CS 468 Lecture 16: Isometry Invariance and Spectral Techniques

Justin Solomon

Scribe: Evan Gawlik

Introduction. In geometry processing, it is often desirable to characterize the shape of an object in a manner that is invariant to isometries – deformations of the object that involve bending without stretching, thereby leaving intrinsic distances undisturbed. Examples where such characterizations are useful include segmentation, symmetry detection, recognition, retrieval, feature extraction, and alignment.

This lecture introduces the mathematical definition of an isometry and describes several *shape descriptors* that are used to characterize geometries in an isometry-invariant manner.

Isometry. Let (X, d_1) and (Y, d_2) be metric spaces. A map $f : X \to Y$ is a *global isometry* if

$$d_1(x, y) = d_2(f(x), f(y))$$

for every $x \in X$ and $y \in Y$.

A related concept applies to the case in which X and Y are Riemannian manifolds with metrics g_1 and g_2, and f is a diffeomorphism. The map f is said to be a *local isometry* if

$$g_1(v, w) = g_2(f_* v, f_* w)$$

for every pair of vector fields v and w on X. Here, $f_* : TX \to TY$ denotes the push-forward.

Shape Descriptors. A *shape descriptor* is an assignment of a real number of tuple of real numbers $h(x) \in \mathbb{R}^n$ to each point x on a surface $S \subset \mathbb{R}^3$, designed in such a way that the tuple stored at each location characterizes the local geometry of the surface and describes the point's "role" on the surface. We have already seen examples of shape descriptors: the Gaussian curvature $K(x) = \kappa_1(x)\kappa_2(x)$ and the mean curvature $H(x) = \kappa_1(x) + \kappa_2(x)$ are two such examples.

Several aims should be kept in mind when designing a good shape descriptor. Clearly, $h(x)$ should provide useful information about the point x. It should be robust against noise in the triangulation and against small deformations, and it should be intrinsic – that is, independent of the manner in which S is embedded in $\mathbb{R}^3$. Finally, it should be invariant under rigid motions and other isometries.

The Hodge Laplacian

$$\Delta = \delta d + d\delta,$$

being an intrinsic operator, provides a useful starting point for the design of many shape descriptors. To intuit its intrinsic nature, note, for example, that solutions to the heat equation

$$u_t = \Delta u$$

on a surface S are unaltered by isometries.

Global Point Signature. An example of a shape descriptor that relies on the Hodge Laplacian is the *Global Point Signature* (GPS). This shape descriptor assigns to each point $x \in S$ the sequence of real numbers

$$GPS(x) = \left(\lambda_1^{-1/2}\phi_1(x), \lambda_2^{-1/2}\phi_2(x), \lambda_3^{-1/2}\phi_3(x), \dots\right),$$

where λ_i are the eigenvalues of Δ and ϕ_i are the corresponding eigenfunctions. Being derived solely from the Laplacian, the GPS is invariant under isometries of S.

Let us also note that the GPS, viewed as a map from S to the space of sequences of real numbers, is injective, provided the surface S does not self-intersect. This follows from the fact the the eigenfunctions $\phi_1, \phi_2, \dots$ form a basis for the space of smooth functions on S. Abstractly, one can think of the image of this map as a surface in infinite-dimensional Euclidean space; injectivity implies that this surface does not self-intersect whenever S does not self-intersect.

The GPS suffers from a few drawbacks. It assumes that the eigenvalues of Δ are unique, and can give rise to abrupt changes in GPS values when a small deformation of the surface leads to a reordering of eigenvalues. Finally, it is a nonlocal feature since the eigenfunctions of the Laplacian generally have global support.

Heat Kernel Signature and Wave Kernel Signature. Two other popular shape descriptors that derive from the Laplacian are the *Heat Kernel Signature* (HKS) and *Wave Kernel Signature* (WKS).

To define the HKS, let $k_t(x, y)$ denote the fundamental solution to the heat equation

$$u_t = \Delta u$$

on S. That is, $k_t(x, y)$ is the value of the solution to the heat equation at time t and position $x \in S$, assuming the initial condition is given by a delta function centered at $y \in S$. In terms of the eigenfunctions ϕ_i and eigenvalues λ_i of $-\Delta$,

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

Fixing a time t, the HKS is then defined as

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

In words, the Heat Kernel Signature measures the amount of heat left at x after t units of time have transpired, assuming the initial heat distribution was concentrated at x.

An example of the HKS at four points on a triangulated surface is shown in Fig. 1. For short times, the four points have nearly identical heat kernel signatures $k_t(x, x)$ since the local geometry (the tips of the dragon's feet) is roughly the same. At a later time, the heat kernel signatures $k_t(x, x)$ capture more global information about the surface's shape and diverge. In this sense, the HKS is a multiscale shape descriptor.

The Wave Kernel Signature (WKS) is a shape descriptor of a similar nature, except that it is based upon solutions to the Schrodinger wave equation

$$u_{tt} = -i\Delta u$$

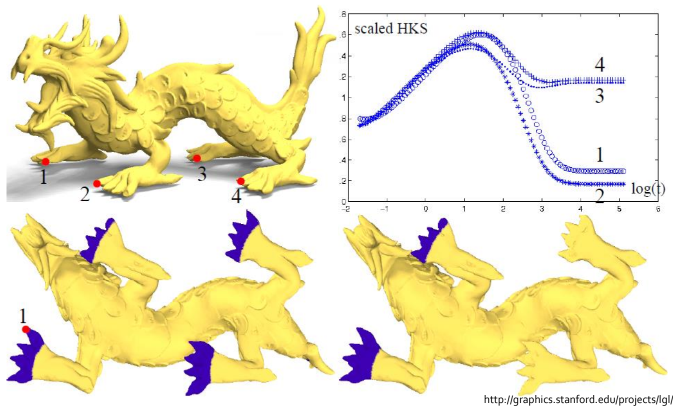

Figure 1: Heat kernel signature.

Upon selecting a family of initial energy distributions $f_E(\lambda)$, $E = 1, 2, \dots$, the WKS is defined as

$$WKS(x) = \left(\sum_i \phi_i(x)^2 f_1(\lambda_i)^2, \sum_i \phi_i(x)^2 f_2(\lambda_i)^2, \sum_i \phi_i(x)^2 f_3(\lambda_i)^2, \dots \right)$$

The entries of this vector correspond to the long-time averages of the squared solution to the Schrodinger wave equation at position x, given the initial energy distributions f_E.

The HKS and WKS have similar advantages and disadvantages. Both are isometry-invariant, easy to compute, and do not suffer from the danger of eigenvalue "switching" under small deformations that we observed for the GPS. Repeated eigenvalues are still an issue, however, and the WKS can sometimes filter out large-scale features that might be worth retaining.

Shape descriptors like those discussed above have applications in a variety of contexts, including feature extraction, correspondence between surfaces, matching surfaces, and detecting discrete symmetries.

Continuous symmetries Much of the machinery developed above is useful for detecting discrete symmetries, such as symmetries under reflection about an axis. A related notion is that of a continuous symmetry – e.g., rotations and translations that leave the geometry invariant.

A *Killing vector field* is a vector field V along which the metric is invariant. Informally, distances between nearby points do not change when transported along the flow of the vector field V.

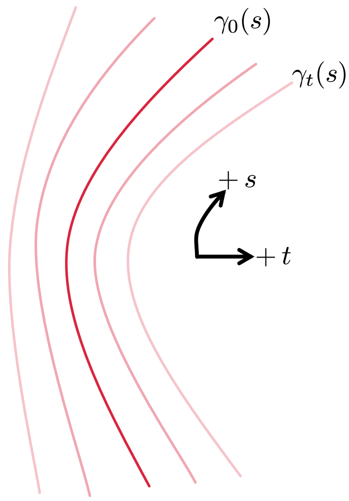

Figure 2: Family of curves $\gamma_t(s)$.

To determine the conditions under which a vector field V qualifies as a Killing vector field, let γ_0 be curve on S parametrized by arclength s, and let $\gamma_t(s)$ denote the location of $\gamma_0(s)$ after being transported along the flow of V by t units of time, as in Fig. 2. If distances are preserved, then the parameter s represents arclength along the deformed curve γ_t; hence

$$\|\gamma_0'(s)\| = \|\gamma_t'(s)\| = 1$$

for every t. Differentiating with respect to time and using the symmetry of mixed partials gives

$$\begin{aligned} 0 &= \frac{\partial}{\partial t}\|\gamma_t'(s)\| \\ &= \frac{1}{\|\gamma_t'(s)\|}\left\langle \gamma_t'(s), \frac{\partial}{\partial t}\gamma_t'(s)\right\rangle \\ &= \left\langle T(s), \frac{\partial}{\partial t}\gamma_t'(s)\right\rangle \\ &= \left\langle T(s), \frac{\partial}{\partial s}\frac{\partial}{\partial t}\gamma_t(s)\right\rangle \end{aligned}$$

Now since $V(s) = \frac{\partial}{\partial t}\gamma_t(s)$, we obtain

$$0 = \left\langle T(s), \frac{\partial}{\partial s}V(s)\right\rangle$$

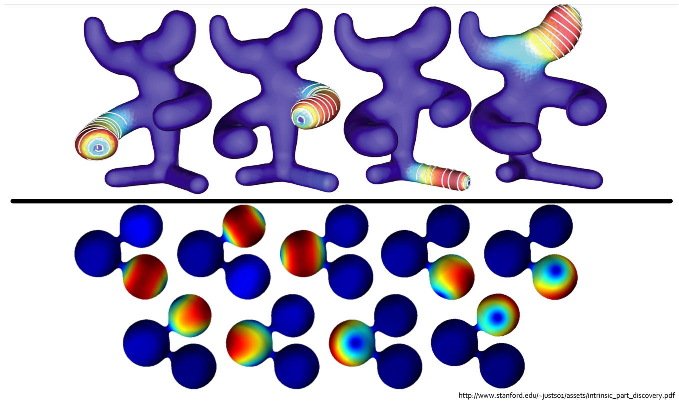

Figure 3: Approximate Killing vector fields.

Equivalently,

$$0 = \langle T, D_T V \rangle \,.$$

Finally, since $D_T V$ has the same component in the T direction as $\nabla_T V$ (the covariant derivative of V in the direction T) we conclude that

$$0 = \langle T, \nabla_T V \rangle \,. \tag{1}$$

For V to be a Killing vector field, this relation must hold for every vector field T on S.

Except on surfaces with high degrees of symmetry, it is often not possible to find a vector field V for which (1) holds exactly for every T. Instead, one can find approximate Killing vector fields via a least squares approach. Denoting

$$PV = \langle T, \nabla_T V \rangle \,,$$

we seek a V which minimizes the *Killing energy*

$$\int_S \|PV\|^2.$$

Some examples of approximate Killing vector fields are shown in Fig. 3.