Computers & Graphics 35 (2011) 549–560

Contents lists available at ScienceDirect

Computers & Graphics

journal homepage: www.elsevier.com/locate/cag

SMI 2011: Full Paper

Diffusion-geometric maximally stable component detection in deformable shapes

Roee Litman[a,*], Alexander M. Bronstein[a], Michael M. Bronstein[b]

[a] School of Electrical Engineering, Tel Aviv University, P.O. Box 39040, Tel Aviv 69978, Israel
[b] Institute of Computational Science, Faculty of Informatics, Universita della Svizzera Italiana, Lugano, Switzerland

ARTICLE INFO

Article history:
Received 11 December 2010
Received in revised form
5 February 2011
Accepted 7 March 2011
Available online 13 April 2011

Keywords:
Deformable shapes
Feature detection
Diffusion geometry
Component tree
Level sets
MSER

ABSTRACT

Maximally stable component detection is a very popular method for feature analysis in images, mainly due to its low computation cost and high repeatability. With the recent advance of feature-based methods in geometric shape analysis, there is significant interest in finding analogous approaches in the 3D world. In this paper, we formulate a diffusion-geometric framework for stable component detection in non-rigid 3D shapes, which can be used for geometric feature detection and description. A quantitative evaluation of our method on the SHREC'10 feature detection benchmark shows its potential as a source of high-quality features.

1. Introduction

Over the past decade, feature-based methods have become a ubiquitous tool in image analysis and a *de facto* standard in many computer vision and pattern recognition problems. More recently, there has been an increased interest in developing similar methods for the analysis of 3D shapes. Feature descriptors play an important role in many shape analysis applications, such as finding shape correspondence [32] or assembling fractured models [11] in computational aracheology. Bags of features [29,24,33] and similar approaches [21] were introduced as a way to construct global shape descriptors that can be efficiently used for large-scale shape retrieval.

Many shape feature detectors and descriptors draw inspiration from and follow analogous methods in image analysis. For example, detection of geometric structures analogous to corners [28] and edges [14] in images has been studied. The histogram of intrinsic gradients used in [36] is similar in principle to the scale invariant feature transform (SIFT) [16] which has recently become extremely popular in image analysis. In [10], the integral invariant signatures [17] successfully employed in 2D shape analysis were extended to 3D shapes.

Examples of 3D-specific descriptors include the popular spin image [12], based on representation of the shape normal field in a local system of coordinates. Recent studies introduced versatile and computationally efficient descriptors based on the heat kernel [31,3] describing the local heat propagation properties on a shape. The advantage of these methods is the fact that heat diffusion geometry is intrinsic and thus deformation-invariant, which makes descriptors based on it applicable in deformable shape analysis.

1.1. Related work

A different class of feature detection methods tries to find stable components or regions in the analyzed image or shape. In the image processing literature, the watershed transform is the precursor of many algorithms for stable component detection [6,34]. In the computer vision and image analysis community, stable component detection is used in the maximally stable extremal regions (MSER) algorithm [18]. MSER represents intensity level sets as a component tree and attempts finding level sets with the smallest area variation across intensity; the use of area ratio as the stability criterion makes this approach affine-invariant, which is an important property in image analysis, as it approximates viewpoint transformations. Alternative stability criteria based on geometric scale-space analysis have been recently proposed in [13].

* Corresponding author. Tel.: +972 3640 8288; fax: +972 3640 7095.
E-mail addresses: roeelitm@post.tau.ac.il (R. Litman), alexbron@ieee.org, bron@eng.tau.ac.il (A.M. Bronstein), michael.bronstein@usi.ch (M.M. Bronstein).

0097-8493/$ - see front matter © 2011 Elsevier Ltd. All rights reserved.
doi:10.1016/j.cag.2011.03.011

In the shape analysis community, shape decomposition into characteristic primitive elements was explored in [22]. Methods similar to MSER have been explored in the works on topological persistence [8]. Persistence-based clustering [4] was used by Skraba et al. [30] to perform shape segmentation. In [7], Digne et al. extended the notion of vertex-weighted component trees to meshes and proposed to detect MSER regions using the mean curvature. The approach was tested only in a qualitative way, and not evaluated as a feature detector.

1.2. Main contribution

The main contribution of our framework is three-fold. First, in Section 2 we introduce a generic framework for stable component detection, which unites vertex- and edge-weighted graph representations (as opposed to vertex weighting used in image and shape maximally stable component detectors [18,7]). Our results (see Section 4) show that the edge-weighted formulation is more versatile and outperforms its vertex-weighted counterpart in terms of feature repeatability. Second, in Section 3 we introduce diffusion geometric weighting functions suitable for both vertex- and edge-weighted component trees. We show that such functions are invariant under a large class of transformations, in particular, non-rigid inelastic deformations, making them especially attractive in non-rigid shape analysis. We also show several ways of constructing scale-invariant weighting functions. Third, in Section 4 we show a comprehensive evaluation of different settings of our method on a standard feature detection benchmark comprising shapes undergoing a variety of transformations (also see Figs. 1 and 2).

2. Diffusion geometry

Diffusion geometry is an umbrella term referring to geometric analysis of diffusion or random walk processes [5]. We models a shape as a compact 2D Riemannian manifold X. In its simplest setting, a diffusion process on X is described by the partial differential equation

$$\left(\frac{\partial}{\partial t}+\Delta\right)f(t,x)=0, \qquad (1)$$

called the *heat equation*, where Δ denotes the positive-semidefinite Laplace–Beltrami operator associated with the Riemannian metric of X. The heat equation describes the propagation of heat on the surface and its solution $f(t,x)$ is the heat distribution at a point x in time t. The initial condition of the equation is some initial heat distribution $f(0,x)$; if X has a boundary, appropriate boundary conditions must be added.

The solution of (1) corresponding to a point initial condition $f(0,x)=\delta(x,y)$ is called the *heat kernel* and represents the amount of heat transferred from x to y in time t due to the diffusion process. The value of the heat kernel $h_t(x,y)$ can also be interpreted as the transition probability density of a random walk of length t from the point x to the point y.

Using spectral decomposition, the heat kernel can be represented as

$$h_t(x,y)=\sum_{i\geq 0}e^{-\lambda_i t}\phi_i(x)\phi_i(y). \qquad (2)$$

Here, ϕ_i and λ_i denote, respectively, the eigenfunctions and eigenvalues of the Laplace–Beltrami operator satisfying $\Delta\phi_i=\lambda_i\phi_i$ (without loss of generality, we assume λ_i to be sorted in increasing order starting with $\lambda_0=0$). Since the Laplace–Beltrami operator is an *intrinsic* geometric quantity, i.e., it can be expressed solely in terms of the metric of X, its eigenfunctions and eigenvalues as well as the heat kernel are invariant under isometric transformations (bending) of the shape. These properties of the Laplacian have been previously exploited in the literature for "natural" parametrizaton of surfaces [15], construction of global shape descriptors [27], and detection of symmetries [25] just to mention a few.

The parameter t can be given the meaning of *scale*, and the family $\{h_t\}_t$ of heat kernels can be thought of as a scale-space of functions on X. By integrating over all scales, a *scale-invariant* version of (2) is obtained:

$$c(x,y)=\sum_{i\geq 1}\frac{1}{\lambda_i}\phi_i(x)\phi_i(y). \qquad (3)$$

This kernel is referred to as the *commute-time kernel* and can be interpreted as the transition probability density of a random walk of any length.

By setting $y=x$, both the heat and the commute time kernels, $h_t(x,x)$ and $c(x,x)$ express the probability density of remaining at a point x, respectively after time t and after any time. The value $h_t(x,x)$, sometimes referred to as the *auto-diffusivity function*, is related to the Gaussian curvature $K(x)$ through

$$h_t(x,x)\approx\frac{1}{4\pi t}\left(1+\frac{1}{6}K(x)t+\mathcal{O}(t^2)\right). \qquad (4)$$

This relation coincides with the well-known fact that heat tends to diffuse slower at points with positive curvature, and faster at points with negative curvature.

For any $t>0$, the values of $h_t(x,y)$ at every x and $y\in B_\varepsilon(x)$ in a small neighborhood around x contain full information about the intrinsic geometry of the shape. Furthermore, Sun et al. [31] show that under mild technical conditions, the set $\{h_t(x,x)\}_{t>0}$ is also fully informative (note that the auto-diffusivity function has to be evaluated at all values of t in order to contain full information about the shape metric).

2.1. Numerical computation

In the discrete setting, we assume that the shape is sampled at a finite number of points $V=\{v_1,\ldots,v_N\}$, upon which a simplicial complex (triangular mesh) with vertices V, edges $E\subset V\times V$ and faces $F\subset V\times V\times V$ is constructed. The computation of the discrete heat kernel $h_t(v_1,v_2)$ and the associated diffusion geometry constructs is performed using formula (2), in which a finite number of eigenvalues and eigenfunctions of the discrete Laplace–Beltrami operator are taken. The latter can be computed directly using the finite elements method (FEM) [27], by discretization of the Laplace operator on the mesh followed by its eigendecomposition. Here, we adopt the second approach according to which the discrete Laplace–Beltrami operator is expressed in the following generic form:

$$(\Delta_X f)_i=\frac{1}{a_i}\sum_j w_{ij}(f_i-f_j), \qquad (5)$$

where $f_i=f(v_i)$ is a scalar function defined on V, w_{ij} are weights, and a_i are normalization coefficients. In matrix notation, (5) can be written as $\Delta_X f=A^{-1}Wf$, where f is an $N\times 1$ vector, $A=\mathrm{diag}(a_i)$ and $W=\mathrm{diag}(\sum_{l\neq i}w_{il})-(w_{ij})$. The discrete eigenfunctions and eigenvalues are found by solving the *generalized eigendecomposition* [15] $W\Phi=A\Phi\Lambda$, where $\Lambda=\mathrm{diag}(\lambda_l)$ is a diagonal matrix of eigenvalues and $\Phi=(\phi_l(v_i))$ is the matrix of the corresponding eigenvectors.

Different choices of A and W have been studied, depending on which continuous properties of the Laplace–Beltrami operator one wishes to preserve [9,35]. For triangular meshes, a popular choice adopted in this paper is the *cotangent weight* scheme