

Analysis of Three-Dimensional Shapes

(IN2238, TU München, Summer 2014)

Intrinsic Metrics
(26.06.2014)


Dr. Emanuele Rodolà
Mathieu Andreux
 {rodola,andreux}@in.tum.de
Room 02.09.058, Informatik IX

Seminar

«Time-discrete geodesics in the space of shells»

Zorah Löhner

Wednesday, July 2nd
14:00 Room 02.09.023




Final exam

When? Second half of July. We'll set up a Doodle.

Where? Probably office **02.09.058**, otherwise room **02.09.023**.

What? Everything we covered in the lecture and exercise classes.



The matching game


Thomas Hörmann
score: 0.947



Wrap-up

Let $T : M \rightarrow N$ be a bijection between two regular surfaces M and N .

Given a scalar function $f : M \rightarrow \mathbb{R}$ on shape M , we can induce a function $g : N \rightarrow \mathbb{R}$ on the other shape by composition:



We can denote this transformation by a functional T_F , such that

$$T_F(f) = f \circ T^{-1}$$

Wrap-up

We called T_F a **functional map** between (scalar functions defined on) the two surfaces. Instead of mapping points to points, **it maps functions to functions**.

Functional maps have some interesting properties:

- They are *linear* maps:

$$T_F(\alpha_1 f_1 + \alpha_2 f_2) = \alpha_1 T_F(f_1) + \alpha_2 T_F(f_2)$$

While T can in general be a very complex transformation between the two shapes, T_F always acts *linearly*.

- Constructing T_F from T is trivial (by definition $T_F(f) = f \circ T^{-1}$)
- Constructing T from T_F is also easy (e.g. use indicator functions)

Wrap-up

These properties suggest that, when we are dealing with correspondences, knowledge of the functional map T_F is equivalent to knowledge of the point-to-point correspondence T .

Together with linearity of T_F , this gives us a way to exploit this knowledge for matching purposes!

Let $\{\phi_i^M\}$ be an **orthogonal** (w.r.t. some inner product) basis for functions f on M , and $\{\phi_j^N\}$ be an orthogonal basis for functions on N . Then we can express the action of T_F in matrix notation:


$$T_F(\mathbf{a}) = C\mathbf{a} = \mathbf{b}$$

$a_i = \langle f, \phi_i^M \rangle$ matrix representation of T_F coefficients in the $\{\phi_j^N\}$ basis

Wrap-up

$$Ca = b$$

This alone does not give us a more convenient way to formulate matching problems. In fact, if we consider *the standard basis* on two shapes of n points, we get to the permutation problem:



$$Pa = b$$

$n \times n$

with this choice of a basis, \mathbf{a} is the usual vector representation of f .

$$f = \sum_i a_i \phi_i^M$$

Wrap-up

Other choices for a basis are possible.


If we use the **eigenfunctions of the Laplace-Beltrami operator**, we can reduce the size of matrix C and still have a good approximation.

$$f = \sum_i a_i \phi_i^M \approx \sum_{i=1}^m a_i \phi_i^M$$

$$C\mathbf{a} = \mathbf{b}$$

Matrix C , which represents our correspondence, is a $m \times m$ matrix. *Its size does not depend on the size of the shapes!*

Typical values for m are 50 or 100



Wrap-up

We have also seen that many common constraints that are used in shape matching problems also become *linear* in the functional map formulation.

Descriptor preservation

$$C \mathbf{a} = \mathbf{b}$$

function on M function on N


For instance, consider *curvature* or other descriptors.

This means that we can set up a linear system (where C is the unknown), and solve it in the least-squares sense:

$$C \begin{pmatrix} | & & | \\ \mathbf{a}_1 & \cdots & \mathbf{a}_n \\ | & & | \end{pmatrix} = \begin{pmatrix} | & & | \\ \mathbf{b}_1 & \cdots & \mathbf{b}_n \\ | & & | \end{pmatrix} \quad \Rightarrow \quad C^* = \arg \min_C \|CA - B\|^2$$

Shapes as metric spaces

As we know, one successful way to model the matching problem is to consider shapes as metric spaces:




We have seen this simple model arising in several different topics, such as:

- **Distance between shapes** (Lipschitz, Gromov-Hausdorff, ...)
- Multi-dimensional scaling (**Euclidean embeddings**, canonical forms, ...)
- Differential geometry ("**natural**" distance on regular surfaces)
- Functional maps (**distance maps** to landmark correspondences)

Gromov-Hausdorff distance

For example, let's look again at our discretization of the Gromov-Hausdorff distance between two metric spaces:


$$d_{GH}(\mathbf{X}, \mathbf{Y}) = \frac{1}{2} \inf_{R \subset \mathbf{X} \times \mathbf{Y}} \sup_{(x,y), (x',y') \in R} |d_X(x, x') - d_Y(y, y')|$$




Gromov-Hausdorff distance

$$d_{GH}(\mathbf{X}, \mathbf{Y}) = \frac{1}{2} \inf_{R \subset \mathbf{X} \times \mathbf{Y}} \sup_{(x,y), (x',y') \in R} |d_X(x, x') - d_Y(y, y')|$$


We already know that the correspondence attaining the infimum will be invariant exactly to the kind of transformations to which the metrics d_X, d_Y are invariant.



geodesic metric \rightarrow isometries




Euclidean metric \rightarrow rigid motions



?

Multi-dimensional scaling


$$f = \arg \min_{f: X \rightarrow \mathbf{R}^m} \sum_{i>j} |d_X(x_i, x_j) - d_{\mathbf{R}^m}(f(x_i), f(x_j))|^2$$



Multi-dimensional scaling

$$f = \arg \min_{f: X \rightarrow \mathbf{R}^m} \sum_{i > j} |d_X(x_i, x_j) - d_{\mathbf{R}^m}(f(x_i), f(x_j))|^2$$

Topological noise can significantly alter distances.




Geodesic distance

We have seen that the first fundamental form on regular surfaces allows us to measure lengths of curves lying on the surface.

We defined the distance $d(p, q)$ between two points of S as

$$d(p, q) = \inf_{\alpha: [0, 1] \rightarrow S} \int_0^1 \|\alpha'(t)\| dt$$

where $\alpha(0) = p$, $\alpha(1) = q$.




This “natural” intrinsic distance on the surface is commonly referred to as **geodesic distance** in the shape analysis literature.

Geodesic distance

$$d(p, q) = \inf_{\alpha: [0,1] \rightarrow S} \int_0^1 \|\alpha'(t)\| dt = \inf_{\alpha: [0,1] \rightarrow S} \int_0^1 \sqrt{I(\alpha'(t))} dt$$

Since isometries preserve the first fundamental form, the *geodesic distance* is preserved under isometries.



Heat diffusion

We have seen how **heat diffusion** on regular surfaces allows to capture their intrinsic geometry. In particular, we studied the following model:

$$\begin{aligned}\frac{\partial u(x, t; u_0)}{\partial t} &= \Delta u(x, t; u_0) \\ u(x, 0) &= u_0(x)\end{aligned}$$

A solution to the heat equation is given by:

$$u(x, t; u_0) = \int_S k_t(x, y) u_0(y) dy$$

The function $k_t : S \times S \rightarrow \mathbb{R}$, called **heat kernel**, describes how much heat is transferred from one point to the other in time t .

Heat kernel

We provided an explicit expression for the heat kernel in \mathbb{R}^n :

$$k_t^{\mathbb{R}^n}(x, y) = \frac{1}{(\sqrt{4\pi t})^n} \exp\left(-\frac{\|x - y\|^2}{4t}\right)$$

as well as in the case of regular surfaces S :

$$k_t^S(x, y) = \sum_k e^{\lambda_k t} \phi_k(x) \phi_k(y) \quad \text{remember that } \lambda_k \leq 0$$

We didn't give any formal proof, but we stated that one can recover the *geodesic distance* on a surface directly from the heat kernel:

$$d_S^2(x, y) = -\lim_{t \rightarrow 0} 4t \log(k_t^S(x, y))$$

A distance based on heat diffusion

Based on these observations, we ask the following question:

Can we define a *new* notion of distance based on the ideas of heat diffusion?

A natural candidate for such a distance is the heat kernel $k_t^S(x, y)$ itself.

However, it is not difficult to see that such a function does *not* satisfy all the metric axioms. In particular, if we look again at the spectral decomposition

$$k_t^S(x, y) = \sum_k e^{\lambda_k t} \phi_k(x) \phi_k(y)$$

we immediately realize that $k_t^S(x, y) = 0 \Leftrightarrow x = y$

Diffusion kernel

The heat kernel $k_t(x, y)$ satisfies the properties of a **diffusion kernel**:

$$k_t(x, y) \geq 0 \quad (\text{non-negativity})$$

$$k_t(x, y) = k_t(y, x) \quad (\text{symmetry})$$

$$\int \int k_t^2(x, y) dx dy < \infty \quad (\text{square integrability})$$

$$\int \int k_t(x, y) f(x) f(y) dx dy \geq 0 \quad (\text{positive semi-definiteness})$$

$$\int k_t(x, y) dy = 1 \quad (\text{conservation})$$




in matrix notation, this corresponds to a stochastic matrix

Exercise

Random walks

A **random walk** is a path modeled as a succession of random steps.

For example, the path traced by a molecule in a liquid, or the path walked by a drunken sailor from the bar to a lamp post.



Brownian motion is the random motion of particles suspended in a fluid. The randomness is the result of the particles colliding with the fluid molecules (or atoms in the case of a gas).

Brownian motion

The physical phenomenon of Brownian motion was modeled mathematically by Einstein in 1905.

In particular, he showed that if $u(x, t)$ is the **density** of Brownian particles at point x and time t , then u satisfies the diffusion equation:

$$\frac{\partial u}{\partial t} = D \Delta u$$

where D is the *mass diffusivity* or *diffusion coefficient*, in general a non-linear function which depends on physical properties such as temperature and viscosity

We already know that a solution to this diffusion equation (with $D = 1$) is given by:

$$u(x, t; u_0) = \int_S k_t(x, y) u_0(y) dy$$

Brownian motion and heat kernel

Thus, the heat diffusion equation provides a model of the time evolution of the **probability density function** $u(x, t)$ associated to the position of a particle undergoing a Brownian motion.

We have seen that, if we start from a δ_z distribution centered around $z \in S$, we get:

$$u(x, t; \delta_z) = \int_S k_t(x, y) \delta_z(y) dy = k_t(x, z)$$


Thus, the probability that a particle is in a small region C around point x after time t , is given by

$$\int_{C \subset S} u(x, t; \delta_z) dx = \int_{C \subset S} k_t(x, z) dx$$

A probabilistic interpretation

This tells us that $k_t(x, y)$ is the **probability density function of transition** from x to y by a **random walk** of length t .

$$u(x, t; u_0) = \int_S k_t(x, y) u_0(y) dy$$



Brownian motion starting at point x , reaching C in time t , with probability given by:

$$\int_C k_t(x, y) dy$$


To emphasize this relationship, some authors denote the heat kernel by $p_t(x, y)$

Diffusion distance

A family of **diffusion distances** $\{d_t\}_{t \in \mathbb{R}_+}$ can be defined by


$$d_t^2(x, y) = \|k_t(x, \cdot) - k_t(y, \cdot)\|^2 = \int_S (k_t(x, z) - k_t(y, z))^2 dz$$

which is nothing but a (weighted) L_2 distance between two probability density functions. Note that the expression above is defining d_t^2 , not d_t .



$$d_t^2(x, t) = 3.6340$$

$t = 10$



$$d_t^2(x, t) = 0.5479$$

Properties

$$d_t^2(x, y) = \|k_t(x, \cdot) - k_t(y, \cdot)\|^2 = \int_S (k_t(x, z) - k_t(y, z))^2 dz$$

- It is a metric.
- Diffusion time t plays the role of a scale parameter.
- It reflects the connectivity of the data at a given scale (denoted by t). If two points x and y are close (in the diffusion sense), there is a large probability of transition from x to y and vice versa.
- The definition involves summing over **all paths** of length $2t$ connecting x to y . As a consequence, this number is very robust to noise perturbation, unlike the geodesic distance (see next slide).

«Lengths of paths»

One useful property of the heat kernel (which we hinted at in the last bullet point of the previous slide) is the following:

$$k_{2T}(x, y) = \int_S k_T(x, z)k_T(z, y)dz$$

To prove this property, we start by imposing:

$$u(x, t) = k_{t+T}(x, y) \quad \text{for some } y$$

Then, applying the heat diffusion model, it must be:

$$\left. \begin{array}{l} u_0(x) = u(x, 0) = k_T(x, y) \\ \frac{\partial u(x, t; u_0)}{\partial t} = \Delta u(x, t; u_0) \end{array} \right\} u(x, t) = \int_S k_t(x, z)u_0(z)dz = \int_S k_t(x, z)k_T(z, y)dz$$

Setting $t = T$ and equating the two expressions for $u(x, t)$, we obtain the desired result.

Diffusion distance in the LB basis

$$\begin{aligned}
 d_t^2(x, y) &= \|k_t(x, \cdot) - k_t(y, \cdot)\|^2 = \left\| \sum_i e^{\lambda_i t} \phi_i(x) \phi_i(\cdot) - \sum_i e^{\lambda_i t} \phi_i(y) \phi_i(\cdot) \right\|^2 \\
 &= \left\| \sum_i e^{\lambda_i t} \phi_i(\cdot) (\phi_i(x) - \phi_i(y)) \right\|^2 = \int_S \left(\sum_i e^{\lambda_i t} \phi_i(z) (\phi_i(x) - \phi_i(y)) \right)^2 dz \\
 &= \int_S \left(\sum_i e^{\lambda_i t} \phi_i(z) (\phi_i(x) - \phi_i(y)) \right) \left(\sum_j e^{\lambda_j t} \phi_j(z) (\phi_j(x) - \phi_j(y)) \right) dz \\
 &= \int_S \sum_{i,j} e^{\lambda_i t} e^{\lambda_j t} (\phi_i(x) - \phi_i(y)) (\phi_j(x) - \phi_j(y)) \phi_i(z) \phi_j(z) dz \\
 &= \sum_{i,j} e^{\lambda_i t} e^{\lambda_j t} (\phi_i(x) - \phi_i(y)) (\phi_j(x) - \phi_j(y)) \langle \phi_i, \phi_j \rangle \\
 &= \sum_i e^{2\lambda_i t} (\phi_i(x) - \phi_i(y))^2 \langle \phi_i, \phi_i \rangle = \sum_i e^{2\lambda_i t} (\phi_i(x) - \phi_i(y))^2
 \end{aligned}$$

1	if $i = j$
0	otherwise

Example: Diffusion distance

$$d_t^2(x, y) = \sum_i e^{2\lambda_i t} (\phi_i(x) - \phi_i(y))^2$$



$t = 1$





$t = 5$



$t = 10$

Pitfall



Diffusion map

$$d_t^2(x, y) = \sum_i e^{2\lambda_i t} (\phi_i(x) - \phi_i(y))^2$$

The definition we gave for the diffusion distance suggests the following Euclidean embedding:

$$p \mapsto (e^{\lambda_1 t} \phi_1(p), e^{\lambda_2 t} \phi_2(p), e^{\lambda_3 t} \phi_3(p) \dots) \quad \text{for a fixed } t \in \mathbb{R}_+$$

We have already seen another similar embedding, which we called GPS:

$$p \mapsto \left(\frac{\phi_1(p)}{\sqrt{-\lambda_1}}, \frac{\phi_2(p)}{\sqrt{-\lambda_2}}, \frac{\phi_3(p)}{\sqrt{-\lambda_3}}, \dots \right)$$

Scale-invariant intrinsic metric

$$p \mapsto (e^{\lambda_1 t} \phi_1(p), e^{\lambda_2 t} \phi_2(p), e^{\lambda_3 t} \phi_3(p) \dots)$$

It is not difficult to see (check it!) that the diffusion map is **not** scale invariant.

However, the previous slides raise the question on whether the following definition is a valid intrinsic metric function:

$$d^2(x, y) = \sum_i \frac{1}{-\lambda_i} (\phi_i(x) - \phi_i(y))^2$$

That is, the L_2 distance between two global point signatures at points x and y .

Commute-time distance

$$d^2(x, y) = \sum_i \frac{1}{-\lambda_i} (\phi_i(x) - \phi_i(y))^2$$

Indeed, it can be proved that this is in fact a metric function! Since we already proved that the GPS embedding is scale-invariant, it is not difficult to see that this metric is also scale-invariant.

The resulting metric is called **commute-time distance**.

Similarly to the diffusion distance, this distance can be rewritten in “kernel notation” as:

$$d^2(x, y) = g(x, x) + g(y, y) - 2g(x, y)$$

where $g(x, y) = \sum_k \frac{1}{-\lambda_k} \phi_k(x) \phi_k(y)$ is the **commute-time kernel**.

Commute-time kernel

At this point, it is interesting to notice the following fact:

$$\begin{aligned}\int_0^\infty k_t(x, y) dt &= \int_0^\infty \sum_k e^{\lambda_k t} \phi_k(x) \phi_k(y) dt && \text{(integrate over all possible times)} \\ &= \sum_k \phi_k(x) \phi_k(y) \int_0^\infty e^{\lambda_k t} dt \\ &= \sum_k \phi_k(x) \phi_k(y) \frac{1}{\lambda_k} e^{\lambda_k t} \Big|_0^\infty && \text{recall that } \lambda_k \leq 0 \\ &= \sum_k \frac{1}{-\lambda_k} \phi_k(x) \phi_k(y) = g(x, y)\end{aligned}$$

In other words, the commute-time kernel corresponds to the probability density function of transition from point x to y by a **random walk of any length**.

Example: Commute-time distance

diffusion, $t=10$



S

commute-time



$\alpha S'$




S




$\alpha S'$



Example: Non-isometries



diffusion, $t=5$



commute-time



Suggested reading

- *Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps.* Coifman et al. PNAS 2005.
- *Über die von der molekularkinetischen Theorie der Wärme geforderte Bewegung von in ruhenden Flüssigkeiten suspendierten Teilchen.* Einstein. Annalen der Physik 2005.