Pattern Recognition 44 (2011) 588–597

Contents lists available at ScienceDirect

Pattern Recognition

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

Artwork 3D model database indexing and classification

Sylvie Philipp-Foliguet [a,*], Michel Jordan [a], Laurent Najman [b], Jean Cousty [b]

[a] ETIS, UMR 8051, CNRS, ENSEA, Université Cergy-Pontoise, F-95000 Cergy-Pontoise, France
[b] Université Paris-Est, LABINFO-IGM, ESIEE, F-93162 Noisy-le-Grand, France

ARTICLE INFO

Article history:
Received 22 October 2009
Received in revised form
23 July 2010
Accepted 29 September 2010

Keywords:
3D database classification
3D shape descriptors
3D surface segmentation
Graph cuts
Watersheds
Artwork database
RETIN

ABSTRACT

This paper presents a framework for the indexing and retrieval of artwork 3D models, allowing global and partial model classification and retrieval. The first part of the paper deals with database classification based on global shape descriptors. A search engine "RETIN-3D", using a SVM classifier coupled with an active learning strategy allows to retrieve categories of similar objects. In a second part, the classification is improved thanks to a local description of the models. A new framework for 3D surface segmentation is proposed. Shape descriptors are adapted to surface regions and kernels on descriptor bags are used to perform the database classification. Our system is designed for classifying and retrieving in ancient artwork 3D databases, and results from this application domain are presented and commented along the paper.

1. Introduction

3D shape modeling and digitizing have received more and more attention for a decade, leading to an increasing amount of 3D model warehouses, either in domain-specific or wide-usage contexts. These 3D model databases require new tools for indexing, classifying, and retrieving the objects, in order to provide the final user an easy access to the models.

Content-based document retrieval (CBDR) has been a very active research field for a few years, and concerns textual documents, images, videos, and more recently 3D models. Usually CBDR is divided into two different steps: (i) an offline step performs the document indexing by computing descriptors and features that are easily and fast compared, and thus builds an efficient summary of each document, called a signature; (ii) an online step, in which the user performs a search in the database thanks to a search engine. By means of signature comparison, the system ranks the database models according to their similarity to a query given as input. A feedback loop based on user interaction refines the results.

In this paper, we focus on 3D model indexing and retrieval and present our search engine "RETIN-3D". The first interactive 3D model search engines appear on the web around 2001–2002. The *Princeton 3D Model Search Engine*, associated to the widely used *Princeton Shape Benchmark* (*PSB*), (http://shape.cs.princeton.edu/benchmark/) allows the user to perform text queries, 2D sketch queries, and to compare 3D models through some 3D shape descriptors [1]. The *3D Search Tool* from the University of Thessaloniki (http://3d-search.iti.gr/3DSearch) is based on the 3D generalized Radon transform and make comparisons within a 2.000 model database [2]; the results are only based on geometric comparisons, without learning, leading to some mis-classifications of the database. The European Network of Excellence *Aim@Shape* (http://www.aimatshape.net) presents a *geometric search engine* which provides content-based retrieval with different matching methods (global or local, *etc.*). The SHREC 3D Shape Retrieval Contests allowed the comparison of 3D shape descriptors and 3D retrieval methods thanks to databases associated with ground-truthes [3]. Ohbuchi et al. [4] proposed a retrieval system based on multiresolution global features, which retrieves object categories from a single example.

We present here a 3D indexing and retrieval search engine dedicated to 3D artwork model databases. Our aim is to provide user-friendly tools for classification, for content-based indexing, for retrieval, and for visualization. These tools are firstly dedicated to historians and archeologists, who will be able to find, display and compare artworks in a few clicks. One can also imagine that museum visitors, provided with their PDA, could have the opportunity to interrogate a database in front of a statue and thus obtain a lot of additional information.

In Section 2, we address the database classification by mean of global shape indexing. Unlike CAD models or artificial models that are often used in 3D model warehouses, artwork models are

* Corresponding author.
E-mail addresses: sylvie.philipp-foliguet@ensea.fr (S. Philipp-Foliguet), michel.jordan@u-cergy.fr (M. Jordan), l.najman@esiee.fr (L. Najman), j.cousty@esiee.fr (J. Cousty).

0031-3203/$ - see front matter © 2010 Elsevier Ltd. All rights reserved.
doi:10.1016/j.patcog.2010.09.016

digitized in a high resolution (between 30,000 and 300,000 vertices) and do not exhibit regular surfaces. We compare several global shape descriptors for a classification task. We introduce the RETIN-3D search engine, which uses these shape descriptors (Section 2.1) to retrieve similar objects thanks to an active learning strategy (Section 2.2). Unlike search engines which are asked by textual requests or 2D/3D sketches, our query consists in a 3D model, and the search engine extracts from the database a category of models similar in a certain way to the query. The user leads the search toward the category, by annotating some objects as relevant or irrelevant for his search (Section 2.3).

Not surprisingly, global shape descriptors are not sufficient to discriminate objects differing by some specific details. In Section 3, we propose to use local shape descriptors computed on regions of the surface. We introduce a new scheme for 3D surface segmentation, based on local curvature computation and watershed cuts (Section 3.1). Shape descriptors are computed for each region of the surface partition (Section 3.2), and the search engine is adapted to 3D region descriptor bags (Section 3.3). Partial matching results are shown and explained in Section 3.4.

Major contributions of this paper are the 3D surface segmentation (Section 3.1) and the active learning for 3D model classification (Sections 2.2 and 3.3), applied to a database of archeological objects. The database consists of high resolution models and includes broken or damaged items.

2. Global shape matching

At first, we compared some of the most used global shape features. We adapted our search engine RETIN [5], which is already used for image and video retrieval to this new modality.

2.1. Global shape features

The literature provides a lot of various 3D shape descriptors, describing geometric as well as topological properties of 3D shapes: global shape descriptors [6–14]; local descriptors [15]; graph based methods [16–18]; geometric methods based on 2D views of 3D models [19,1,20]. Interested readers can refer to some recent review papers [21–24] for more details.

Before computing shape descriptors for each model of the database, we perform a spatial alignment preprocessing which aims to put the 3D models in a canonical coordinate system. The origin of the coordinate system is set at the center of gravity of the model, and the spatial alignment is achieved thanks to a principal component analysis transform.

Our 3D models are VRML meshes, featuring 3D vertices V (the model geometry) and 3D facets (triangles). We implemented most of the global shape descriptors [5] described above. We briefly describe below the ones which gave exploitable results for our high resolution models.

Cord histograms: A cord [25] is defined as the vector from a vertex to the center of the model, and is characterized by three features: (i) the length of the cord; (ii) its angle with the first principal axis; (iii) its angle with the second principal axis. We built a descriptor named "Cord2D" based upon two normalized histograms of the two first features (*cf.* Fig. 1(a)).

Extended Gaussian images (*EGI*): The 3D model is projected on a Gaussian sphere, and each point of the sphere is attributed with the total area of the faces having the same orientation [6]; for each facet of the Gaussian sphere, of orientation n_k:

$$P_{n_k} = \sum_{l=1}^{N_k} A_{l,n_k} \qquad (1)$$

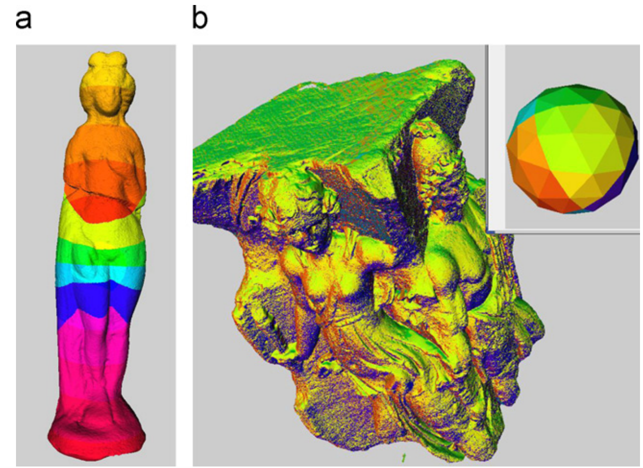

Fig. 1. Color representation of shape descriptors on a Venus figurine: (a) Cord2D descriptor; each color represents one cordlength; (b) EGI descriptor with the corresponding face orientations.

where N_k is the number of faces of the model in direction n_k and A_{l,n_k} is the area of the l-th face of orientation n_k (*cf.* Fig. 1(b)). The main drawback of this index is that it is unable to differentiate convex and concave parts of the objects.

Complex Extended Gaussian Images (*CEGI*): In order to deal with this drawback, the CEGI feature [7] describes an object in the complex space through two attributes: the face orientation, and the distance between the center of the face and the center of the object. With the same notation as Eq. (1)

$$P_{n_k} = \sum_{l=1}^{N_k} A_{l,n_k} e^{id_{l,k}} \qquad (2)$$

where $d_{l,k}$ is the distance between the center of the model and the center of the face l (this distance is negatively signed if the face is directed towards the model center, and positively else). The CEGI descriptor is composed of the module and the phase of P_{n_k}.

3D Hough transform: The 3D Hough feature [8] is an extension of the Hough transform consisting in accumulating the spherical coordinates (s,θ,ϕ) which uniquely define the plane containing a model face. A 3D histogram is then computed, where each face contributes proportionally to its area.

Spherical harmonic representation: Spherical harmonic transforms have been used to compute shape descriptors based on a voxel representation of 3D models [1]. More recently, spherical harmonic transform has been directly computed on 3D star-shape triangulated models [26,27]. We use a descriptor which is based on the spherical harmonic transform computation using 32 concentric spheres centered in the center of gravity of each object.

2.2. Global shape indexing and retrieval

We present in this paper an extension of our search engine RETIN, which was originally built for image retrieval. We consider browsing a database or retrieving a category as a supervised classification problem. The user gives one or several examples of the objects he is looking for and the system returns the objects the most similar to these examples. In classification tasks, there are many ways to form the classes, depending on the user expectations. For example is he looking for greek vases, for vases with one handle, for vases with a particular shape or painting? Another problem is to initialize the search: if the database is large, how extract enough examples?