Analysis of 3D Shapes (IN2238)

Frank R. Schmidt
Matthias Vestner

Summer Semester 2017

11. Euclidean Embeddings

Multidimensional Scaling

Shape matching

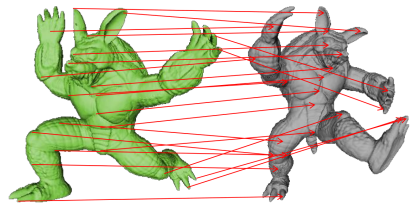

Our goal is to assign each point on the source shape a corresponding point on the target shape. Although a diffeomorphic (bijective and diff'able in both directions) mapping is desired, most of the approaches we discuss will not even guarantee injective mappings (remember nearest neighbors from ICP). Eventually we deal with discretized shapes, mostly triangular meshes. The correspondence will then be a mapping between the vertices.

Isometries

A mapping $\Phi : \mathcal{M} \to \mathcal{N}$ between two shapes (manifolds) is an isometry if

$$d_{\mathcal{M}}(x, y) = d_{\mathcal{N}}(\Phi(x), \Phi(y)) \quad \text{for all points } x, y \in \mathcal{M}.$$

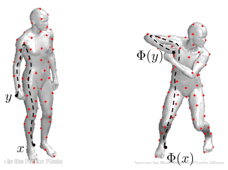

If such a mapping exists $\mathcal{M}$ and $\mathcal{N}$ are called isometric. Many shape matching approaches assume that the shapes to be matched are (nearly) isometric. The task then becomes to find the (almost-)isometry Φ.

Euclidean isometry

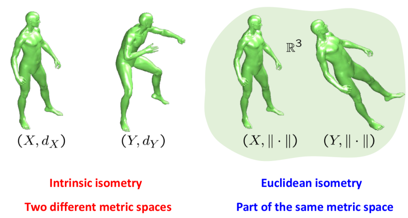

Last week we have seen two methods (ICP and PCA) that can be used to find a correspondence between shapes that are isometric with respect to the Euclidean metric (rigid alignment).
Today we discuss a way to transform the more difficult problem of intrinsic isometries into a rigid alignment problem.

Canonical forms

The main idea is to transform the two shapes to be matched into **canonical forms** such that the two canonical forms are isometric with respect to the euclidean metric.

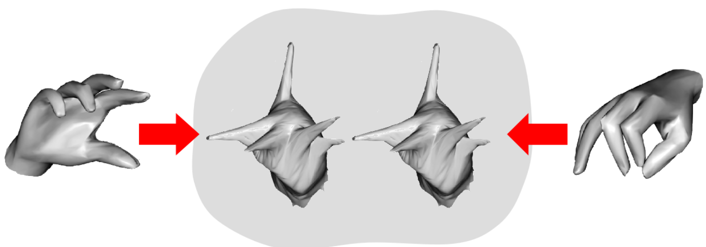

Isometric embedding

The immediate approach for isometric shapes is to find mappings $X \to \mathbb{R}^k$ that preserve all distances.

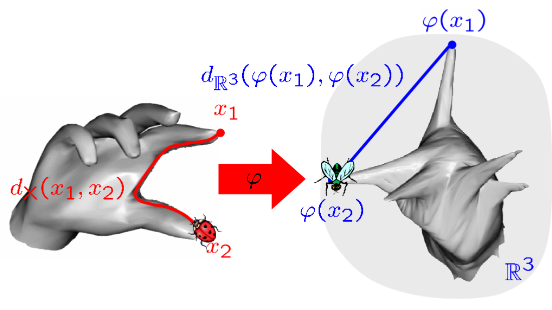

Map makers problem

$d_S = d_{\mathbb{R}^2}$

No isometric embedding

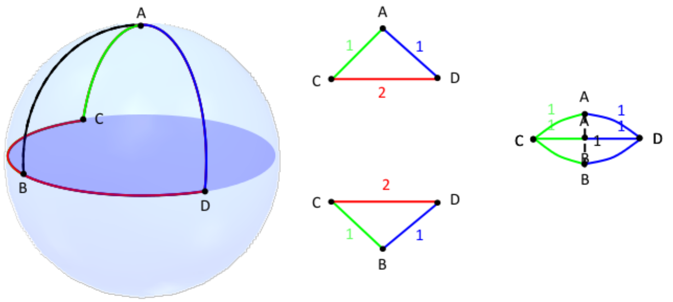

Conclusion: There is no isometric embedding in any euclidean space.

Multidimensinal Scaling

If we cannot find an isometric embedding, try to find a mapping $\Phi : (X, d_X) \to \mathbb{R}^k$ that distorts the distances the least.
We assume X is sampled at the vertices $\{v_1, \dots, v_V\}$ and we have a way to calculate the geodesic distances $d_X(v_i, v_j)$.
Notation:

- $\mathbf{D} \in \mathbb{R}^{V \times V}$ is a matrix storing all the pairwise distances on X: $\mathbf{D}_{ij} = d_X(v_i, v_j)$
- $\mathbf{Z} \in \mathbb{R}^{V \times k}$ is representing the embedding: $\mathbf{Z}_{j\cdot} = \Phi(v_j)$
- $\tilde{\mathbf{D}}(Z) \in V \times V$ is storing the euclidean distances of the mappes vertices: $\tilde{\mathbf{D}}_{ij}(Z) = \|z_i - z_k\|_{\mathbb{R}^k}$

We want to find the matrix Z minimizing the distortion (stress)

$$\sigma(\mathbf{Z}) = \sum_{i=1}^{V} \sum_{j=i+1}^{V} |\tilde{\mathbf{D}}_{ij}(Z) - \mathbf{D}_{ij}|^2$$

In matrix notation

One can show that

$$\sigma(\mathbf{Z}) = \sum_{i=1}^{V} \sum_{j=i+1}^{V} |\tilde{\mathbf{D}}_{ij}(Z) - \mathbf{D}_{ij}|^2$$
$$= \operatorname{tr}(\mathbf{Z}^T \mathbf{W} \mathbf{Z}) - 2 \operatorname{tr}(\mathbf{Z}^T \mathbf{B}(\mathbf{Z}) \mathbf{Z}) + \sum_{i=1}^{V} \sum_{j=i+1}^{V} \mathbf{D}_{ij}^2$$

with

$$\mathbf{W} = \begin{pmatrix} V-1 & -1 & & -1 \\ -1 & V-1 & & -1 \\ \vdots & & \ddots & \vdots \\ -1 & -1 & \dots & V-1 \end{pmatrix} \quad \mathbf{B}_{ij}(\mathbf{Z}) = \begin{cases} -\frac{\mathbf{D}_{ij}}{\tilde{\mathbf{D}}_{ij}(Z)} & i \neq j, \tilde{\mathbf{D}}_{ij}(Z) \neq 0 \\ 0 & i \neq j, \tilde{\mathbf{D}}_{ij}(Z) = 0 \\ -\sum_{k \neq i} \mathbf{B}_{ij}(\mathbf{Z}) & i = j \end{cases}$$

MDS objective

We want to solve the optimization problem

$$\underset{\mathbf{Z} \in \mathbb{R}^{n \times k}}{\operatorname{argmin}} \operatorname{tr}(\mathbf{Z}^T \mathbf{W} \mathbf{Z}) - 2 \operatorname{tr}(\mathbf{Z}^T \mathbf{B}(\mathbf{Z}) \mathbf{Z}) + \sum_{i=1}^{V} \sum_{j=i+1}^{V} \mathbf{D}_{ij}^2$$

- last term independent of $\mathbf{Z}$
- not linear
- not even quadratic (second term)
- gradient descent very expensive (a lot of evaluations of stress and gradient) and therefore slow

Iterative majorization

Assume we can find a family of functions $h(\mathbf{Z}; \mathbf{Q})$ satisfying

- $h(\mathbf{Q}; \mathbf{Q}) = \sigma(\mathbf{Q})$
- $h(\mathbf{Z}; \mathbf{Q}) \geqslant \sigma(\mathbf{Z}) = \operatorname{tr}(\mathbf{Z}^T \mathbf{W} \mathbf{Z}) - 2 \operatorname{tr}(\mathbf{Z}^T \mathbf{B}(\mathbf{Z}) \mathbf{Z}) + \sum_{i=1}^{V} \sum_{j=i+1}^{V} \mathbf{D}_{ij}^2$
- $h(\cdot; \mathbf{Q})$ is convex for all $\mathbf{Q}$

$h(\cdot, Q)$ are then called majorizing functions of σ.

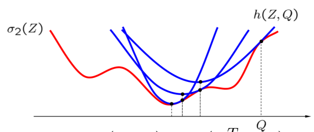

Iterative minorization:

- $\mathbf{Z}^{(k+1)} = \operatorname{argmin}_Z h(\mathbf{Z}; \mathbf{Q}^{(k)})$
- $\mathbf{Q}^{(k+1)} = Z^{(k)}$

$$h(\mathbf{Z}, \mathbf{Q}) = \operatorname{tr}(\mathbf{Z}^T \mathbf{W} \mathbf{Z}) - 2 \operatorname{tr}(\mathbf{Z}^T \mathbf{B}(\mathbf{Q}) \mathbf{Q}) + \sum_{i=1}^{V} \sum_{j=i+1}^{V} \mathbf{D}_{ij}^2$$

Quadratic majorizer

The majorizing function

$$h(\mathbf{Z}, \mathbf{Q}) = \operatorname{tr}(\mathbf{Z}^T \mathbf{W} \mathbf{Z}) - 2 \operatorname{tr}(\mathbf{Z}^T \mathbf{B}(\mathbf{Q}) \mathbf{Q}) + \sum_{i=1}^{V} \sum_{j=i+1}^{V} \mathbf{D}_{ij}^2$$

is **quadratic** in $\mathbf{Z}$. A necessary condition for a minimizer is that the gradient vanishes:

$$\nabla_{\mathbf{Z}} h(\mathbf{Z}, \mathbf{Q}) = 2\mathbf{W}\mathbf{Z} - 2\mathbf{B}(\mathbf{Q})\mathbf{Q} = 0$$

The matrix $\mathbf{W} = V \cdot \mathbf{I} - \mathbb{1}$ does not have full rank (interpretation?), so we make use of its pseudo inverse $\mathbf{W}^\dagger = \frac{1}{V}(\mathbf{I} - \frac{1}{V}\mathbb{1})$ to get

$$\mathbf{Z} = \mathbf{W}^\dagger \mathbf{B}(\mathbf{Q})\mathbf{Q} = \frac{1}{V}(\mathbf{I} - \frac{1}{V}\mathbb{1})\mathbf{B}(\mathbf{Q})\mathbf{Q} = \frac{1}{V}\mathbf{B}(\mathbf{Q})\mathbf{Q}$$

SMACOF algorithm

The SMACOF algorithm (Scaling by minimizing a convex function) thus reads as follows:
While not converged

$$\mathbf{Z}^{(k+1)} = \frac{1}{V}\mathbf{B}(Z^{(k)})Z^{(k)}$$

- no guarantee of global convergence
- decreasing stress at each iteration
- Complexity: $O(kV^2)$

Complexity

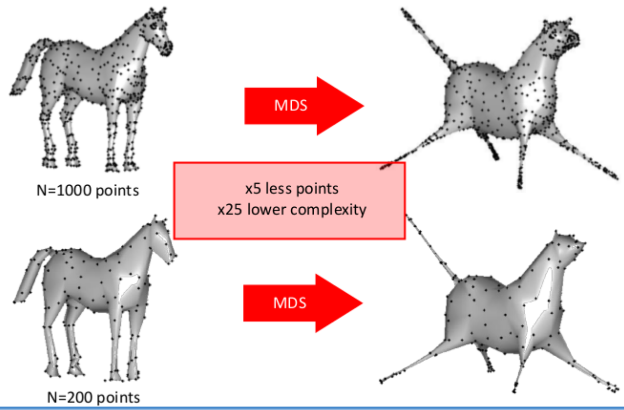

Gradient/Dirichlet energy

Dirichlet energy

As an alternative to MDS a popular approach to embed a shape $\mathcal{M}$ into a Euclidean space is by finding functions $\varphi_i : \mathcal{M} \to \mathbb{R}$ that are orthonormal *i.e.*$\langle \varphi_i, \varphi_j \rangle_{L^2(\mathcal{M})} = \delta_{ij}$ and minimize the **Dirichlet energy**

$$E_D(\varphi_i) = \int_{\mathcal{M}} \langle \nabla \varphi_i, \nabla \varphi_i \rangle dp = \int_{\mathcal{M}} \|\nabla \varphi_i\|^2 \, dp.$$

The Dirichlet energy measures how *variable* a function is. Let $\mathcal{M} = (-\pi, \pi)$, then

$$E_D(\cos(kx)) = \int_{-\pi}^{\pi} \|\nabla \cos(kx)\|^2 \, dx = k^2 \int_{-\pi}^{\pi} \sin^2(kx) dx = k^2 \left[\frac{x}{2} - \frac{\sin(2kx)}{4k} \right]_{-\pi}^{\pi} = \pi k^2$$

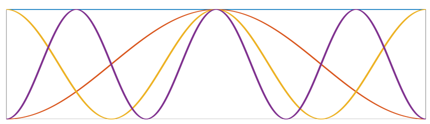

Gradient

We have yet not defined what the gradient of a function $f : \mathcal{M} \to \mathbb{R}$ is. We do know gradients of functions defined on Euclidean domains. For a function $\tilde{f} : \mathbb{R}^2 \to \mathbb{R}$ the gradient is given by

$$\nabla \tilde{f} = \begin{pmatrix} \frac{\partial \tilde{f}}{\partial u_1} \\ \frac{\partial \tilde{f}}{\partial u_1} \end{pmatrix}$$

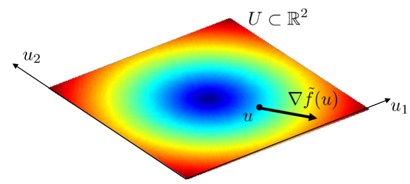

Geometric meaning of the gradient

Geometric meaning of the gradient

- the vector that points in the **direction of steepest increase** of $\tilde{f}$
- its length measures the strength of increase
- relationship with the differential of $\tilde{f}$:

$$\begin{aligned} d\tilde{f}(u)(\vec{v}) &= \lim_{h \to 0} \frac{\tilde{f}(u + h\vec{v}) - \tilde{f}(u)}{h} \\ &= \frac{d}{dh} \tilde{f}(u + h\vec{v})|_{h=0} \\ &= \langle \nabla \tilde{f}(u), \vec{v} \rangle \end{aligned}$$

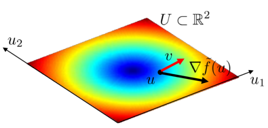

Riesz representation

Dual space:
Let X be a vector space. Then we denote by

$$X^* = \{\psi : X \to \mathbb{R} | \psi \text{ is linear}\}$$

the **dual space** of X.

Riesz Representation theorem:
Let $(H, \langle \cdot, \cdot \rangle)$ be a Hilbertspace (inner product, complete). Then for each continuous $\psi \in H^*$ there exists a unique $y \in H$ such that

$$\psi(x) = \langle y, x \rangle \qquad \forall x \in H$$

Differential of f

The differential of a differentiable function $f : \mathcal{M} \to \mathbb{R}$ at a point $p \in \mathcal{M}$ is the linear mapping $df(p) : T_p\mathcal{M} \to \mathbb{R}$ satisfying

$$df(p)[\vec{v}] = \lim_{t \to 0} \frac{f(c(t)) - f(p)}{t}$$

for all curves $c : (-\varepsilon, \varepsilon) \to \mathcal{M}$ with $c(0) = p$ and $\dot{c}(0) = \vec{v}$.

Uniqueness and Linearity

Let $c_1(0) = c_2(0) = p = x(u)$ and $\dot{c}_1(0) = \dot{c}_2(0) = \vec{v}$. By defining the preimages $\gamma_i(t) = x^{-1} \circ c_i(t)$ and as usual $\tilde{f} = f \circ x$ we get

$$\begin{aligned} df(p)[v] &= \lim_{t \to 0} \frac{f(c_i(t)) - f(p)}{t} = \lim_{t \to 0} \frac{\tilde{f}(\gamma_i(t)) - \tilde{f}(u)}{t} = \frac{d}{dt} \tilde{f}(\gamma_i(t)) \\ &= \langle \nabla \tilde{f}(u), \dot{\gamma}_i(0) \rangle = \langle \nabla \tilde{f}(u), (Dx)^{-1} \dot{c}_i(0) \rangle = \langle \nabla \tilde{f}(u), (Dx)^{-1} \vec{v} \rangle \end{aligned}$$

This is independent of the choice of c and linear in $\vec{v}$.

Gradient on manifold

Let $f : \mathcal{M} \to \mathbb{R}$ be a differentiable function. The **gradient** $\nabla f(p)$ at $p \in \mathcal{M}$ is the unique element of $T_p\mathcal{M}$ such that

$$\langle \nabla f(p), \vec{v} \rangle = df(p)[v]$$

In local coordinates

Let $p = x(u)$. Given $\nabla \tilde{f}(u) \in \mathbb{R}^2$ and the first fundamental form $g(u) \in \mathbb{R}^{2 \times 2}$, the coefficients $\alpha \in \mathbb{R}^2$ (local coordinates) of $\nabla f = Dx \cdot \alpha \in T_p\mathcal{M}$ are given by

$$\alpha = g^{-1}(u) \nabla \tilde{f}(u)$$

. Let $\beta \in \mathbb{R}^2$ be the coefficients of $\vec{v} \in T_p\mathcal{M}$. Then

$$df(p)[\vec{v}] = \langle \nabla \tilde{f}(u), \beta \rangle = \langle \alpha, \beta \rangle_{g(u)} = \langle \nabla f, \vec{v} \rangle$$

Notice that this in general is a different vector then $\nabla \tilde{f}(u)$!

Example

Let $\mathcal{M}$ be a sphere parametrized via $x(u) = \begin{pmatrix} \cos(u_1)\cos(u_2) \\ \sin(u_1)\cos(u_2) \\ \sin(u_2) \end{pmatrix}$ and $f(p) = x_2(p)$.

Then $\tilde{f}(u) = \sin(u_1)\cos(u_2)$, $\nabla \tilde{f}(u) = \begin{pmatrix} \cos(u_1)\cos(u_2) \\ -\sin(u_1)\sin(u_2) \end{pmatrix}$ and

$g^{-1}(u) = \begin{pmatrix} \frac{1}{\cos^2(u_2)} & 0 \\ 0 & 1 \end{pmatrix}$. For the local coordinates α of ∇f this yields

$\alpha = \begin{pmatrix} \frac{\cos(u_1)}{\cos(u_2)} \\ -\sin(u_1)\sin(u_2) \end{pmatrix}$