3D SHAPE MATCHING AND TEICHMÜLLER SPACES OF POINTED RIEMANN SURFACES

CLAUDIO FONTANARI AND LETIZIA PERNIGOTTI

Abstract. Shape matching represents a challenging problem in both information engineering and computer science, exhibiting not only a wide spectrum of multimedia applications, but also a deep relation with conformal geometry. After reviewing the theoretical foundations and the practical issues involved in this fashinating subject, we focus on two state-of-the-art approaches relying respectively on local features (landmark points) and on global properties (conformal parameterizations). Finally, we introduce the Teichmüller space $T_{g,n}$ of n-pointed Riemann surfaces of genus g into the realm of multimedia, showing that its beautiful geometry provides a natural unified framework for three-dimensional shape matching.

As incisively summarized in the survey paper [1], shape matching deals with transforming a shape and measuring the resemblance with another one via some similarity measure, thus providing an essential ingredient in shape retrieval and registration. By the way, there seems to be no universal and standard definition of what a shape is (see [1] and [2]): [1] goes even back to Plato's Meno, where Socrates claims in "terms employed in geometrical problems" that "figure is limit of solid" ([3]), while [2] relies on Kendall's definition of shape as "all the geometrical information that remains when location, scale, and rotational effects (Euclidean transformations) are filtered out from an object" ([4]).

The authors of [2] also emphasize the progressive shift of interest in the last decade from 2D (where shapes are just silhouettes) to 3D (where shapes are embedded surfaces), motivated by spectacular advances in the field of three-dimensional computer graphics. Unluckily, as remarked in [5], most 2D methods do not generalize directly to 3D model matching, thus inducing a flurry of recent research to focus on the specific problem of 3D shape retrieval. It is worth stressing that, unlike text documents, 3D models do not allow the simplest form of searching by keyword and it is now a rather common belief (see for instance [6]) that an efficient 3D shape retrieval system should take into account the principles of the human visual system as disclosed by cognitive neurosciences.

Before turning to the description of a couple of current approaches moving along these lines, we wish to mention at least another challenging application of three-dimensional shape matching. The so-called computational anatomy arises in medical imaging, in particular neuroimaging, and according to [7] it involves comparison of the shape of anatomic structures between two individuals, and development of a statistical theory which allows shape to be studied across populations.

A main component in this analysis, after obtaining the individual model representations for the subjects being studied, is the establishment of correspondence of anatomically homologous substructures between the subjects. For example, if we are interested in comparing shape differences between faces of two individuals in images, we would like to ensure that the coordinates of the left eye in one image correspond to the left eye in the other image. On a finer scale, we would like to ensure that the left corner of the left eye corresponds appropriately ([7]). More generally, following [8], we will define the landmarks of an object as the points of interest of the object that have important shape attributes. Examples of landmarks are corners, holes, protrusions, and high curvature points.

Landmark-based shape recognition is motivated by such a concept of dominant points. It uses landmarks as shape features to recognize objects in a scene or to establish correspondences between objects, by extending in an optimal way over the entire structure the correspondence at a finite subset. This last process is called landmark matching ([7]) and originates from the viewpoint of the human visual system, which suggests that some dominant points along an object contour are rich in information content and are sufficient to characterize the shape of the object (see [8] and the references therein). In particular, the paper [7] presents a methodology and algorithm for generating diffeomorphisms of the sphere onto itself, given the displacement of a finite set of template landmarks.

The restriction to a spherical domain is quite natural from the point of view of brain imaging, where a basic assumption is that the topology of the brain surface is the same as that of a crumpled sheet and, in particular, does not have any holes or self intersections ([9]). It then follows by Riemann Uniformization (see for instance [10], Theorem 4.4.1 (iii)) that there exists a conformal diffeomorphism of such a surface of genus zero onto a sphere. More generally, every orientable compact embedded surface can be made into a Riemann surface with conformal structure (see for instance [11], Theorem IV.1.1).

As pointed out in [12] (see also references therein), the visual field is represented in the brain by mappings which are, at least approximately, conformal. Thus, to simulate the imaging properties of the human visual system conformal image mapping is a necessary technique. Moreover, beside its theoretical soundness, the application of conformal geometry to the 3D shape classification problem presents several practical advantages: according to [13], the conformal structures are independent of triangulation, insensitive to resolution, and robust to noises.

As far as the genus zero case is concerned, explicit conformal flattenings can be obtained by numerically solving the Laplace-Beltrami equation (the heart of this procedure consists in a finite element reduction to a system of linear equations, see [9] and [14]) or by deforming a homeomorphism in order to minimize the harmonic energy (both convergence of the algorithm and uniqueness of the solution are ensured by imposing further constraints, see [15] and [16]).

A natural generalization involves quasi-conformal mappings, which do not distort angles arbitrarily (as it is well-known, conformal mappings are angle-preserving): in particular, the least-squares conformal mappings (introduced in [17] via a least-squares approximation of the Cauchy-Riemann equations and in [16] through a discrete version of the harmonic energy minimization method) provide a natural solution to 3D nonrigid surface alignment and stitching ([16]), at least in genus zero.

In higher genus, a parameterization method based on Riemann surface structure has been developed in a series of papers by Shing-Tung Yau and collaborators (see in particular [18]). The basic idea is to segment the surface according to its conformal structure, parameterize the patching using a holomorphic one form, and finally glue them together via harmonic maps ([20]). A comprehensive and self-contained account of 3D applications of conformal geometry is now available in [19].

Our (perhaps tendentious) survey of the subject has presented three-dimensional shape matching as a classification problem for Riemann surfaces carrying (land)marked points. Indeed, following [21], let us fix nonnegative integers g and n such that $2g-2+n>0$, a compact connected oriented reference surface S_g of genus g and a sequence of n distinct points $(x_1, \dots, x_n)$ on S_g. By definition, the Teichmüller space $T_{g,n}$ is the space of conformal structures on S_g up to isotopies that fix $\{x_1, \dots, x_n\}$ pointwise. It is, in a natural way, a complex manifold of dimension $3g-3+n$.

The idea of classifying 3D shapes according to their conformal structure goes back to [13], where the general principle is stated that in nature it is highly unlikely for different shapes to share the same conformal structure. Of course by Riemann Uniformization this fails in genus $g=0$, but the choice of a suitable number $n >> 0$ of landmark points allows to address in a uniform way also such an exceptional case. Indeed, we believe that the Teichmüller space $T_{g,n}$ can provide a solid mathematical framework to three-dimensional shape matching, by supporting a unified geometric theory of landmark matching and conformal parameterizations.

The case $n=0$ has recently been considered in [22] and [23], where the difference between two shapes of the same genus g is measured by the Euclidean distance between the corresponding Fenchel-Nielsen coordinates in the Teichmüller space $T_{g,0}$. In order to compute the Fenchel-Nielsen coordinates of a shape, the associated surface is conformally deformed along the surface Ricci flow, until its Gaussian curvature is -1 everywhere. The surface is then decomposed into several pairs of hyperbolic pants, i.e., genus zero surfaces with three geodesic boundaries. Each pair of hyperbolic pants is uniquely described by the lengths of its boundaries and the way of glueing different pairs of pants is encoded by the twisting angles between two adjacent pairs of pants which share a common boundary.

Here we show how to adapt to the case $n > 0$ all steps of the algorithm described in [23], §4:

(1) Compute topological pants decomposition.
(2) Compute the hyperbolic metric using Ricci flow.
(3) Compute hyperbolic pants decomposition.
(4) Compute the Fenchel-Nielsen coordinates.

In order to approximate a punctured Riemann surface by a triangular mesh, we consider the corresponding compact surface with boundaries obtained by removing open discs of sufficiently small but strictly positive radius instead of points.

Step (1) follows the approach of [24], which relies on [25] and is explicitly presented in the case of surfaces with boundary.

As far step (2) is concerned, we recall that the exponential convergence of the solution of the Ricci flow equation to a complete metric with prescribed constant scalar curvature has been proved in [26] for compact surfaces with negative Euler characteristic. The case of interest for us, namely, an open surface conformal to a punctured compact Riemann surface, has been recently addressed in [27] (see in particular Theorem 3). On the other hand, the combinatorial version of the Ricci flow is introduced in [28] for triangulations on compact surfaces with boundary (see in particular Theorem 5.1). Since the description of step (2) in [23], §4, relies on [28], it directly extends to our setting (for implementation issues in the case of surfaces with boundary we refer to [19], §12).

Finally, steps (3) and (4) in [23], §4, rely on [29], which addresses also surfaces with boundary, and can be repeated verbatim to produce (up to renumbering) a sequence of lenghts and twisting angles

$$\{(l_1, \theta_1), (l_2, \theta_2), \ldots, (l_{3g-3+n}, \theta_{3g-3+n}), l_{3g-3+n+1}, \ldots, l_{3g-3+2n}\}.$$

Indeed, there are $3g - 3 + 2n$ length parameters, one for each curve of the pants decomposition and one for each boundary curve, and $3g - 3 + n$ twist parameters, one for each curve of the pants decomposition. By setting the length parameters $l_{3g-3+n+1}, \ldots, l_{3g-3+2n}$ to be zero, we can turn boundary components into punctures (see for instance [30], §10.6.3). The corresponding Fenchel-Nielsen coordinates are given by

$$\{(l_1, \theta_1), (l_2, \theta_2), \ldots, (l_{3g-3+n}, \theta_{3g-3+n})\}.$$

Full implementation details and numerical results will be presented elsewhere.

All the above undoubtedly points towards a potentially fruitful interaction between two such apparently unrelated fields as moduli spaces and information technologies. We hope to see deeper into this in the next future.

References

[1] R. C. Veltkamp: Shape matching: similarity measures and algorithms. SMI 2001 International Conference on Shape Modeling and Applications, 7–11 May 2001, 188–197.

[2] N. Iyer, S. Jayanti, K. Lou, Y. Kalyanaraman, K. Ramani: Three-dimensional shape searching: state-of-the art review and future trends. Elsevier Computer-Aided Design 37 (2005), 509–530.

[3] Plato: Meno. W. Lamb (ed.): Plato in Twelve Volumes. Vol. 3, Harvard University Press, 1967.

[4] D. G. Kendall: The diffusion of shape. Adv. Appl. Prob. 9 (1977), 428–430.

[5] J.W.H. Tangelder and R.C. Veltkamp: A survey of content based 3D shape retrieval methods. Shape Modeling Applications, 2004, 145–156.

[6] H.Y. Sean Lin, J. C. Lin, H.-Y. Mark Liao: 3D shape retrieval using cognitive psychology-based principles. Seventh IEEE International Symposium on Multimedia, 12–14 December 2005, 8 pp.

[7] J. Glaunès, M. Vaillant, M. I. Miller: Landmark matching via large deformation diffeomorphisms on the sphere. Journal of Mathematical Imaging and Vision 20 (2004), 179–200.

[8] N. Ansari and E.J. Delp: Partial shape recognition: a landmark-based approach. IEEE Transactions on Pattern Analysis and Machine Intelligence, Volume 12, Issue 5, May 1990, 470–483.

[9] S. Angenent, S. Haker, A. Tannenbaum, R. Kikinis: On the Laplace-Beltrami operator and brain surface flattening. IEEE Transactions on Medical Imaging, Volume 18, Issue 8, August 1999, 700–711.

[10] J. Jost: Compact Riemann surfaces. An introduction to contemporary mathematics. Translated from the German manuscript by R. R. Simha. Universitext. Springer-Verlag, Berlin, 1997. xiv+291 pp.

[11] O. Lehto: Univalent functions and Teichmüller spaces. Graduate Texts in Mathematics, 109. Springer-Verlag, New York, 1987. xii+257 pp.

[12] C. Frederick and E. L. Schwartz: Conformal image warping. IEEE Computer Graphics and Applications, Volume 10, Issue 2, March 1990, 54–61.

[13] X. Gu and S.-T. Yau: Surface classification using conformal structures. Ninth IEEE International Conference on Computer Vision, 2003, 701–708 (Vol.1)

[14] S. Haker, S. Angenent, A. Tannenbaum, R. Kikinis, G. Sapiro, M. Halle: Conformal surface parameterization for texture mapping. IEEE Transactions on Visualization and Computer Graphics, Volume 6, Issue 2, April-June 2000, 181–189.

[15] X. Gu and S.-T. Yau: Computing conformal structures of surfaces. Communications in Information and Systems 2 (2002), 121–146.

[16] X. Gu, Y. Wang, T. F. Chan, P.M. Thompson, S.-T. Yau: Genus zero surface conformal mapping and its application to brain surface mapping. IEEE Transactions on Medical Imaging, Volume 23, Issue 8, August 2004, 949–958.

[17] B. Lévy, S. Petitjen, N. Ray, J. Maillot: Least Squares Conformal Maps for Automatic Texture Atlas Generation. ACM SIGGRAPH Conference (2002), 362–371.

[18] Y.Wang, L. M. Lui, X. Gu, K. M. Hayashi, T. F. Chan, A. W. Toga, P. M. Thompson, S.-T. Yau: Brain Surface Conformal Parameterization Using Riemann Surface Structure. IEEE Transactions on Medical Imaging, Volume 26, Issue 6, June 2007, 853–865.

[19] X. D. Gu and S.-T. Yau: Computational Conformal Geometry, Advanced Lectures in Mathematics Volume III, International Press, Sommerville, Massachusetts, U.S.A. and Higher Education Press, Beijing, China, 2008.

[20] Y.Wang, X. Gu, K. M. Hayashi, T. F. Chan, P. M. Thompson, S.-T. Yau: Surface parameterization using Riemann surface structure. ICCV 2005. Tenth IEEE International Conference on Computer Vision, 17-21 October 2005, 1061–1066 (Vol. 2).
[21] R. Hain and E. Looijenga: Mapping class groups and moduli spaces of curves. Algebraic geometry–Santa Cruz 1995, 97–142, Proc. Sympos. Pure Math., 62, Part 2, Amer. Math. Soc., Providence, RI, 1997.
[22] M. Jin, W. Zeng, F. Luo, and X. Gu: Computing Teichmüller Shape Space. IEEE Transactions on Visualization and Computer Graphics, Volume 15, Issue 3, May-June 2009, 504–517.
[23] M. Jin, W. Zeng, N. Ding, X. D. Gu and S.-T. Yau: Computing Fenchel-Nielsen Coordinates in Teichmüller Shape Space, Commun. Inf. Syst. Volume 9, Number 2 (2009), 213-234.
[24] X. Li, X. D. Gu and H. Qin: Surface Mapping Using Consistent Pants Decomposition, IEEE Transactions on Visualization and Computer Graphics, Volume 15, Issue 4, July 2009, 558–571.
[25] T. K. Dey, K. Li and J. Sun: On Computing Handle and Tunnel Loops. International Conference Cyberworlds (2007), 357–366.
[26] B. Chow: The Ricci flow on the 2-sphere. J. Differ. Geom. 33 (1991), 325–334.
[27] L. Ji, R. Mazzeo, and N. Sesum: Ricci flow on surfaces with cusps, Math. Ann. 345 (2009), 819–834.
[28] B. Chow and F. Luo: Combinatorial Ricci flows on surfaces, J. Differential Geometry 63 (2003), 97–129.
[29] M. Jin, J. Kim, F. Luo, and X. Gu: Discrete Surface Ricci Flow. IEEE Transactions on Visualization and Computer Graphics, Volume 14, Issue 5, September-October 2008, 1030–1043.
[30] B. Farb and D. Margalit: A primer on mapping class groups. Version 5.0 (2011). Available online at http://www.math.utah.edu/m̃argalit/primer/.

Dipartimento di Matematica, Università degli Studi di Trento, Via Sommarive 14, 38123 Povo (Italy)

E-mail address: fontanar@science.unitn.it, pernigotti@science.unitn.it