Abstract Surfaces in Image Processing

E. Meinhardt-Llopis, G. Facciolo

March 27, 2014

Abstract

In this article we review the typical uses of Riemannian geometry in image processing and we propose new applications arising from natural geometric operations. Abstract surfaces are a common tool in image processing, where a metric in the plane is built using the gradient of an image, thus defining a Riemannian manifold. The most famous examples are active contours, which are are closed geodesics, and anisotropic diffusion, which is the heat equation on the manifold. In this context, we investigate Gaussian curvature, Ricci flow, and the spectral theory of the Laplace-Beltrami operator. Each of these constructions leads to a new tool for image processing.

1 Introduction

The purpose of this article is to explore the applications of Riemannian geometry to image processing. Some of these applications, such as active contours [1], image-guided diffusion [2] or Cheeger cuts [3] are already standard tools in image processing. However, there are still many standard ideas from Riemannian geometry [4] that are seldom, if ever, used in the context of image processing. For example, Gaussian curvature, the Laplace-Beltrami operator or Ricci flow. Thus the main goal of this article is to investigate how useful these ideas are as image processing tools.

The focus of this article is geometric rather than analytical. Thus, we do not worry about questions of convergence. However, we have implemented all the proposed techniques and the source code of the experiments is available as supplementary material.

In this article, an image determines an *abstract* surface; not the graph $(x, y, I(x, y))$, but an abstract manifold (M, g) where M is a rectangle and g is a metric field, typically defined from the gradient of I. In particular our surfaces need *not* be embedded in any higher-dimensional space. All the described geometric techniques are already in use for surfaces embedded in 3D. The novelty of the present article is to use them for the abstract surfaces that are defined by images.

Plan of the paper: §2 Summary of Riemannian geometry. §3 Examples of metric tensors. §4 Common uses of Riemannian geometry in image processing.

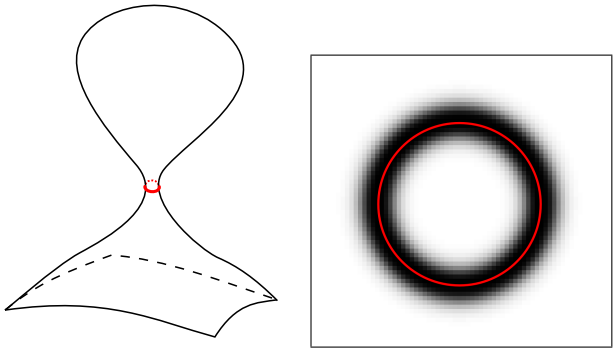

Figure 1: Two interpretations of the same Riemannian manifold: a complicatedly curved submanifold of a flat space, or a flat manifold with complicated metric field. In both cases, a closed geodesic is shown. The second interpretation is useful in image processing, where the geodesics are called active contours.

§5 New applications.

2 Summary of Riemannian geometry

On this section we remind the basic constructions of Riemannian geometry [4].

Let M be a flat surface, for example a rectangle or the whole plane. A *Riemannian metric* assigns to each point $p \in M$ a symmetric and positive definite matrix $g_p = \left(\begin{smallmatrix} E & F \\ F & G \end{smallmatrix}\right)$. This matrix defines an inner product on the space of all vectors emanating from point p. Thus, if x_p and y_p are vectors at point p, their inner product is

$$g_p(x_p, y_p) := x_p^T \cdot g_p \cdot y_p$$

and the norm of a vector is

$$\|x_p\|_g := \sqrt{g_p(x_p, x_p)}.$$

The inner product allows to compute the angle between two vectors at the same point, but the angle between two vectors at different points, is not defined. This is **the most important difference** with respect to the Euclidean plane, where vectors based at different points can be compared directly. Henceforth, the set M equipped with the metric g, will be called *the manifold*. It is a Riemannian manifold of dimension 2, also called Riemannian surface (but not a Riemann surface, which is a different thing requiring a complex structure).

2.1 Length, area and distance

The *length* of a curve $\gamma : [0, 1] \to M$ is

$$\text{length}(\gamma) := \int_0^1 \|\gamma'(t)\|_g \, \mathrm{d}t. \tag{1}$$

The length is independent of the parametrization. The *action* of the curve is

$$\text{action}(\gamma) := \int_0^1 \|\gamma'(t)\|_g^2 \, \mathrm{d}t$$

The length is independent of the parametrization, but the action depends on the parametrization. Given two points $p, q \in M$ their *distance* is defined as

$$\text{dist}(p, q) := \inf_{\gamma} \ \text{length}(\gamma)$$

where the inf is taken among all the curves γ joining p and q. With this distance the manifold has the structure of a metric space. Finally, the *area* of a region $B \subseteq M$ is defined as

$$\text{area}(B) := \int_B \sqrt{EG - F^2} \, \mathrm{d}x\mathrm{d}y.$$

With this area measure, the manifold has the structure of a measure space.

These seemingly innocent expressions encode a lot of notation. The length of a curve $(x(t), y(t))$ is, fully expanded,

$$\int_0^1 \sqrt{E(x(t), y(t))\dot{x}(t)^2 + 2F(x(t), y(t))\dot{x}(t)\dot{y}(t) + G(x(t), y(t))\dot{y}(t)^2} \, \mathrm{d}t.$$

A moderately verbose notation, useful for practical computations, is the following

$$\text{length} = \int_0^1 \sqrt{E\dot{x}^2 + 2F\dot{x}\dot{y} + G\dot{y}^2}$$

or even

$$\text{length} = \int \mathrm{ds}^2$$

where

$$\mathrm{ds}^2 = E\mathrm{d}x^2 + 2F\mathrm{d}x\mathrm{d}y + G\mathrm{d}y^2$$

2.2 Geodesics

Geodesics are curves of locally minimal length. This means that a curve is a geodesic when it can not be shortened by making a small deformation. Geodesics parametrized by arc-length satisfy the following differential equations:

$$\begin{cases} \ddot{x} &= \Gamma^x_{xx}\,\dot{x}^2 \ + \ 2\Gamma^x_{xy}\,\dot{x}\dot{y} \ + \ \Gamma^x_{yy}\,\dot{y}^2 \\ \ddot{y} &= \Gamma^y_{xx}\,\dot{x}^2 \ + \ 2\Gamma^y_{xy}\,\dot{x}\dot{y} \ + \ \Gamma^y_{yy}\,\dot{y}^2 \end{cases}$$

where the Christoffel symbols Γ^i_{jk} are expressions containing E, F, G and their partial derivatives with respect to x and y. There may be many geodesics of different length joining the same pair of points; but the shortest curve that joins two points, when it exists, is always geodesic.

2.3 The Laplace-Beltrami operator

The differential operators of vector calculus have to be adapted to the non-Euclidean settings. The differential operators of the gradient, divergence and Laplacian need to be adapted to the non-Euclidean setting. For example, we want the gradient of a function f to be a vector perpendicular to the level lines of f (where perpendicularity is defined using the metric tensor). There are many equivalent definitions, but for computational purposes we content ourselves with their coordinate expressions:

$$\operatorname{grad}(f) := \left(\frac{Gf_x - Ff_y}{EG - F^2},\ \frac{Ef_y - Ff_x}{EG - F^2} \right)$$

$$\operatorname{div}(p, q) := \frac{\left(p\sqrt{EG - F^2}\right)_x + \left(q\sqrt{EG - F^2}\right)_y}{\sqrt{EG - F^2}}$$

and the Laplacian, which is called the Laplace-Beltrami operator, is defined as the divergence of the gradient:

$$\Delta_g f := \frac{1}{\sqrt{EG - F^2}} \left[\left(\frac{f_x G - f_y F}{\sqrt{EG - F^2}} \right)_x + \left(\frac{f_y E - f_x F}{\sqrt{EG - F^2}} \right)_y \right]$$

These operators describe many geometric constructions on the manifold. For example, the distance function to any subset of the manifold satisfies the Eikonal equation $|\operatorname{grad}(u)| = 1$. The diffusion of information in the manifold is governed by the heat equation $u_t = \Delta_g u$. Similarly, the wave, Poisson and Schrdinger equations on the manifold can be defined, and so on.

The operator $-\Delta_g$ is linear and, assuming Neumann boundary conditions, self-adjoint with respect to the area element ($\int p \Delta_g q \, \mathrm{d}A = \int q \Delta_g p \, \mathrm{d}A$). Moreover, it is positive-definite. Thus, it has a sequence of eigenvalues $0 = \lambda_1 < \lambda_2 < \cdots$ which, unless the manifold has symmetries, are all different. The eigenfunctions $\{\varphi_k\}$ satisfy $\Delta_g \varphi_k = \lambda_k \varphi_k$ and are an orthogonal basis of smooth functions on M. When the metric is conformal to the Euclidean (meaning that $F = 0$ and $E = G$), then $\Delta_g f = \frac{1}{E}(f_{xx} + f_{yy})$.

2.4 Gaussian curvature

Gaussian curvature measures how non-Euclidean the surface is at each point. Typically, a surface has regions of positive curvature (called elliptical points) and regions of negative curvature (called hyperbolic points) separated by curves of zero curvature (called parabolic curves). Some degenerate cases are possible, for example isolated parabolic points or regions of vanishing curvature, which are called flat.

There are many equivalent definitionsof curvature which provide different insights to its meaning [5]. For example, the relative difference of the perimeter of a circle of radius ϵ with $2\pi\epsilon$, as ϵ tends to zero. For our purposes, a coordinate expression is more manageable. Due to the many symmetries of the Christoffel

symbols there are a lot of different coordinate expressions. The shortest one that we could find is the following one: A possible coordinate expression is the following:

$$K = \operatorname{div}\left(\frac{-\Gamma^y_{xy}}{E},\ \frac{\Gamma^y_{xx}}{E}\right)$$

In the particular case of a metric conformal to the Euclidean, Gaussian curvature has the following form

$$K = -\frac{1}{2}\Delta \log E \tag{2}$$

Points where $K = 0$ are called *parabolic points*.

2.5 Geodesic curvature

The geodesic curvature of a curve is a measure of how far the curve is to being a geodesic. If the curve is parametrized by arc-length, the geodesic curvature is its acceleration. In general coordinates, the geodesic curvature of a parametrized curve $\gamma(t) = (x(t), y(t))$ is given by

$$\kappa(\gamma) := \frac{g(\nabla_{\gamma'}\gamma'\ ,\ J\gamma')}{\|\gamma'\|^3_g}$$

where $\nabla_{\gamma'}\gamma'$ is the intrinsic acceleration, given by

$$\nabla_{\gamma'}\gamma' := \begin{pmatrix} \ddot{x} + \Gamma^x_{xx}\,\dot{x}^2 \ + \ 2\Gamma^x_{xy}\,\dot{x}\dot{y} \ + \ \Gamma^x_{yy}\,\dot{y}^2 \\ \\ \ddot{y} + \Gamma^y_{xx}\,\dot{x}^2 \ + \ 2\Gamma^y_{xy}\,\dot{x}\dot{y} \ + \ \Gamma^y_{yy}\,\dot{y}^2 \end{pmatrix}$$

Notice that the geodesic equations say exactly $\nabla_{\gamma'}\gamma' = 0$, thus geodesics have vanishing geodesic curvature. The geodesic curvature of the coordinate axis has a particularly simple form:

$$\kappa(\text{horizontal}) = \Gamma^y_{xx}\frac{\sqrt{EG - F^2}}{E\sqrt{E}}$$

$$\kappa(\text{vertical}) = -\Gamma^x_{yy}\frac{\sqrt{EG - F^2}}{G\sqrt{G}}.$$

2.6 Ricci flow

So far we have worked on a manifold with a single, fixed metric. The Ricci flow is a *geometric flow* by which the metric itself evolves according to a diffusion equation, and ultimately converges to a metric of constant curvature. It is defined by the following system of four PDE:

$$g_t = -Kg$$

Where K is the Gaussian curvature of the metric g. To obtain this simple form [6] we have to assume periodic boundary conditions (so that M is a flat

torus). The matrix $\mathrm{Ric} = Kg$ is called the Ricci tensor, hence the name. It is known [6] that Ric is approximated to first order by $-\Delta g$, where the Laplace-Beltrami acts componentwise on g. Thus the Ricci flow is well approximated by the heat equation acting on the metric $g_t = \Delta g$. However, this interpretation serves only as a heuristic guide, since the approximation is not stable over time. For the actual computations, it is better to stick to the initial description. Yet, the current numerical methods for solving this highly nonlinear equation are expectedly unstable [7].

3 Metrics

In this section we summarize some metrics that can be defined from a base image I.

The simplest metrics are **conformal** to the Euclidean, which have the form $g = h(\|\nabla I\|)\left(\begin{smallmatrix}1 & 0\\ 0 & 1\end{smallmatrix}\right)$, where h is a decreasing function, such as $h(s) = \frac{1}{s}$ or $h(s) = e^{-s^2/2\sigma_2}$. Another example are **anisotropic** metrics, like

$$g = \nabla I \otimes \nabla I \;+\; \lambda \nabla I^{\perp} \otimes \nabla I^{\perp}$$

Here λ is intended to be a small number, or $h(\|\nabla I\|)$.

Any embedding $\phi : M \hookrightarrow \mathbf{R}^n$ of the image domain defines a metric, **induced** by the Euclidean metric of $\mathbf{R}^n$:

$$g = \begin{pmatrix} \phi_x \cdot \phi_x & \phi_x \cdot \phi_y \\ \phi_x \cdot \phi_y & \phi_y \cdot \phi_y \end{pmatrix} \tag{3}$$

where the dots represent the Euclidean scalar product in $\mathbf{R}^n$. For example, the graph of the image intensity is an embedding into $\mathbf{R}^3$, and the RGB colors are an embedding into $\mathbf{R}^5$. Local patches also define an embedding. For example, patches of size 7×7 give an embedding into $\mathbf{R}^{49}$. The matrix of the induced metric coincides with the **local structure tensor** computed using 7×7 square neighborhoods.

Finally, the **affine invariant gradient** [8] is defined as

$$\nabla_{\mathrm{aff}} I := \frac{|\det(\mathrm{HI})|}{|\nabla I^T \cdot HI \cdot \nabla I|}$$

where HI denotes the Hessian matrix of I. This is an affine invariant quantity that can be used to replace the gradient magnitude $\|\nabla I\|$ in the previous formulas.

4 Explicit uses of Riemannian geometry in image processing

In this section we recall four common and well-established uses of Riemannian metrics in image processing.

(a)  (b) 

Figure 2: Application of Riemannian thin plates to the problem of colorization. (a) Color traces drawn on a grayscale image. (b) Solution of the problem described in equation (5)

4.1 Active contours

Snakes [9, 1] or active contours were introduced as energy-minimizing curves on an abstract surface. In the original description, the manifold is vividly called a *snake pit*, where a user is supposed to draw a curve (the "snake") which will be the starting point for minimizing the energy. The general snake energy $E(\gamma) = \int L(\gamma, \dot{\gamma}, \ddot{\gamma})$, includes our geodesic length (1) as a particular case when $L = \|\dot{\gamma}\|_g$.

Independently, it was realised the convenience to represent the evolution of curves implicitly; thus, instead of evolving a curve expressed parametrically, the equations describe the evolution of a function whose zero-level set is the curve. This was first done for curves evolving under mean-curvature motion [10]. Finally, *geodesic active contours* [1] were introduced as a general settings to describe the evolution of implicit curves in an arbitrary Riemannian manifold towards a closed geodesic. Notice however that closed geodesics are much less general than arbitrary geodesics, and implicit methods can only find closed geodesics. In the context of image processing, active contours are used to find smooth curves that follow the edges of an image I. Thus, curves following the edges need to have a much smaller energy that curves that cross through flat regions. This means that *the metric has to be small on the parts of high gradient of* I. An ideal candidate is then the function $\frac{1}{\|\nabla I\|}$, which has the desirable feature of being infinite on constant regions. In practice, it is better to use a bounded decreasing function of the gradient (see section 3).

4.2 Anisotropic diffusion

Anisotropic diffusion [11, 2] smooths an image without blurring the edges. Many kinds of anisotropic diffusion correspond [12] to the heat equation on a certain manifold (M, g), the metric tensor g being the diffusion tensor:

$$\begin{cases} u_t = \Delta_g u \\ u(0) = I \end{cases} \tag{4}$$

In the conformal case, $\Delta_g = \frac{1}{E}\Delta$ (notice that E appears *outside* the Euclidean Laplacian). Notice that the metric remains constant along all the evolution, so that the original Perona-Malik diffusion, which uses a time-varying metric, can not yet be described in this framework.

The purpose of anisotropic diffusion is to avoid diffusion through edges. Thus, the metric *has to be large on the parts of high gradient of* I. This is the opposite as for the computation of geodesics. In fact, looking at the units of each term in a physical setting, for diffusion it makes sense to use the *dual* metric, whose matrix is the inverse of g.

4.3 Cheeger cuts

Normalized cuts [3] are a popular graph-based image segmentation technique. They can be interpreted as a discretization [13] of an isoperimetric problem on a manifold (M, g). The resulting techniques for this and related problems have been widely used for image processing [14].

Weights are assigned to the edges of a graph according to the color difference joined pixels. Then, a relaxed discrete optimization is used to find the partition of the vertex whose cut (sum of edge weights) has the minimal cost (appropriately normalized by the volume of each part).

Thanks to the discretization technique used in graph-cuts [13], any Riemannian metric can be discretized into a graph so that cuts of the graph have a weight that approximates geodesic length. Thus, normalized cuts are readily interpretable in a continuous setting (M, g) as an isoperimetric problem

$$\min_{A \subseteq M} \frac{\text{length}(\partial A)}{\min(\text{area}(A), \text{area}(M \backslash A))}$$

The resulting techniques of this and related problems have been widely used for image processing [14].

4.4 Geodesic Voronoi segmentations

A typical tool in surface segmentation are Voronoi diagrams [15]. The same techniques are used on the abstract surfaces of images [16, 17]. The resulting fast marching algorithms are modeled upon the Eikonal equation $\|\nabla_g u\| = 1$.

(a)  (b) 

Figure 3: Parabolic curves drawn on (a) the surface of the Apollo Belvedere, and (b) the abstract surface of the Lena image. In the first case they have no apparent structure. In the second case they follow the edges of the image.

5 New Possibilities

In this section we propose six new applications of Riemannian geometry to image processing: Riemannian thin-plate splines, curvature edges, spectral descriptors, Chladni figures, eigenfunction approximation and Ricci flow.

5.1 Riemannian thin plates

Thin plate splines [18] interpolate data given at scattered points. They are determined by the equation $\Delta\Delta u = 0$, with boundary conditions given by the data points.

This definition is readily mapped to our Riemannian setting, The advantage is that now, biharmonic functions are slowly varying on the Riemannian surface, but they can be forced to be almost discontinuous in the Euclidean coordinates. leading to a form of *guided interpolation* [19]. Thus given some data values $f\colon B \to \mathbf{R}$ on a set $B \subseteq M$, we consider the following equation on $\Omega = M \backslash B$:

$$\begin{cases} \Delta_g \Delta_g u = 0 & \text{in } \Omega \\ u = f & \text{in } \partial\Omega \end{cases} \tag{5}$$

This linear equation is discretized and solved by standard methods such as Gauss-Seidel. A possible application is image colorization (see Fig. 2), where we solve this equation three times, one for each color channel f, and finally we impose the intensities of the original image. Thanks to the metric, most of the variation of the solutions f happens along the edges of the image.

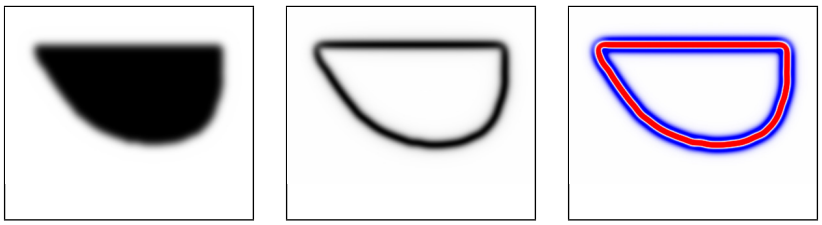

Figure 4: Partition of the image domain according to the sign of Gaussian curvature. Left: Smooth synthetic image I. Middle: Conformal metric given by $E = e^{-\|\nabla I\|^2/2\sigma^2}$. Right: Curvature, positive (blue), negative (red), or very close to zero (white). The points of negative curvature identify the position and width of the edges.

5.2 Curvature edges

A famous experiment by Felix Klein [5], [20] consisted in drawing the curves of parabolic points on the surface of Apollo's Belvedere (see Fig. 3). These curves do not have a particularly informative structure. On the other hand, the curves of parabolic points on the abstract surface defined from an image follow fairly well the edges of the image. In practice, they seem to offer no advantage over simpler edge detection operators, besides the fact that they can be made invariant to affine transforms when using an affine invariant metric.

The sign of Gaussian curvature is also a useful edge descriptor. We observe (see Fig. 4) that along each blurred edge there is a region of points of negative curvature $K < 0$, surrounded by two curves of parabolic points $K = 0$. The rest of the image domain consists in points of positive curvature $K > 0$, albeit very close to zero except near the edges.

For the conformal metric given by $E = e^{-\|\nabla I\|/2\sigma^2}$ the curvature can be computed explicitly by formula (2), giving the following expression

$$K = \frac{1}{4\sigma^2}\Delta\|\nabla I\|^2$$

which is possibly the simplest third order elliptic operator.

By Gauss-Bonnet theorem, the integral of K over the whole image domain is a constant that does not depend on the metric, but only on the topology of the boundary conditions. For periodic boundary conditions, the integral of K always vanishes.

5.3 Laplace-Beltrami spectral descriptors

The simplest structure of spectral theory is the sequence of eigenvalues $0 = \lambda_1 < \lambda_2 < \lambda_3 < \cdots$. They correspond to the frequencies of the vibration modes of the surface. The number λ_2 is called the fundamental tone, and the rest are called the harmonics. In the case of embedded surfaces this sequence has been called

the *shape DNA* [21, 22] because it is a useful descriptor of the shape, invariant to arbitrary isometries of the surface. This name makes sense because, while isospectral surfaces exist, they are relatively rare [23].

The Laplace-Beltrami spectrum is closely related [24] to the *geodesic length spectrum* $\gamma_1 < \gamma_2 < \gamma_3 < \cdots$, which is the sequence of lengths of periodic geodesics on the manifold. The first values of γ_i identify the "necks" of the surface, while further values are the lengths of geodesics that surround several necks in different combinations of increasing length.

The definition of the shape DNA applies also to our abstract surfaces, where it provides a *global* descriptor of the image. By construction, it is invariant to arbitrary isometric diffeomerphisms of the image domain; thus it inherits all the invariance properties of the metric. However, its usefulness is limited because it depends strongly on the boundary conditions that are used to define the eigenfunctions. The interpretation of the geodesic length spectrum in this context is interesting: the first values correspond to the perimeters of the blobs of the image, and the following values encode the relative distances between the blobs.

5.4 Laplace-Beltrami eigenvector segmentations

The eigenfunctions φ_k represent the shape of each vibration mode. For a physical surface, for example the board of a guitar, the sets $[\varphi_k = 0]$ can be observed by pouring sand on the surface and letting it vibrate with frequency λ_k. The sand will move away from the vibrating parts of the surface and it will accumulate on the parts that remain static. The resulting patterns are called *Chladni figures*, or nodal sets. They were first observed by Robert Hooke in 1680, and later extensively studied by Ernst Chladni in 1787. These sets have limited complexity because [24] the set $M \backslash [\varphi_k = 0]$ has at most k connected components.

The nodal sets are a common tool in the modern study of three-dimensional shapes [25, 26]. In the case of our abstract surfaces, they adapt well to the boundaries of the objects (see Fig. 5). Due to their increasing complexity, they give rise to a hierarchy of segmentations of the image domain.

5.5 Laplace-Beltrami eigenvector approximation

The eigenfunctions φ_k are an orthogonal basis of $L^2(M)$, allowing to perform harmonic analysis. Thus, we can represent any function $f : M \to \mathbf{R}$ as a linear combination of φ_k. Most linear PDE can be solved explicitly using this representation. For example, the Cauchy problem for the heat equation is solved by means of the heat kernel [27]. First we express the initial condition $u(0)$ as a linear combination of eigenfunctions: $u(0) = \sum c_k \varphi_k$ and then the solution is given by

$$u(t) = \sum c_k e^{-\lambda_k t} \varphi_k.$$

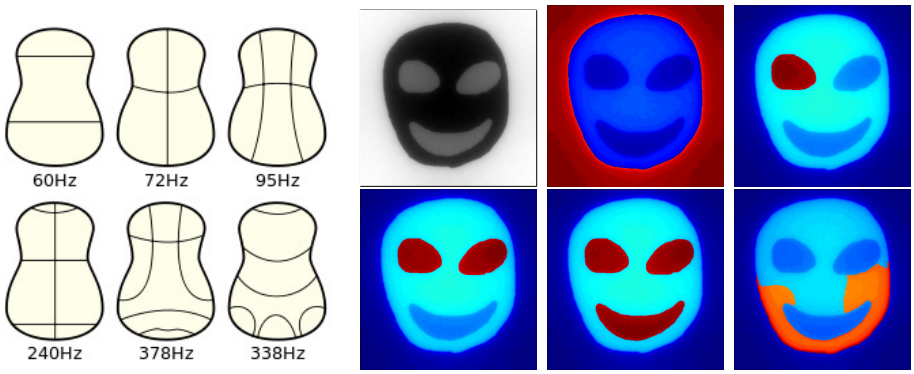

Figure 5: Left: First nodal sets of a guitar baseplate. Right: A synthetic image and the first five nodal sets of its abstract surface (the color represents the sign of each eigenfunction).

Similarly, the solution of Poisson equation $\Delta_g u = f$ is $u = \sum \frac{c_k}{\lambda_k} \varphi_k$ where c_k are the coefficients of f.

This is not a practical method of solving equations because finding the eigenfunctions is harder than solving the equation directly. Still, these functions are a convenient basis of smooth functions on the surface [28], analogous to sinusoidal functions on the plane or spherical harmonics on the sphere.

5.6 Ricci Scale space

Ricci flow is already an important tool in the theoretical [29] and numerical [30] study of embedded surfaces. Numerically, it is problematic because the surface has to be re-embedded in R^3 after each evolution step. For abstract surfaces, the situation is much easier because we can directly evolve the metric field using the equation $g_t = \mathrm{Ric}_g$. Notice that this describes the evolution of the metric, not of the image. But at any time the image can be easily recovered from the metric field by solving an Eikonal equation in the Euclidean plane. For example, if E is a function of the gradient $E = h(\|\nabla I\|)$, then we solve the Eikonal equation $\|\nabla u\| = h^{-1}(E(t))$ to recover the image from the metric coefficient E at a certain moment t. The resulting scale space has the same visual effect as the Perona-Malik diffusion.

6 Conclusion

We have proposed several new applications of Riemannian geometry to the context of image processing. Some of these applications provide a new insight to classical methods, while others (e.g., interpolation by Riemannian thin plate splines) are new techniques.

The source code (written in C and Octave/Matlab) for reproducing the experiments described in this paper is available as supplementary material.

References

[1] V. Caselles, R. Kimmel, and G. Sapiro, "Geodesic active contours," *International journal of computer vision*, vol. 22, no. 1, pp. 61–79, 1997.

[2] J. Weickert, *Anisotropic diffusion in image processing*, vol. 1, Teubner Stuttgart, 1998.

[3] J. Shi and J. Malik, "Normalized cuts and image segmentation," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 22, no. 8, pp. 888–905, 2000.

[4] M. Berger, *A panoramic view of Riemannian geometry*, Springer, 2003.

[5] D. Hilbert and S. Cohn-Vossen, *Geometry and the Imagination*, 1952.

[6] B. Chow and D. Knopf, "The Ricci flow: an introduction," *American Mathematical Society, Providence, RI*, 2004.

[7] J.-H. Rubinstein and R. Sinclair, "Visualizing ricci flow of manifolds of revolution," *Experimental Mathematics*, vol. 14, no. 3, pp. 285–298, 2005.

[8] P.-J. Olver, G. Sapiro, and A. Tannenbaum, "Affine invariant detection: edge maps, anisotropic diffusion, and active contours," *Acta Applicandae Mathematica*, vol. 59, no. 1, pp. 45–77, 1999.

[9] M. Kass, A. Witkin, and D. Terzopoulos, "Snakes: Active contour models," *International journal of computer vision*, vol. 1, no. 4, pp. 321–331, 1988.

[10] S. Osher and J.-A. Sethian, "Fronts propagating with curvature-dependent speed: algorithms based on hamilton-jacobi formulations," *Journal of computational physics*, vol. 79, no. 1, pp. 12–49, 1988.

[11] P. Perona and J. Malik, "Scale-space and edge detection using anisotropic diffusion," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 12, no. 7, pp. 629–639, 1990.

[12] D. Eberly, "A differential geometric approach to anisotropic diffusion," in *Geometry-Driven Diffusion in Computer Vision*, pp. 371–392. Springer, 1994.

[13] Y. Boykov and V. Kolmogorov, "Computing geodesics and minimal surfaces via graph cuts," in *Computer Vision, 2003. Proceedings. Ninth IEEE International Conference on*. IEEE, 2003, pp. 26–33.

[14] V. Caselles, G. Facciolo, and E. Meinhardt-Llopis, "Anisotropic cheeger sets and applications," *SIAM Journal on Imaging Sciences*, vol. 2, no. 4, pp. 1211–1254, 2009.

[15] R. Kunze, F.-E. Wolter, and T. Rausch, "Geodesic voronoi diagrams on parametric surfaces.," in *Computer Graphics International*, 1997, vol. 97, pp. 230–237.

[16] G. Peyre and L. Cohen, "Surface segmentation using geodesic centroidal tesselation," in *3D Data Processing, Visualization and Transmission, 2004. 3DPVT 2004. Proceedings. 2nd International Symposium on*. IEEE, 2004, pp. 995–1002.

[17] G. Facciolo and V. Caselles, "Geodesic neighborhoods for piecewise affine interpolation of sparse data," in *Image Processing (ICIP), 2009 16th IEEE International Conference on*. IEEE, 2009, pp. 365–368.

[18] K. Rohr, H.-S. Stiehl, R. Sprengel, T.-M. Buzug, J. Weese, and M.-H. Kuhn, "Landmark-based elastic registration using approximating thin-plate splines," *Medical Imaging, IEEE Transactions on*, vol. 20, no. 6, pp. 526–534, 2001.

[19] P. Pérez, M. Gangnet, and A. Blake, "Poisson image editing," in *ACM Transactions on Graphics (TOG)*. ACM, 2003, vol. 22, pp. 313–318.

[20] J.-J. Koenderink, *Solid shape*, vol. 2, Cambridge Univ Press, 1990.

[21] M. Reuter, F.-E. Wolter, and N. Peinecke, "Laplace–beltrami spectra as shape-dnaof surfaces and solids," *Computer-Aided Design*, vol. 38, no. 4, pp. 342–366, 2006.

[22] A. Dubrovina and R. Kimmel, "Matching shapes by eigendecomposition of the laplace-beltrami operator," in *Proc. 3DPVT*, 2010, vol. 2.

[23] M. Kac, "Can one hear the shape of a drum?," *Amer. Math. Monthly*, vol. 73, pp. 1–23, 1966.

[24] Y.-C. De Verdiere, "Spectrum of the laplace operator and periodic geodesics: thirty years after," in *Annales-Institut Fourier*, 2007, vol. 57, p. 2429.

[25] B. Lévy, "Laplace-beltrami eigenfunctions towards an algorithm that understands geometry," in *Shape Modeling and Applications, 2006. SMI 2006. IEEE International Conference on*. IEEE, 2006, pp. 13–13.

[26] A. Bronstein and R. Kimmel, *Numerical geometry of non-rigid shapes*, Springer, 2008.

[27] S. Seo, M. Chung, and H. Vorperian, "Heat kernel smoothing of anatomical manifolds via laplace-beltrami eigenfunctions," 2010.

[28] A. Qiu, D. Bitouk, and M. Miller, "Smooth functional and structural maps on the neocortex via orthonormal bases of the laplace-beltrami operator," *Medical Imaging, IEEE Transactions on*, vol. 25, no. 10, pp. 1296–1306, 2006.

[29] R.-S. Hamilton, "The ricci flow on surfaces," *Contemp. Math*, vol. 71, no. 1, 1988.

[30] G. Zou, J. Hua, Z. Lai, X. Gu, and M. Dong, "Intrinsic geometric scale space by shape diffusion," *Visualization and Computer Graphics, IEEE Transactions on*, vol. 15, no. 6, pp. 1193–1200, 2009.