn a collection of group elements ρ_{ij} that would transform object o_i to o_j.

If it is possible to find a sequence of group elements that allow for an accurate transformation between objects then this set can be synchronised.

Synchronisation, informally

The goal of our procedure is to adjust the pairwise correspondences (often in the presence of noise or incomplete measurements) to obtain a globally consistent arrangement of transformations.

Previous work in this direction assumed that data are sampled from an underlying manifold, but we show that this is not necessary.

Our procedure is based more on a kind of rigidity than smoothness.

Rotational consistency

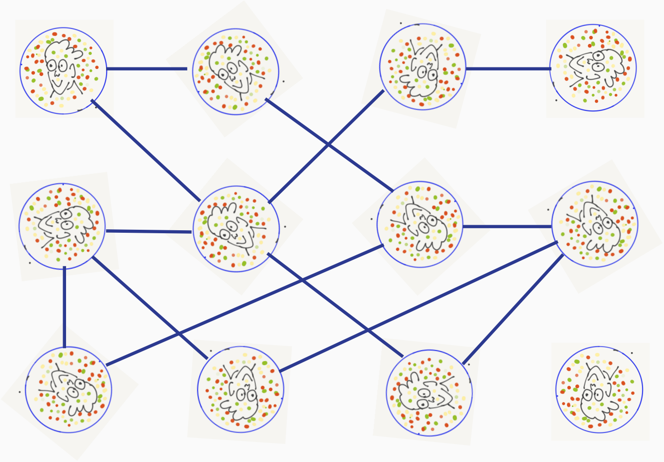

Cryo-Electron Microscopy

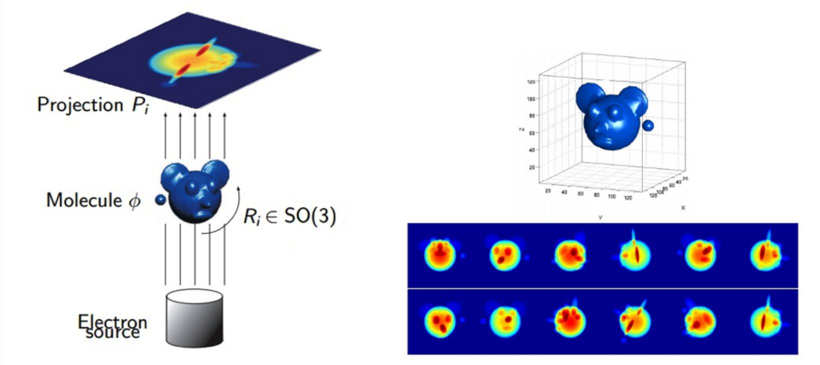

Picture from Singer, Shkolnisky via Bandeira *Ten lectures and forty two problems ...*

Shape space

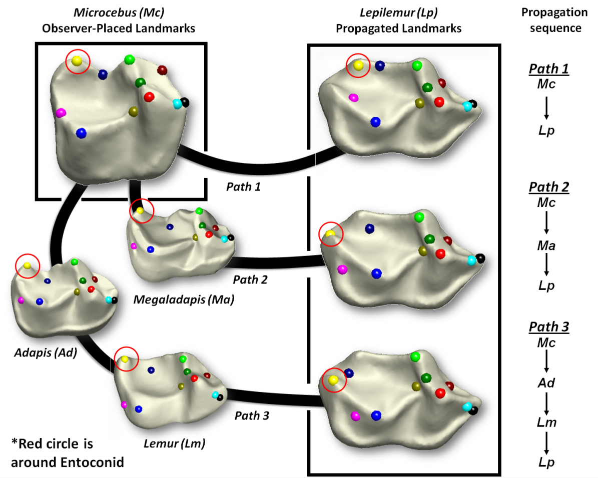

Microcebus (Mc)
Observer-Placed Landmarks
Lepilemur (Lp)
Propagated Landmarks
Propagation sequence
Path 1
Mc
Lp
Path 2
Mc
Ma
Lp
Path 3
Mc
Ad
Lm
Lp
Path 1
Path 2
Megaladapis (Ma)
Adapis (Ad)
Path 3
Lemur (Lm)
*Red circle is around Entoconid

Synchronisation problems

- Group actions
- Bundle representation
- Twisted de Rham complex

Vertex and edge potentials

Let $\mathcal{G} = (V, E)$ be a graph with vertex set V and edge set E. Let G be a group.

- A *vertex potential* is a map $f : V \to G$.
- An *edge potential* is a function $\rho : E \to G$ which is symmetric: for every edge $(i, j) \in E$ we have $\rho(j, i) = \rho(i, j)^{-1}$

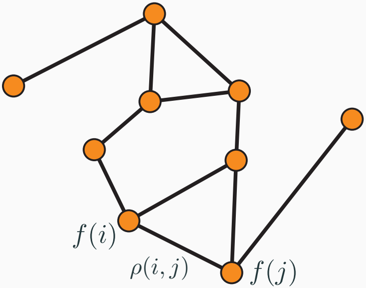

Compatible potentials

- We say that a group potential f and an edge potential ρ are *compatible across an edge* (i, j) iff

$$f(i) = \rho(i, j) f(j).$$

- We say that the two potentials ν and ρ are *compatible* iff they are compatible across every edge $(i, j) \in E$.
- This definition also makes sense if f takes values in a G-module M.

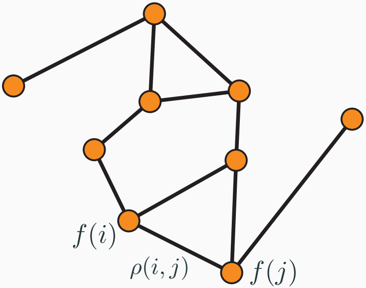

Synchronisation problem

Bandeira, Singer and Spielman (2013) posed the following question:

Given an edge potential $\rho : E \to G$, does there exits a vertex potential $f : V \to G$ which is compatible with ρ?

Note that in this formulation the converse problem is easy. Given a vertex potential $f : V \to G$, one can define a compatible edge potential ρ by

$$\rho(i,j) = f(i)f(j)^{-1}.$$

Paths

Let v and w be vertices in the graph G. A *path* γ connecting the vertices v and w is a sequence of edges

$$\gamma = (e_1, e_2, \ldots, e_n)$$

where for $k = 1, \ldots, n$, $e_k = (i_k, j_k)$, $v = i_1$, $j_k = i_{k+1}$ and $w = j_n$.

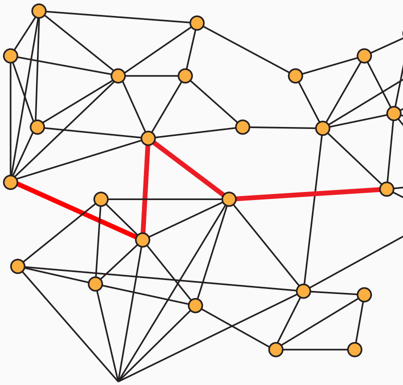

An edge potential ρ gives rise to a map

$$\{\text{Paths in } \mathcal{G}\} \longrightarrow G, \qquad R : \gamma \mapsto \rho(e_1) \dots \rho(e_n).$$

Reversal property For $e = (i, j)$, let $e^{-1} = (j, i)$, and

$$\gamma^{-1} = (e_n^{-1}, \dots, e_1^{-1}).$$

Because the edge potential ρ is symmetric, we have that

$$R(\gamma^{-1}) = \rho(e_n^{-1}) \dots \rho(e_1^{-1}) = \rho(e_n)^{-1} \dots \rho(e_1)^{-1} = R(\gamma)^{-1}.$$

Concatenation

$$R(\gamma_1 \circ \gamma_2) = R(\gamma_1) R(\gamma_2).$$

Holonomy groups

Let ω is a based loop, so that $v = w$ and the vertex v is a selected base point on the loop ω.

Let Ω_v denote the space of loops based at v together with the trivial loop $0 = (v, v)$. We assume that $R(0) = 1$.

Then we have a natural map from the space Ω_v of loops based at v to G defined by

$$\omega \mapsto R(\omega)$$

Lemma

For every vertex v, $H(v) = R(\Omega_v)$ is a subgroup of G, called the holonomy group at v. If the graph $\mathcal{G}$ is connected, for every two vertices v and w, the holonomy groups Ω_v and Ω_w are isomorphic.

Restriction to paths

For a path γ connecting a vertex v to a vertex w, let $V(\gamma)$ be the set of vertices along γ, and let $E(\gamma)$ denote the set of edges comprising γ.

If $f : V \to G$, then f_γ is the restriction of f to $V(\gamma)$. Similarly, ρ_γ is the restriction of ρ to $E(\gamma)$.

We say that f_γ is compatible with ρ_γ along γ if and only if for every edge $e \in E(\gamma)$, $e = (i, j)$, $f_\gamma(i) = \rho(i, j) f_\gamma(j)$.

It follows that

$$f_\gamma(v) = R(\gamma) f_\gamma(w).$$

Synchronisation along loops

Lemma

Let $R(\omega) = 1$ for a based loop ω. Then there exists an edge potential $f_\omega : V_\omega \to G$ along ω which is compatible with the edge potential ρ_ω along ω.

Proof.

Let $\omega = (e_1, \ldots, e_n)$, where $i_1 = j_n = v$. We put $f_\omega(v) = 1$, and $f_\omega(j_k) = f(i_{k+1}) = \rho(e_1) \ldots \rho(e_k)$.

Then for every edge $e_k = (i_k, j_k)$, $f_\omega(i_k) = \rho(i_k, j_k) f(j_k)$.

When $k = n$,

$$f_\omega(j_n) = \rho(e_1) \ldots \rho(e_n) = R(\omega) = 1$$

so

$$f_\omega(j_n) = f_\omega(i_1) = 1$$

Theorem

Let $\mathcal{G}(V, E)$ be a connected graph and let ρ be an edge potential on $\mathcal{G}$. Then there exists a vertex potential $f : V \to G$ compatible with ρ if and only if the holonomy group $H(v)$ of some vertex v is trivial. It follows that the holonomy group of every other vertex is also trivial, as is the holonomy of every loop in $\mathcal{G}$.

Sketch of proof

Assume that the holonomy of every loop is trivial. Choose a base vertex v and put $f(v) = 1$. Then for every vertex w define

$$f(w) = R(\gamma)$$

where γ is a path connecting v to w. If γ_1 is a different path connecting v to w, then $\omega = \gamma \circ \gamma_1^{-1}$ is a loop based at v, so that

$$1 = R(\omega) = R(\gamma)R(\gamma_1)^{-1}$$

which shows that our definition does not depend on the choice of γ. The group potential f is compatible with ρ by construction.

Corollary

Let $\rho \in P(E, G)$ be an edge potential for which the synchronization problem has a solution, that is, there exists a vertex potential $f \in P(V, G)$ compatible with ρ. Then ρ satisfies the following condition. For every triangle (v_0, v_1, v_2) in $\mathcal{G}$

$$\rho(v_0, v_1)\rho(v_1, v_2)\rho(v_2, v_0) = 1.$$

Cycle consistency

In summary, the triviality of holonomy groups for a symmetric edge potential implies the following well-known cycle consistency conditions:

$$\begin{aligned}
\rho(i,i) &= 1 \quad \text{for all } i \in V, \\
\rho(i,j)\rho(j,i) &= 1 \quad \text{for all } (i,j) \in E, \\
\rho(i,j)\rho(j,k)\rho(k,i) &= 1 \quad \text{if } (i,j),(j,k),(k,i) \in E.
\end{aligned}$$

Eurographics Symposium on Geometry Processing 2013
Yaron Lipman and Hao Zhang
(Guest Editors)

Volume 32 (2013), Number 5

Consistent Shape Maps via Semidefinite Programming

Qi-Xing Huang and Leonidas Guibas

Computer Science Department, Stanford University, Stanford, CA

Definition 2.1 Given a shape collection $\mathcal{S} = \{S_1, \cdots, S_n\}$ of n shapes where each shape consists of the same number of samples, we say a map collection $\Phi = \{\phi_{ij} : S_i \rightarrow S_j | 1 \leq i, j \leq n\}$ of maps between all pairs of shapes is cycle consistent if and only if the following equalities are satisfied:

$$\begin{aligned} \phi_{ii} &= id_{S_i}, \quad 1 \leq i \leq n, && \text{(1-cycle)} \\ \phi_{ji} \circ \phi_{ij} &= id_{S_i}, \quad 1 \leq i < j \leq n, && \text{(2-cycle)} \\ \phi_{ki} \circ \phi_{jk} \circ \phi_{ij} &= id_{S_i}, \quad 1 \leq i < j < k \leq n, && \text{(3-cycle)} \end{aligned} \tag{1}$$

where id_{S_i} denotes the identity self-map on S_i.

Fibre Bundles and Consistency Conditions

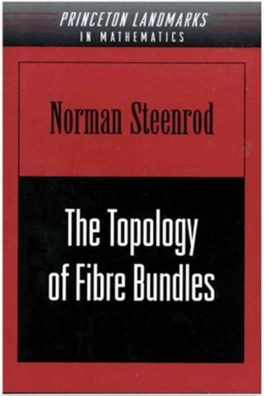

Bundle Existence Theorem. If G is a topological group acting on Y, and $\{U_i\}$, $\{g_{ij}\}$ is a system of coordinate transformations in the space X satisfying the cycle-consistency conditions, then there exists a fibre bundle $\mathcal{B}$ with base space X, fibre Y, group G, and coordinate transforms $\{\rho_{ij}\}$.

In other words we have constructed a bundle with fibre G over each vertex of the graph. The graph can be regarded as the nerve of the cover in the above theorem.

Bundle representation

Theorem

Let G be a topological group, $\Gamma = (V, E)$ a connected undirected graph, and $\rho : E \to G$ a map satisfying $\rho_{ij} = \rho_{ji}^{-1}$ for all $(i, j) \in E$. Write $\mathfrak{U} = \{U_i \mid 1 \leq i \leq |V|\}$ for an open cover of Γ in which U_i is the union of the singe vertex set $\{i\}$ with the interior of all edges adjacent to the vertex i. Then ρ defines a flat principal G-bundle $\mathcal{B}_\rho$ over Γ with a system of local trivializations defined on the open sets in $\mathfrak{U}$ with constant bundle transition functions ρ_{ij} on non-empty $U_i \cap U_j$. Furthermore, ρ is synchronizable if and only if $\mathcal{B}_\rho$ is trivial.

Synchronisation bundle

The fibre bundle $\mathcal{B}_\rho$ associated with the connected graph Γ and edge potential ρ is called a *synchronization bundle associated with the edge potential* ρ *over* Γ.

Obstruction to synchronisation

The compatibility condition on edges is

$$f(i) = \rho(i,j)f(j)$$

To construct a useful way to measure the failure to satisfy this condition, we now consider a G-module M (equipped with an inner product), and M-valued functions $f : V \to M$ on the vertex set V.

We want to interpret the difference

$$f(i) - \rho(i,j)f(j).$$

Frustration

Given a graph $\Gamma = (V, E)$ and a G-valued edge potential ρ, the success of finding an M-valued vertex potential f compatible with ρ is measured by

$$\min_{f:V\to F} \sum_{(i,j)\in E} \mathrm{Cost}_F\left(\rho_{ij} f_j, f_i\right), \tag{1}$$

where $\mathrm{Cost}_F : F \times F \to [0, \infty)$ is a cost function on M (e.g. derived from a distance or a norm).

Optimisation

The case $M = G$ corresponds to the optimisation problem

$$\min_{f:V\to G} \sum_{(i,j)\in E} \mathrm{Cost}_G\left(\rho_{ij} f_j, f_i\right). \tag{2}$$

which, in the case the Cost_G is G-invariant, is equivalent to

$$\min_{f:V\to G} \sum_{(i,j)\in E} \mathrm{Cost}_G\left(\rho_{ij}, f_i f_j^{-1}\right). \tag{3}$$

If ρ is synchronizable, a minimizer of (2) (or (1)) can be geometrically realized as a global section of the synchronization bundle $\mathcal{B}_\rho$; such a minimizer implies the triviality of the principal bundle $\mathcal{B}_\rho$

Examples

The following examples have been studied so far:

- $G = M = O(d)$, and $G = O(d), M = S^{d-1}$, Bandeira, Singer, Spielman 2013;
- Synchronisation of rotations, $G = M = SO(d)$, Boumal, Singer, Absil, Blondel 2014;
- Orientation detection $G = M = O(1)$, Singer, Wu 2011;
- Cryo-electron microscopy $G = M = SO(2)$, Singer, Wu 2012;
- Global alignment of three-dimensional scans $G = M = SO(3)$, Tzveneva, Singer, Rusinkiewicz 2011.

The standard graph Laplacian

Given a graph $\mathcal{G}(V, E)$ we define the gradient $d : \mathbb{C}[V] \to \mathbb{C}[E]$ by the formula

$$df(i,j) = f(i) - f(j).$$

The divergence operator $\delta : \mathbb{C}[E] \to \mathbb{C}[V]$ is given by

$$\delta\phi(i) = \sum_{j \sim i} \phi(i,j).$$

The graph Laplacian is the operator $\Delta : \mathbb{C}[V] \to \mathbb{C}[V]$ defined by $\Delta_\rho = \delta d$:

$$\Delta f(i) = \sum_{j : j \sim i} (f(i) - f(j))$$

Twisted differential

M is a G-module equipped with an inner product, and ρ a symmetric potential as before.

$$d_\rho : \mathbb{C}(V, M) \longrightarrow \mathbb{C}(E, M)$$

$$d_\rho f(i,j) = f(i) - \rho(i,j) f(j)$$

for every $f \in \mathbb{C}(V, M)$.

A solution to the synchronisation problem is an element of the kernel of d_ρ.

Twisted Laplacian

The usual divergence:

$$\delta : \mathbb{C}(E, M) \longrightarrow \mathbb{C}(V, M)$$

$$\delta\phi(i) = \sum_{j:j\sim i} \phi(i, j)$$

together with d_ρ give a twisted Laplacian

$$\Delta_\rho = \delta d_\rho = \sum_{j:j\sim i} (f(i) - \rho(i, j)f(j))$$

The Raleigh quotient

$$\eta_\rho(f) = \frac{\langle f, \Delta_\rho f \rangle}{\langle f, f \rangle} = \frac{\sum_{(i,j)\in E} w_{ij} \left\| f(i) - \rho(i,j) f(j) \right\|^2}{\sum_{i \in V} d_i \left\| f(i) \right\|^2}$$

$\eta_\rho(f)$ is defined in Bandeira et al. (2013) as the *frustration*

Synchronisation and cohomology

We have defined a simple de Rham-type complex

$$\mathcal{C} : 0 \longrightarrow \mathbb{C}(V, M) \xrightarrow{d_\rho} \mathbb{C}_\rho(E, M) \longrightarrow 0.$$

The cohomology of this complex is

$$H^0(\mathcal{C}, d) = \ker d_\rho$$

in degree zero. In degree one, we have that

$$H^1(\mathcal{C}, d) = \mathbb{C}(E, M)/d_\rho(\mathbb{C}(V, M)).$$

Thus solutions to the synchronisation problem form the zeroth cohomology group, which is the same as the kernel of d_ρ.

Theorem

1. *The space of solutions of the synchronisation problem given by a unitary edge potential ρ is isomorphic to the space of harmonic functions f, i.e., functions f with the property $\Delta_\rho f = 0$. Moreover, we have the following orthogonal decomposition:*

$$\mathbb{C}(V, M) = \ker d_\rho \oplus \mathrm{Im}\ \delta.$$

2. *The Laplace operator Δ_ρ is self-adjoint and positive.*

Learning group actions

Given a group G acting on a set X, simultaneously learn a new action of G on X and a partition of X into disjoint subsets $X_1, \cdots, X_K$, such that the new action is as close as possible to the given action and cycle-consistent on each X_i $(1 \leq i \leq K)$.

Example: representations

If the set X is a vector space and we seek a direct sum decomposition $X = \bigoplus_{i=1}^{K} X_i$, the LGA problem reduces to the search for all irreducible G-subrepresentations of X.

Example: spins

$X = \{x_1, \cdots, x_n\}$ equipped with $S : X \to \{\pm 1\}$ that assigns to each x_i either value $+1$ or -1.

Let $G = \{\pm 1\}$ act on X transitively as $(g_{ji}, x_i) \mapsto x_j$, $g_{ji} = S(x_j)S(x_i)$.

Suppose the spin of each point in X (i.e. the label map S) is unknown, but we have full access to the group actions $\{g_{ij}\}$, we can reconstruct S — up to flipping labels ± 1 — by spectral clustering the dataset X, viewed as vertices of a complete graph Γ with weight $w_{ij} = g_{ij}$ on the edge connecting x_i and x_j.

Algorithm

Algorithm 1 SYNCHRONIZATION CUT: Learning Group Actions by Synchronization

procedure SYNCUT(Γ, ρ, K) ▷ weighted graph $\Gamma = (V, E, w)$, $\rho \in C^1(\Gamma; G)$, number of partitions K
  $t = 0$
  $\epsilon = w$
  while not converge **do**
    $f^{(t)} \in C^0(\Gamma; G) \leftarrow$ SYNCHRONIZE(Γ, ρ, ϵ)
    for $(i, j) \in E$ **do** ▷ calculate weights ϵ on graph Γ for spectral clustering
      $\epsilon_{ij} \leftarrow w_{ij} \exp\left(-\left\| f_i^{(t)} - \rho_{ij} f_j^{(t)} \right\|_{\mathrm{F}}^2\right)$
    end for
    $\{S_1, \cdots, S_K\} \leftarrow$ SPECTRALCLUSTERING(Γ, ϵ)
    for $\ell = 1, \cdots, K$ **do**
      $g^{(\ell)} \in \Omega^0(S_\ell; G) \leftarrow$ SYNCHRONIZE($S_\ell, \rho|_{S_\ell}, \epsilon|_{S_\ell}$)
    end for
    $f^* \in \Omega^0(\Gamma; G) \leftarrow$ COLLAGE$\left(\{S_\ell\}_{\ell=1}^K, \left\{g^{(\ell)}\right\}_{\ell=1}^K\right)$
    for $(i, j) \in E$ **do** ▷ update weights $w^{(t)}$ on graph Γ for next iteration
      $\epsilon_{ij} \leftarrow w_{ij} \exp\left(-\left\| f_i^* - \rho_{ij} f_j^* \right\|_{\mathrm{F}}^2\right)$
    end for
    $\{S_1, \cdots, S_K\} \leftarrow$ SPECTRALCLUSTERING(Γ, ϵ)
    $t \leftarrow t + 1$
  end while
  return $\{S_1, \cdots, S_K\}$, f^* ▷ f^* defines a cycle-consistent edge potential on each partition
end procedure

Synthetic data

In a graph, create two connected components S_1 and S_2 of equal size.

(1) Randomly generate a vertex potential $g \in C^0(\Gamma; G)$ for the entire graph Γ;

(2) Set the value of ρ on edge (i, j) according to

$$\rho_{ij} = \begin{cases} g_i g_j^{-1} & \text{if both } i, j \in S_1 \text{ or } i, j \in S_2, \\ \text{a random matrix in } O(d) & \text{otherwise.} \end{cases}$$

Synthetic data

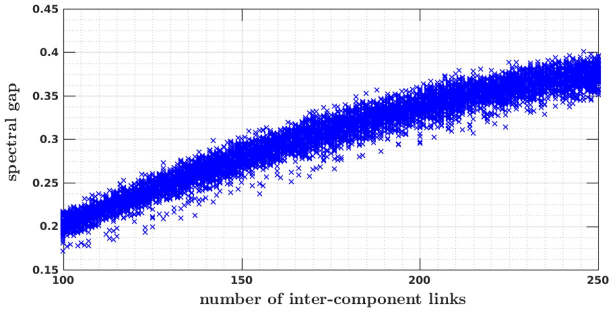

A scatter plot of the correlation between the number of inter-component links and the spectral gap in our random graph model, with $N = 100$ vertices and the (integer) number of inter-component links uniformly distributed between 100 and 250.

Simulation

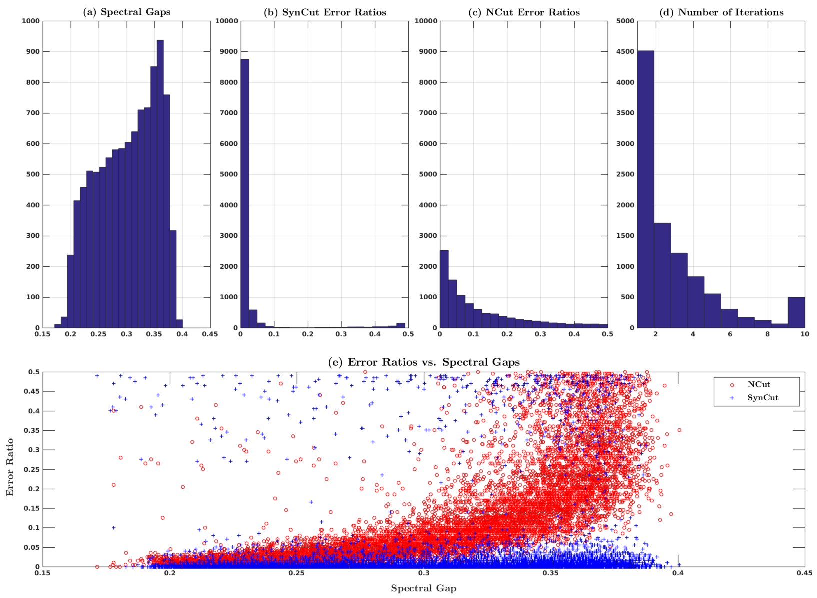
(a) Spectral Gaps
(b) SynCut Error Ratios
(c) NCut Error Ratios
(d) Number of Iterations
(e) Error Ratios vs. Spectral Gaps
NCut
SynCut
Error Ratio
Spectral Gap

Spin point cloud

- A simple example: a set of particles with 'spin' pointing up or down
- First approximation: similarity between points is zero if spins point in the same direction, 1 otherwise.
- Spectral clustering using the associated Laplacian reveals the original spin clusters.

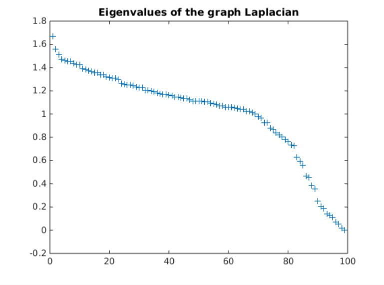

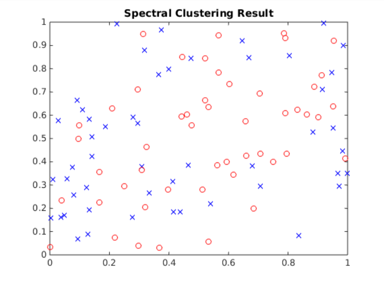

Lemurs

- We use a dataset that contains 116 second mandibular molars of simian and prosimian mammals, each discretized into a triangular mesh with 5,000 vertices and 10,000 faces.
- Molars in this dataset come with landmarks specified manually by evolutionary anthropologists at Duke University.
- Each tooth comes with 16 such landmarks, of which the order indicates correspondence across the entire dataset.

Lemurs

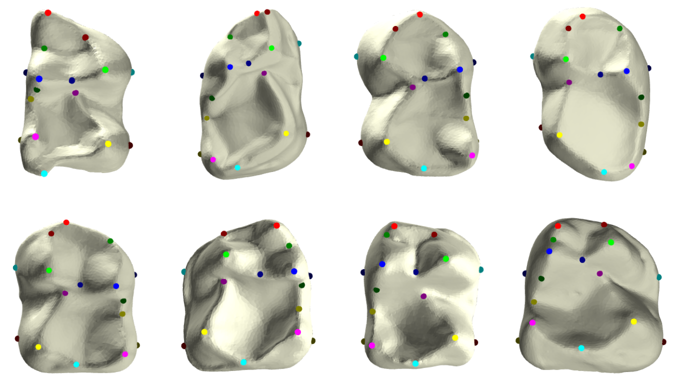

Geometric clustering

- It is clear that there are left molars and right molars; as indicated by the orientation of the landmarks. Our first step is to the teeth into two distinct (left/right) groups based on pairwise comparisons.
- Extending the idea of Procrustes similarity, for each pair of teeth we compute an orthogonal transformation in $O(3)$ that best aligns this pair;
- The determinant of each transformation is a measure of similarity. This corresponds to synchronisation with $G = \mathbb{Z}_2$ as in the spin example.
- Using other groups, like the Euclidean group $E(3)$ reveals more information.

Lemurs

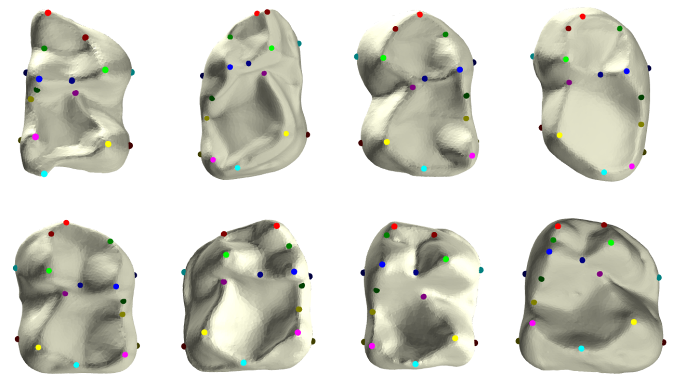

As expected, teeth 1, 2, 7, 8 belong to one cluster, while 3, 4, 5, 6 belong to the other cluster.

Higher dimensions: Oriented simplicial complexes

Let K be a simplicial complex of dimension n. We denote by K^p the set of p-simplices in K. Let G be a group and let M be a G-module equipped with an inner product $\langle -, - \rangle$. We assume that the simplices in K are oriented, and an orientation of a p-simplex S in K is represented by an ordering of its vertices:

$$[v_0, v_1, \ldots, v_p].$$

Two presentations are equivalent if they differ by an even permutation of the vertices of K.

Forms

A p-form ω on K with values in M is a map $\omega : K^p \to M$ which is skew-symmetric in the sense that $\omega(-S) = -\omega(S)$, where $-S$ is the simplex S with opposite orientation.

Alternatively, ω can be viewed as a skew-symmetric M-valued function $\omega : V^{p+1} \to M$.

Given an edge potential ρ we construct a differential complex

$$\Omega^0(K, M) \xrightarrow{d_\rho} \Omega^1(K, M) \xrightarrow{d_\rho} \cdots \xrightarrow{d_\rho} \Omega^n(K, M)$$

equipped with a twisted differential d_ρ associated with the edge potential ρ.

First note that there is a very natural operation $\rho\wedge : \Omega^p(K, M) \to \Omega^{p+1}(K, M)$ defined by

$$(\rho \wedge \omega)[v_0, v_1, \dots, v_{p+2}] \\ = \sum_{\sigma\in\Sigma_{p+2}} \mathrm{sgn}\,(\sigma)\rho(v_{\sigma(0)}, v_{\sigma(1)})\omega[v_{\sigma(1)}, \dots, v_{\sigma(p+2)}]$$

Definition

Let K be a simplicial complex of dimention n, and let $\rho : K^1 \to G$ be a symmetric edge potential. Then the twisted differential d_ρ is a map

$$d_\rho : \Omega^p(K, M) \to \Omega^{p+1}(K, M)$$

defined by

$$d_\rho(\omega) = d\omega + \rho \wedge \omega.$$

Here d is the usual simplicial differential defined for every $\omega \in \Omega^p(K, M)$ by

$$d\omega[v_0, \ldots, v_{p+2}] = \sum_{i=1}^{p+1} (-1)^i \omega[v_0, \ldots, \widehat{v_i}, \ldots, v_p].$$

Theorem

The operator $d_\rho : \Omega^p(K, M) \to \Omega^{p+1}(K, M)$ is a differential, that is, $d_\rho^2 = 0$, if and only if the edge potential ρ satisfies the equation

$$d(\rho \wedge \omega) + \rho \wedge d\omega + \rho \wedge \rho \wedge \omega = 0$$

for every p-form ω.

Remark

Using the Leibniz rule purely formally we can write

$$d(\rho \wedge \omega) = d\rho \wedge \omega - \rho \wedge d\omega.$$

Hence the expression $d(\rho \wedge \omega) + \rho \wedge d\omega$ plays the role of the term $d\rho \wedge \omega$.

Using this formal analogy, the condition derived in the theorem can be written as

$$d\rho + \rho \wedge \rho = 0$$

which is a familiar condition defining a flat connection in the theory of vector bundles.

Higher twisted differential

Theorem

Let K be a simplicial complex, and G a group. Assume that ρ is a G-valued edge potential with the property that $\rho[v_0, v_1] = \rho[v_1, v_0]^{-1}$ for every edge $[v_0, v_1]$ in K. Let us assume that and that the synchronization problem has a solution for ρ. So in particular, ρ satisfies the cocycle condition:

$$\rho(i,j)\rho(j,k) = \rho(i,k).$$

for every three edges $[i,j]$, $[j,k]$, and $[i,k]$ that form a triangle in K. Then $d_\rho^2 = 0$.