Analysis of Three-Dimensional Shapes (IN2238, TU München, Summer 2014)

Intrinsic Shape Descriptors
(05.06.2014)

Dr. Emanuele Rodolà
Matthias Vestner
Thomas Windheuser
{rodola,vestner,windheus}@in.tum.de
Room 02.09.058, Informatik IX

Seminar

«An image processing approach to surface matching»

Felix Schmid

Wednesday, June 11th

14:00 Room 02.09.023

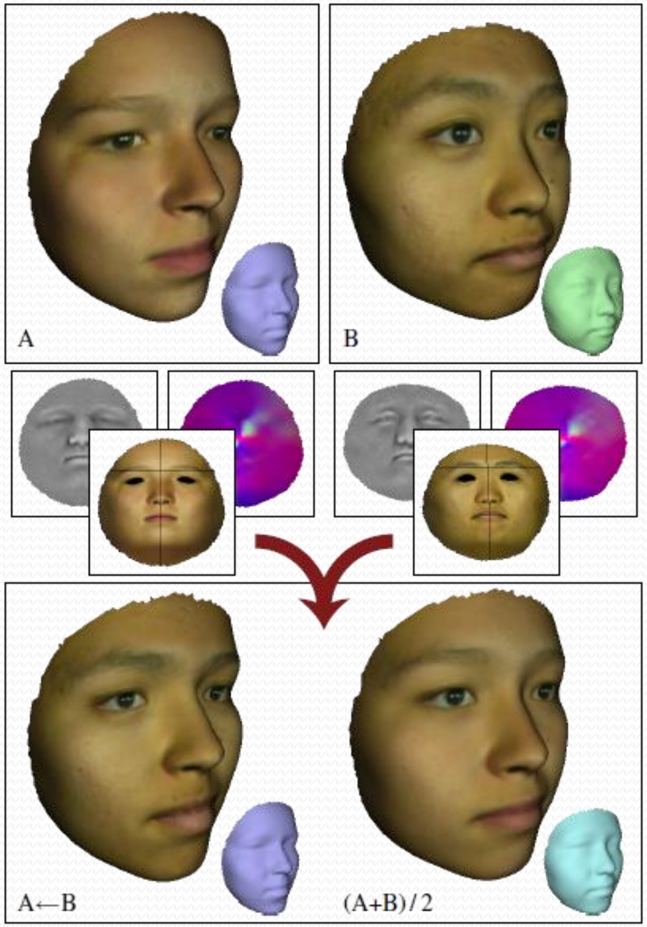

Wrap-up

We introduced two other *linear* operators for our manifolds, namely:

The **divergence** of a vector field *V* (e.g. $V = \nabla f$):

$$\text{div } V = \frac{1}{\sqrt{\det g(x)}} \sum_i \frac{\partial}{\partial x_i} \left(V_i(x) \sqrt{\det g(x)} \right)$$

The **Laplacian**, or **Laplace-Beltrami** operator of a scalar function *f* defined on the surface:

$$\Delta f = \frac{1}{\sqrt{\det g}} \sum_{i,j} \frac{\partial}{\partial x_i} \left(g^{ij} \frac{\partial \tilde{f}(x)}{\partial x_j} \sqrt{\det g(x)} \right)$$

Wrap-up

Then we considered the Helmholtz equation, involving the Laplace-Beltrami operator:

$$\Delta f = \lambda f$$

- The eigenvalues are real, non-positive, and countable.

$$0 = \lambda_0 > \lambda_1 \geq \lambda_2 \geq \ldots \to -\infty \qquad -\lambda_j \sim \frac{\pi}{|S|} j \quad \text{for } j \to \infty$$

Weyl's law

- There is exactly one **zero eigenvalue**, and its corresponding eigenfunction is constant

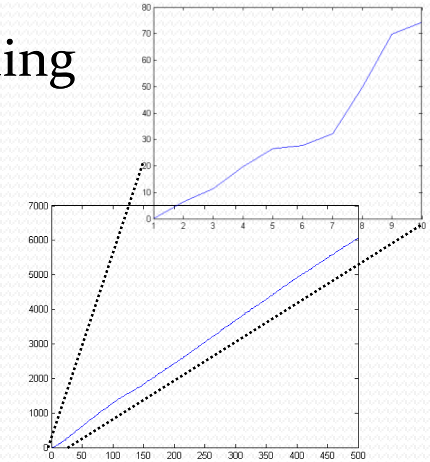

- **Eigenvectors** to distinct eigenvalues are **orthogonal**

Wrap-up

$$\Delta f = \frac{1}{\sqrt{\det g}} \sum_{i,j} \frac{\partial}{\partial x_i} \left(g^{ij} \frac{\partial \tilde{f}(x)}{\partial x_j} \sqrt{\det g(x)} \right)$$

The Laplace-Beltrami operator only depends on the metric tensor g. It is therefore **invariant under isometric deformations** of the surface.

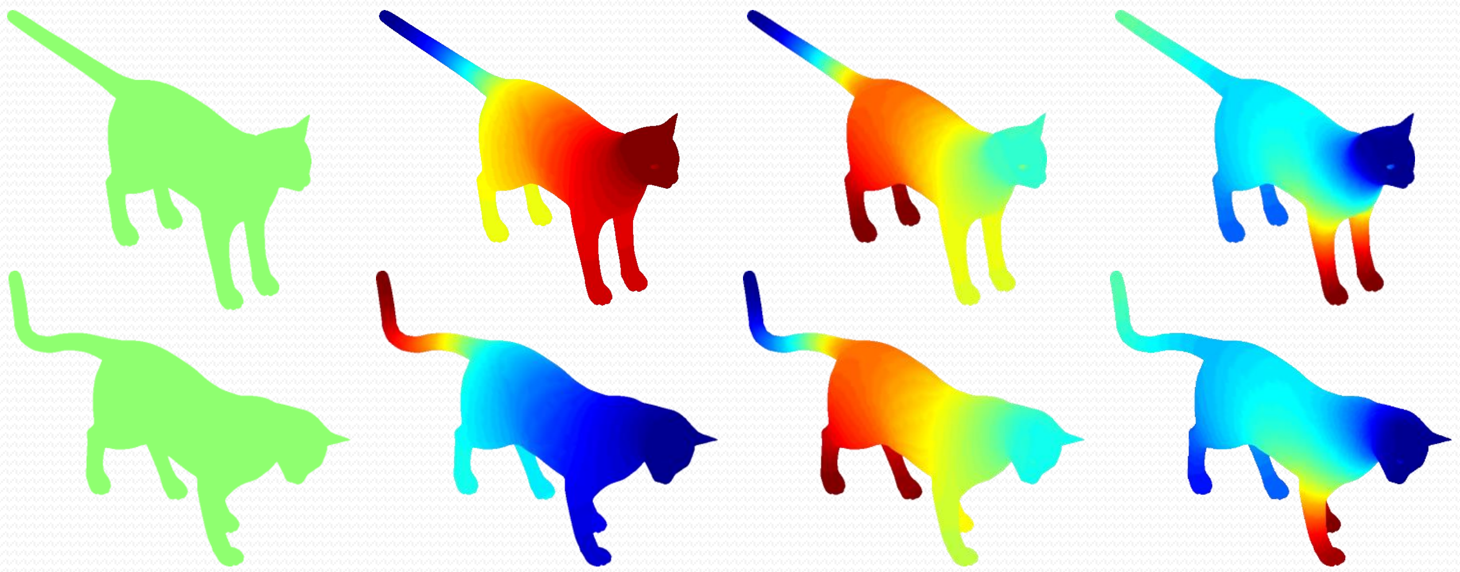

Wrap-up

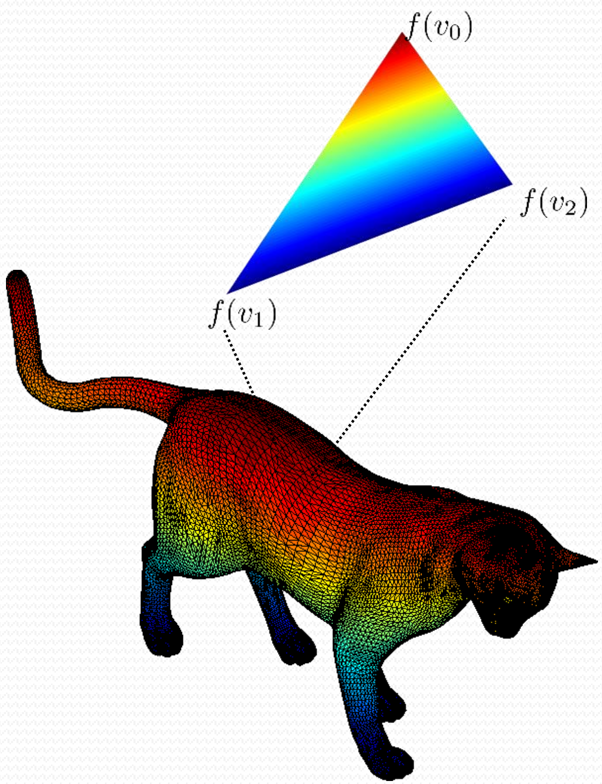

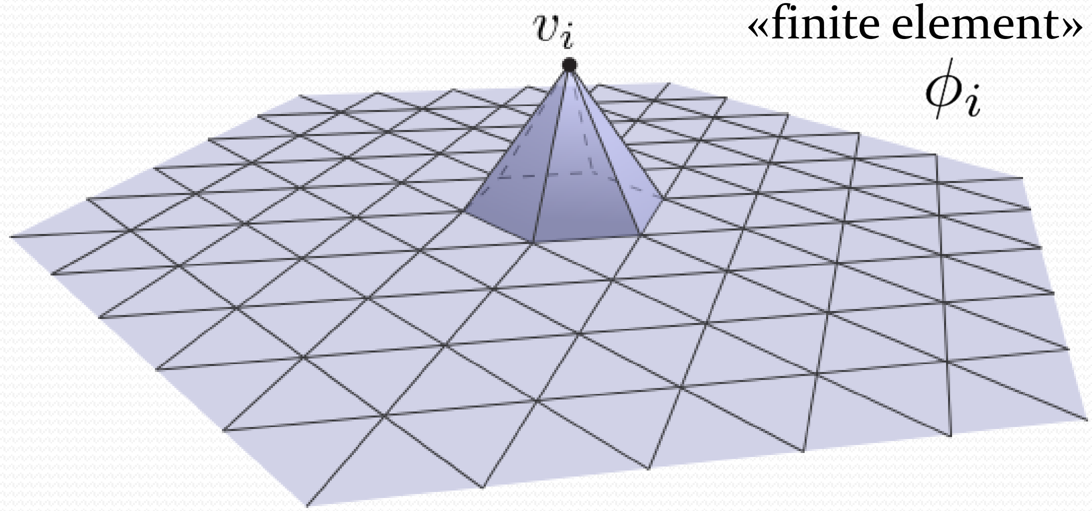

«finite element»

ϕ_i

$$h(x) := \Delta f(x) = \sum_i h_i \phi_i(x)$$

$$h = Lf$$

$$f(x) = \sum_i f(v_i)\phi_i(x)$$ piecewise linear

Wrap-up

Main idea: Test with ϕ_j

$$\langle h, \phi_j \rangle = \langle \Delta f, \phi_j \rangle$$

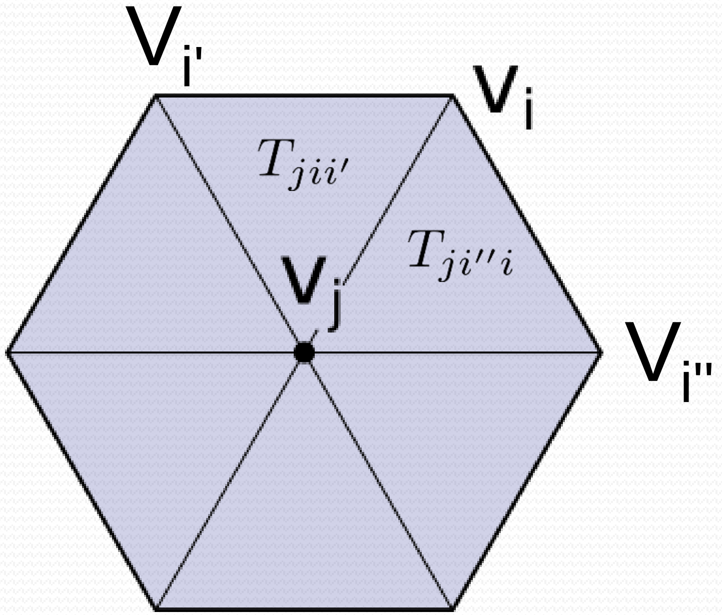

$\langle h, \phi_j \rangle = (Mh)_j$ mass matrix

$\langle \Delta f, \phi_j \rangle = (Cf)_j$ stiffness matrix

⇩

$$\boxed{h = M^{-1}Cf = Lf}$$

$$M_{ij} = \frac{|T_{jii'}|}{12} + \frac{|T_{ji''i}|}{12} \qquad M_{jj} = \sum_{i \in N(j)} \frac{|T_{jii'}|}{6}$$

$$C_{ij} = \frac{1}{2}(\cot \alpha_{ij} + \cot \beta_{ij}) \qquad C_{jj} = -\sum_{i \neq j} C_{ij}$$

Wrap-up

Note that, since $h = M^{-1}Cf = Lf$, we can rewrite the Helmholtz equation as an equivalent *generalized* eigenvalue problem:

$$\Delta f = \lambda f \Rightarrow M^{-1}Cf = \lambda f \Rightarrow \boxed{Cf = \lambda M f}$$

The eigenvalues and eigenvectors (eigenfunctions) are the same as in the original case. In particular, since *C* is symmetric and *M* is symmetric positive-definite, the generalized eigenvectors *f* are still **orthonormal** *with respect to the M-inner product*:

$$\langle f, g \rangle_M = f^{\mathrm{T}} M g$$

Exercise

In other words, we are approximating the continuous inner product as follows:

$$\int_S f(x)g(x)dx \approx f^{\mathrm{T}} M g = \langle f, g \rangle_M$$

Wrap-up

For a regular surface *S* the diffusion of heat is described by the heat equation:

$$\frac{\partial u(x,t)}{\partial t} = \Delta u(x,t)$$

$$u(x,0) = u_0(x)$$

$$u(\partial U, t) = \ldots$$

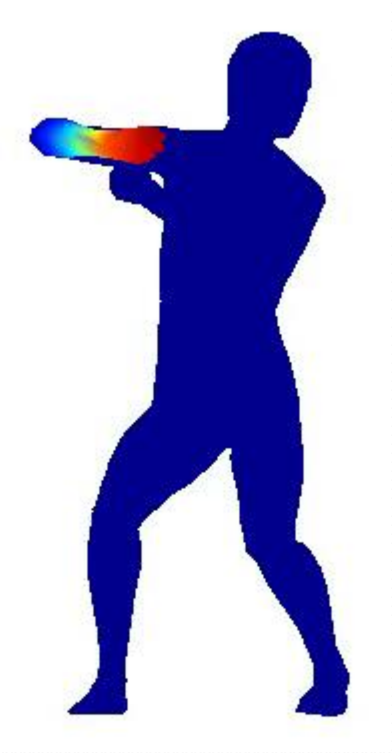

Our newly derived expressions for the Laplacian allow us to study heat diffusion on surfaces from a practical point of view.

Euclidean embeddings

In the previous lectures we have seen how to translate a general, non-rigid matching problem to a **rigid** matching problem.

We did so by finding maps $f:(X,d_X)\rightarrow(\mathbf{R}^m,\|\cdot\|)$ minimizing a *quadratic stress*:

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

multi-dimensional scaling

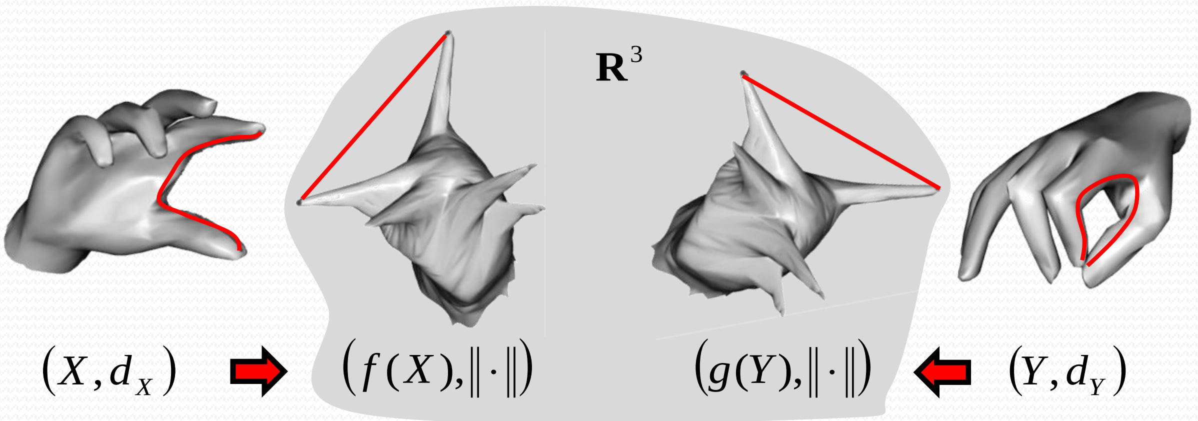

Euclidean embeddings

We referred to the minimizing *f* as a **minimum-distortion embedding** of the shape into Euclidean space.

The minimum-distortion embedding is defined in terms of *pairwise* quantities on the shape (namely, evaluations of a distance function).

Can we define alternative embeddings by making use of the new differential-geometric tools we have introduced?

The embedding should be:

- deformation-invariant
- robust to discretization process
- defined using intrinsic properties of the shape (i.e. metric tensor)
- easy to deal with

General approach

Construct an embedding that relies on the Laplace-Beltrami operator. Two important properties are immediately evident:

- The operator is isometry invariant
- Its eigenfunctions have a global nature, and are thus more stable to local changes

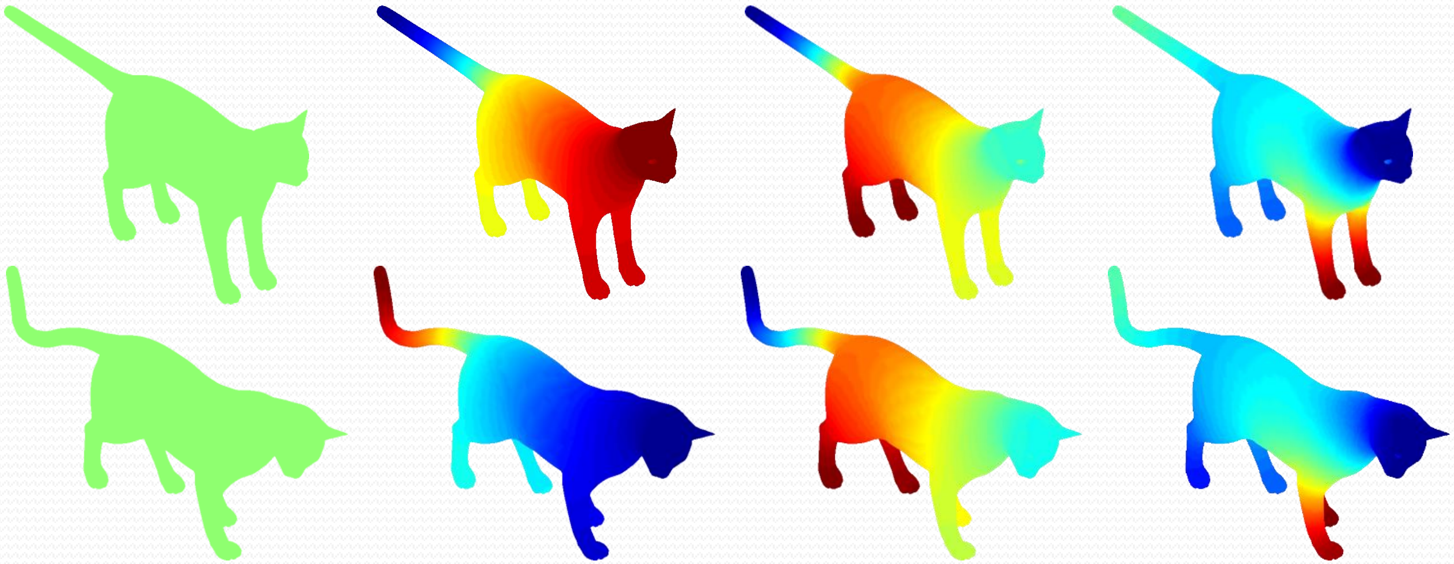

Euclidean embedding via Δ_S

The most straightforward approach is to map each point $p \in S$ to an infinite-dimensional vector according to the eigenfunctions of Δ_S:

$$\Delta_S \varphi = \lambda \varphi \qquad 0 = \lambda_0 > \lambda_1 \geq \lambda_2 \geq \ldots \to -\infty$$

$$p \mapsto (\varphi_0(p), \varphi_1(p), \varphi_2(p), \ldots) \in \mathbb{R}^\infty$$

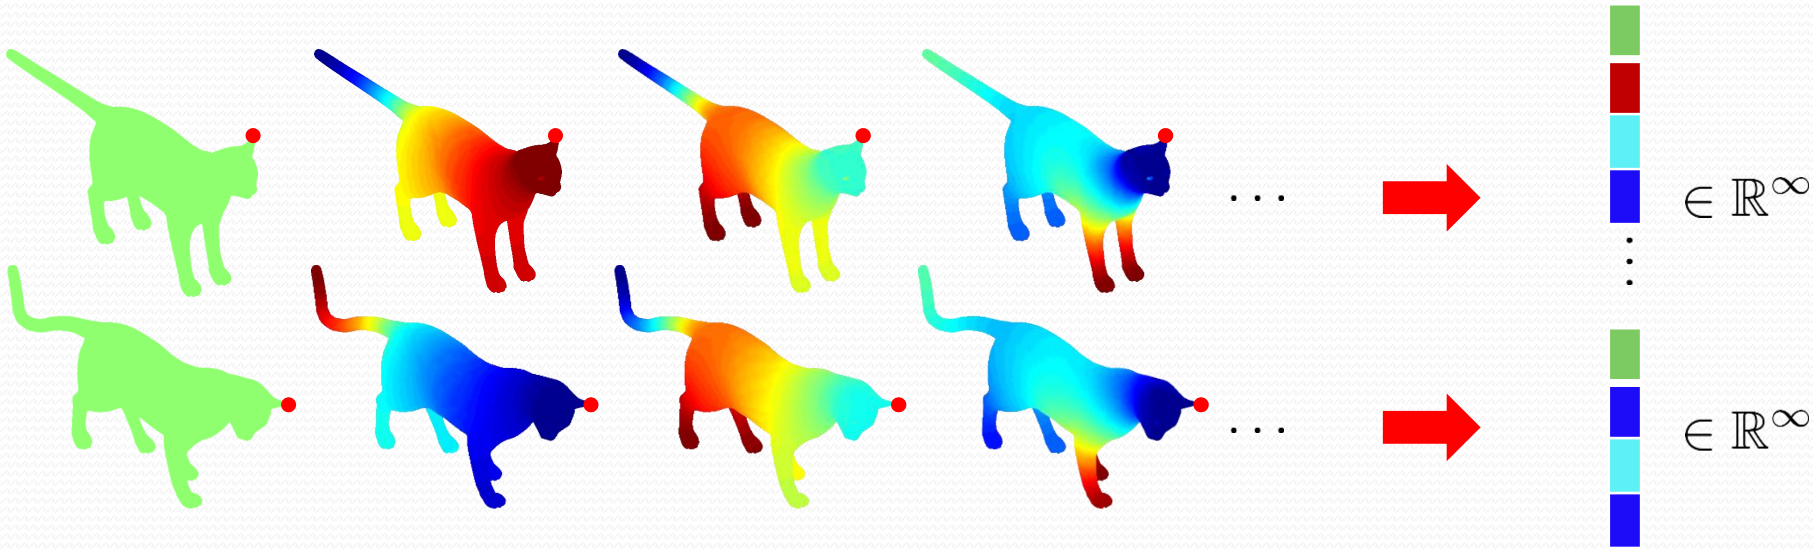

Euclidean embedding via Δ_S

$$p \mapsto (\varphi_0(p), \varphi_1(p), \varphi_2(p), \dots) \in \mathbb{R}^\infty$$

Is this a meaningful embedding?

In general, we can not expect to be given two *exactly* isometric shapes.
Two main issues we can directly deal with:

- The eigenfunctions have different *signs*
- The eigenfunctions have different *scales*

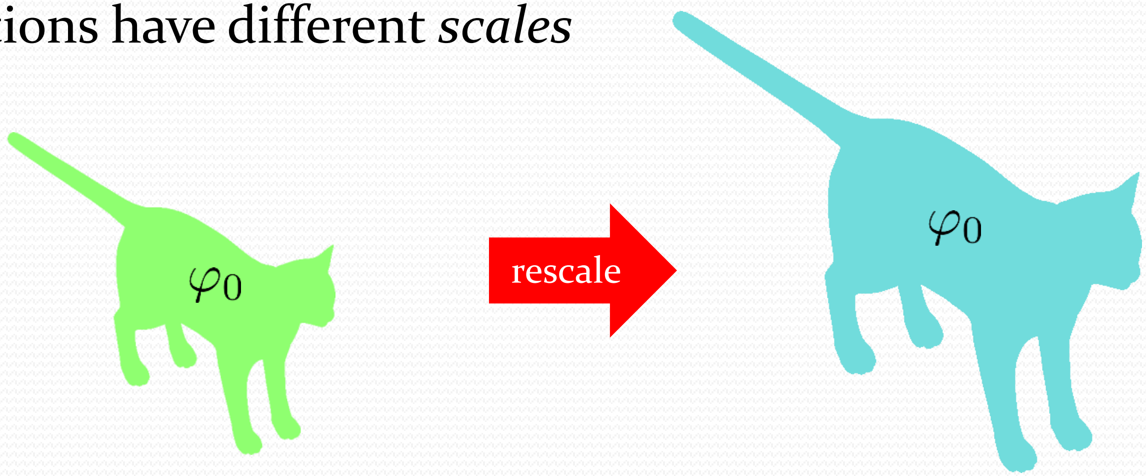

Changes in scale

Observe that for φ_0 we have:

$$\langle \varphi_0, \varphi_0 \rangle_M = \int_S K \cdot K dx = K^2 \int_S dx = K^2 |S|$$

$$\langle \varphi_0, \varphi_0 \rangle_M = 1$$ by orthonormality of $\{\varphi_k\}$

⇩

$$\varphi_0 = K = \frac{1}{\sqrt{|S|}}$$

In general, the scale of the eigenfunctions *depends on the size of the shape*.

This should come as no surprise, remember for instance Weyl's law:

$$-\lambda_j \sim \frac{\pi}{|S|} j \quad \text{for } j \to \infty$$

Rescaling areas

Let us be given a shape S and its scaled version $S' = \alpha S$

$$\text{Area}(\mathbf{x}(U)) = \int_U \sqrt{\det g}\, du dv$$

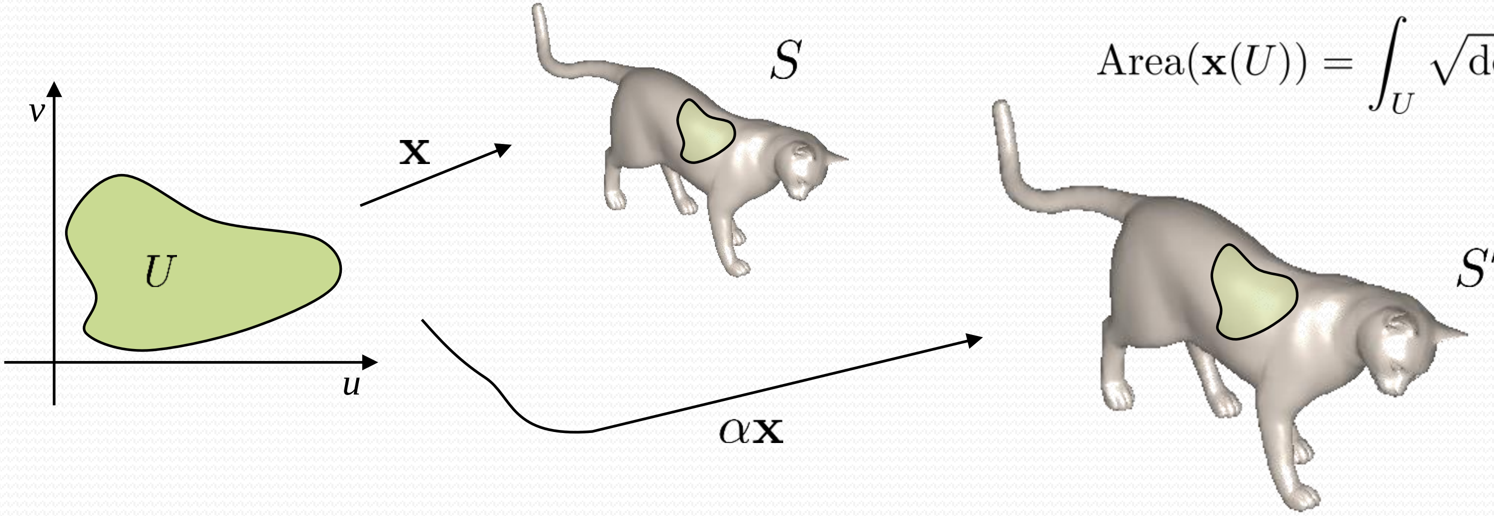

$$g' = \begin{pmatrix} \langle \alpha\mathbf{x}_u, \alpha\mathbf{x}_u \rangle & \langle \alpha\mathbf{x}_u, \alpha\mathbf{x}_v \rangle \\ \langle \alpha\mathbf{x}_v, \alpha\mathbf{x}_u \rangle & \langle \alpha\mathbf{x}_v, \alpha\mathbf{x}_v \rangle \end{pmatrix} = \alpha^2 g$$

$$\text{Area}(\alpha\mathbf{x}(U)) = \int_U \sqrt{\det g'}\, du dv = \int_U \sqrt{\det \alpha^2 g}\, du dv = \int_U \sqrt{\alpha^4 \det g}\, du dv = \alpha^2 \int_U \sqrt{\det g}\, du dv$$

$$= \alpha^2 \text{Area}(\mathbf{x}(U))$$

Rescaling eigenvalues

Let us be given a shape S and its scaled version $S' = \alpha S$, and let us consider the generalized eigenvalue problem for the first shape:

$$C\varphi = \lambda M\varphi$$

For the second shape, we have:

$$S \mapsto \alpha S$$
$$\lambda \mapsto \frac{1}{\alpha^2}\lambda$$

$$C\varphi' = \lambda' M' \varphi' = \lambda'(\alpha^2 M)\varphi' = (\alpha^2 \lambda')M\varphi'$$

$C' \equiv C$ since cotangents do not change with scale

the areas in M scale up with α^2

One could pre-process the shapes by normalizing their eigenvalues. For instance, pick an eigenvalue $\lambda^\star$ and rescale the given shape S as

$$S' = \sqrt{|\lambda^\star|}S$$

for example, choose $\lambda^\star = \max\{|\lambda_k|\}$

Rescaling eigenfunctions

What happens to the eigenfunctions?

Let us have a look at what happens to the *first* (constant) eigenfunction:

$$\langle {\varphi'}_0, {\varphi'}_0 \rangle_{\alpha^2 M} = {K'}^2 \int_{\alpha S} dx = {K'}^2 |\alpha S| = {K'}^2 \alpha^2 |S|$$

$$ {\varphi'}_0 = K' = \frac{1}{\alpha \sqrt{|S|}}$$

It looks like eigenfunctions are rescaled as $\varphi \mapsto \frac{1}{\alpha}\varphi$. We are going to prove this statement for arbitrary eigenfunctions in the following slides.

Rescaling eigenfunctions

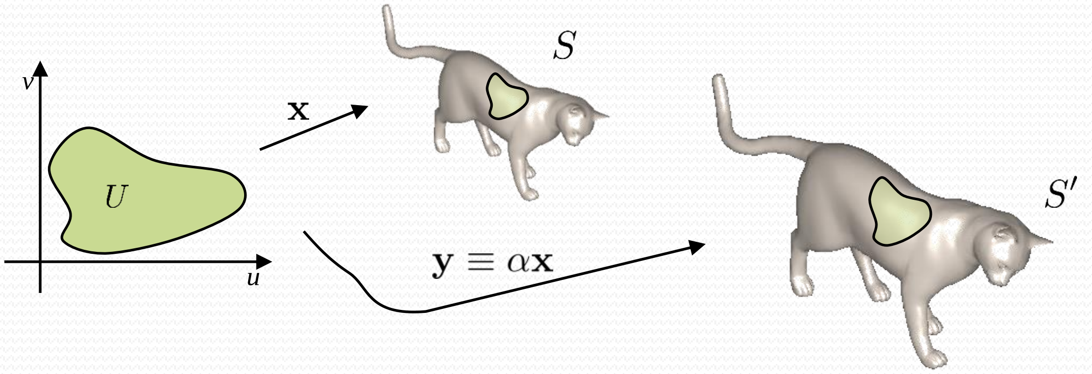

A few slides ago we showed that
$\mathrm{Area}(\alpha\mathbf{x}(U)) = \alpha^2 \mathrm{Area}(\mathbf{x}(U))$

In particular, since $y \equiv \alpha x$, we can compute the area element on *S'* as
$dy = \alpha^2 dx$

Let us consider the generic eigenfunction φ on *S*. How is it transformed by the rescaling $S \mapsto \alpha S$?

$$\varphi'(y) = \kappa \varphi\left(\frac{y}{\alpha}\right) \quad \text{unknown}$$

$$\int_{\alpha S} \varphi'^2(y) dy = \int_{\alpha S} \kappa^2 \varphi^2 \left(\frac{y}{\alpha}\right) dy = \kappa^2 \int_S \varphi^2(x) \alpha^2 dx = \alpha^2 \kappa^2 \int_S \varphi^2(x) dx = \alpha^2 \kappa^2$$

By orthonormality of φ': $\alpha^2 \kappa^2 = 1 \Rightarrow \kappa = \frac{1}{\alpha}$

Since this holds for any eigenfunction, we have proved that $\varphi \mapsto \frac{1}{\alpha} \varphi$

Scale-invariant embedding

These results allow us to act directly at the descriptor level, i.e. when the embedding is performed:

$$p \mapsto (\varphi_0(p), \varphi_1(p), \varphi_2(p), \ldots)$$

⇩

$$p \mapsto \left(\frac{\varphi_0(p)}{\sqrt{|\lambda_0|}}, \frac{\varphi_1(p)}{\sqrt{|\lambda_1|}}, \frac{\varphi_2(p)}{\sqrt{|\lambda_2|}}, \ldots \right)$$

The resulting embedding is **scale-invariant**. Indeed:

$$\frac{\varphi_j(p)}{\sqrt{|\lambda_j|}} \quad \xrightarrow{\text{rescale}} \quad \frac{\varphi'_j(p)}{\sqrt{|\lambda'_j|}} = \frac{\frac{1}{\alpha}\varphi_j(p)}{\sqrt{\frac{|\lambda_j|}{\alpha^2}}} = \frac{\frac{1}{\alpha}\varphi_j(p)}{\frac{1}{\alpha}\sqrt{|\lambda_j|}} = \frac{\varphi_j(p)}{\sqrt{|\lambda_j|}}$$

Global Point Signature

$$p \mapsto \left(\frac{\varphi_1(p)}{\sqrt{|\lambda_1|}}, \frac{\varphi_2(p)}{\sqrt{|\lambda_2|}}, \frac{\varphi_3(p)}{\sqrt{|\lambda_3|}}, \dots \right)$$

This new, scale-invariant embedding defines a descriptor known as the **Global Point Signature** (**GPS**).

The GPS embedding of a shape is an *isometry-invariant* Euclidean embedding. Differently, multi-dimensional scaling was determined only up to rigid motions!

In practice, GPS is truncated to the first *m* eigenfunctions.

<u>Main issues:</u>

1. The signs of eigenvectors are undefined
2. Two eigenvectors may be swapped

Example: Segmentation

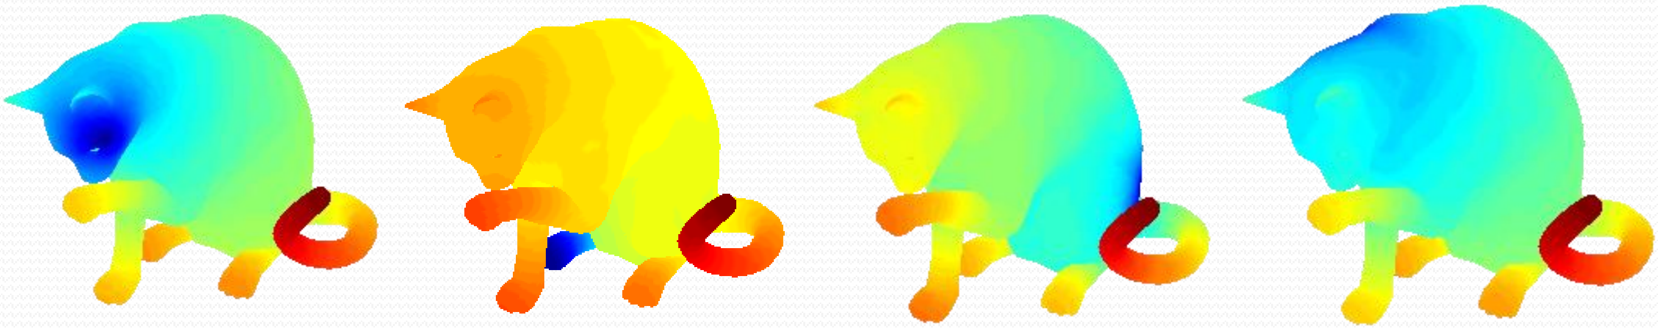

Distance maps (standard Euclidean metric on GPS descriptors) from different source points. Distance goes from blue to red.

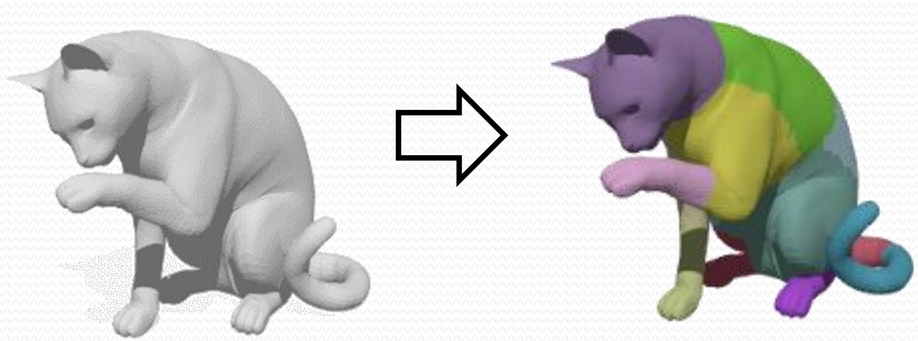

Segmentation obtained via *k*-means clustering of the GPS embedding.

Robust to isometric deformations!

Example: Matching

Recall the Gromov-Hausdorff formulation we gave for the matching problem:

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

all possible correspondences!

We can use descriptors to *reduce* the set of correspondences over which to optimize:

Just consider as good «candidate» matches the ones among points with similar descriptors.

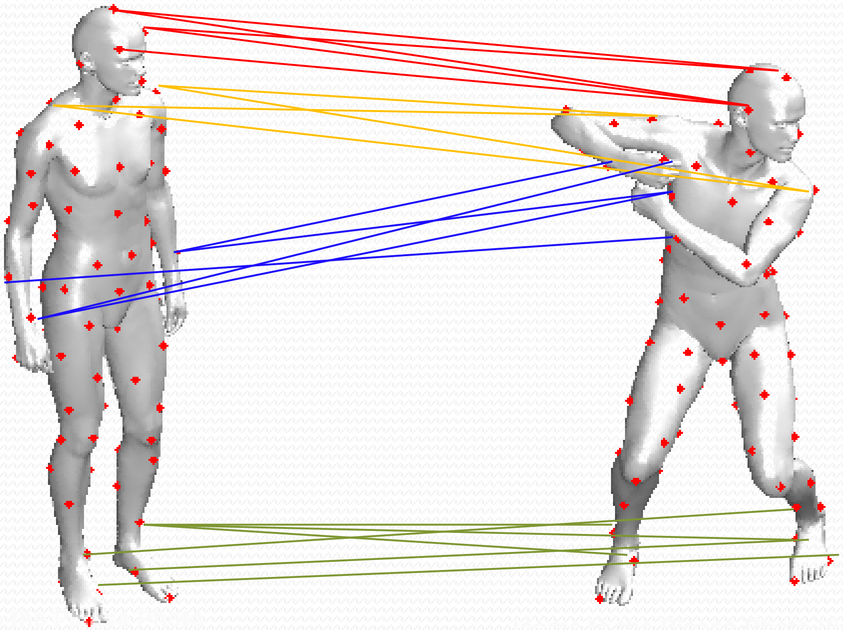

Point descriptors

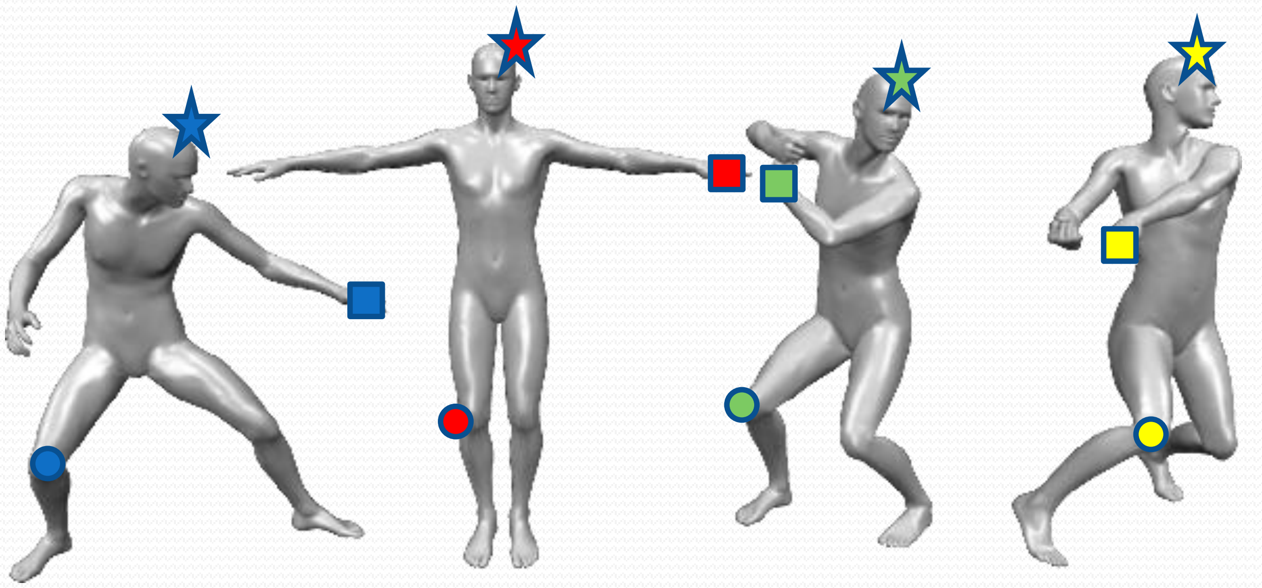

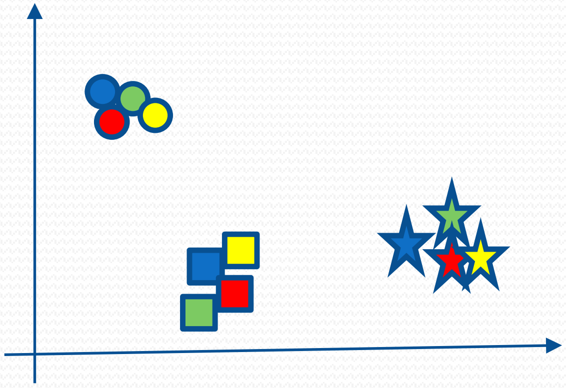

Goal:

- invariant under isometric deformations
- **multiscale** (from local to global)

Multiscale property

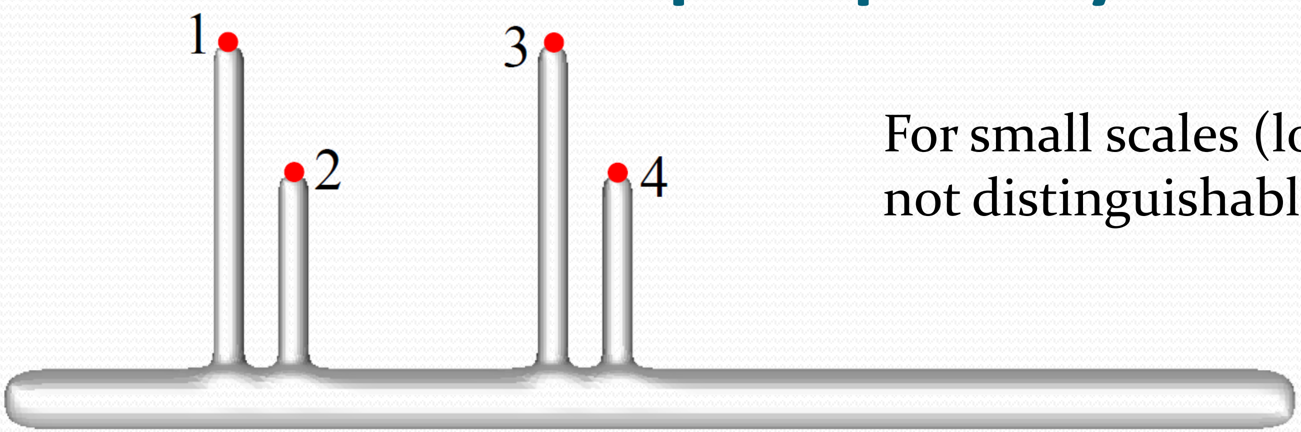

For small scales (locally) 1 and 3 are not distinguishable

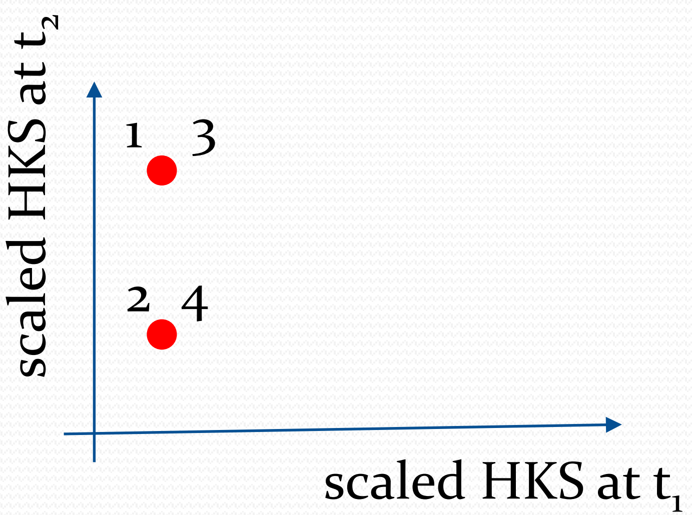

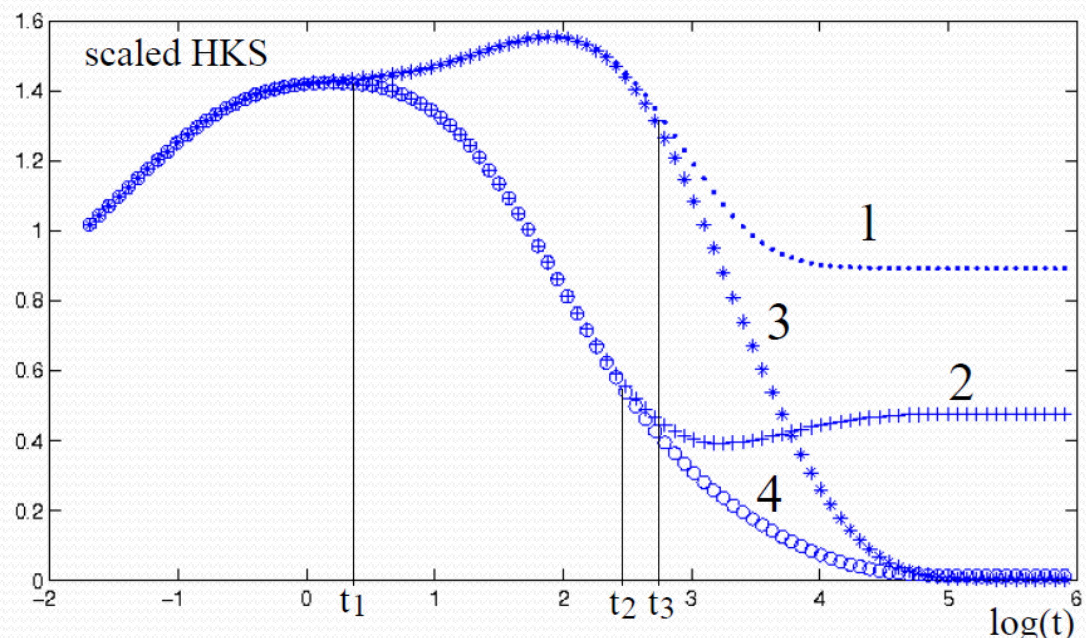

Heat diffusion on surfaces

For a regular surface *S* the diffusion of heat can be described by the heat equation:

$$\frac{\partial u(x,t;u_0)}{\partial t} = \Delta u(x,t;u_0)$$

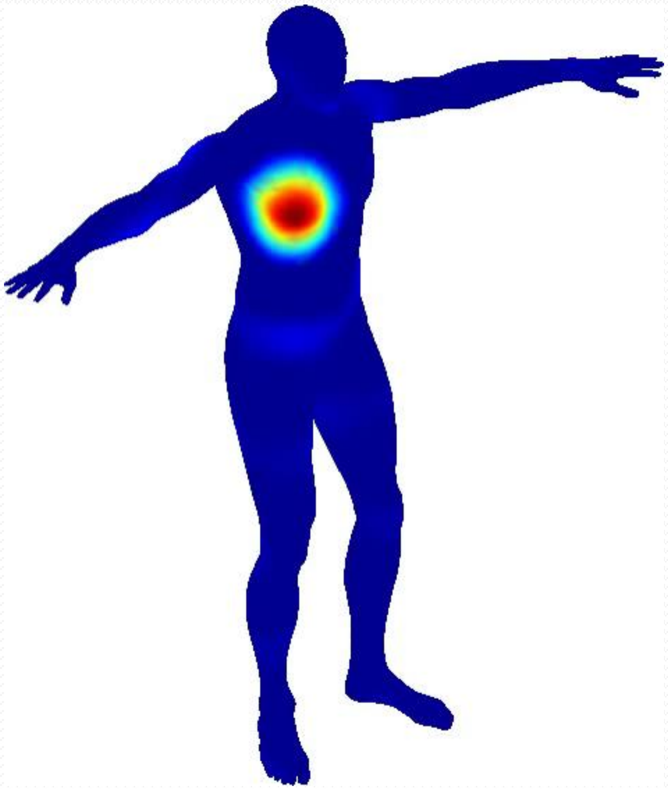

We write $u(x,t;u_0)$ for the amount of heat at point *x* after time *t*, when at time zero the distribution of heat is given by

$$u(x,0) = u_0(x)$$

Heat kernel

A solution to the heat equation is given by

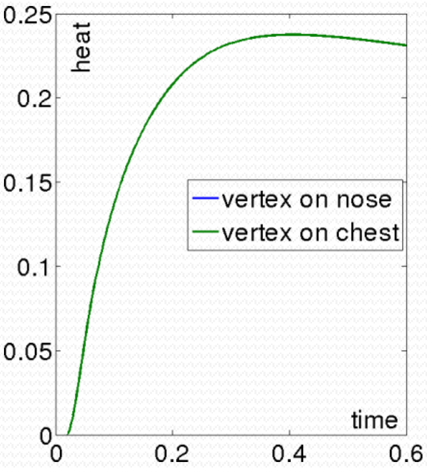

$$u(x, t; u_0) = \int_S k_t(x, y) u_0(y) dy$$

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

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

The dirac-"function" δ_z satisfies

$$\int f(x) \delta_z(x) dx = f(z)$$

k_t is called the *heat kernel* and depends on the geometry of S.

Heat kernel

The heat kernels in $\mathbb{R}^n$ are given by:

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

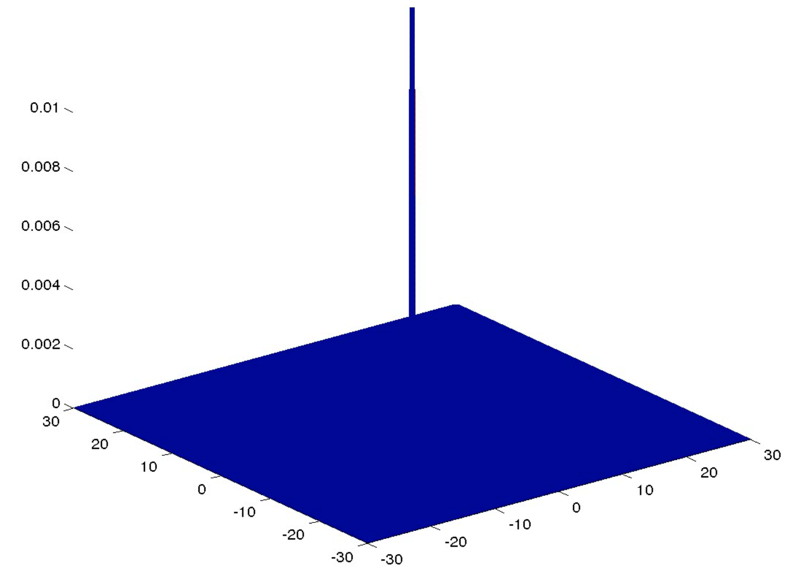

The dirac-"function" δ_z can be seen as

$$\lim_{t\to 0} k_t^S(z,\cdot)$$

One can recover the geodesic distances on general surfaces from the respective heat kernel:

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

Informative property

A surjective $T : S \to S'$ is an isometry iff $k_t^S(x,y) = k_t^{S'}(T(x),T(y))$

Solving the heat equation

We know that the eigenfunctions $\{\psi_k(x)\}_{k=0}^{\infty}$ of the Laplace-Beltrami operator form a basis, thus for every *t* we can write

$$u(t, x; u_0) = \sum_{k=0}^{\infty} c_k(t)\psi_k(x) \approx \sum_{k=0}^{m} c_k(t)\psi_k(x)$$

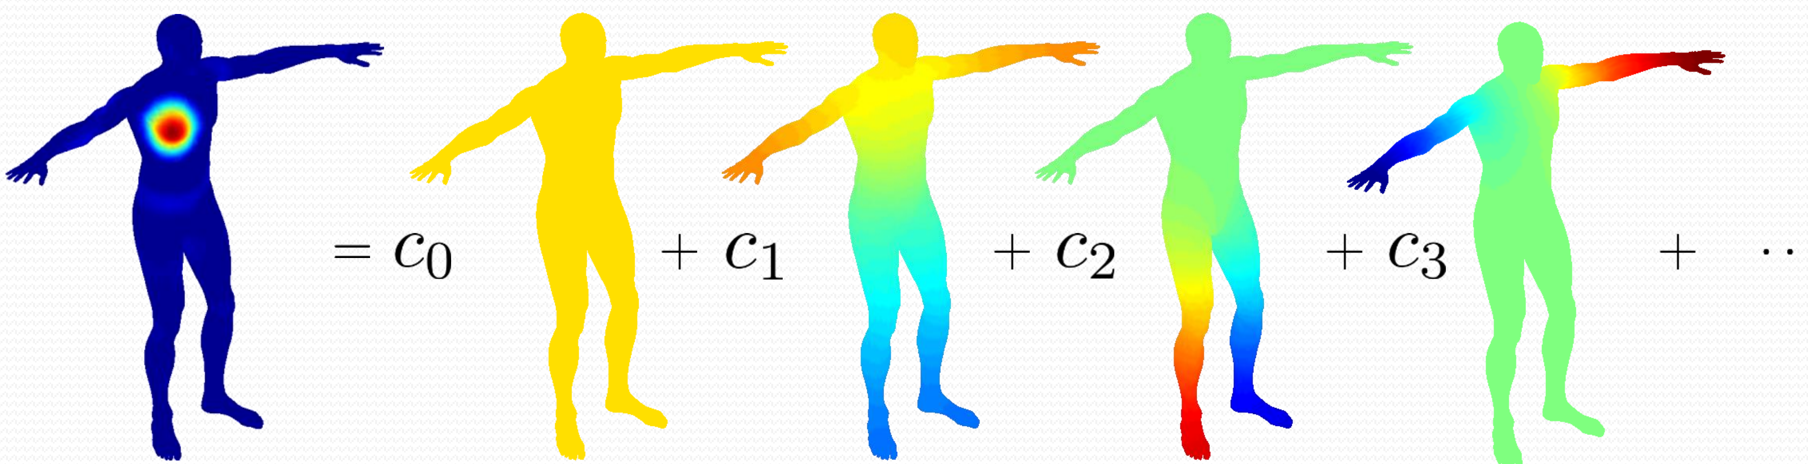

Solving the heat equation

$$u(t,x;u_0) = \sum_{k=0}^{\infty} c_k(t)\psi_k(x) = \sum_{k=0}^{\infty} d_k e^{\lambda_k t}\psi_k(x)$$

$$\Delta u(t,x;u_0) = \sum_{k=0}^{\infty} c_k(t)\Delta\psi_k(x) = \sum_{k=0}^{\infty} c_k(t)\lambda_k\psi_k(x)$$

$$\frac{\partial u(t,x;u_0)}{\partial t} = \sum_{k=0}^{\infty} \dot{c}_k(t)\psi_k(x) \qquad \Rightarrow c_k(t) = d_k e^{\lambda_k t}$$

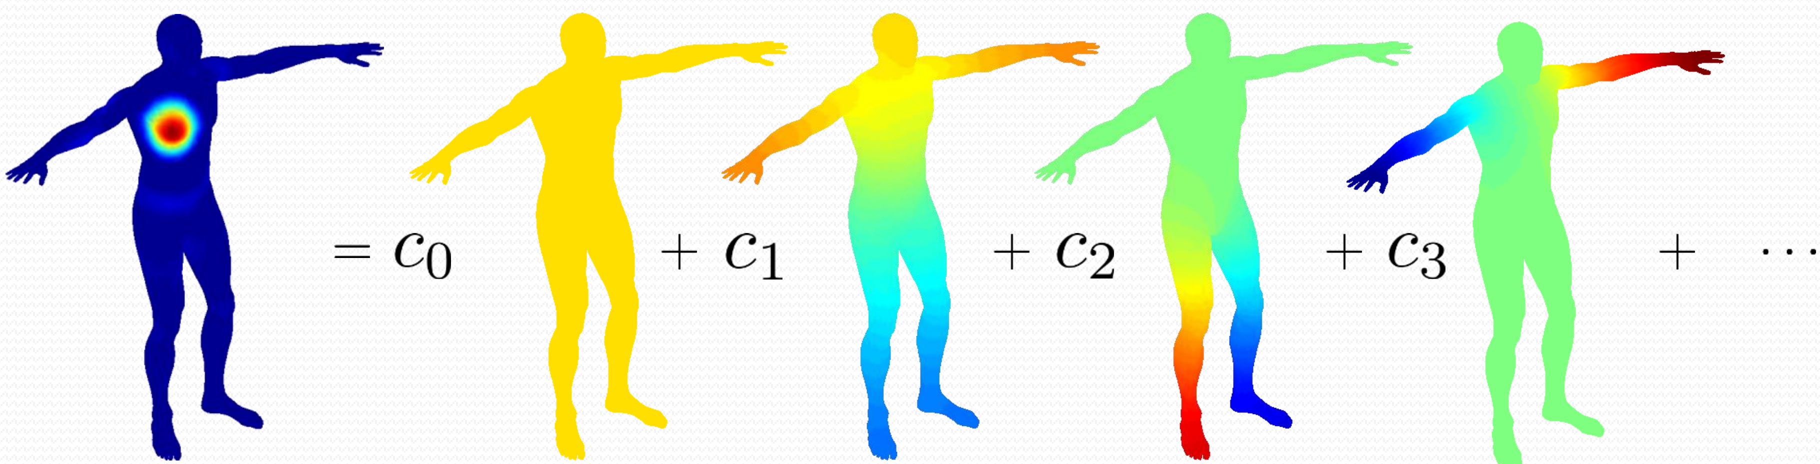

Solving the heat equation

$$u(t, x; u_0) = \sum_{k=0}^{\infty} d_k e^{\lambda_k t} \psi_k(x)$$

$$u(0, x; u_0) = \sum_{k=0}^{\infty} d_k \psi_k(x) \overset{!}{=} u_0(x)$$

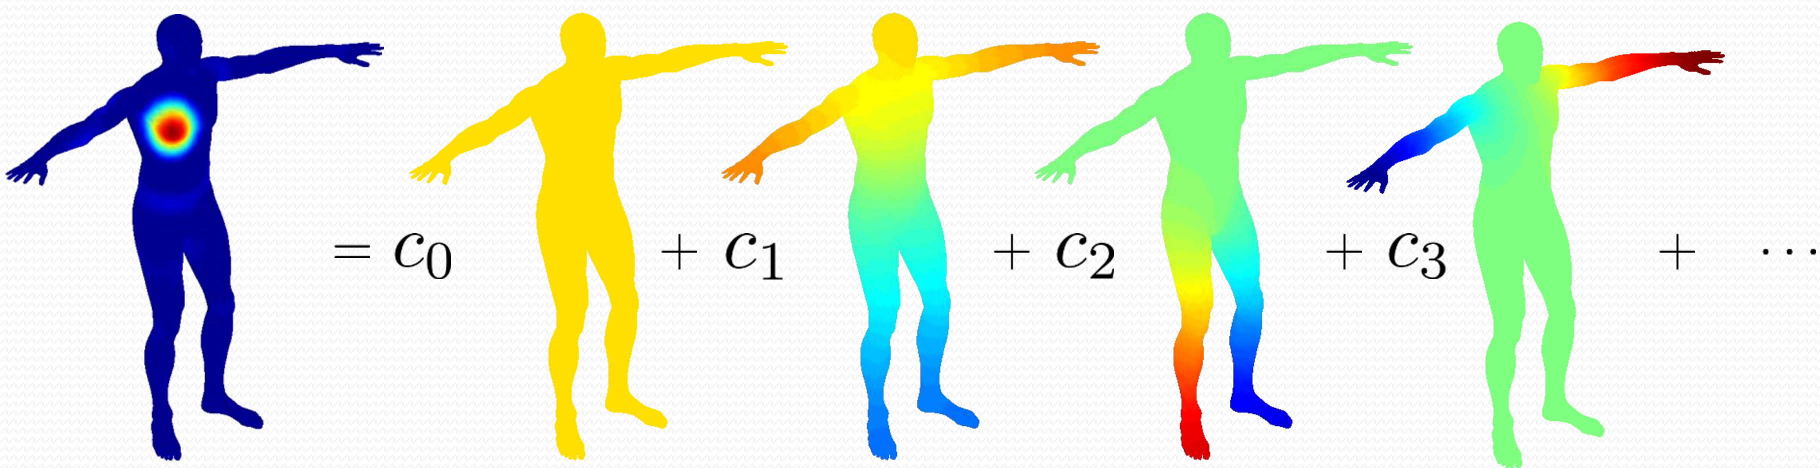

Heat diffusion on discrete surface

In the discrete setting functions are represented as vectors and linear operators as matrices. In particular we can collect all the eigenvectors of the Laplacian matrix $L = M^{-1}C$ in a matrix

$$\Psi = \begin{pmatrix} | & & | \\ \psi_0 & \dots & \psi_{n-1} \\ | & & | \end{pmatrix}$$

and write

$$u(0,\cdot;u_0) = \sum_{k=0}^{n-1} d_k \psi_k = \Psi \begin{pmatrix} d_0 \\ \vdots \\ d_{n-1} \end{pmatrix} = \begin{pmatrix} | \\ u_0 \\ | \end{pmatrix} \qquad \Rightarrow d = \Psi^{-1} \begin{pmatrix} | \\ u_0 \\ | \end{pmatrix}$$

Note that due to our discretization of the Laplacian Ψ is not necessarily orthogonal!

Heat kernel signature

Informative property
A surjective $T : S \to S'$ is an isometry iff $k_t^S(x, y) = k_t^{S'}(T(x), T(y))$

One could use the functions $k_t(x, \cdot) : S \times \mathbb{R}_{>0} \to \mathbb{R}_{>0}$ as descriptors.

But

- to compare two heat kernels at points on different shapes, one needs information about the ordering of the vertices
- a lot of information in the heat kernel is redundant

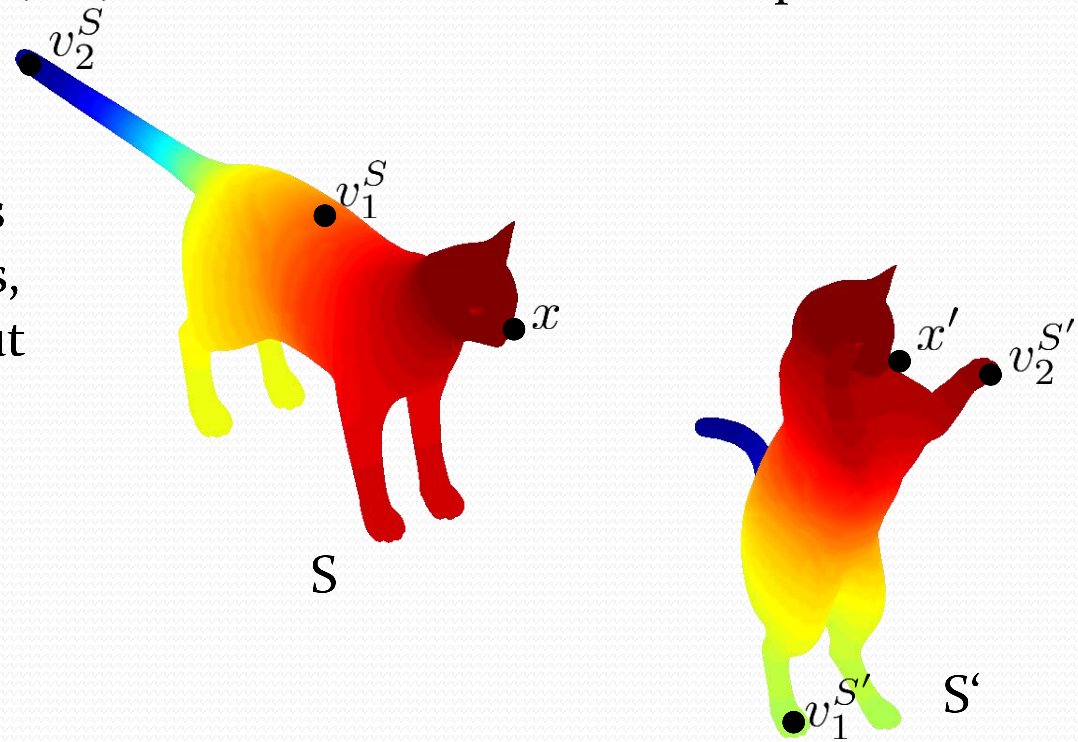

Heat kernel signature

Given a point *x* on a surface its Heat Kernel Signature *HKS(x)* is a function over the temporal domain $HKS(x) : \mathbb{R}_{>0} \to \mathbb{R}_{>0}$:

$$HKS(x,t) = k_t(x,x) = u(t,x;\delta_x)$$
$$= \sum_{k=0}^{n-1} e^{\lambda_k t} \psi_k^2(x)$$

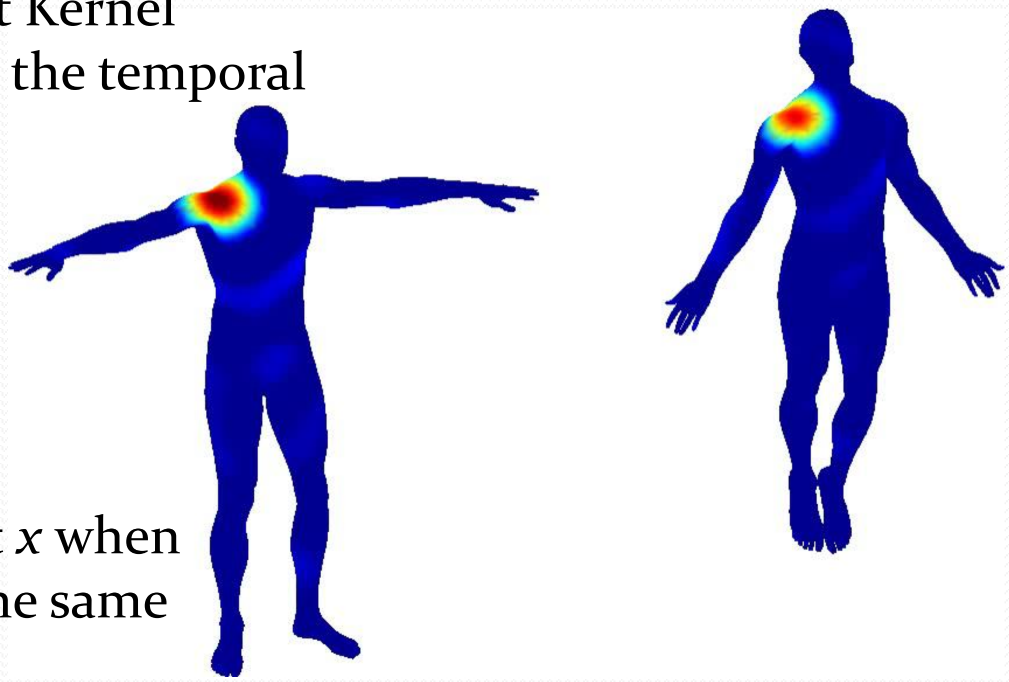

The amount of heat staying at point *x* when starting with a unit heat source at the same point.

Informative property
If the eigenvalues of the Laplacian of *S* and *S*' are not repeated then a homeomorphism $T : S \to S'$ is an isometry iff $k_t^S(x,x) = k_t^{S'}(T(x), T(x))$.

Heat kernel signature

Similarly to GPS, in practice HKS is truncated to the first *m* eigenfunctions.

$$HKS(x,t) = \sum_{k=0}^{m-1} e^{\lambda_k t} \psi_k^2(x)$$

<u>Main issues of GPS:</u>

1. The signs of eigenvectors are undefined ⇨ We are now squaring!
2. Two eigenvectors may be swapped ⇨ We are taking sums!

Heat kernel signature

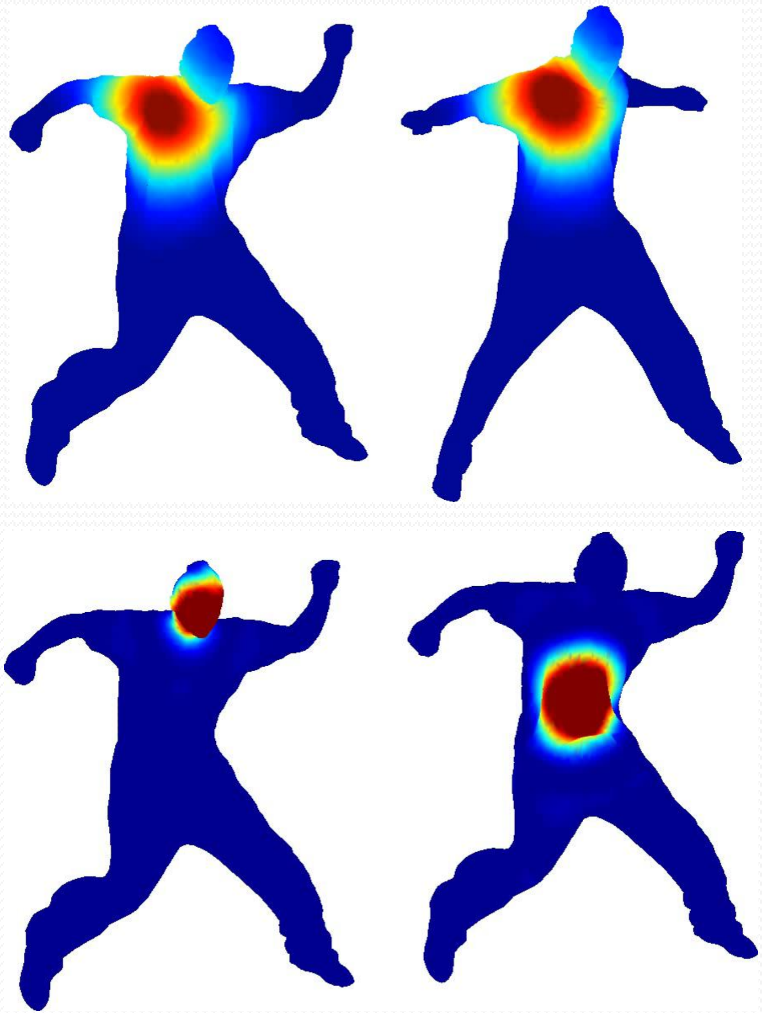

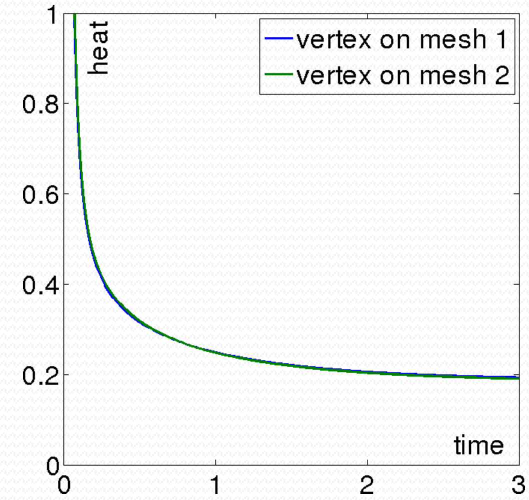

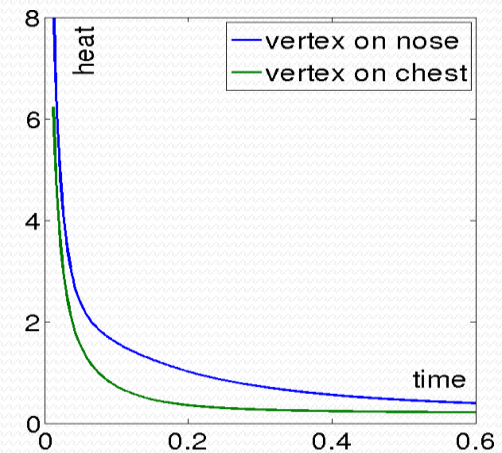

Multiscale property

$\{k_t(x,x)\}$ encodes information about neighborhood in a multiscale way

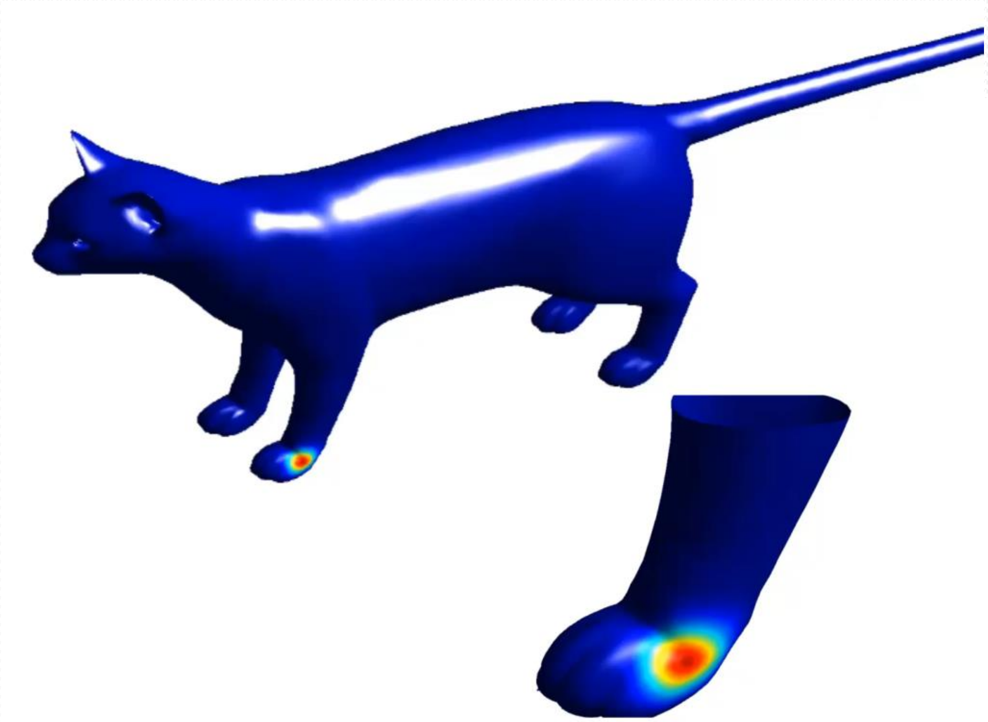

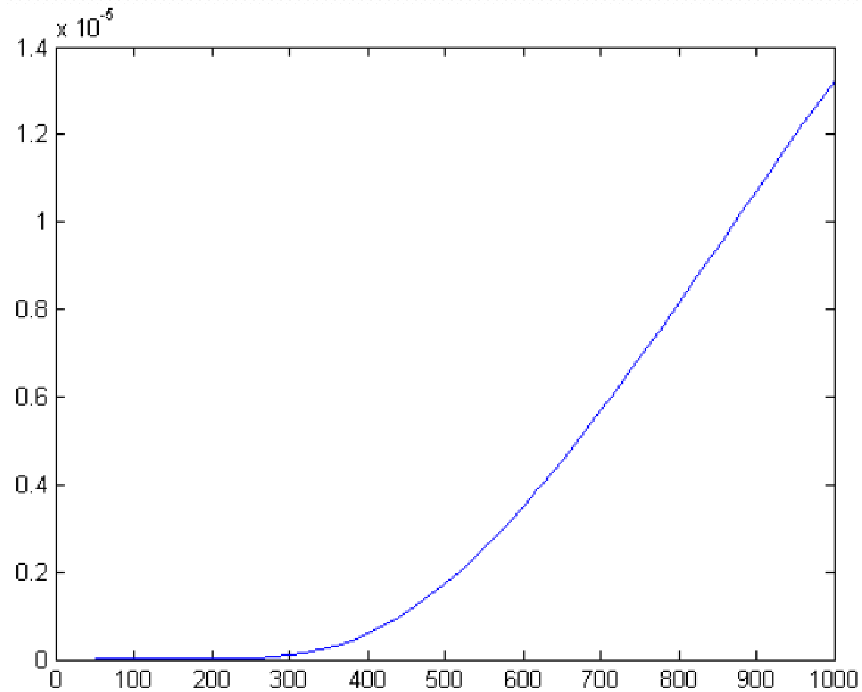

Scaled Heat Kernel Signature

Difference $|k_t(x,x) - k_t(x',x')|$ decreases exponentially as t increases.

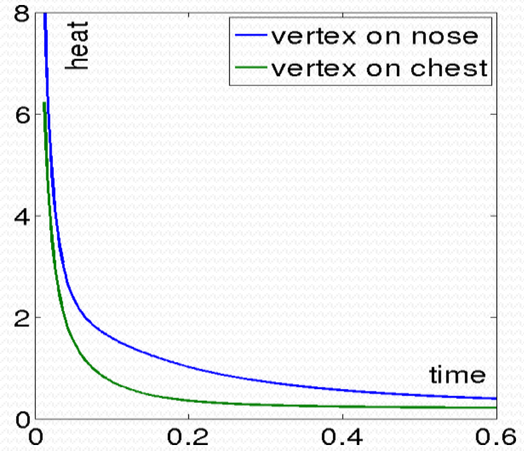

Large scales have minor influence.

Workaround
Consider scaled heat kernel signatures

$$sHKS(x,t) = \frac{k_t(x,x)}{\int_S k_t(y,y)dy}$$

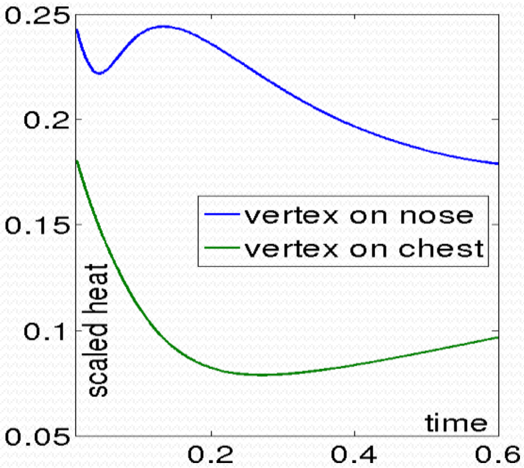

Differences between two signatures at different time scales contribute approximately equally.

Heat kernel signature

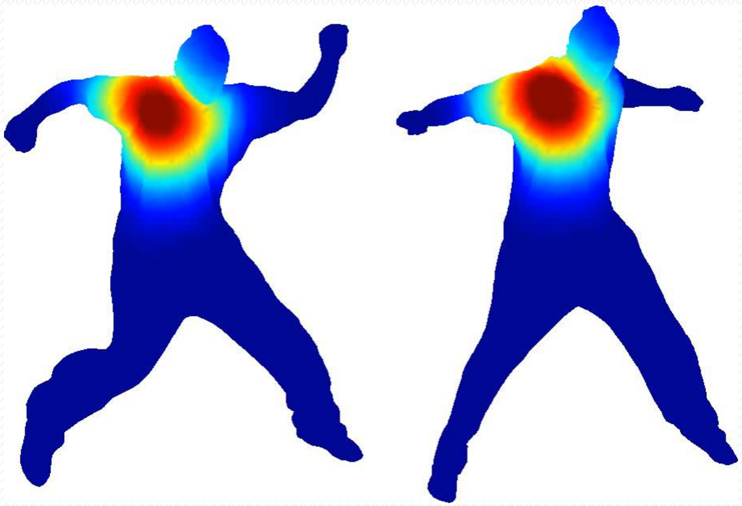

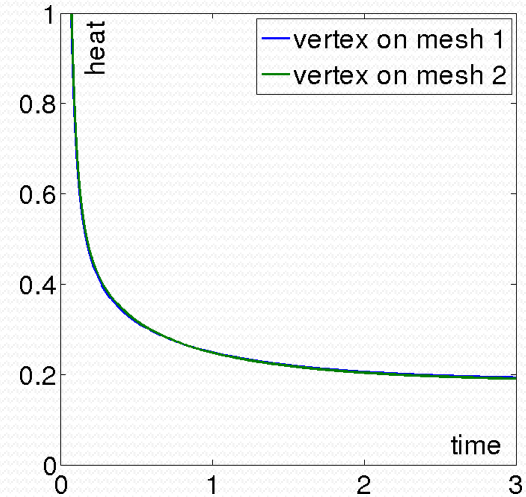

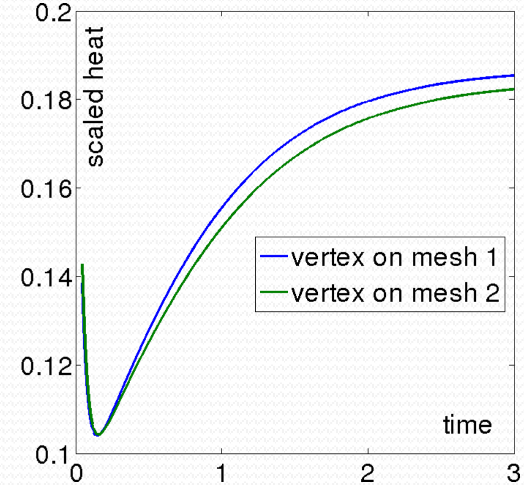

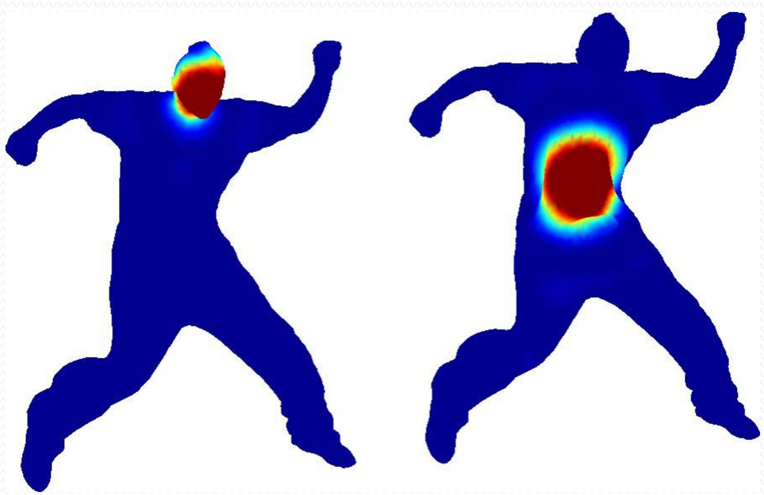

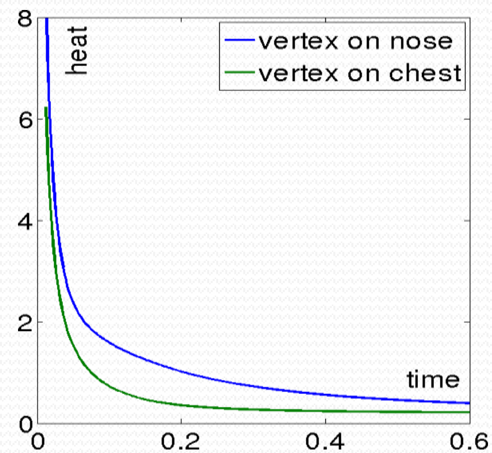

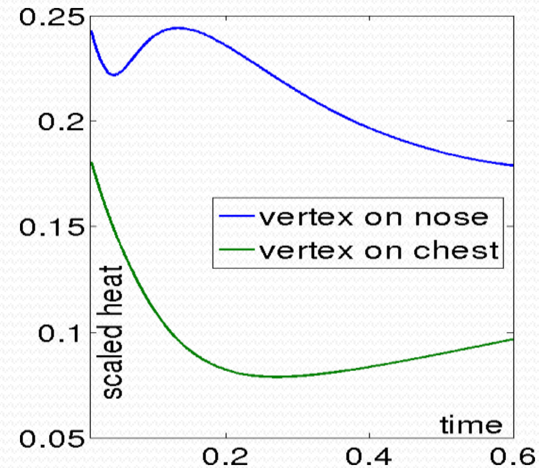

Distance between signatures

$$d^2_{[t_1,t_2]}(x,x') = \int_{t_1}^{t_2} (sHKS(x,t) - sHKS(x',t))^2 d\log t$$

sample *sHKS* uniformly over the logarithmic scaled temporal domain

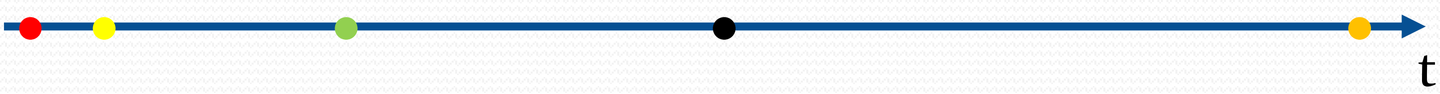

$$d^2_{[t_1,t_2]}(x,x') = \left\| \begin{pmatrix} sHKS(x,t_1) \\ \ldots \\ sHKS(x,t_k) \end{pmatrix} - \begin{pmatrix} sHKS(x',t_1) \\ \ldots \\ sHKS(x',t_k) \end{pmatrix} \right\|$$

Example: Feature detection

Feature points can be selected as the local maxima of $k_t(x,x)$ for a fixed *t*, or as the *persistent* maxima across different time steps.

Suggested reading

- *Laplace-Beltrami eigenfunctions for deformation invariant shape representation*. Rustamov. Proc. SGP 2007.
- *A concise and provably informative multi-scale signature based on heat diffusion*. Sun, Ovsjanikov, Guibas. Proc. SGP 2009.