

Equi-affine Invariant Geometries of Articulated Objects

Dan Raviv, Alexander M. Bronstein, Michael M. Bronstein,
Ron Kimmel, and Nir Sochen

Technion, Computer Science Department, Israel
Tel Aviv University, School of Electrical Engineering, Israel
Università della Svizzera Italiana, Faculty of Informatics, Switzerland
Technion, Computer Science Department, Israel
Tel Aviv University, Department of Applied Mathematics

Abstract. We introduce an (equi-)affine invariant geometric structure by which surfaces that go through squeeze and shear transformations can still be properly analyzed. The definition of an affine invariant metric enables us to evaluate a new form of geodesic distances and to construct an invariant Laplacian from which local and global diffusion geometry is constructed. Applications of the proposed framework demonstrate its power in generalizing and enriching the existing set of tools for shape analysis.

1 Introduction

Shape analysis has been one of the principal research fields in computer vision for many years. Numerous methods are based on modeling shapes as Riemannian manifolds, from which it is possible to derive many geometric invariances. Differential geometry and diffusion geometry have been bold players in this growing field. Schwartz *et al.* [22] proposed to embed a non-rigid shape in an Euclidean domain both conformal and isometric, followed by Elad *et al.* [14] that discussed embeddings in higher dimensions, and presented a practical representation of shapes referred to as *canonical forms*. Later on Elad *et al.* [13] and Bronstein *et al.* [5] showed that for some surfaces, such as faces, a spherical domain better captures intrinsic properties. In 2005 Memoli *et al.* [17] pointed the importance of *Gromov-Hausdorff* distance for shape analysis, followed by Bronstein *et al.* [6] who introduced a variational framework that minimizes the *Gromov-Hausdorff* distance by a direct embedding between two non-rigid shapes which does not suffer from an unbounded distortion of an intermediate ambient space. Diffusion geometry, referred to as spectral geometry, based on heat diffusion on manifolds and the properties of the Laplace Beltrami operator have become growingly popular in shape analysis in the past years. Driving inspiration from Berard *et al.* 1994 work [2], Lafon *et al.* [10] proposed in 2006 a probabilistic analysis of algorithms using graph Laplacians. In 2007, Rustamov [21] showed how shapes can be analyzed using the eigen-functions of the Laplace Beltrami operator, and later on Gebal *et al.* [15] discussed auto diffusion functions. Sun *et al.* [24] used

the decay of heat as a feature, known as *Heat Kernel Signatures*, which was further used by [18] as volumetric descriptors. Diffusion geometric constructs in general were found to be more robust than their geodesic counterparts [7], hence they have found successful applications in many shape analysis tasks, such as [19].

However, all of these constructions depend on the definition of the Riemannian metric tensor. So far, the default choice of the metric induced by the Euclidean embedding of the shape has been used. Such a metric and all the related constructions is invariant to inelastic deformations of the shape and global Euclidean transformations (rotations, reflections and translations). In this paper, we show a different construction of a metric that has a wider class of invariance, being also invariant to equi-affine transformations. It contains the metric evaluation we presented in [29] and [30] for both diffusion and differential geometry.

The rest of the paper is organized as follows. In Section 2 we provide the mathematical background of Euclidean and diffusion geometry, followed by Section 3 where we elaborate on the equi-affine metric. Section 4 is dedicated to numerical aspects, and several applications are presented in Section 5. We conclude the paper in Section 6.

2 Mathematical Background

2.1 Differential Geometry

We model a surface (X, g) as a compact complete two dimensional Riemannian manifold X with a metric tensor g , evaluated on the tangent plane $T_x X$ of point x in the natural basis using the inner product $\langle \cdot, \cdot \rangle_x : T_x X \times T_x X \rightarrow \mathbb{R}$. We further assume that X is embedded into $\mathbb{E} = \mathbb{R}^3$ by means of a regular map $\mathbf{x} : U \subseteq \mathbb{R}^2 \rightarrow \mathbb{R}^3$, so that the metric tensor can be expressed in coordinates as

$$g_{ij} = \left\langle \frac{\partial \mathbf{x}}{\partial u_i}, \frac{\partial \mathbf{x}}{\partial u_j} \right\rangle, \quad (1)$$

where the u_i 's are the coordinates of U , which yields the infinitesimal displacement dp

$$dp^2 = g_{11}du_1^2 + 2g_{12}du_1du_2 + g_{22}du_2^2. \quad (2)$$

Minimal geodesics, or shortest paths, are the minimizers of all path length

$$d_X(x, x') = \min_{C \in \Gamma(x, x')} \ell(C) \quad (3)$$

over the set of all admissible paths $\Gamma(x, x')$ between the points x and x' on the surface X , where due to completeness assumption, a minimizer always exists (not necessary unique). Many algorithms have been proposed for the computation of geodesic distances. They differ by accuracy and complexity. In this paper we focus on the family of algorithms simulating wavefront propagation known as *fast marching methods* [16].

2.2 Differential Operators

Laplace Beltrami operator (LBO), named after Eugenio Beltrami, is the generalization of the Laplace operator. It is a linear operator, defined as the divergence of the gradient of a scalar function $f : X \rightarrow \mathbb{R}$ on a manifold

$$\Delta_g f = \operatorname{div}_g \operatorname{grad}_g f. \tag{4}$$

The operator can be extended to tensors, but it is beyond the scope of this note.

In local coordinates u of a chart [11], the LBO assumes the form of

$$\Delta_g f = \frac{1}{\sqrt{|g|}} \frac{\partial}{\partial u^\alpha} \left(\sqrt{|g|} g^{\alpha\beta} \frac{\partial}{\partial u^\beta} f \right), \tag{5}$$

where $X(u^1, u^2, \dots, u^n) = (X^1, X^2, \dots, X^n)$ is the embedding of an n -dimensional manifold. Since our focus will be two dimensional affine invariants, we constrain ourself to two dimensions

$$X(u^1, u^2) = (x(u^1, u^2), y(u^1, u^2), z(u^1, u^2)). \tag{6}$$

2.3 Diffusion Geometry

The Laplace-Beltrami operator gives rise to the partial differential equation

$$\left(\frac{\partial}{\partial t} + \Delta_g \right) f(t, x) = 0, \tag{7}$$

called the *heat equation*. The heat equation describes the propagation of heat on the surface and its solution $f(t, x)$ is the heat distribution at a point x in time t . The initial condition of the equation is some initial heat distribution $f(0, x)$; if X has a boundary, appropriate boundary conditions must be added. The solution of (7) corresponding to a point initial condition $f(0, x) = \delta(x - x')$, is called the *heat kernel* and represents the amount of heat transferred from x to x' in time t by the diffusion process. Using spectral decomposition, the heat kernel can be represented as

$$h_t(x, x') = \sum_{i \geq 0} e^{-\lambda_i t} \phi_i(x) \phi_i(x') \tag{8}$$

where ϕ_i and λ_i are, respectively, the eigenfunctions and eigenvalues of the Laplace-Beltrami operator satisfying $\Delta \phi_i = \lambda_i \phi_i$ (without loss of generality, we assume λ_i to be sorted in increasing order starting with $\lambda_0 = 0$). Since the Laplace-Beltrami operator is an *intrinsic* geometric quantity, i.e., it can be expressed solely in terms of the metric of X , its eigenfunctions and eigenvalues as well as the heat kernel are invariant under isometric transformations of the manifold.

The value of the heat kernel $h_t(x, x')$ can be interpreted as the transition probability density of a random walk of length t from the point x to the point

x' . This allows to construct a family of intrinsic metrics known as *diffusion metrics*,

$$d_t^2(x, x') = \int (h_t(x, y) - h_t(x', y))^2 dy = \sum_{i>0} e^{-\lambda_i t} (\phi_i(x) - \phi_i(x'))^2, \tag{9}$$

which measure the “connectivity rate” of the two points by paths of length t .

The parameter t can be given the meaning of *scale*, and the family $\{d_t\}$ can be thought of as a scale-space of metrics. By integrating over all scales, a *scale-invariant* version of (9) is obtained,

$$d_{CT}^2(x, x') = 2 \int_0^\infty d_t^2(x, x') dt = \sum_{i>0} \frac{1}{\lambda_i} (\phi_i(x) - \phi_i(x'))^2. \tag{10}$$

This metric is referred to as the *commute-time distance* and can be interpreted as the connectivity rate by paths of any length. We will broadly call constructions related to the heat kernel, diffusion and commute time metrics as *diffusion geometry*.

3 Equi-affine Metric

An *affine transformation* $\mathbf{x} \mapsto \mathbf{A}\mathbf{x} + \mathbf{b}$ of the three-dimensional Euclidean space can be parametrized by a regular 3×3 matrix \mathbf{A} and a 3×1 vector \mathbf{b} . since all constructions discussed here are trivially translation invariant, we will omit the vector \mathbf{b} . The transformation is called *special affine* or *equi-affine* if it is volume-preserving, i.e., $\det \mathbf{A} = 1$.

As the standard Euclidean metric is not affine-invariant, the Laplace-Beltrami Operators associated with X and $\mathbf{A}X$ are generally distinct, and so are the resulting diffusion geometries. In what follows, we are going to substitute the Euclidean metric by its equi-affine invariant counterpart. That, in turn, will induce an equi-affine-invariant Laplace-Beltrami Operator and define equi-affine-invariant diffusion geometry.

The equi-affine metric can be defined through the parametrization of a curve [8,23]. Let C be a curve on X parametrized by p . By the chain rule,

$$\begin{aligned} \frac{dC}{dp} &= \mathbf{x}_1 \frac{du_1}{dp} + \mathbf{x}_2 \frac{du_2}{dp} \\ \frac{d^2C}{dp^2} &= \mathbf{x}_1 \frac{d^2u_1}{dp^2} + \mathbf{x}_2 \frac{d^2u_2}{dp^2} + \mathbf{x}_{11} \left(\frac{du_1}{dp}\right)^2 + \\ &\quad 2\mathbf{x}_{12} \frac{du_1}{dp} \frac{du_2}{dp} + \mathbf{x}_{22} \left(\frac{du_2}{dp}\right)^2, \end{aligned} \tag{11}$$

where, for brevity, we denote $\mathbf{x}_i = \frac{\partial \mathbf{x}}{\partial u_i}$ and $\mathbf{x}_{ij} = \frac{\partial^2 \mathbf{x}}{\partial u_i \partial u_j}$. As volumes are preserved under the equi-affine group of transformations, we define the invariant arclength p through

$$\det(\mathbf{x}_1, \mathbf{x}_2, C_{pp}) = 1. \quad (12)$$

Plugging (11) into (12) yields

$$dp^2 = \det(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_{11} du_1^2 + 2\mathbf{x}_{12} du_1 du_2 + \mathbf{x}_{22} du_2^2), \quad (13)$$

from where we readily have an equi-affine-invariant pre-metric tensor

$$\hat{g}_{ij} = \tilde{g}_{ij} |\tilde{g}|^{-1/4}, \quad (14)$$

where $\tilde{g}_{ij} = \det(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_{ij})$. The pre-metric tensor (14) defines a true metric only on strictly convex surfaces [8]; in more general cases, it might cease from being positive definite. In order to deal with arbitrary surfaces, we extend the metric definition by restricting the eigenvalues of the tensor to be positive. Representing \hat{g} as a 2×2 matrix admitting the eigendecomposition $\hat{\mathbf{G}} = \mathbf{U}\mathbf{\Gamma}\mathbf{U}^T$, where \mathbf{U} is orthonormal and $\mathbf{\Gamma} = \text{diag}\{\gamma_1, \gamma_2\}$, we compose a new first fundamental form for non-vanishing Gaussian curvature matrix $\mathbf{G} = \mathbf{U}|\mathbf{\Gamma}|\mathbf{U}^T$. The metric tensor g is positive definite and is equi-affine invariant.

4 Numerical Considerations

4.1 Local Fitting

In order to compute the equi-affine metric we need to evaluate the second-order derivatives of the surface with respect to some parametrization coordinates. While this can be done practically in any representation, here we assume that the surface is given as a triangular mesh. For each triangular face, the metric tensor elements are calculated from a quadratic surface patch fitted to the triangle itself and its three adjacent neighbor triangles. The four triangles are unfolded to the plane, to which an affine transformation is applied in such a way that the central triangle becomes a unit simplex. The coordinates of this planar representation are used as the parametrization \mathbf{u} with respect to which the first fundamental form coefficients are computed at the barycenter of the simplex (Figure 1). This step is performed for every triangle of the mesh and is summarized in [30].

4.2 Affine Geodesics

Calculating geodesic distances was intensively explored in past decades. Several fast and accurate numerical schemes [27,16,25,26] can be used for this purpose. We use the FMM technique, after locally rescaling each edge according to the equi-affine metric. The (affine invariant) length of each edge is defined by $L^2(dx, dy) = g_{11}dx^2 + 2g_{12}dxdy + g_{22}dy^2$. Specifically, for our canonical triangle with vertices at $(0, 0)$, $(1, 0)$ and $(0, 1)$ we have $L_1^2 = g_{11}$, $L_2^2 = g_{22}$ and $L_3^2 = g_{11} - 2g_{12} + g_{22}$. Each edge may appear in more than one triangle. In our experiments we use the average length as an approximation, while verifying that the triangle inequality holds.

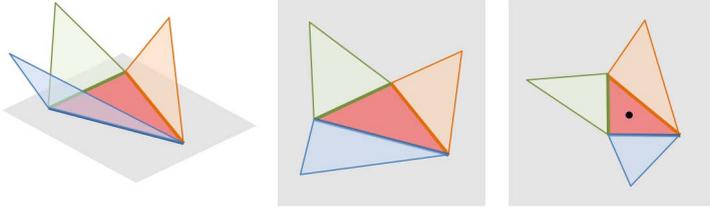


Fig. 1. Left to right: part of a triangulated surface about a specific triangle. The three neighboring triangles together with the central one are unfolded flat to the plane. The central triangle is canonized into a right isosceles triangle; three neighboring triangles follow the same planar affine transformation. Finally, the six surface coordinate values at the vertices are used to interpolate a quadratic surface patch from which the metric tensor is computed.

4.3 Finite Elements Method (FEM)

Having the discretized first fundamental form coefficients, our next target is to discretize the Laplace-Beltrami Operator. Since our final goal is not the operator itself but its eigendecomposition, we skip the explicit construction of the Laplacian and discretize its eigenvalues and eigenfunctions directly. This is achieved using the finite elements method (FEM) proposed in [12] and used in shape analysis in [20]. For that purpose, we translate the eigendecomposition of the Laplace-Beltrami Operator $\Delta\phi = \lambda\phi$ into a *weak form*

$$\int \psi_k \Delta\phi \, da = \lambda \int \psi_k \phi \, da \tag{15}$$

with respect to some basis $\{\psi_k\}$ spanning a (sufficiently smooth) subspace of $L^2(X)$. Specifically, we choose the ψ_k 's to be the first-order finite element functions obtaining a value of one at a vertex k and decaying linearly to zero in its 1-ring (the size of the basis equals to the number of vertices in the mesh). Substituting these functions into (15), we obtain

$$\int \psi_k \Delta\phi \, da = \int \langle \nabla \psi_k, \nabla \phi \rangle_x \, da = \int g^{ij} (\partial_i \phi) (\partial_j \psi_k) \, da = \lambda \int \psi_k \phi \, da. \tag{16}$$

Next, we approximate the eigenfunction ϕ in the finite element basis by $\phi = \sum_{l=1} \alpha_l \psi_l$. This yields

$$\int g^{ij} \left(\partial_i \sum_l \alpha_l \psi_l \right) (\partial_j \psi_k) \, da = \lambda \int \psi_k \sum_l \alpha_l \psi_l \, da,$$

or, equivalently,

$$\sum_l \alpha_l \int g^{ij} (\partial_i \psi_l) (\partial_j \psi_k) \, da = \lambda \sum_l \alpha_l \int \psi_k \psi_l \, da.$$

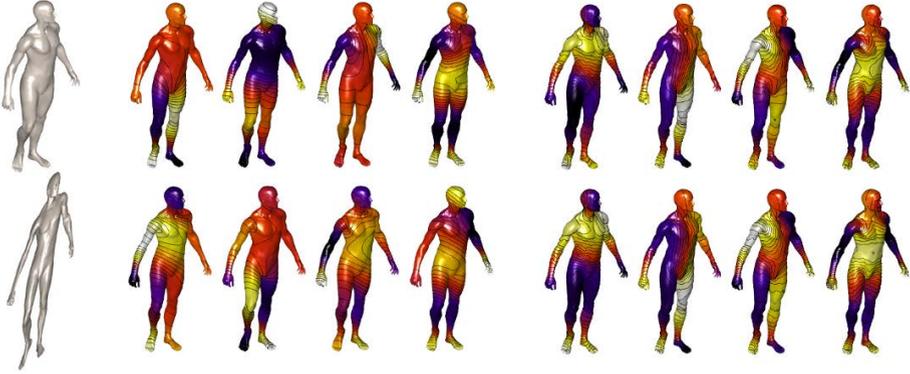


Fig. 2. Four eigenfunctions of the standard (second through fifth columns) and the proposed equi-affine-invariant (four rightmost columns) Laplace-Beltrami operator. Two rows show a shape and its equi-affine transformation. For convenience of visualization, eigenfunctions are textured mapped onto the original shape.

The last equation can be rewritten in matrix form as a generalized eigendecomposition problem $\mathbf{A}\alpha = \lambda\mathbf{B}\alpha$ solved for the coefficients α_l , where

$$a_{kl} = \int g^{ij} (\partial_i \psi_l) (\partial_j \psi_k) da,$$

$$b_{kl} = \int \psi_k \psi_l da,$$

and the local surface area is expressed in parametrization coordinates as $da = \sqrt{g} du_1 du_2$. The resulting eigendecomposition can be used to define an equi-affine-invariant diffusion geometry. Eigenfunctions, heat kernels, and diffusion distances remain invariant under volume-preserving affine transformations of the shape (Figures 2–3).

Evaluating the proposed metric is bounded by the number of adjacent neighbors of each vertex, from which we conclude that the new metric is evaluated in linear time with relation of the number of vertices. Spectral decomposition is performed using the power method, implemented in MATLAB, and in practice we only need few (below 200) eigenvectors.

5 Applications

To evaluate the performance of the proposed approach for the construction of local descriptors, we used the Shape Google framework [28] based on standard and affine-invariant Heat Kernel Signatures. HKS and AI-HKS were computed at six arbitrary scales ($t = 1024, 1351.2, 1782.9, 2352.5, \text{ and } 4096$). Bags of features were computed using soft vector quantization with variance taken as twice the median of all distances between cluster centers. Approximate nearest neighbor method [1] was used for vector quantization. Both the standard and the affine-invariant Laplace-Beltrami Operator discretization were computed using finite

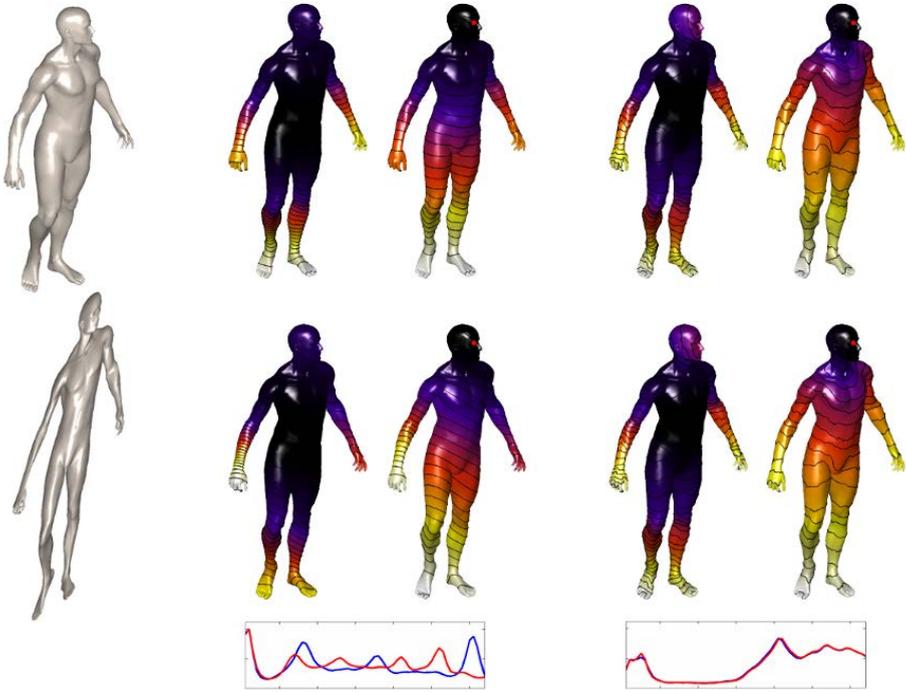


Fig. 3. Heat kernel signature $h_t(x, x)$ and diffusion metric ball (second and third columns, respectively), and their equi-affine invariant counterparts (fourth and fifth columns, respectively). Two rows show a shape and its transformation. For convenience of visualization, the kernel and the metric are overlaid onto the original shape. Plots under the figure show the corresponding metric distributions before and after the transformation.

elements. Heat kernels were approximated using the first 100 eigenpairs of the discrete Laplacian. The geometric vocabulary size was set to 64.

Evaluation was performed using the SHREC 2010 robust large-scale shape retrieval benchmark methodology [4]. The dataset consisted of two parts: 793 shapes from 13 shape classes with simulated transformation of different types (Figure 4) and strengths (60 per shape) used as queries, and additional 521 shapes from a large variety of objects. The total dataset size was 1314. Retrieval was performed by matching 780 transformed queries to shape classes. Each query had one correct corresponding null shape in the dataset. Performance was evaluated using precision/recall characteristic. *Precision* $P(r)$ is defined as the percentage of relevant shapes in the first r top-ranked retrieved shapes. *Mean average precision* (mAP), defined as $mAP = \sum_r P(r) \cdot rel(r)$, where $rel(r)$ is the relevance of a given rank, was used as a single measure of performance. Intuitively, mAP is interpreted as the area below the precision-recall curve. Ideal performance retrieval performance results in first relevant match with mAP=100%. Performance results were broken down according to transformation class and strength.

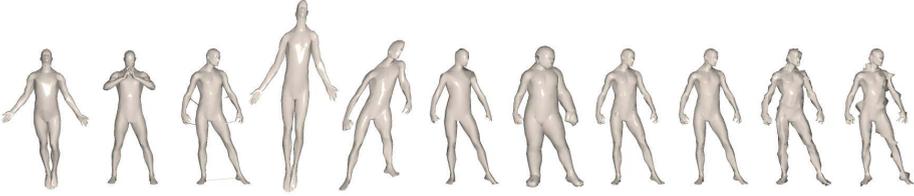


Fig. 4. Examples of query shape transformations used in the shape retrieval experiment (left to right): null, isometry, topology, affine, affine+isometry, sampling, local dilation, holes, microholes, Gaussian noise, shot noise

Tables 2–1 show that in contrast to the Euclidean metric, the equi-affine metric preserves the high accuracy rate of shape retrieval for all deformations, including equi-affine. In some deformations we can see an improvement, which we attribute to the smoothing effect of the second order interpolation. As this metric is based on second derivatives it is less robust to noise than its Euclidean adversary. Yet, since the numeric is based on the weak form (FEM) of the LBO, the integration improves robustness. Adding that to the usage of low frequencies from the eigendecomposition, explains the competitive results even without performing noise reduction and/or resampling as a preprocessing step.

The equi-affine metric can be used in many existing methods that compute geodesic distances. In what follows, we show several examples for using the new metric in known applications such as Voronoi tessellation and non-rigid matching.

Voronoi tessellation is a partitioning of (X, g) into disjoint open sets called Voronoi cells. A set of k points $(x_i \in X)_{i=1}^k$ on the surface defines the Voronoi cells $(V_i)_{i=1}^k$ such that the i -th cell contains all points in X closer to x_i than to any other x_j in the sense of the metric g . Voronoi tessellations created with the equi-affine metric commute with equi-affine transformations as visualized in Figure 5.

Table 1. Performance (mAP in %) of Shape Google with HKS descriptors

Transform.	Strength				
	1	≤ 2	≤ 3	≤ 4	≤ 5
<i>Isometry</i>	100.00	100.00	100.00	100.00	100.00
<i>Equi-Affine</i>	100.00	86.89	73.50	57.66	46.64
<i>Iso.+Equi-Affine</i>	94.23	86.35	76.84	70.76	65.36
<i>Topology</i>	100.00	100.00	98.72	98.08	97.69
<i>Holes</i>	100.00	96.15	92.82	88.51	82.74
<i>Micro holes</i>	100.00	100.00	100.00	100.00	100.00
<i>Local scale</i>	100.00	100.00	97.44	87.88	78.78
<i>Sampling</i>	100.00	100.00	100.00	96.25	91.43
<i>Noise</i>	100.00	100.00	100.00	99.04	99.23
<i>Shot noise</i>	100.00	100.00	100.00	98.46	98.77

Table 2. Performance (mAP in %) of Shape Google with equi-affine-invariant HKS descriptors

Transform.	Strength				
	1	≤2	≤3	≤4	≤5
<i>Isometry</i>	100.00	100.00	100.00	100.00	99.23
<i>Affine</i>	100.00	100.00	100.00	100.00	97.44
<i>Iso. + Equi-Affine</i>	100.00	100.00	100.00	100.00	100.00
<i>Topology</i>	96.15	94.23	91.88	89.74	86.79
<i>Holes</i>	100.00	100.00	100.00	100.00	100.00
<i>Micro holes</i>	100.00	100.00	100.00	100.00	100.00
<i>Local scale</i>	100.00	100.00	94.74	82.39	73.97
<i>Sampling</i>	100.00	100.00	100.00	96.79	86.10
<i>Noise</i>	100.00	100.00	89.83	78.53	69.22
<i>Shot noise</i>	100.00	100.00	100.00	97.76	89.63

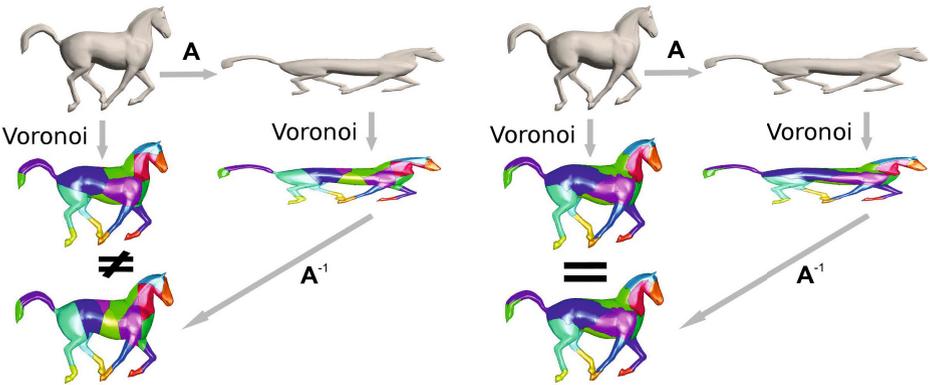


Fig. 5. Voronoi cells generated by a fixed set of 20 points on a shape undergoing an equi-affine transformation. The standard geodesic metric (left) and its equi-affine counterpart (right) were used. Note that in the latter case the tessellation commutes with the transformation.

Two non-rigid shapes X, Y can be considered similar if there exists an isometric correspondence $\mathcal{C} \subset X \times Y$ between them, such that $\forall x \in X$ there exists $y \in Y$ with $(x, y) \in \mathcal{C}$ and vice-versa, and $d_X(x, x') = d_Y(y, y')$ for all $(x, y), (x', y') \in \mathcal{C}$, where d_X, d_Y are geodesic distance metrics on X, Y . In practice, no shapes are perfectly isometric, and such a correspondence rarely exists; however, one can attempt finding a correspondence minimizing the metric *distortion*,

$$\text{dis}(\mathcal{C}) = \max_{\substack{(x,y) \in \mathcal{C} \\ (x',y') \in \mathcal{C}}} |d_X(x, x') - d_Y(y, y')|. \tag{17}$$

The smallest achievable value of the distortion is called the *Gromov-Hausdorff distance* [9] between the metric spaces (X, d_X) and (Y, d_Y) ,

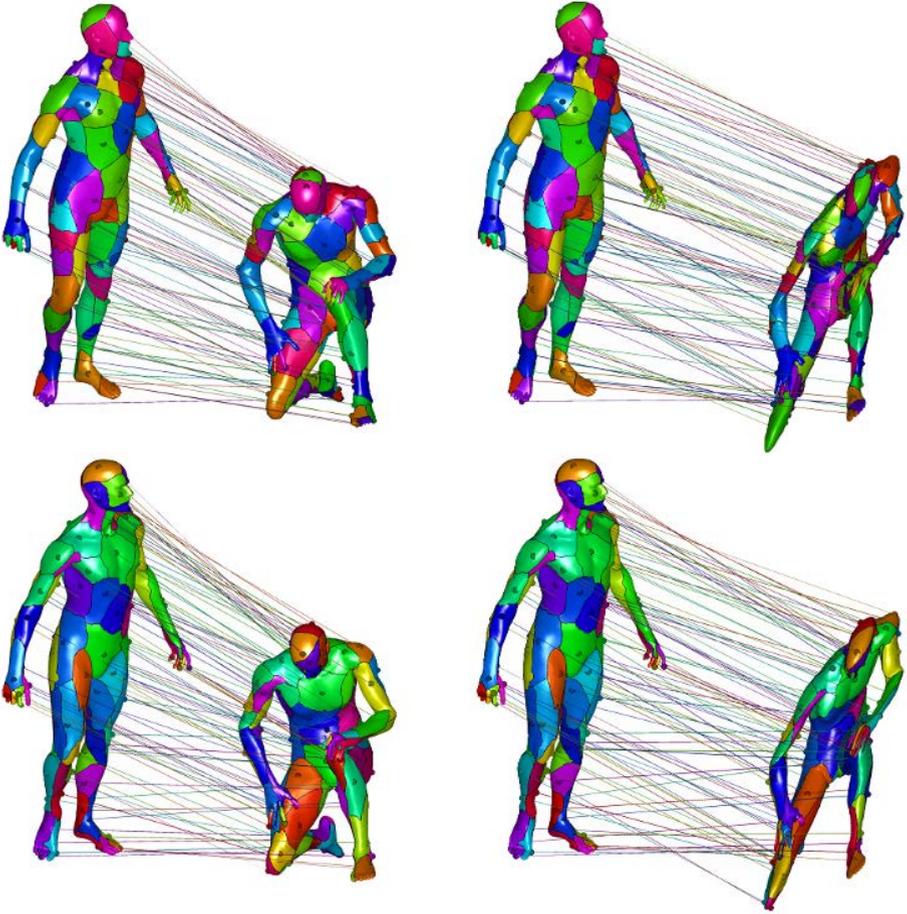


Fig. 6. The GMDS framework is used to calculate correspondences between a shape and its isometry (left) and isometry followed by an equi-affine transformation (right). Matches between shapes are depicted as identically colored Voronoi cells. Standard distance (first row) and its equi-affine-invariant counterpart (second row) are used as the metric structure in the GMDS algorithm. Inaccuracies obtained in the first case are especially visible in the legs and arms.

$$d_{\text{GH}}(X, Y) = \frac{1}{2} \inf_C \text{dis}(C), \quad (18)$$

and can be used as a criterion of shape similarity.

The choice of the distance metrics d_X, d_Y defines the invariance class of this similarity criterion. Using geodesic distances, the similarity is invariant to inelastic deformations. Here, we use geodesic distances induced by our equi-affine Riemannian metric tensor, which gives additional invariance to affine transformations of the shape.

Bronstein *et al.* [3] showed how (18) can be efficiently approximated using a convex optimization algorithm in the spirit of multidimensional scaling (MDS), referred to as generalized MDS (GMDS). Since the input of this numeric framework are geodesic distances between mesh points, all that is needed to obtain an equi-affine GMDS is one additional step where we substitute the geodesic distances with their equi-affine equivalents. Figure 6 shows the correspondences obtained between an equi-affine transformation of a shape using the standard and the equi-affine-invariant versions of the geodesic metric.

6 Conclusion

We introduced an equi-affine-invariant metric that can cope with surfaces that do not have vanishing Gaussian curvature. We showed a wide range of applications, from shape retrieval through Voronoi tessellation to correspondence search, based on differential geometry tools and spectral analysis. The limitation of the method is the fixed scale restriction that will be solved in the future.

Acknowledgments. This research was supported by European Community's FP7- ERC program, grant agreement no. 267414. MB was supported by the Swiss High-Performance and High-Productivity Computing (HP2C).

References

1. Arya, S., Mount, D.M., Netanyahu, N.S., Silverman, R., Wu, A.Y.: An optimal algorithm for approximate nearest neighbor searching. *J. ACM* 45, 891–923 (1998)
2. Bérard, P., Besson, G., Gallot, S.: Embedding riemannian manifolds by their heat kernel. *Geometric and Functional Analysis* 4, 373–398 (1994)
3. Bronstein, A., Bronstein, M., Kimmel, R.: Efficient computation of isometry-invariant distances between surfaces. *SIAM J. Scientific Computing* 28(5), 1812–1836 (2006)
4. Bronstein, A.M., Bronstein, M.M., Castellani, U., Falcidieno, B., Fusiello, A., Godil, A., Guibas, L.J., Kokkinos, I., Lian, Z., Ovsjanikov, M., Patané, G., Spagnuolo, M., Toldo, R.: SHREC 2010: robust large-scale shape retrieval benchmark. In: *Proc. 3DOR* (2010)
5. Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Expression-invariant face recognition via spherical embedding. In: *Proc. Int'l Conf. Image Processing (ICIP)*, vol. 3, pp. 756–759 (2005)
6. Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Efficient computation of isometry-invariant distances between surfaces. *SIAM J. Scientific Computing* 28, 1812–1836 (2006)
7. Bronstein, M.M., Bronstein, A.M.: Shape recognition with spectral distances with spectral distances. *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)* 33, 1065–1071 (2011)

8. Buchin, S.: Affine differential geometry. Science Press, Beijing (1983)
9. Burago, D., Burago, Y., Ivanov, S.: A course in metric geometry. Graduate studies in mathematics, vol. 33. American Mathematical Society (2001)
10. Coifman, R.R., Lafon, S.: Diffusion maps. *Applied and Computational Harmonic Analysis* 21, 5–30 (2006)
11. Dierkes, U., Hildebrandt, S., Kuster, A., Wohlrab, O.: *Minimal Surfaces I*. Springer, Heidelberg (1992)
12. Dziuk, G.: Finite elements for the Beltrami operator on arbitrary surfaces. In: Hildebrandt, S., Leis, R. (eds.) *Partial Differential Equations and Calculus of Variations*, pp. 142–155 (1988)
13. Elad, A., Keller, Y., Kimmel, R.: Texture mapping via spherical multi-dimensional scaling. In: *Proc. Scale-Space Theories in Computer Vision*, pp. 443–455 (2005)
14. Elad, A., Kimmel, R.: On bending invariant signatures for surfaces. *Trans. on Pattern Analysis and Machine Intelligence (PAMI)* 25, 1285–1295 (2003)
15. Gebal, K., Bærentzen, J.A., Aanæs, H., Larsen, R.: Shape analysis using the auto diffusion function. In: *Proc. of the Symposium on Geometry Processing*, pp. 1405–1413 (2009)
16. Kimmel, R., Sethian, J.A.: Computing geodesic paths on manifolds. *Proc. National Academy of Sciences (PNAS)* 95, 8431–8435 (1998)
17. Mémoli, F., Sapiro, G.: A theoretical and computational framework for isometry invariant recognition of point cloud data. *Foundations of Computational Mathematics* 5, 313–346 (2005)
18. Raviv, D., Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Volumetric heat kernel signatures. In: *Proc. 3D Object recognition (3DOR)*, part of ACM Multimedia (2010)
19. Raviv, D., Bronstein, A.M., Bronstein, M.M., Kimmel, R., Sapiro, G.: Diffusion symmetries of non-rigid shapes. In: *Proc. International Symposium on 3D Data Processing, Visualization and Transmission (3DPVT)* (2010)
20. Reuter, M., Biasotti, S., Giorgi, D., Patanè, G., Spagnuolo, M.: Discrete Laplace–Beltrami operators for shape analysis and segmentation. *Computers & Graphics* 33, 381–390 (2009)
21. Rustamov, R.M.: Laplace-Beltrami eigenfunctions for deformation invariant shape representation. In: *Proc. Symposium on Geometry Processing (SGP)*, pp. 225–233 (2007)
22. Schwartz, E.L., Shaw, A., Wolfson, E.: A numerical solution to the generalized mapmaker’s problem: flattening nonconvex polyhedral surfaces. *Trans. on Pattern Analysis and Machine Intelligence (PAMI)* 11, 1005–1008 (1989)
23. Sochen, N.: Affine-invariant flows in the Beltrami framework. *Journal of Mathematical Imaging and Vision* 20, 133–146 (2004)
24. Sun, J., Ovsjanikov, M., Guibas, L.J.: A concise and provably informative multi-scale signature based on heat diffusion. In: *Proc. Symposium on Geometry Processing (SGP)* (2009)
25. Surazhsky, V., Surazhsky, T., Kirsanov, D., Gortler, S., Hoppe, H.: Fast exact and approximate geodesics on meshes. In: *Proc. ACM Transactions on Graphics (SIGGRAPH)*, pp. 553–560 (2005)
26. Yatziv, L., Bartesaghi, A., Sapiro, G.: $O(N)$ implementation of the fast marching algorithm. *J. Computational Physics* 212, 393–399 (2006)

27. Tsitsiklis, J.N.: Efficient algorithms for globally optimal trajectories. *IEEE Trans. Automatic Control* 40, 1528–1538 (1995)
28. Ovsjanikov, M., Bronstein, A.M., Bronstein, M.M., Guibas, L.J.: Shape Google: a computer vision approach to invariant shape retrieval. In: *Proc. NORDIA* (2009)
29. Raviv, D., Bronstein, A.M., Bronstein, M.M., Kimmel, R., Sochen, N.: Affine-invariant diffusion geometry of deformable 3D shapes. In: *Proc. Computer Vision and Pattern Recognition (CVPR)* (2011)
30. Raviv, D., Bronstein, A.M., Bronstein, M.M., Kimmel, R., Sochen, N.: Affine-invariant geodesic geometry of deformable 3D shapes. *Computers & Graphics* 35, 692–697 (2011)