J Math Imaging Vis
DOI 10.1007/s10851-013-0467-y

Equi-affine Invariant Geometry for Shape Analysis

Dan Raviv · Alexander M. Bronstein · Michael M. Bronstein · Dan Waisman · Nir Sochen · Ron Kimmel

Abstract Traditional models of bendable surfaces are based on the exact or approximate invariance to deformations that do not tear or stretch the shape, leaving intact an intrinsic geometry associated with it. These geometries are typically defined using either the shortest path length (geodesic distance), or properties of heat diffusion (diffusion distance) on the surface. Both measures are implicitly derived from the metric induced by the ambient Euclidean space. In this paper, we depart from this restrictive assumption by observing that a different choice of the metric results in a richer set of geometric invariants. We apply equi-affine geometry for analyzing arbitrary shapes with positive Gaussian curvature. The potential of the proposed framework is explored in a range of applications such as shape matching and retrieval, symmetry detection, and computation of Voroni tessellation. We show that in some shape analysis tasks, equi-affine-invariant intrinsic geometries often outperform their Euclidean-based counterparts. We further explore the potential of this metric in facial anthropometry of newborns. We show that intrinsic properties of this homogeneous group are better captured using the equi-affine metric.

Keywords Equi-affine · Affine · Metric invariant · Intrinsic geometry · Shape analysis

D. Raviv (✉)
Media Lab, Massachusetts Institute of Technology (MIT), Cambridge, USA
e-mail: darav@mit.edu

A.M. Bronstein
School of Electrical Engineering, Tel Aviv University, Tel Aviv, Israel

M.M. Bronstein
Institute of Computational Science, Faculty of Informatics, Università della Svizzera Italiana, Lugano, Switzerland

D. Waisman
Department of Neonatology, Carmel Medical Center and Faculty of Medicine, Technion, Haifa, Israel

N. Sochen
Dept. of Applied Mathematics, Tel Aviv University, Tel Aviv, Israel

R. Kimmel
Dept. of Computer Science, Technion, Haifa, Israel

1 Introduction

Many methods of shape analysis in computer vision, graphics, and pattern recognition, model shapes as Riemannian manifolds. Differential geometry provides a broad arsenal of tools allowing to describe local and global invariant properties of the shape and compute similarity and correspondence between two given shapes. Moreover, it is possible to represent metric structures in different spaces that are convenient to work with.

One of the early efforts in this domain is of Schwartz *et al.* [58], who proposed to study the surface of the brain cortex by embedding it into the plane. This approach was further developed by Elad and Kimmel [27], who showed that by finding the most isometric embedding of a non-rigid 3D shape into a Euclidean space with more than two dimensions, it is possible to undo the deformations. The embedding was computed using multidimensional scaling (MDS). The resulting *canonical form* representation allowed to compute shape similarity invariant to inelastic deformations. Canonical forms were used for different applications in deformable shape analysis, ranging from texture mapping [26] to face recognition [11, 14], while embeddings into non-Euclidean spaces (e.g., a sphere) were discussed in [10, 70].

Published online: 02 October 2013

The introduction of canonical forms led to a number of follow-up papers, where shapes were represented as metric spaces with some intrinsic (e.g. geodesic) distance metric, and their comparison involved the similarity of the underlying metric spaces. Mémoli and Sapiro [39] and Bronstein *et al.* [12, 13] computed shape similarity by discretizing the Gromov-Hausdorff metric [31]. In [12], an MDS-like optimization algorithm (branded generalized MDS or GMDS) was introduced for the computation of the Gromov-Hausdorff metric, and a different optimization based on hierarchical graph labeling was used in [69]. Different metrics were studied in [16, 38] in the Gromov-Hausdorff framework for shape matching and later on in [48, 50] for the detection of intrinsic symmetries of deformable shapes, modeled as self-isometries.

In [7] Bérard *et al.* introduced the concept of embedding the structure of a Riemannian manifold into the eigenstructure of its Laplace-Beltrami operator. This theoretical framework was later exploited by Rustamov [55] who used the fact that it is possible to isometrically embed surfaces into Euclidean spaces defined by the eigenfunctions of the Laplace-Beltrami operator. Belkin and Niyogi [6] realized that heat kernels extracted from the Laplace-Beltrami operator can be used for quantifying the distance between points in abstract manifolds. In this new geometry, the distance between two points is expressed as the integrated difference between two heat kernels each of which is centered about one of the points. The *diffusion geometry* was later refined and applied to manifold learning problems by Coifman and Lafon [21] while Lévy [36] showed that significant geometric information of a shape can be captured by considering diffusion operators on its Riemannian manifold. Properties of Laplacian eigenfunctions were later used for intrinsic symmetry detection [45], and shape matching [59].

The diffusion framework was found useful also for the construction of local shape descriptors. Sun *et al.* [65], and Gebal *et al.* [29] used the heat kernel of the diffusion equation to construct multi-scale intrinsic descriptors referred to as *heat kernel signatures* (HKS). HKS descriptors were extended to textured shapes in [34] and to volumetric data in [49]. Scale-invariant HKS were proposed in [18]. These descriptors performed successfully in large-scale shape retrieval applications [17, 43], using the bag of features [60] paradigm. Ovsjanikov *et al.* [44] used HKS to find dense correspondence, while Castellani *et al.* were using HKS [20] in brain analysis applications.

A different approach for alignment was adopted within the medical community, where a smooth velocity field is found. Early attempts such as [3] addressed this problem but faced self folding of the grid. A later approach, which became the standard for non rigid mapping, was presented in [5]. In their scheme, a numerical procedure called Large Deformation Diffeomorphic Metric Mapping (LDDMM) optimizes for the best diffeomorphism between images by combining smoothness of the velocity field with a data term from the images. Another approach named SPHARM-PDM [63] rigidly aligns shapes according to their spherical parameterization. It showed good results on brain images, but is only relevant for those with a spherical topology. In an entirely different path, we encounter skeletal representations such as [56], which can handle larger deformations, but limited to coarse matching. An improved alignment which better handles reparameterization, as noted in [32], was later addressed in [35].

The starting point of all the aforementioned methods is the model of a shape or shape space as a Riemannian manifold. For example, MDS and GMDS frameworks that are based on diffusion geometry rely on a Laplace-Beltrami operator that is expressed through the Riemannian metric. Local descriptors are also expressed through Riemannian structures defined on the manifold. In the vast majority of cases, the standard choice of a metric is the one induced by the embedding of the shape. Such a metric is invariant to inelastic deformations of the shape and to global Euclidean transformations, that is, rotations, reflections, and translations. In this paper we show how to construct a metric that is invariant to a larger group of transformations. Besides theoretical importance, we argue that this invariance has many important applications in shape analysis (see, e.g., [30, 34]).

Affine invariance was explored in the field of image processing [2] more than twenty years ago, and re-appeared in state of the art descriptors as ASIFT *affine-scale invariant feature transform* [41], that are used to successfully locate repeatable informative (invariant) features in images. We can find affine measures in planar curves [19], flows [57, 61], and structures [22] while stability was the main topic of [4]. Here, we apply the known equi-affine invariant quadratic form for surfaces with positive Gaussian curvature to less trivial geometric structures. Using the resulting metric, we show its application to spectral shape analysis by deriving an associated Laplace-Beltrami operator that is not just invariant to bendings but also to equi-affine transformations. This, in turn, allows us to derive global affine-invariant measures such as diffusion (Fig. 1) and commute-time distances, as well as local equi-affine ones such as HKS.

The equi-affine metric was introduced by Blaschke [8, 64] as a theoretical framework for dealing with volume preserving linear transformations of R^3 in which a given two dimensional surface is embedded. We first applied equi-affine geometry to surfaces in [51, 52], where computational and numerical considerations were explored. Here, we present a somewhat more mature view on using the equi-affine geometry for surface analysis, new numerical procedures are discussed, a compact derivation of the equi-affine quadratic form is provided, as well as an alternative construction of the invariant measures. We follow our brief exploration of

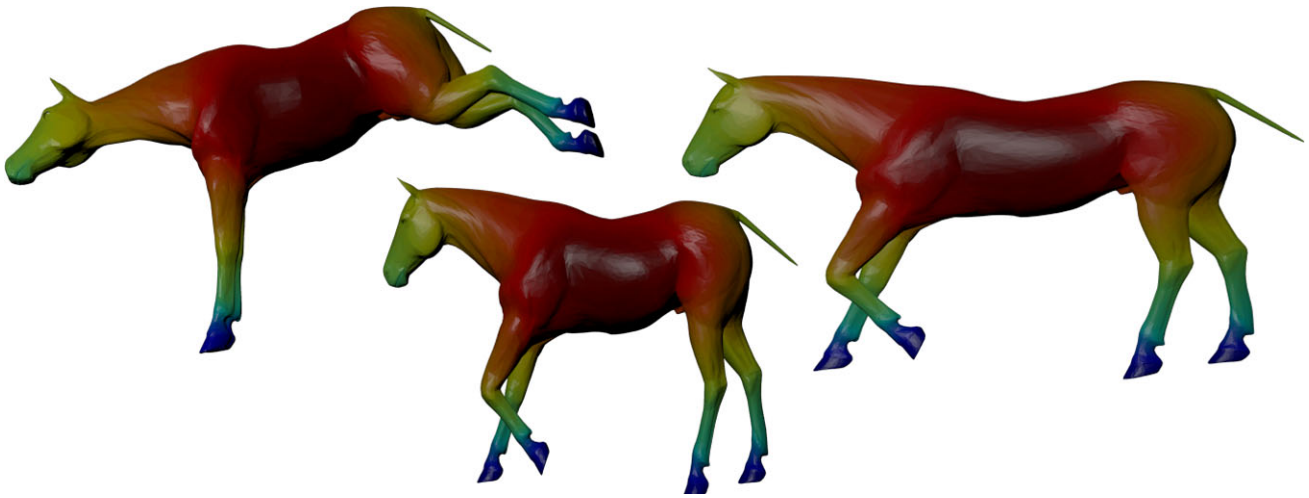

Fig. 1 Diffusion distances from the horse's abdomen (*colored red*) are invariant under equi-affine stretching (*middle*) and non-rigid deformations (*left*). (Color figure online)

exploiting the diffusion geometry that was presented in a conference [51] for the construction of stable equi-affine invariant measures on which we elaborate in this journal publication.

The rest of the paper is organized as follows. We begin in Sect. 2 by describing the contribution of this paper. In Sect. 3 we provide the mathematical background of Euclidean, Riemmannian, and Diffusion geometries, followed by Sect. 4 where we elaborate on the equi-affine metric. Section 5 is dedicated to numerical aspects of the proposed framework, with several synthetic experiments presented in Sect. 6. We explore the potential of this metric for facial morphometry in Sect. 7, and conclude the paper in Sect. 8.

2 Main Contributions

This note contains the mathematical aspects of the equi-affine metric first shown in [51, 52]. The novelty of this paper can be summarized as follows:

- New and self contained proof of the equi-affine quadratic form in matrix formulation.
- Analysis of degenerate cases, when the metric is undefined or vanishes.
- New synthetic experiments, showing the advantages of the equi-affine metric over the Euclidean one.
- Analysis of facial measurements in newborns. We show the potential of the equi-affine metric for statistical geometric measurements in medical imaging.

3 Background

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 base using the inner product $\langle \cdot, \cdot \rangle_x : T_x X \times T_x X \to \mathbb{R}$. We further assume that X is embedded into $\mathbb{E} = \mathbb{R}^3$ by means of a regular map $\mathbf{x} : U \subseteq \mathbb{R}^2 \to \mathbb{R}^3$, so that the metric tensor can be expressed in coordinates as the coefficients of the first fundamental form

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

where u_i are the coordinates of U. The metric tensor yields the arclength

$$dp^2 = g_{11} {du_1}^2 + 2 g_{12} du_1 du_2 + g_{22} {du_2}^2. \tag{2}$$

Given a smooth scalar field f on the manifold, its gradient $\operatorname{grad} f$ is the vector field satisfying $f(x + dr) = f(x) + \langle \operatorname{grad} f(x), dr \rangle_x$ for every infinitesimal tangent vector $dr \in T_x X$. The inner product $\langle \operatorname{grad} f(x), v \rangle_x$ can be interpreted as the directional derivative of f in the direction v. A directional derivative of f whose direction at every point on X is defined by the values of a vector field V on the manifold is called the Lie derivative of f along V. The Lie derivative of the manifold volume (area) form along a vector field V is called the divergence of V, $\operatorname{div} V$. The negative divergence of the gradient of a scalar field f,

$$\Delta f = -\operatorname{div} \operatorname{grad} f, \tag{3}$$

is called the *Laplacian* of f. The operator Δ is called the *Laplace-Beltrami* operator, and it generalizes the standard notion of the Laplace operator to manifolds. Note that we define the Laplacian with the negative sign to conform to the computer graphics and computational geometry convention.

Evaluating Δ in local coordinates for any metric g holds an important part in this paper, as it bridges between local metric evaluation and global features of a shape. A well

known formulation of (3) in local coordinates $\mathbf{u} = (u_1, u_2)^{\mathrm{T}}$ is [23]

$$\Delta f = \frac{1}{\sqrt{\det \mathbf{G}}} \frac{\partial^{\mathrm{T}}}{\partial \mathbf{u}} \left(\sqrt{\det \mathbf{G}} \mathbf{G}^{-1} \frac{\partial}{\partial \mathbf{u}} f \right), \tag{4}$$

where the sign is flipped to impose a semi-positive structure, and $\frac{\partial}{\partial \mathbf{u}} = (\frac{\partial}{\partial u_1}, \frac{\partial}{\partial u_2})^{\mathrm{T}}$ and $\mathbf{G} = (g_{ij}) \in \mathbb{R}^{2\times 2}$. Being a positive self-adjoint operator, the Laplacian admits an eigendecomposition

$$\Delta \phi = \lambda \phi \tag{5}$$

with non-negative eigenvalues λ and corresponding orthogonormal eigenfunctions ϕ. Furthermore, due to the assumption that our domain is compact, the spectrum is discrete, $0 = \lambda_1 < \lambda_2 < \cdots$. In physics, (5) is known as the *Helmholtz equation* representing the spatial component of the wave equation. Thinking of our domain as of a vibrating membrane (with appropriate boundary conditions), the ϕ_i's can be interpreted as natural vibration modes of the membrane, while the λ_i's assume the meaning of the corresponding vibration frequencies. In fact, in this setting the eigenvalues have inverse area or squared spatial frequency units.

3.1 Extrinsic and Intrinsic Geometry

There exist two natural ways to measure distances on X: First, one can simply consider X a subset of $\mathbb{E}$ equipped with the standard Euclidean metric $d_{\mathbb{E}}$, and measure the distance on X using the *restricted Euclidean metric* $d_{\mathbb{E}}|_{X\times X}$. Under this metric, shortest paths between points on X are straight lines, possibly not entirely contained in X. Quantities expressed with it are often referred to as the *extrinsic geometry* of the shape.

The second choice is to consider a *path* $\gamma : [0, 1] \to X$ on the surface. By subdividing the interval $[0, 1]$ into n points $0 = t_1 < t_2 < \cdots < t_n = 1$, one can bound below the path length by the sum of the lengths $d_{\mathbb{E}}(\gamma(t_i), \gamma(t_{i+1}))$. In this way, a *length structure*

$$L(\gamma) = \sup_{n, \{t_1, \ldots, t_n\}} \sum_{i=1}^{n-1} d_{\mathbb{E}}\big(\gamma(t_i), \gamma(t_{i+1})\big) \tag{6}$$

is constructed; a path is said *rectifiable* if the supremum exists. If further $\gamma = \mathbf{x}(\beta)$ is the image of some path β in the parametrization domain, its length on the surface can be expressed in terms of the Riemannian structure as

$$L(\gamma) = \int_0^1 \sqrt{g_{11}\beta_1^2(t) + 2g_{12}\beta_1(t)\beta_2(t) + g_{22}\beta_2^2(t)}\, dt, \tag{7}$$

where β_i is the partial derivative of β with respect to u_i. Being totally expressible in terms of the metric tensor, the length structure $L(\gamma)$ is said *intrinsic*; such quantities are collectively called the *intrinsic geometry* of X.

Using the length structure, one can define a metric between two points x and x' on X as the length of the shortest path connecting between them,

$$d_L(x, x') = \inf_{\substack{\gamma \\ \gamma(0)=x \\ \gamma(1)=x'}} L(\gamma), \tag{8}$$

which is called the *induced* or the *geodesic* metric. Geodesic distances can be obtained as the viscosity solution to the *eikonal equation*

$$\|\nabla d\|_2 = 1 \tag{9}$$

(i.e., the largest d satisfying $\|\nabla d\|_2 \leq 1$) with boundary condition at the source point $d(x) = 0$, where the solution represents $d(x') = d_L(x, x')$. In the past two decades, several algorithms have been proposed for the computation of geodesic distances, differing in accuracy and complexity. In this study, we focus on a family of simulated wavefront propagation algorithms called *fast marching* [33]. While the complexity of the computation of geodesic distance from a single point to the rest of the points on the surface using fast marching is $\mathcal{O}(N \log N)$, N being the number of points, it was shown that linear complexity can be achieved without sacrificing the linear order of approximation [66]. On parametric surfaces, fast marching can be carried out by means of a raster scan and efficiently parallelized, which makes it especially attractive for GPU-based computation [62, 68].

We conclude that a shape can be modeled as a metric space equipped either with an extrinsic (restricted Euclidean) or intrinsic (geodesic) metrics [15, 37]. The choice of the metric impacts dramatically on the resulting geometric properties. While the extrinsic geometry is invariant only to rigid motion, its intrinsic counterpart remains intact under inelastic deformations that do not stretch or tear the surface. However, a major disadvantage of the geodesic metric is its extreme sensitivity to topological changes. Even when localized to a single point, topological noise greatly affects the shortest paths on the manifold and, consequently, the intrinsic geometry they define. A way counter this phenomenon is by replacing the notion of a shortest path by some notion of an "average" path length. In what follows, we briefly overview a construction of a family of intrinsic geometries based on this principle.

3.2 Diffusion Geometries

As a starting point, let us consider the *heat equation* describing diffusion processes on manifolds, which in our notation can be expressed as

$$\left(\Delta + \frac{\partial}{\partial t} \right) f(x, t) = 0 \tag{10}$$

where $f(x,t)$ is the distribution of heat on the manifold at point x at time t. The initial condition is some initial heat distribution $f_0(x)$ at time $t=0$, and boundary conditions are applied in case the manifold has a boundary.

The solution of the heat equation at time t can be expressed as the application of the *heat operator*

$$f(x,t) = \int h_t(x,y) f_0(y) da(y) \tag{11}$$

to the initial distribution. The kernel $h_t(x,y)$ of this integral operator is called the *heat kernel* and it corresponds to the solution of the heat equation at point x at time t with the initial distribution being a delta function at point y. From the signal processing perspective, the heat kernel can be interpreted as a non shift-invariant "impulse response". It also describes the amount of heat transferred from point x to point y after time t, as well as the transition probability density from point x to point y by a random walk of length t.

According to the spectral decomposition theorem, the heat kernel can be expressed as

$$h_t(x,y) = \sum_{i\geq 0} \exp(-\lambda_i t)\phi_i(x)\phi_i(y), \tag{12}$$

where $\exp(-\lambda t)$ can be interpreted as its "frequency response" (note that with a proper selection of units in (11), the eigenvalues λ_i assume inverse time or frequency units). The bigger is the time parameter, the lower is the cut-off frequency of the low-pass filter described by this response and, consequently, the bigger is the support of h_t on the manifold. Since the Laplace-Beltrami operator is an *intrinsic* geometric quantity, 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*,

$$\begin{aligned} d_t^2(x,x') &= \int \big(h_t(x,\cdot) - h_t(x',\cdot)\big)^2 da \\ &= \sum_{i>0} e^{-2\lambda_i t}\big(\phi_i(x) - \phi_i(x')\big)^2, \end{aligned} \tag{13}$$

which measure the diffusion distance of the two points for a given time 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. It appears that by integrating d_t^2 over all scales, a *scale-invariant* version of (13) is obtained,

$$d_{\mathrm{CT}}^2(x,x') = \int_0^\infty d_t^2(x,x')dt$$

$$= \sum_{i>0} \frac{1}{2\lambda_i}\big(\phi_i(x) - \phi_i(x')\big)^2. \tag{14}$$

This metric is referred to as the *commute-time distance* [46] and can be interpreted as the connectivity rate by paths of any length.

The quantity

$$h_t(x,x) = \sum_{i\geq 0} \exp(-\lambda_i t)\phi_i^2(x), \tag{15}$$

sometimes referred to as the *autodiffusivity function*, describes the amount of heat remaining at point x after time t. Furthermore, for small values of t is it related to the manifold curvature according to

$$h_t(x,x) = \frac{1}{4\pi t} + \frac{K(x)}{12\pi} + \mathcal{O}(t), \tag{16}$$

where $K(x)$ denotes the Gaussian (in general, sectional) curvature at point x.

In [65], Sun *et al.* showed that under mild technical conditions, the sequence $\{h_t(x,x)\}_{t>0}$ contains *full* information about the metric of the manifold. The authors proposed to associate each point x on the manifold with a vector $\mathrm{HKS}(x) = (h_{t_1}(x,x),\ldots,h_{t_n}(x,x))$ of the autodiffusivity functions sampled at some finite set of times $t_1,\ldots,t_n$. The authors dubbed such a feature descriptor as the *heat kernel signature* (HKS).

Several follow-ups extended these concepts and include scale-invariance [1, 18], geometric-photometric relation [34], and volumetric analysis [49], to say a few.

3.3 Isometries and Symmetries

The metric model allows casting the problem of shape similarity as similarity of metric structures. Two metric spaces (X,d_X) and (Y,d_Y) are said *isometric* if there exists bijection $g: X \to Y$ such that $d_Y \circ (g\times g) = d_X$. Such a distance-preserving map is called an *isometry*. Isometries form a group with the function composition operation, denoted by $\mathrm{Iso}(d)$.

Clearly, the notion of isometry is metric-dependent. Considering X and Y equipped with the extrinsic metric, we will that they are isometric if they are *congruent*, i.e., are related by a rigid motion $X = i(Y)$, $i \in \mathrm{Iso}(d_{\mathbb{E}})$. When an intrinsic metric d_X on X is considered, the resulting isometry group is usually richer, as a congruence is, naturally, also a d_X-isometry (in other words, the group of congruences of X is a subgroup of the group of its d_X-isometries). However, for some objects these two classes coincide, meaning that they have no incongruent isometries. Such shapes are called *rigid*, and their extrinsic geometry is completely defined by the intrinsic one.

True isometries are merely a mathematical idealization not existing in practice due to imperfections in measurement, representation, and deviations from the perfectly inelastic model. A generalization of the notion of isometry can be obtained by first defining a *correspondence* $C \subset X \times Y$ as a set of pairs of points from X and Y satisfying: 1) for every $x \in X$, there exists (at least one) $y \in Y$ such that $(x, y) \in C$; and, vice versa, 2) for every $y \in Y$, there exists $x \in X$ such that $(x, y) \in C$. This notion generalizes the notion of a bijective map between the space.

Suppose two pairs of points (x, y) and (x', y') are in correspondence. Then, we can quantify the quality of the correspondence by measuring to which extent the distance between x and x' measured on X using d_X matches the distance between the corresponding points y and y' measured on Y using d_Y,

$$\epsilon\big(x, y, x', y'\big) = \big|d_X\big(x, x'\big) - d_Y\big(y, y'\big)\big|. \tag{17}$$

This allows assigning a correspondence C the distortion

$$\mathrm{dis}(C) = \|\epsilon\|_{L^p(C\times C)}, \tag{18}$$

where, for example, the L^∞ norm

$$\|\epsilon\|_{L^\infty(C\times C)} = \sup_{(x,y),(x',y')\in C} \epsilon\big(x, y, x', y'\big) \tag{19}$$

can be used.

Minimizing the distortion over all possible correspondences between X and Y yields a distance

$$D(X, Y) = \frac{1}{2}\inf_C \mathrm{dis}(C) \tag{20}$$

between X and Y called the *Gromov-Hausdorff distance*. If the infimum is realized by some C^*, the latter is called a *minimum distortion correspondence* (note that more than one minimum distortion correspondence might exist if the shape possesses intrinsic symmetries). Bronstein *et al.* [12] showed that the above distance can be efficiently approximated using a generalized multidimensional scaling (GMDS) algorithm.

A particular case of shape similarity is the similarity of a shape to itself. The collection of *self-isometries* or *symmetries* of (X, d) also forms a group under function composition. A metric space (X, d) is self-similar if there exists a *self-isometry* on (X, d) (an isometry from (X, d) to itself). Since this definition depends on the choice of the metric, we distinguish between the group of *extrinsic symmetries* $\mathrm{Sym}(X, d_{\mathbb{E}})$ and that of *intrinsic symmetries* $\mathrm{Sym}(X, d_X)$, where d_X is an intrinsic metric on X [48].

4 Affine Invariance

In this section, we develop an equi-affine invariant metric on arbitrary surfaces with positive Gaussian curvature. We arrive at this result by first briefly reminding the classical construction of the equi-affine arclength for planar curves, and developing an equi-affine arclength on convex surfaces in $\mathbb{R}^3$, as first shown by Blaschke [8].

4.1 Planar Curves

An affine transformation of the plane can be expressed by six parameters as

$$\mathbf{Ax} + \mathbf{b} = \begin{pmatrix} a & b \\ c & d \end{pmatrix}\begin{pmatrix} x \\ y \end{pmatrix} + \begin{pmatrix} c \\ d \end{pmatrix}, \tag{21}$$

where $\det\mathbf{A} > 0$. The area preserving affine group (also called *special* or *equi-affine*) is defined by demanding $\det\mathbf{A} = 1$, which leaves five degrees of freedom.

Let $C : I \subset \mathbb{R} \to \mathbb{R}^2$ be a curve parameterized by a parameter p over some interval I. One way to construct an equi-affine arclength p is by asserting the conservation of the area of the parallelogram formed by the velocity and acceleration vectors of $C(p)$,

$$|C_p \times C_{pp}| \;=\; \det(C_p, C_{pp}) = \text{const.} \tag{22}$$

We construct the arclength as $f(\det(C_p, C_{pp}))dp$, where the function f is determined by demanding reparametrization invariance of p. For that goal, let ξ be a reparametrization of p such that $p = \xi(q)$. We require

$$f\big(\det(C_p, C_{pp})\big)dp = f\big(\det(C_q, C_{qq})\big)dq. \tag{23}$$

Using the chain rule,

$$\begin{aligned} C_q &= C_p\dot{\xi} \\ C_{qq} &= C_{pp}\dot{\xi}^2 + C_p\ddot{\xi}. \end{aligned} \tag{24}$$

Hence,

$$\begin{aligned} \det(C_q, C_{qq}) &= \det\left(\begin{pmatrix} \dot{\xi} & \\ \ddot{\xi} & \dot{\xi}^2 \end{pmatrix}(C_p, C_{pp})\right) \\ &= \det\begin{pmatrix} \dot{\xi} & \\ \ddot{\xi} & \dot{\xi}^2 \end{pmatrix} \cdot \det(C_p, C_{pp}) \\ &= \dot{\xi}^3 \cdot \det(C_p, C_{pp}). \end{aligned} \tag{25}$$

Combining (23) with $dp = \dot{\xi}dq$ yields

$$f\big(\det(C_p, C_{pp})\big)dp = f\big(\det(C_p, C_{pp})\dot{\xi}^3\big)\frac{dp}{\dot{\xi}}, \tag{26}$$

from where the equi-affine arclength is defined by

$$dv = \det(C_p, C_{pp})^{\frac{1}{3}}dp. \tag{27}$$

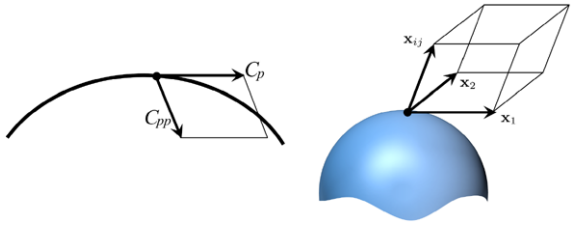

Fig. 2 An equi-affine local measurement for curves' length is the area of the parallelogram (*left*). For surfaces (*right*) we obtain the metric elements as the volume of the parallelepiped constructed by $\mathbf{x}_1$ and $\mathbf{x}_2$, and the corresponding second derivatives. $h_{ij} = \det(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_{12})$

4.2 Quadratic Form for Curved Surfaces

Let us now be given a two-dimensional surface X embedded into $\mathbb{R}^3$ by a regular embedding $\mathbf{x} : U \subseteq \mathbb{R}^2 \to \mathbb{R}^3$. We repeat the construction of the equi-affine arclength on the surface [64]. Let $C(p)$ be a curve on X; by the chain rule,

$$\begin{aligned} C_p &= \mathbf{x}_1 \frac{du_1}{dp} + \mathbf{x}_2 \frac{du_2}{dp} \\ C_{pp} &= \mathbf{x}_1 \frac{d^2u_1}{dp^2} + \mathbf{x}_2 \frac{d^2u_2}{dp^2} + \mathbf{x}_{11}\left(\frac{du_1}{dp}\right)^2 \\ &\quad + 2\mathbf{x}_{12}\frac{du_1}{dp}\frac{du_2}{dp} + \mathbf{x}_{22}\left(\frac{du_2}{dp}\right)^2, \end{aligned} \tag{28}$$

where, for brevity, we denote $\mathbf{x}_i = \frac{\partial \mathbf{x}}{\partial u_i}$ and $\mathbf{x}_{ij} = \frac{\partial^2 \mathbf{x}}{\partial u_i \partial u_j}$.

Once again, we assert the conservation of the volume of the parallelepiped formed by the vectors $\mathbf{x}_1$, $\mathbf{x}_2$ on the surface and the acceleration vector C_{pp} (Fig. 2),

$$C_{pp} \cdot (\mathbf{x}_1 \times \mathbf{x}_2) \;=\; \det(\mathbf{x}_1, \mathbf{x}_2, C_{pp}) = \text{const.} \tag{29}$$

Plugging (28) and using the fact that $\det(\mathbf{x}_i, \mathbf{x_j}, \mathbf{x_i}) = \mathbf{0}$, yields the squared arclength

$$\begin{aligned} &\det(\mathbf{x}_1, \mathbf{x}_2, C_{pp})dp^2 \\ &\quad = \det\left(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_{11}du_1^2 + 2\mathbf{x}_{12}du_1du_2 + \mathbf{x}_{22}du_2^2\right) \\ &\quad = \left(h_{11}du_1^2 + 2h_{12}du_1u_2 + h_{22}du_2^2\right) \\ &\quad = d\mathbf{u}^{\mathrm{T}}\mathbf{H}d\mathbf{u}, \end{aligned} \tag{30}$$

where $h_{ij} = \det(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_{ij})$.

We again assert reparametrization invariance, this time to a reparametrization of the surface. Let $V \subset \mathbb{R}^2$ be some other parametrization domain, so that $\mathbf{u} = \boldsymbol{\xi}(\mathbf{v})$, where $\boldsymbol{\xi} : V \to U$, $\boldsymbol{\xi}(\mathbf{v}) = (\xi(\mathbf{v}), \eta(\mathbf{v}))^{\mathrm{T}}$, is a reparametrization of the surface. Using the chain rule, we obtain

$$\mathbf{y}_i = \frac{\partial \mathbf{x}}{\partial v_i} \;=\; \mathbf{x}_1\xi_i + \mathbf{x}_2\eta_i, \tag{31}$$

where $\xi_i = \frac{\partial u_1}{\partial v_i}$ and $\eta_i = \frac{\partial u_2}{\partial v_i}$. Similarly,

$$\begin{aligned} \mathbf{y}_{ij} = \frac{\partial^2 \mathbf{x}}{\partial v_i \partial v_j} &= \mathbf{x}_1\xi_{ij} + \mathbf{x}_2\eta_{ij} + \mathbf{x}_{11}\xi_i\xi_j \\ &\quad + \mathbf{x}_{12}(\xi_i\eta_j + \xi_j\eta_i) + \mathbf{x}_{22}\eta_i\eta_j, \end{aligned} \tag{32}$$

where $\xi_{ij} = \frac{\partial^2 u_1}{\partial v_i \partial v_j}$ and $\eta_{ij} = \frac{\partial^2 u_2}{\partial v_i \partial v_j}$. Substituting the latter into the squared arclength (30) yields

$$\begin{aligned} \det(\mathbf{y}_1, \mathbf{y}_2, C_{pp})dp^2 &= d\mathbf{v}^{\mathrm{T}}\mathbf{H}'d\mathbf{v} \\ &= d\mathbf{u}^{\mathrm{T}}\mathbf{J}^{-\mathrm{T}}\mathbf{H}'\mathbf{J}^{-1}d\mathbf{u}, \end{aligned} \tag{33}$$

where $\mathbf{J}$ is the Jacobian of $\boldsymbol{\xi}$ and the elements of $\mathbf{H}'$ are given by $h'_{ij} = \det(\mathbf{y}_1, \mathbf{y}_2, \mathbf{y}_{ij})$. In order to achieve reparametrization arclength invariance we introduce a function f operating on the determinant, for which $\mathbf{Q} = \mathbf{H}f(\det \mathbf{H})$, and consequently

$$\begin{aligned} \det \mathbf{Q} = \det \mathbf{H} \cdot f^2(\det \mathbf{H}) &= \frac{\det \mathbf{H}' \cdot f^2(\det \mathbf{H}')}{\det^2 \mathbf{J}} \\ &= \det\left(\mathbf{J}^{-\mathrm{T}}\mathbf{Q}'\mathbf{J}^{-1}\right). \end{aligned} \tag{34}$$

Using the fact that $(\mathbf{y}_1, \mathbf{y}_2) = \mathbf{J}(\mathbf{x}_1, \mathbf{x}_2)$, and the identities

$$\begin{aligned} \det\left(\mathbf{J}(\mathbf{x}_1, \mathbf{x}_2), \mathbf{x}_{ij}\right) &= \det \mathbf{J} \cdot \det(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_{ij}) \\ \det(\mathbf{x}_1, \mathbf{x}_2, a\mathbf{x} + b\mathbf{y}) &= (\mathbf{x}_1 \times \mathbf{x}_2) \cdot (a\mathbf{x} + b\mathbf{y}) \\ &= a(\mathbf{x}_1 \times \mathbf{x}_2) \cdot \mathbf{x} + b(\mathbf{x}_1 \times \mathbf{x}_2) \cdot \mathbf{y} \\ &= a\det(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}) + b\det(\mathbf{x}_1, \mathbf{x}_2, \mathbf{y}), \end{aligned} \tag{35}$$

we have

$$\begin{aligned} h'_{ij} &= \det \mathbf{J}\big(\xi_i\xi_j \det(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_{11}) \\ &\quad + (\xi_i\eta_j + \xi_j\eta_i)\det(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_{12}) \\ &\quad + \eta_i\eta_j \det(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_{22})\big) \\ &= \det \mathbf{J} \cdot \boldsymbol{\xi}_i^{\mathrm{T}}\mathbf{H}\boldsymbol{\xi}_j, \end{aligned} \tag{36}$$

where $\boldsymbol{\xi}_i^{\mathrm{T}} = (\xi_i, \eta_i)$. Hence,

$$\begin{aligned} \det \mathbf{H}' &= \det^2 \mathbf{J} \cdot \left(\boldsymbol{\xi}_1^{\mathrm{T}}\mathbf{H}\boldsymbol{\xi}_1\boldsymbol{\xi}_2^{\mathrm{T}}\mathbf{H}\boldsymbol{\xi}_2 - \boldsymbol{\xi}_1^{\mathrm{T}}\mathbf{H}\boldsymbol{\xi}_2\boldsymbol{\xi}_1^{\mathrm{T}}\mathbf{H}\boldsymbol{\xi}_2\right) \\ &= \det^2 \mathbf{J} \cdot \left(\boldsymbol{\xi}_1^{\mathrm{T}}\mathbf{H}\big(\boldsymbol{\xi}_1\boldsymbol{\xi}_2^{\mathrm{T}} - \boldsymbol{\xi}_2\boldsymbol{\xi}_1^{\mathrm{T}}\big)\mathbf{H}\boldsymbol{\xi}_2\right) \\ &= \det^3 \mathbf{J} \cdot \boldsymbol{\xi}_1^{\mathrm{T}}\mathbf{H}\begin{pmatrix} & 1 \\ -1 & \end{pmatrix}\mathbf{H}\boldsymbol{\xi}_2 \\ &= \det^3 \mathbf{J} \cdot \det(\mathbf{H}\boldsymbol{\xi}_1, \mathbf{H}\boldsymbol{\xi}_2) \\ &= \det^4 \mathbf{J} \cdot \det \mathbf{H}. \end{aligned} \tag{37}$$

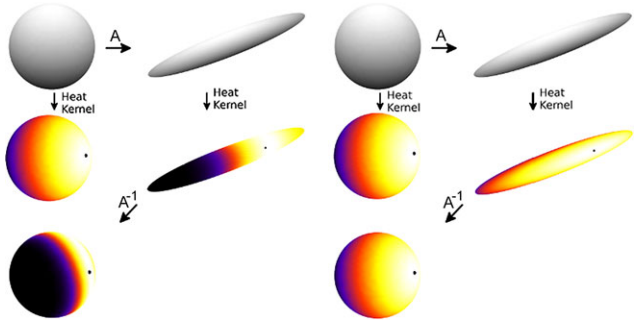

Fig. 3 Heat diffusion originated from one source on a sphere. *On the left* we used the standard intrinsic metric induced by the Euclidean metric of the embedding space, while *on the right* the equi-affine metric was used. Heat propagation is shown on the original mesh and after an affine transformation. For convenience of visualization, heat kernels are overlaid onto the untransformed shape

Substituting the latter into (34) yields

$$\det \mathbf{H} \cdot f^2(\det \mathbf{H}) \cdot \det^2 \mathbf{J} = \det^4 \mathbf{J} \cdot \det \mathbf{H} \cdot f^2\big(\det^4 \mathbf{J} \cdot \det \mathbf{H}\big), \tag{38}$$

which reduces to

$$f(\det \mathbf{H}) = \det \mathbf{J} \cdot f\big(\det^4 \mathbf{J} \cdot \det \mathbf{H}\big) \tag{39}$$

yielding an optional solution $f(t) = t^{-1/4}$. We conclude that the equi-affine arclength is given by the quadratic fundamental form [61, 64]

$$\mathbf{Q} = \mathbf{H} \cdot \det^{-1/4} \mathbf{H}, \tag{40}$$

To avoid complex normalization, Su [64] suggested to define $f(t) = |t|^{-1/4}$, for which

$$\mathbf{Q} = \mathbf{H} \cdot |\det \mathbf{H}|^{-1/4}, \tag{41}$$

and normalize the form by -1 if $\det \mathbf{H} < 0$, which leads to an invariance in case the new parameters have a different orientation. In Fig. 3, we depict a heat signature measured on a sphere originated from a single source point, using the induced Euclidean metric and the one created using $\mathbf{Q}$ defined above.

4.3 Equi-affine Metric

In order to construct a metric we propose to project the quadratic form (40) onto the space of positive definite matrices. A valid (pseudo-)metric is obtained in all cases except where the surface has a positive Gaussian curvature.

From the eigendecomposition of

$$\mathbf{H} \cdot |\det \mathbf{H}|^{-1/4} = \mathbf{U} \Gamma \mathbf{U}^{\mathrm{T}}, \tag{42}$$

Fig. 4 Three surfaces with similar equi-affine metrics

where $\mathbf{U}$ is orthonormal and $\Gamma = \mathrm{diag}\{\gamma_1, \gamma_2\}$, we compose a new metric $\bar{\mathbf{Q}}$, such that

$$\bar{\mathbf{Q}} = \mathbf{U} |\Gamma| \mathbf{U}^{\mathrm{T}}, \tag{43}$$

is positive semi-definite and equi-affine invariant, where $|\Gamma| = \mathrm{diag}\{|\gamma_1|, |\gamma_2|\}$. Note that Γ can be complex. One can think of this projection as a change of orientation.

For points with negative curvature the local metric has one positive and one negative eigenvalues. In this case we set the metric tensor to be zero. That is, such regions would practically be ignored by the proposed geometry. In the next section we elaborate on regularization of hyperbolic and parabolic points.

Under the proposed approach, non-isometric surfaces can now be intrinsically indistinguishable, and considered isometric. Let us provide a simple example for clarity. We denote by X^a, X^b and X^c, three surfaces

$$\begin{aligned} X^a(u,v) &= \big(u, v, u^2 + v^2\big) \\ X^b(u,v) &= \big(u, v, -u^2 - v^2\big) \\ X^c(u,v) &= \big(2u, v, u^2 + v^2\big), \end{aligned} \tag{44}$$

depicted in Fig. 4.

A simple calculation reveals that the Euclidean metrics of the three surfaces are

$$\mathbf{G}^a = \begin{bmatrix} 1+4u^2 & 4uv \\ 4uv & 1+4v^2 \end{bmatrix},$$

$$\mathbf{G}^b = \begin{bmatrix} 1+4u^2 & 4uv \\ 4uv & 1+4v^2 \end{bmatrix},$$

$$\mathbf{G}^c = \begin{bmatrix} 4+4u^2 & 4uv \\ 4uv & 1+4v^2 \end{bmatrix},$$

respectively. Not surprising X^a and X^b have similar metrics, and consequently they are isometric in a conventional manner, while X^c is not isometric to the rest. Evaluating the equi-affine quadratic forms involves the following normalization

$$\mathbf{Q}^a = \begin{bmatrix} \sqrt{2} & 0 \\ 0 & \sqrt{2} \end{bmatrix},$$

$$\mathbf{Q}^b = \begin{bmatrix} -\sqrt{2} & 0 \\ 0 & -\sqrt{2} \end{bmatrix},$$

$$\mathbf{Q}^c = \begin{bmatrix} \sqrt{2} & 0 \\ 0 & \sqrt{2} \end{bmatrix}.$$

Since we wish all three to be valid metrics while maintaining meaningful geometric features, we project the metrics onto semi-positive space, which yields

$$\bar{\mathbf{Q}}^a = \bar{\mathbf{Q}}^b = \bar{\mathbf{Q}}^c = \begin{bmatrix} \sqrt{2} & 0 \\ 0 & \sqrt{2} \end{bmatrix}. \tag{45}$$

We conclude the section by summarizing the main result: *For surfaces with positive Gaussian curvature, $\bar{\mathbf{G}}$ is equi-affine invariant, positive-definite, and parameterization invariant metric.*

4.4 Surfaces with Parabolic and Hyperbolic Points

Projecting the metric onto a positive semi-definite subspace, as shown in (43), is insufficient for parabolic and hyperbolic points, as we consider the metric to be zero for both. At such points, the metric (43) should be redefined. Numerically, we assign $(\bar{q}_{ij}) = \epsilon I = \epsilon \left(\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}\right)$ metric, where ϵ is a small constant.

In order to justify this regularization, let us examine the area measurement (determinant) of the proposed metric near parabolic points. Since the surface normal is

$$\bar{n} = \frac{\mathbf{x}_1 \times \mathbf{x}_2}{\|\mathbf{x}_1 \times \mathbf{x}_2\|}, \tag{46}$$

and the second fundamental form coefficients are

$$b_{ij} = \langle \mathbf{x}_{ij}, \bar{n} \rangle = \frac{\det(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_{ij})}{\sqrt{\det g}}, \tag{47}$$

we can write the equi-affine quadratic form as

$$\mathbf{Q} = \left(\frac{\det g}{\det b}\right)^{\frac{1}{4}} b. \tag{48}$$

This explains the problematic metric definition once the Gaussian curvature vanishes, as indeed

$$K = \frac{\det b}{\det g} = 0. \tag{49}$$

The local area of the equi-affine quadratic form can be written as

$$\det \mathbf{Q} = \left(\left(\frac{\det g}{\det b}\right)^{\frac{1}{4}}\right)^2 \det b = \sqrt{\det g \det b}, \tag{50}$$

and consequently the determinant of the equi-affine pseudo-metric becomes

$$\det \bar{\mathbf{Q}} = \sqrt{\det g |\det b|}. \tag{51}$$

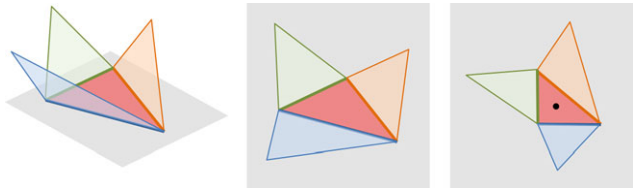

Fig. 5 The three neighboring triangles together with the central one are unfolded onto the plane. The *central triangle* is canonized into a right unit isosceles triangle while the rest of its three neighboring triangles follow the same planar affine transformation. Finally, the six surface coordinate values at the vertices are used to interpolate a quadratic surface patch from which the metric tensor is computed

We infer that assuming b and g are bounded, the local area is bounded and is well defined near points with zero Gaussian curvature. Thus, regularization would solve the degenerate cases introducing numerical inaccuracy at the order of ϵ.

5 Numerical Computation

5.1 Affine Invariant Metric

We represent the surface X as a triangular mesh, and consider each triangular face with the three neighboring faces. The four triangles are unfolded to the plane and transformed so that the central triangle becomes the canonical simplex in the plane. This constructs a local system of coordinates (Fig. 5), in which three quadratic functions describing a second-order patch are fitted to the $\mathbb{R}^3$ coordinates of the six triangle vertices, and evaluated at the barycenter of the central triangle (Algorithm 1). From the coefficients of these

Algorithm 1: Equi-affine invariant metric discretization

Input: 3×6 matrix $\mathbf{P}$ of triangle vertex coordinates in $\mathbb{R}^3$ (each column $\mathbf{P}_i$ represents the coordinates of a vertex, the first three columns belonging to the central triangle).

Output: 6×3 matrix of coefficients $\mathbf{D}$

1 Flatten the triangles to a plane, such that each vertex $\mathbf{P}_i$ becomes $\mathbf{Q}_i \in \mathbb{R}^2$, and (i) the first vertex becomes the origin, $\mathbf{C}_1 = [0\ 0]^T$; (ii) edge lengths are preserved, $d(\mathbf{C}_i, \mathbf{C}_j) = d(\mathbf{P}_i, \mathbf{P}_j)$ for all i and j; and (iii) the orientation is unchanged, $\text{sign}(\mathbf{C}_i^T \mathbf{C}_j) = \text{sign}(\mathbf{P}_i^T \mathbf{P}_j)$.

2 Construct a new parameterization $\hat{\mathbf{C}}_i = \mathbf{M}\mathbf{C}_i$, where $\mathbf{M} = [\mathbf{C}_2\ \mathbf{C}_3]^{-1}$.

3 Calculate the coefficients $\mathbf{D} = \mathbf{N}^{-1}\mathbf{P}^T$ of each coordinate polynomial, where $u = \hat{\mathbf{C}}_{i1}$, $v = \hat{\mathbf{C}}_{i2}$, and $\mathbf{N}$ is a 6×6 matrix with each row defined as $\mathbf{N}_i = [1\ u\ v\ uv\ u^2\ v^2]$.

Algorithm 2: Equi-affine pre-metric coefficients evaluation

Input: Coefficient matrix $\mathbf{D} = (\mathbf{d}_1, \ldots, \mathbf{d}_6)^{\mathrm{T}}$ from Algorithm 1.

Output: 2×2 matrix $\mathbf{Q}$ with the equi-affine quadratic form value at the barycenter of the central triangle.

1 Construct

$$\mathbf{x}_1 = (\mathbf{d}_2, 2\mathbf{d}_5, \mathbf{d}_4)(1, u, v)^{\mathrm{T}}$$
$$\mathbf{x}_2 = (\mathbf{d}_3, 2\mathbf{d}_6, \mathbf{d}_4)(1, u, v)^{\mathrm{T}}$$
$$\mathbf{x}_{11} = 2\mathbf{d}_5$$
$$\mathbf{x}_{22} = 2\mathbf{d}_6$$
$$\mathbf{x}_{12} = \mathbf{d}_4.$$

2 Evaluate the pre-metric tensor at $(u, v) = (0.5, 0.5)$

$$h_{ij} = \det(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_{ij})$$

3 Calculate the equi-affine pre-metric tensor as $\mathbf{Q} = |h_{11}h_{22} - h_{12}^2|^{-\frac{1}{4}} \cdot \mathbf{H}$.

second-order polynomials, the coefficients of the equi-affine quadratic form are computed (Algorithm 2), and projected onto the positive semi-definite cone.

5.2 Geodesics and Fast Marching

The computation of geodesic distances on a surface is used by discretizing the viscosity solution to the eikonal equation (9). Here, we adopt the family of fast marching algorithms simulating the propagation of a wavefront on the surface and recording its time of arrival at different vertices. Fast marching methods resemble in structure and properties the Dijkstra shortest path algorithm, with the important difference that unlike in graphs, the path is not restricted to the edges and can pass anywhere on the triangular faces of the mesh [33]. The numerical core of fast marching is the *update* step by which the values of the distance map known at two vertices of a triangle are propagated to the third one.

Typically, the geometry of the triangle itself is used to simulate propagation of the wavefront with unit speed. In our case, in order to enforce a different metric, we modify the geometry of the triangle being updated such that the Euclidean metric defined on it represents our equi-affine invariant metric at that point.

Since the equi-affine invariant arclength is given by $dp^2 = \mathbf{du}^{\mathrm{T}}\mathbf{Qdu}$, we can evaluate the equi-affine edges length of each triangle with respect to the local coordinates $(0, 0)$, $(1, 0)$ and $(0, 1)$, yielding $L_1^2 = g_{11}$, $L_2^2 = g_{22}$, and $L_3^2 = g_{11} - 2g_{12} + g_{22}$.

5.3 Laplace-Beltrami Operator Discretization

Many approaches to the discretization of the Laplace-Beltrami operators have been proposed in the literature [67]. Here, we adopt the finite elements method (FEM) discretization proposed in [25]. Given N points and their triangulation we translate the eigendecomposition of the Laplace-Beltrami operator (5) into a *weak form*

$$\int \psi_k \Delta\phi \, da = \lambda \int \psi_k \phi \, da, \tag{52}$$

with respect to some basis $\{\psi_k\}$ spanning a (sufficiently smooth) subspace of $L^2(X)$. Though, the basis functions can be of any order, here we limit our attention to linear ψ_k's obtaining a zero value in k's 1-ring and one at k itself.

Substituting these functions into (52) and performing derivation by parts yields [23]

$$\int_{\partial X} \psi_k \frac{\partial \phi}{\partial n} ds - \int \psi_k \Delta\phi \, da = \int \langle \nabla\psi_k, \nabla\phi \rangle_x \, da, \tag{53}$$

where n is the direction of the exterior normal. Inserting a vanishing Neumann boundary condition $\frac{\partial\phi(x)}{\partial n} = 0, \forall x \in \partial X$ into (53) yields the well known formulation

$$\int \psi_k \Delta\phi \, da = -\int g^{ij} (\partial_i \phi)(\partial_j \psi_k) \, da = \lambda \int \psi_k \phi \, da.$$

Note, that ψ_k is no longer required to satisfy the boundary condition, since it has become a part of the weak formulation. For that reason, the Neumann boundary condition is called a *natural boundary condition*, in contrast to the Dirichlet counterpart which is refereed to as an *essential boundary condition*.

Next, we approximate the eigenfunction ϕ as a combination of the ψ_l's,

$$\phi = \sum_{l=1} \alpha_l \psi_l, \tag{54}$$

which yields

$$\int g^{ij} \left(\partial_i \sum_l \alpha_l \psi_l \right) (\partial_j \psi_k) \, da = -\lambda \int \psi_k \sum_l \alpha_l \psi_l \, da, \tag{55}$$

or due to compactness,

$$\sum_l \alpha_l \int g^{ij} (\partial_i \psi_l)(\partial_j \psi_k) \, da = -\lambda \sum_l \alpha_l \int \psi_k \psi_l \, da. \tag{56}$$

Note that the latter integrals depend only on the basis functions ψ_k and the surface and can be therefore precomputed.

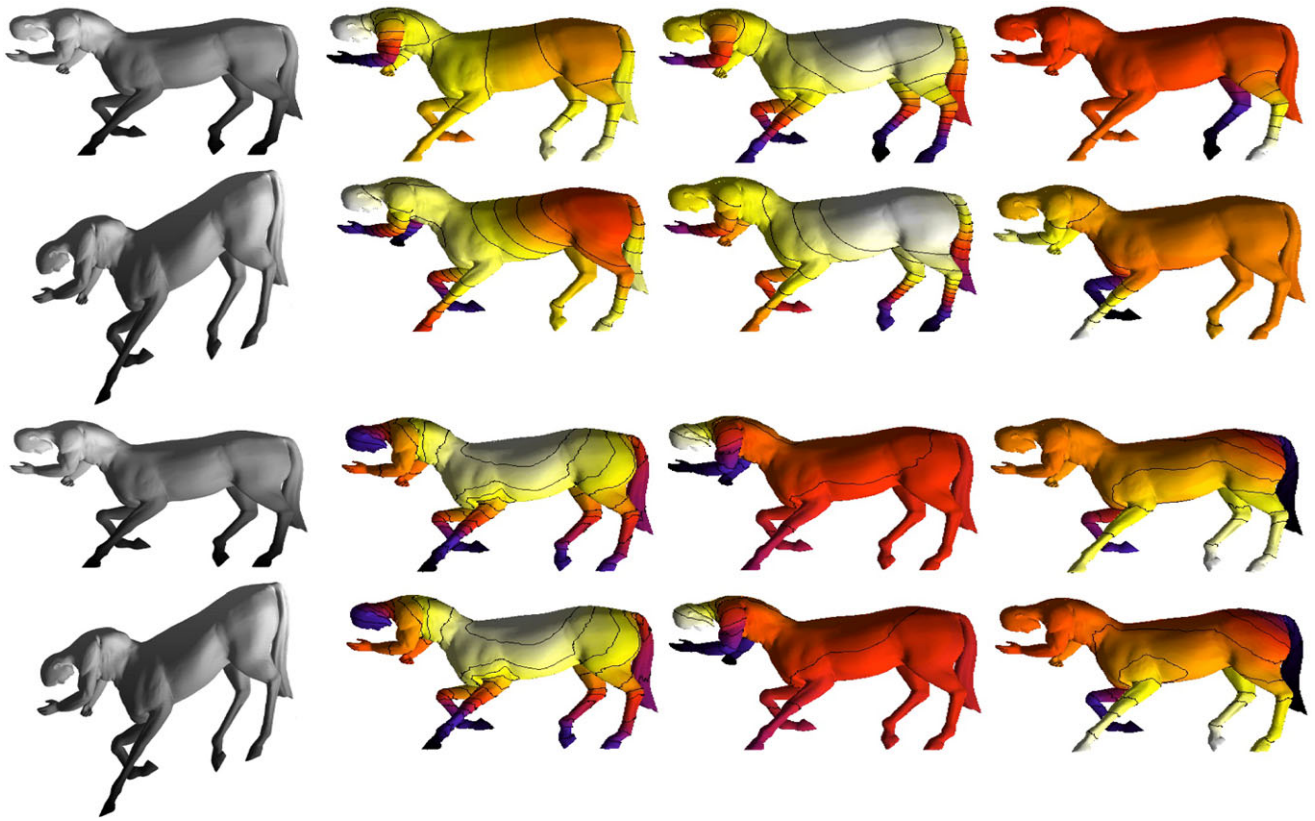

Fig. 6 Three eigenfunctions (2nd, 3rd, and 5th) of the standard (*first two rows*) and the proposed equi-affine invariant (*last two rows*) Laplace-Beltrami operators. The *first and the third rows* relate to the original shape, while the *second and forth rows* to its equi-affine transformation. For convenience of visualization, eigenfunctions are overlaid onto the untransformed shape

The last equation can be rewritten in a matrix form as the generalized eigendecomposition $\mathbf{A}\alpha = \lambda\mathbf{B}\alpha$ solved for the coefficients α_l [54], where

$$\begin{aligned} a_{kl} &= \int g^{ij}(\partial_i \psi_l)(\partial_j \psi_k)\, da, \\ b_{kl} &= -\int \psi_k \psi_l\, da, \end{aligned} \tag{57}$$

are the elements of the matrices $\mathbf{A}$ and $\mathbf{B}$, respectively. A visualization of several eigenfunctions can be seen in Fig. 6.

6 Synthetic Data

In this section, we show a few applications of our construction in non-rigid shape retrieval, correspondence, and partitioning. In all the applications, we compare the proposed equi-affine-invariant metric and the standard Riemannian metric induced by the Euclidean embedding.

6.1 Shape Retrieval

The bag of features approach for shape retrieval (dubbed as Shape Google) was shown in [43]. In this approach, local shape descriptors (HKS) were quantized into a "geometric vocabulary", creating for each shape a distribution of words from the vocabulary (bag of features). Comparing two bags of features allows to measure the similarity of the underlying shapes.

Here, we compared how Shape Google performs with HKS descriptors obtained using different metrics. The SHREC 2010 large-scale shape retrieval benchmark methodology [9] was used for evaluation. The dataset in the benchmark consisted of two parts: 793 shapes from 13 shape classes with simulated transformation of different types (Fig. 8) and strengths (total 60 transformations per shape, used as queries) and the remaining 521 shapes used as the queried corpus. Transformations classes *affine* and *isometry+affine* were added to the original SHREC query set, representing, respectively, equi-affine transformations of different strengths of the null shape and its approximate isometries.

Retrieval was performed by matching 780 transformed queries to the null shapes. Each query had one correct corresponding null shape in the dataset. Scale-space was sampled at six scales $t = 1024, 1351, 1783, 2353, 3104, 4096$. Bags of features were computed using soft vector quantization with variance taken as twice the median of all distances

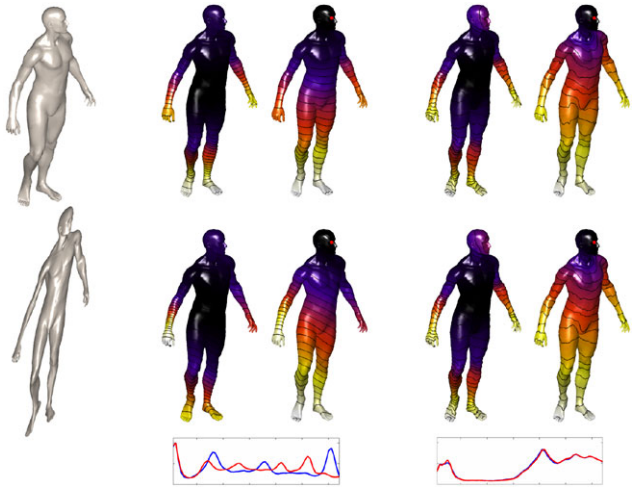

Fig. 7 Heat kernel signature $h_t(x,x)$ and diffusion metric ball around a point on the head obtained using the standard (*second and third columns*, respectively) and the equi-affine invariant (*fourth and fifth columns*, respectively) metrics. Two rows show a shape and its equi-affine transformation. For convenience of visualization, the heat kernels are overlaid onto the untransformed shape. *Plots below the figure* show the corresponding metric distributions before and after the transformation

between cluster centers in a vocabulary of 64 entries. Both the standard and the affine-invariant Laplace-Beltrami operator discretization were computed using finite elements, Heat kernels were approximated using the first smallest 100 eigenvectors and eigenvalues. Examples of heat kernel signatures obtained using two different metrics are shown in Fig. 7.

Performance was evaluated using precision/recall characteristic. *Precision* $\mathrm{P}(r)$ is defined as the percentage of relevant shapes in the first r top-ranked retrieved shapes. *Mean average precision* (mAP), defined as $\mathrm{mAP} = \sum_r \mathrm{P}(r) \cdot \mathrm{rel}(r)$, where $\mathrm{rel}(r)$ 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 retrieval performance ($mAP = 100\ \%$) is achieved when all queries return relevant first matches. Performance results were broken down according to transformation class and strength.

Tables 1–2 show that the equi-affine version of Shape Google is competitive to the one using the standard (embedding-induced) metric, and exhibits nearly perfect retrieval results in the equi-affine transformations class.

6.2 Voronoi Tessellation

The second application we show is the partitioning of non-rigid shapes. 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 on X closer to x_i than to any other x_j in the sense of the metric g. In other words, a point $x \in X$ belongs the i'th cell if $d_X(x,x_i) \leq d_X(x,x_j)\ j \neq i$, where $d_X : X \times X \to \mathbb{R}$ is the distance map induced by the Riemannian metric tensor g. Boundary points which have a similar distance to more than one centroid are considered to be in none or all the relevant cells. Voronoi tessellations are ubiquitous in geometry processing applications, and their definition is related to the metric used.

In our experiment, we compared the standard metric with the equi-affine metric for Voronoi tessellation generation. For each choice of the metric, we computed the associated Laplace-Beltrami operator and used its eigenvector and eigenvalues to compute the diffusion distance with a time scale $t = 10$. Thirteen points were picked on the shape using the farthest point sampling strategy, and tessellation was performed before and after applying an equi-affine transformation to the shape. Figure 9 clearly shows the invariance of the resulting tessellation in the case our metric is used.

Table 1 Performance of Shape Google with equi-affine HKS descriptors

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

Table 2 Performance of ShapeGoogle with Euclidean HKS descriptors

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

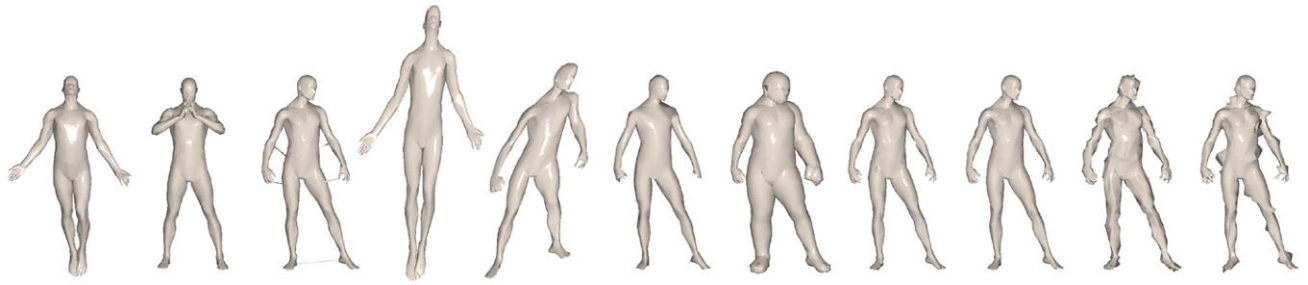

Fig. 8 Examples of query shape transformations used in the shape retrieval experiment (*left to right*): null, isometry, topology, affine, affine+isometry, sampling, local scale, holes, microholes, Gaussian noise, shot noise

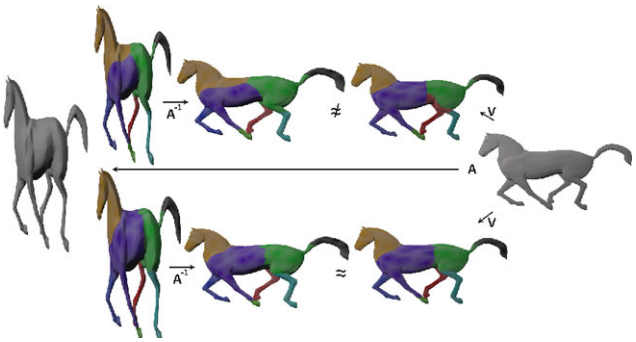

Fig. 9 Voronoi cells generated by a fixed set of 13 points on a shape undergoing an equi-affine transformation. We used diffusion distances with $t = 10$. The standard (*top*) and the equi-affine (*bottom*) metrics were used. Note that in the latter case, the tessellation commutes with the transformation

6.3 Non-rigid Matching and Intrinsic Symmetry

As the last application, we show an archetypal problem in shape analysis: intrinsic correspondence of two non-rigid shapes, and its particular case of self-correspondence (symmetry detection).

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 distance metrics on X, Y. In practice, no shapes are truly isometric, and such a correspondence rarely exists; however, one can attempt finding a correspondence minimizing the metric distortion.

We used the GMDS framework [12] to minimize (20). In our experiment, we initialized the algorithm manually (several heuristic algorithm can be used instead [24, 53, 69]). As distances in the GMDS algorithm, we used the geodesic distances computed using the standard and our equi-affine Riemannian metric.

Figures 10 and 11 show the correspondence between different poses of an armadillo with and without equi-affine stretching. It can be seen that correspondence based on the equi-affine metric performs better than the Euclidean one in the presence of bendings as well as equi-affine stretching.

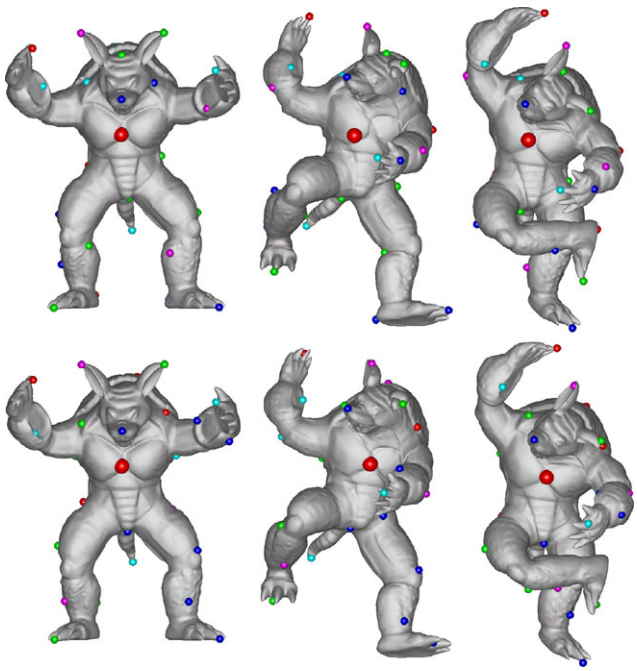

Fig. 10 Correspondence between different poses using the standard (*first row*) and the equi-affine (*second row*) metrics. It is clear that the equi-affine metric performs better than the Euclidean one in the presence of bendings. See for example the *enlarged red ball in the center of the body* that represents a corresponding point mapped from the original surface. (Color figure online)

An intrinsic symmetry is a particular case of shape correspondence with itself, defined as a minimum distortion self-embedding of a shape [47]. We used the same GMDS algorithm with different initialization in order to find the intrinsic reflective symmetry of the armadillo shape (Fig. 12).

7 Facial Morphometry

Facial morphometry is important in a variety of applications, such as diagnosis of generic disorders [28], statistical measurements of child growth [40], and construction of facial masks [42]. Evaluating geometric statistics of bendable and

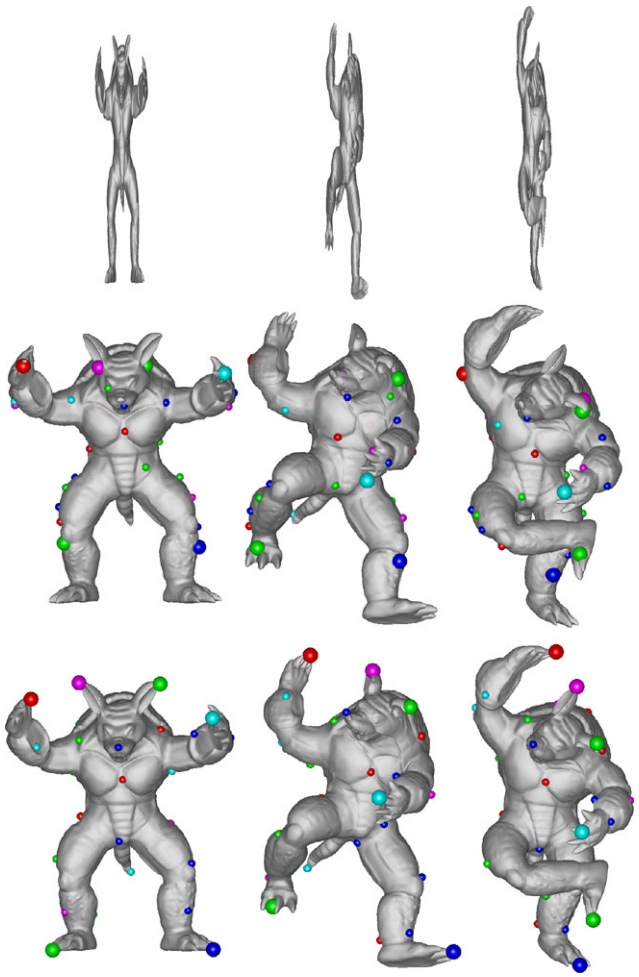

Fig. 11 Correspondence between different poses after stretching using the standard (*second row*) and the equi-affine (*third row*) metrics. Results based on the equi-affine metric are similar to Fig. 10 while the Euclidean metric suffers decrease in accuracy. See for example the *enlarged balls* that represent corresponding points in the limbs. For clarity we depicted the correspondence on the original shape, while a half size model of the *stretched shapes* appear in the *first row*

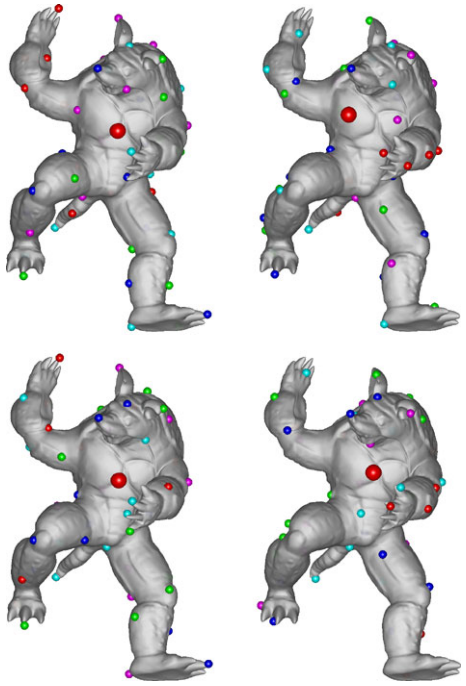

Fig. 12 Intrinsic symmetry of the armadillo shape using the Euclidean (*first row*) and the equi-affine (*second row*) metrics. It is clear that the equi-affine metric performs better than the Euclidean one in the presence of bendings. See for example the *enlarged red ball in the center of the body* that should keep its location under perfect symmetry. (Color figure online)

Fig. 13 Structured light scanner used to capture the facial geometry of newborns in the hospital

stretchable objects such as the face, forces us to align them within their ambient space.

A coarse correspondence between matching parts can be found based on local features. Usually, this alignment is performed using the albedo (texture) of the facial images, as it is easier than using geometry alone. Many models were proposed to provide dense geometric correspondence between surfaces. Most of these methods assume that the shapes, and in our case faces, are almost isometric with respect to the Euclidean ambient space. While this assumption could hold when considering facial surfaces of the same subject under various expressions [11], it would most probably fail when trying to align faces of different subjects. To overcome this problem, feature points are first matched, and then a smoothness term compensates for the stretches while interpolating the dense correspondence.

In this section, we argue that the equi-affine metric is more robust to local stretches, hence it is better suited for analyzing intrinsic geometric properties of faces, specifically when dense alignment is required. We further claim that diffusion distances, which are widely used in the computer vision community, are also applicable in medical applications when coupled with the equi-affine metric.

In the following experiments we scanned 50 newborn infants with a structured light camera (Fig. 13) and constructed a triangulated mesh with 5000 vertices for each child. We manually marked different facial features (Fig. 14) and evaluated the eigenvalues and eigenfunctions of the facial surface given the Euclidean metric and the equi-affine one.

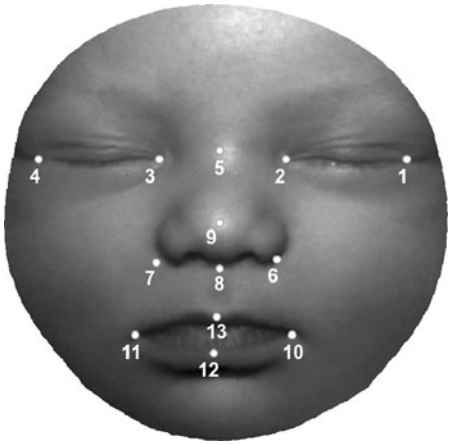

Fig. 14 Facial fiducials were marked by hand for evaluation of local stretches between faces. The inner canthis (P2-3) distance is used for global scaling

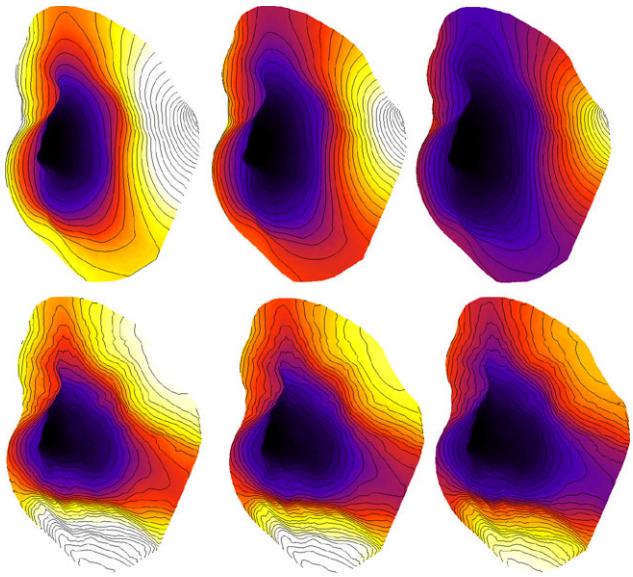

Fig. 15 Diffusion distances from the tip of the nose given three different times: *from left to right* 0.10, 0.15 and 0.20. In the *first row* we used the Euclidean metric and in the *second row* the equi-affine metric. We can see that the equi-affine metric is much more robust to small time shifts. Quantitative results appear in Fig. 16

We evaluated the diffusion distances from the tip of the nose to the entire face in several time scales as can be seen in Fig. 15. As each choice of time provides a different measurement of length, it is unclear which is the best time for a given application. The equi-affine invariant distances remain similar for small time changes, while the Euclidean invariant distances change dramatically. We provide a quantitative assessment in Fig. 16 where we show the distance distribution for three different times. In order to validate the affine invariant stability claim, we evaluated the distances between six different manually marked facial features. For each pair (50 × 49) we measured the ratio between those distances

Table 3 Mean and standard deviation of 50 × 49 distance pairs ratios after a global affine alignment was performed

	P2-3	P1-4	P6-7	P5-9	P1-6	P2-6
Mean	1.17	1.13	1.12	1.15	1.08	1.09
STD	0.39	0.65	0.23	0.21	0.19	0.10

Table 4 Time depended diffusion distances ratios between 50 × 49 pairs

Time	P2-3	P1-4	P6-7	P5-9	P1-6	P2-6
Diffusion distances ratios based on the Euclidean metric						
0.10	1.14	1.20	1.18	1.74	1.40	1.41
0.15	1.74	1.92	1.91	3.63	2.66	2.83
0.20	3.60	4.15	4.34	10.90	7.26	8.40
Diffusion distances ratios based on the equi-affine metric						
0.10	0.98	0.97	1.05	1.19	0.99	0.97
0.15	1.03	1.03	1.12	1.25	1.05	0.99
0.20	1.15	1.16	1.30	1.38	1.19	1.06

in various times that correspond to different scales. Results are presented in Fig. 17 and Table 4. The selected timescale for analysis is the dynamic range for which diffusion distances represent well the facial geometry. Smaller time values provide unstable results, while larger times fail to maintain meaningful information about the geometric structure. In Fig. 16 we see that in the effective dynamic range, the equi-affine distances contain meaningful information and are better spread than the Euclidean ones.

If the faces would have been transformed by a single affine transformation, then a global normalization would have been sufficient, which is rarely the case. In Table 3 we show the pairwise distortion error of straight Euclidean distances between the facial features after we performed a global affine alignment. We witness different distortions for different distances. Hence, faces are difficult to align using a single global affine transformation, while using local deformations, such as those supported by the equi-affine metric, provides a better alignment tool.

The advantages of the equi-affine metric are twofolds. First, it is much more robust to the time chosen in the diffusion geometry setting. Second, the ratio between the computed distances of corresponding points, based on the equi-affine invariant metric is closer to one. It means that any alignment algorithm which minimizes intrinsic properties will provide better results using the equi-affine metric.

Next, we compared the quality of facial alignment using the GMDS framework [13] given diffusion distances and commute time distances based on the Euclidean and Equi-affine metric. In Table 5 we show the distortion from ground truth of 45 comparisons between 10 faces. We define the error as the average Euclidean intrinsic (geodesics)

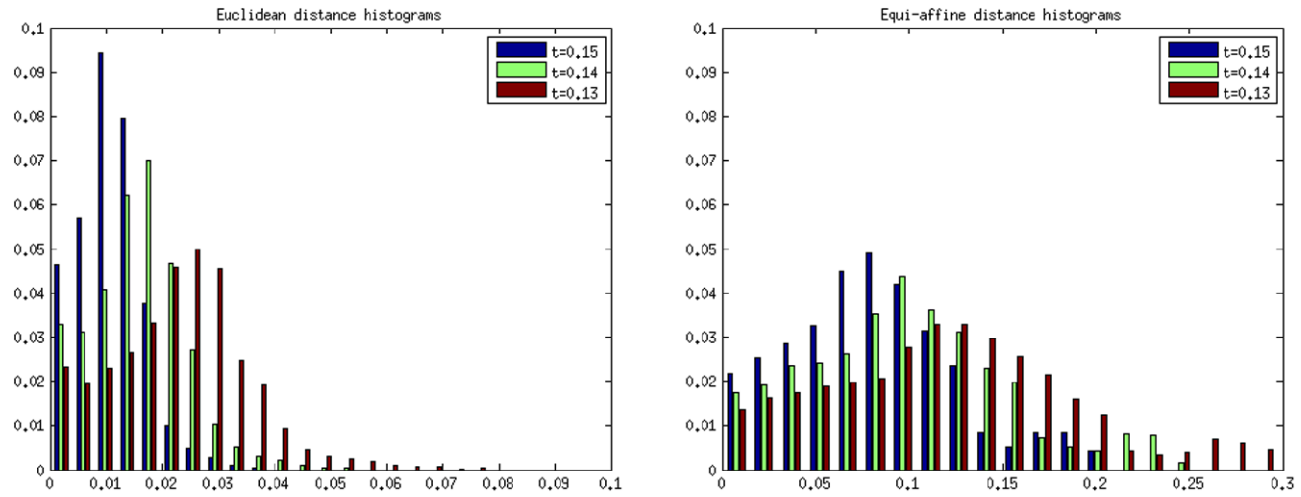

Fig. 16 Histogram of diffusion distances measured from the tip of the nose (Fig. 15) for three different times. We can see that the equi-affine metric is more robust to time changes

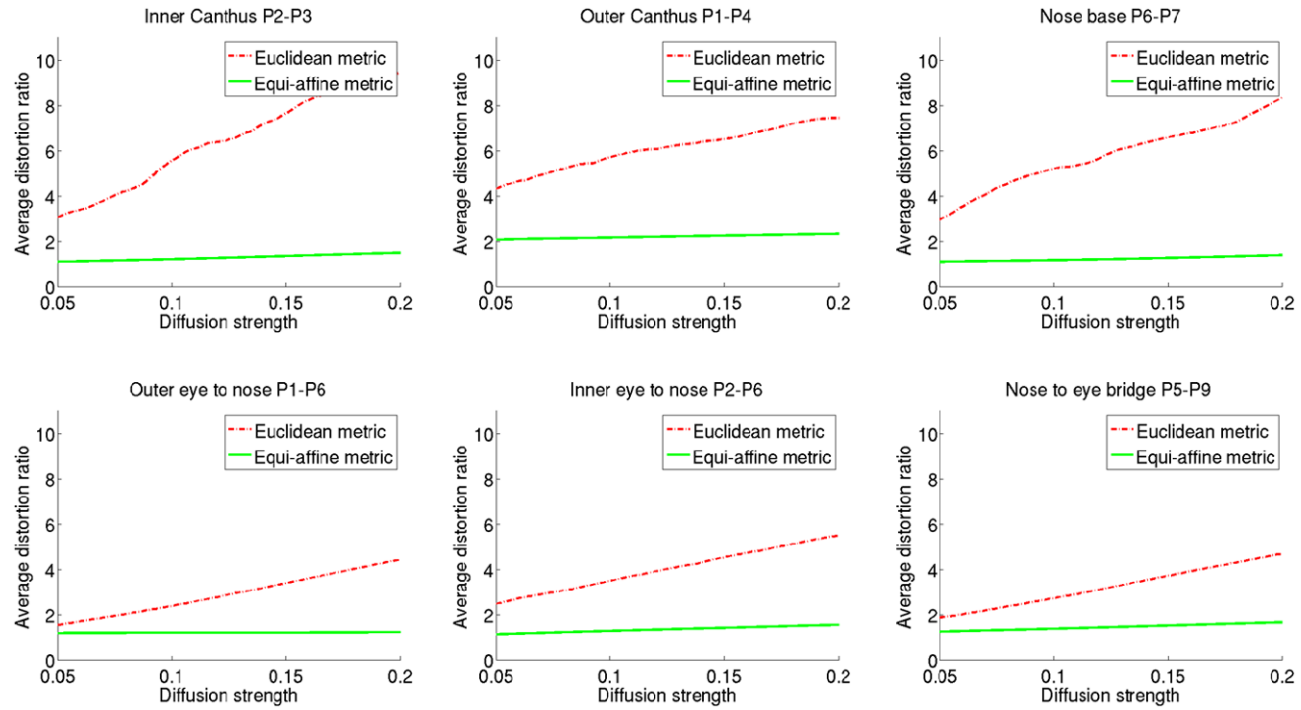

Fig. 17 A pairwise average ratios of diffusion distances between facial fiducials. We compared 6 different distances for every pair from our 50 newborns. A ratio of one means that the distances are almost the same. Due to local stretches the ratios are different between graphs. As can be seen, the Euclidean based distances are less robust to the choice of time, are different from pair to pair and have a higher variance then those measured with the equi-affine metric. We provide quantitative results in Table 4

shifts of marked points (1 to 9 from Fig. 14). In addition, we scaled all models such that the inner canthus distance will be 1. In the majority of our experiments the equi-affine metric significantly outperformed the Euclidean one. We summarize the results in Table 6. On average, the Euclidean metric showed more than 30 % distortions compared to the equi-affine one using commute time distance, and more than 120 % distortion for diffusion distances with time 0.1. Choosing a longer time introduced even larger distortions.

8 Conclusions

A computational framework for constructing equi-affine invariant distances for surfaces was presented. It was shown to

Table 5 Alignment distortion from ground truth using GMDS framework for 10 different faces (45 comparisons). *On the left* we present the Euclidean based distortion, and on the right the Equi-affine. The faces were normalized such that the inner canthus distance (points 2–3) is 1. *On the top table* we used diffusion distances, and on the bottom commute time distances

Euc/EA	1	2	3	4	5	6	7	8	9	10
Diffusion distances $t = 0.1$										
1	0/0	0.55/0.22	0.33/0.17	0.63/0.11	0.37/0.08	0.35/0.13	0.61/0.14	0.45/0.13	0.63/0.44	0.37/0.14
2		0/0	0.30/0.24	0.62/0.12	0.36/0.14	0.36/0.16	0.61/0.15	0.46/0.28	0.61/0.34	0.37/0.18
3			0/0	0.43/0.11	0.36/0.23	0.35/0.17	0.61/0.17	0.46/0.13	0.64/0.37	0.37/0.17
4				0/0	0.45/0.07	0.35/0.11	0.61/0.09	0.46/0.12	0.57/0.29	0.36/0.18
5					0/0	0.33/0.15	0.61/0.20	0.42/0.24	0.59/0.39	0.38/0.29
6						0/0	0.07/0.16	0.46/0.13	0.62/0.31	0.35/0.17
7							0/0	0.41/0.28	0.60/0.37	0.37/0.22
8								0/0	0.13/0.40	0.37/0.12
9									0/0	0.46/0.37
10										0/0
Commute time distances										
1	0/0	0.20/0.14	0.26/0.16	0.27/0.18	0.24/0.17	0.25/0.21	0.21/0.14	0.19/0.13	0.43/0.37	0.26/0.18
2		0/0	0.16/0.19	0.23/0.14	0.25/0.17	0.25/0.17	0.24/0.16	0.22/0.18	0.35/0.22	0.30/0.20
3			0/0	0.24/0.18	0.26/0.17	0.26/0.26	0.36/0.18	0.31/0.17	0.37/0.34	0.22/0.20
4				0/0	0.18/0.17	0.36/0.23	0.31/0.13	0.32/0.16	0.31/0.27	0.26/0.20
5					0/0	0.22/0.21	0.20/0.17	0.20/0.20	0.26/0.36	0.26/0.18
6						0/0	0.18/0.16	0.18/0.12	0.36/0.31	0.27/0.13
7							0/0	0.17/0.17	0.41/0.38	0.31/0.16
8								0/0	0.23/0.24	0.23/0.12
9									0/0	0.22/0.26
10										0/0

Table 6 Mean/median distortions of GMDS alignment from Table 5. The faces were normalized according to the inner canthus

Mean/median	Diffusion $t = 0.1$	Commute time
Euclidean	0.449/0.430	0.261/0.254
Equi-affine	0.203/0.170	0.199/0.182

gracefully handle degenerate cases. The proposed method allowed us to overcome semi-local equi-affine distortions without the need to explicitly estimate them during the analysis of non-rigid surfaces. The equi-affine transformations we can handle could be either semi-local or global, and act on either rigid or non-rigid shapes. In order to prove robustness we showed how to apply the equi-affine invariant metric to various analysis methods designed for non-rigid shapes. Finally, we explored the potential of the proposed metric in statistical analysis of facial morphometry. A limitation of the proposed metric is its sensitivity to scale. In a follow-up paper we plan to marry the proposed metric with the scale invariant one [1], and construct an affine invariant metric that can also handle scales.

Acknowledgements This research was supported by European Community's FP7-ERC program, grant agreement no. 267414. This research was mostly done while D.R. was a PhD candidate in the Technion, Israel.

References

1. Aflalo, Y., Kimmel, R., Raviv, D.: Scale invariant geometry for non-rigid shapes. SIAM J. Imaging Sci. (2013). Accepted
2. Alvarez, L., Guichard, F., Lions, P.L., Morel, J.M.: Axioms and fundamental equations of image processing. Arch. Ration. Mech. Anal. **123**, 199–257 (1993)
3. Amit, y., Let, F.: A non-linear variational problem for image matching. SIAM J. Sci. Comput. **15**, 207–224 (1994)

4. Andrade, M., Lewiner, T.: Affine-invariant curvature estimators for implicit surfaces. Comput. Aided Geom. Des. **29**(2), 162–173 (2012)
5. Beg, M.F.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. Int. J. Comput. Vis. **61**(2), 139–157 (2005)
6. Belkin, M., Niyogi, P.: Laplacian eigenmaps for dimensionality reduction and data representation. Neural Comput. **15**(6), 1373–1396 (2003)
7. Bérard, P., Besson, G., Gallot, S.: Embedding Riemannian manifolds by their heat kernel. Geom. Funct. Anal. **4**(4), 373–398 (1994)
8. Blaschke, W.: Vorlesungen uber Differentialgeometrie und geometrische Grundlagen von Einsteins Relativitatstheorie (1923)
9. 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. Workshop on 3D Object Retrieval (3DOR) (2010)
10. 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)
11. Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Three-dimensional face recognition. Int. J. Comput. Vis. **64**(1), 5–30 (2005)
12. Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Efficient computation of isometry-invariant distances between surfaces. SIAM J. Sci. Comput. **28**(5), 1812–1836 (2006)
13. Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. Proc. Natl. Acad. Sci. USA **103**(5), 1168–1172 (2006)
14. Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Robust expression-invariant face recognition from partially missing data. In: Proc. European Conf. Computer Vision (ECCV), pp. 396–408 (2006)
15. Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Rock, paper, and scissors: extrinsic vs. intrinsic similarity of non-rigid shapes. In: Proc. International Conference on Computer Vision (ICCV) (2007)
16. Bronstein, A.M., Bronstein, M.M., Kimmel, R., Mahmoudi, M., Sapiro, G.: A Gromov-Hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching. Int. J. Comput. Vis. **89**(2–3), 266–286 (2010)
17. Bronstein, A.M., Bronstein, M.M., Ovsjanikov, M., Guibas, L.J.: Shape Google: geometric words and expressions for invariant shape retrieval. ACM Trans. Graph. **30**(1) (2011)
18. Bronstein, M.M., Kokkinos, I.: Scale-invariant heat kernel signatures for non-rigid shape recognition. In: Proc. Computer Vision and Pattern Recognition (CVPR) (2010)
19. Bruckstein, A.M., Holt, R.J., Netravali, A.N., Richardson, T.J.: Invariant signatures for planar shape recognition under partial occlusion. CVGIP, Image Underst. **58**, 49–65 (1993)
20. Castellani, U., Mirtuono, P., Murino, V., Bellani, M., Rambaldelli, G., Tansella, M., Brambilla, P.: A new shape diffusion descriptor for brain classification. In: Proc. Medical Image Computing and Computer Assisted Intervention (MICCAI), pp. 426–433 (2011)
21. Coifman, R.R., Lafon, S.: Diffusion maps. Appl. Comput. Harmon. Anal. **21**, 5–30 (2006)
22. Craizer, M., Anciaux, H., Lewiner, T.: Discrete affine minimal surfaces with indefinite metric. Differ. Geom. Appl. **28**(2), 158–169 (2010)
23. Dierkes, U., Hildebrandt, S., Kuster, A., Wohlrab, O.: Minimal Surfaces I. Springer, Heidelberg (1992)
24. Dubrovina, A., Kimmel, R.: Matching shapes by eigendecomposition of the Laplace Belrami operator. In: Proc. Symposium on 3D Data Processing Visualization and Transmission (3DPVT) (2010)
25. Dziuk, G.: Finite elements for the Beltrami operator on arbitrary surfaces. In: Partial Differential Equations and Calculus of Variations, pp. 142–155 (1988)
26. 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)
27. Elad, A., Kimmel, R.: On bending invariant signatures for surfaces. IEEE Trans. Pattern Anal. Mach. Intell. **25**(10), 1285–1295 (2003)
28. Farkas, L.G.: Surface anatomy of the face in Down's syndrome: anthropometric proportion indices in the craniofacial regions. J. Craniofac. Surg. **12**(6), 519–524 (2001)
29. Gebal, K., Baerentzen, A.J., Aanaes, H., Larsen, R.: Shape analysis using the auto diffusion function. In: Proc. of the Symposium on Geometry Processing (SGP), pp. 1405–1413 (2009)
30. Ghosh, D., Amenta, N., Kazhdan, M.: Closed-form blending of local symmetries. In: Proc. Symposium on Geometry Processing (SGP) (2010)
31. Gromov, M.: Structures Métriques Pour les Variétés Riemanniennes. Textes Mathématiques, vol. 1 (1981)
32. Jermyn, I.H.: Invariant Bayesian estimation on manifolds. Ann. Stat. **33**(2) (2005)
33. Kimmel, R., Sethian, J.A.: Computing geodesic paths on manifolds. Proc. Natl. Acad. Sci. USA **95**(15), 8431–8435 (1998)
34. Kovnatsky, A., Bronstein, M.M., Bronstein, A.M., Kimmel, R.: Photometric heat kernel signatures. In: Proc. Scale Space and Variational Methods (SSVM) (2011)
35. Kurtek, S., Klassen, E., Ding, Z., Jacobson, S., Jacobson, J.L., Avison, M.J., Srivastava, A.: Parameterization-invariant shape comparisons of anatomical surfaces. Trans. Med. Imaging **30**(3) (2011)
36. Lévy, B.: Laplace-Beltrami eigenfunctions towards an algorithm that "understands" geometry. In: Proc. Shape Modeling and Applications (2006)
37. Mémoli, F.: Gromov-Hausdorff distances in Euclidean spaces. In: Proc. Non-rigid Shape Analysis and Deformable Image Registration (NORDIA) Workshop. Part of Computer Vision and Pattern Recognition (CVPR) (2008)
38. Mémoli, F.: Spectral Gromov-Wasserstein distances for shape matching. In: Proc. Non-rigid Shape Analysis and Deformable Image Registration (NORDIA) Workshop (2009)
39. Mémoli, F., Sapiro, G.: A theoretical and computational framework for isometry invariant recognition of point cloud data. Found. Comput. Math. **5**, 313–346 (2005)
40. Merlob, P.: Anthropometric measurements of the newborn infant (27 to 41 gestational weeks). Birth Defects Orig. Artic. Ser. **20**(7), 1–52 (1984)
41. Morel, J.M., Yu, G.: ASIFT: a new framework for fully affine invariant image comparison. SIAM J. Imaging Sci. **2**, 438–469 (2009)
42. Newhouse, M., Luder, A., Chalamish, A., Raviv, D., Kimmel, R., Amirav, I.: Design of aerosol facemasks for children using computerized 3d face analysis. J. Aerosol Med. Pulm. Drug Deliv. (2013). Accepted
43. Ovsjanikov, M., Bronstein, A.M., Bronstein, M.M., Guibas, L.J.: Shape Google: a computer vision approach to invariant shape retrieval. In: Proc. Non-rigid Shape Analysis and Deformable Image Alignment (NORDIA) (2009)
44. Ovsjanikov, M., Mérigot, Q., Mémoli, F., Guibas, L.J.: One point isometric matching with the heat kernel. In: Proc. Symposium on Geometry Processing (SGP), vol. 29, pp. 1555–1564 (2010)
45. Ovsjanikov, M., Sun, J., Guibas, L.J.: Global intrinsic symmetries of shapes. Comput. Graph. Forum **27**(5), 1341–1348 (2008)
46. Qiu, H., Hancock, E.R.: Commute times, discrete Green's functions and graph matching. In: International Conference on Image Analysis and Processing (ICIAP), pp. 454–462 (2005)

47. Raviv, D., Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Symmetries of non-rigid shapes. In: Proc. Workshop on Non-rigid Registration and Tracking Through Learning (NRTL). Part of International Conference on Computer Vision (ICCV) (2007)
48. Raviv, D., Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Full and partial symmetries of non-rigid shapes. Int. J. Comput. Vis. **89**(1) (2010)
49. 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)
50. 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)
51. 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)
52. Raviv, D., Bronstein, A.M., Bronstein, M.M., Kimmel, R., Sochen, N.: Affine-invariant geodesic geometry of deformable 3D shapes. Comput. Graph. **35**(3), 692–697 (2011)
53. Raviv, D., Dubrovina, A., Kimmel, R.: Hierarchical framework for shape correspondence. Numer. Math. Theor. Meth. Appl. **6**(1) (2013)
54. Reuter, M., Biasotti, S., Giorgi, D., Patanè, G., Spagnuolo, M.: Discrete Laplace–Beltrami operators for shape analysis and segmentation. Comput. Graph. **33**(3), 381–390 (2009)
55. Rustamov, R.M.: Laplace-Beltrami eigenfunctions for deformation invariant shape representation. In: Proc. Symposium on Geometry Processing (SGP), pp. 225–233 (2007)
56. Sabastian, T.M., Tek, H., Crisco, J.J., Kimia, B.B.: Segmentation of carpal bones from CT images using skeletally coupled deformable models. Med. Image Anal. **7**(1) (2003)
57. Sapiro, G.: Affine invariant shape evolutions. Ph.D. thesis, Technion—IIT (1993)
58. Schwartz, E.L., Shaw, A., Wolfson, E.: A numerical solution to the generalized mapmaker's problem: flattening nonconvex polyhedral surfaces. IEEE Trans. Pattern Anal. Mach. Intell. **11**, 1005–1008 (1989)
59. Sharma, A., Horaud, R.P., Cech, J., Boyer, E.: Topologically-robust 3D shape matching based on diffusion geometry and seed growing. In: Proc. Computer Vision and Pattern Recognition (CVPR) (2011)
60. Sivic, J., Zisserman, A.: Video Google: a text retrieval approach to object matching in videos. In: Proc. Computer Vision and Pattern Recognition (CVPR) (2003)
61. Sochen, N.: Affine-invariant flows in the Beltrami framework. J. Math. Imaging Vis. **20**(1), 133–146 (2004)
62. Spira, A., Kimmel, R.: An efficient solution to the eikonal equation on parametric manifolds. Interfaces Free Bound. **6**(3), 315–327 (2004)
63. Styner, M., Oguz, I., Xu, S., Brechbühler, C.: Framework for the statistical shape analysis of brain structures using SPHARM-PDM. Insight **1071** (2006)
64. Su, B.: Affine differential geometry. Science Press, Beijing (1983)
65. 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)
66. Tsitsiklis, J.N.: Efficient algorithms for globally optimal trajectories. IEEE Trans. Autom. Control **40**(9), 1528–1538 (1995)
67. Wardetzky, M., Mathur, S., Kälberer, F., Grinspun, E.: Discrete Laplace operators: no free lunch. In: Proc. Symposium on Geometry Processing (SGP) (2007)
68. Weber, O., Devir, Y.S., Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Parallel algorithms for approximation of distance maps on parametric surfaces. In: Proc. ACM Transactions on Graphics (SIGGRAPH) (2008)
69. Zhang, C., Bronstein, M.M., Bronstein, A.M., Paragios, N.: Discrete minimum distortion correspondence problems for non-rigid shape matching. In: Proc. Scale Space and Variational Methods (SSVM) (2011)
70. Zigelman, G., Kimmel, R., Kiryati, N.: Texture mapping using surface flattening via multi-dimensional scaling. IEEE Trans. Vis. Comput. Graph. **9**(2), 198–207 (2002)

Dan Raviv is a postdoctoral fellow in the Camera Culture Group in the Media Lab, MIT, Massachusetts Institute of Technology. He earned his Ph.D. in Computer Science Department, Technion, Israel institute of Technology working in GIP laboratory on invariant metrics of non-rigid shapes. He has a bachelor degree, Summa cum Laude, in both Mathematics and Computer Science, and is an alumnus of Technion s Excellence Program. Dan's research was granted several excellence prizes including the Gutwirth and Intel awards. Dan was elected best lecturer on several occasions including the prestigious Technion directorship prize for outstanding adjunct teacher.

Alexander M. Bronstein received the B.Sc. and M.Sc. from the Department of Electrical Engineering in 2002 and 2005, and Ph.D. from the Department of Computer Science, Technion in 2007. Since 2010, he is with the School of Electrical Engineering at Tel Aviv University. His main research interests are theoretical and computational methods in metric geometry and their application to problems in computer vision, pattern recognition, shape analysis, computer graphics, image processing, and machine learning. He has authored over 100 publications in leading journals and conferences, over a dozen of patents and patent applications, a monograph, and coedited several books. Dr. Bronstein's research was recognized by numerous awards, including the Hershel Rich Technion Innovation award, the Krill Prize by Wolf Foundation (2012), and the European Research Council (ERC) Startup Grant (2013). Highlights of his research were featured in CNN, Wired, and SIAM News. He co-chaired the IEEE International Workshop on Non-rigid shapes and deformable image alignment (NORDIA) in 2008–2011, the International Conference on n Scale Space and Variational Methods in Computer Vision (SSVM) in 2011, served as the program chair of the Eurographics Workshop on 3D Object Retrieval (3DOR) in 2012, area chair of the IEEE Asian Conference on Computer Vision (ACCV) in 2010, and participated in program committees of major conferences in his field. Dr. Bronstein held visiting appointments in Politecnico di Milano (2008), Stanford university (2009), Verona University (2010), and Duke University (2013). In addition to his academic activities, he was a co-founder of a Silicon Valley startup Novafora, Inc., where he served from 2004 till 2009 as a scientist and a Vice President of video technology, leading a group of researchers and engineers in developing novel Internet-scale video analysis technologies. Dr. Bronstein was a principal technologist and an inventor of the 3D acquisition technology in the foundation of the Israeli startup company Invision that was acquired by Intel Corporation in 2011.

Michael M. Bronstein is an assistant professor in the Institute of Computational Science at the Faculty of Informatics, University of Lugano (USI), Switzerland. Before joining USI, he held a visiting appointment at Stanford university. Michael Bronstein received the B.Sc. summa cum laude (2002) from the Department of Electrical Engineering and Ph.D. with distinction (2007) from the Department of Computer Science, Technion—Israel Institute of Technology. He has authored a book and over 70 publications in leading journals and conferences, mostly on theoretical and computational methods in metric and spectral geometry and their application to problems in computer vision, pattern recognition, and shape analysis. His research was featured in international news such as CNN, and recognized by numerous awards, including the Hershel Rich Technion Innovation award (2003), Adams fellowship (2005) and most recently the ERC starting grant (2012). Besides academic work, Prof. Bronstein is actively involved in the development and transfer of technology and is the inventor of over a dozen of patents. His industrial track record includes the start-ups Novafora (2004–2009 as co-founder and VP of video technology) and Invision (2009–2012 as one of the principal technologists). Following the acquisition of Invision by Intel in 2012, Michael Bronstein currently also serves as research scientist at Intel.

Dan Waisman is head of the Newborn Unit, Department of Neonatology, Carmel Medical Center and Clinical Assistant Professor in the Faculty of Medicine, Technion—Israel Institute of Technology. Involved in basic translational research, and in development and application of new technologies in the area of experimental and clinical medicine.

Nir Sochen finished his B.Sc studies in Physics 1986 and his M.Sc. studies in theoretical physics 1988 both in the University of Tel-Aviv. Received his PhD in Theoretical physics 1992 from the Universite de Paris-Sud while conducting his research in the Service de Physique Theorique at the Centre d'Etude Nucleaire at Saclay, France. He continued with a one year research in the Ecole Normale Superieure in Paris on the Haute Etude Scientifique fellowship, and a three years NSF fellowship in the physics department of the University of California at Berkeley. It is in Berkeley that his interests shifted from quantum field theories and integrable models, related to high-energy physics and string theory, to computer vision and image processing. He spent one year in the physics department at the University of Tel-Aviv and two years in the Faculty of Electrical Engineering in the Technion-Israel Institute of Technology. Currently he is a senior lecturer in the Department of Applied Mathematics, University of Tel-Aviv. He is also a member of the Ollendorf center in the Technion. His main research interests are the applications of differential geometry and statistical physics in image processing and computational vision.

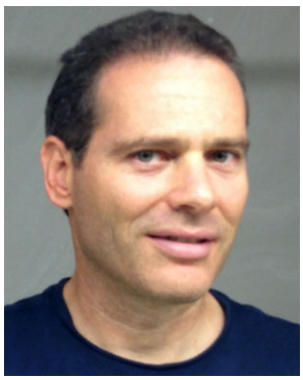

Ron Kimmel is a Professor of Computer Science at the Technion where he holds the Montreal Chair in Sciences. He held a post-doctoral position at UC Berkeley and a visiting professorship at Stanford University. He has worked in various areas of image and shape analysis in computer vision, image processing, and computer graphics. Kimmel's interest in recent years has been non-rigid shape processing and analysis, medical imaging and computational biometry, numerical optimization of problems with a geometric flavor, and applications of metric geometry and differential geometry. He is an IEEE Fellow for his contributions to image processing and non-rigid shape analysis. He is an author of two books, an editor of more than a couple, and an author of numerous articles. He is the founder of the Geometric Image Processing Lab as well several successful shape acquisition, processing and analysis companies.