Graphical Models 71 (2009) 63–76

Contents lists available at ScienceDirect

Graphical Models

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

Alignment of 3D models

Mohamed Chaouch *, Anne Verroust-Blondet

INRIA, Paris – Rocquencourt, 78153 Le Chesnay, France

ARTICLE INFO

Article history:
Received 23 July 2008
Received in revised form 2 December 2008
Accepted 22 December 2008
Available online 14 January 2009

Keywords:
3D alignment
Principal component analysis
Symmetry detection
3D shape retrieval

ABSTRACT

In this paper we present a new method for alignment of 3D models. This approach is based on two types of symmetries of the models: the reflective symmetry and the local translational symmetry along a direction. Inspired by the work on the principal component analysis (PCA), we select the best optimal alignment axes within the PCA-axes, the plane reflection symmetry being used as a selection criterion. This pre-processing transforms the alignment problem into an indexing scheme based on the number of the retained PCA-axes. In order to capture the local translational symmetry of a shape along a direction, we introduce a new measure we call the local translational invariance cost (*LTIC*). The mirror planes of a model are also used to reduce the number of candidate coordinate frames when looking for the one which corresponds to the user's perception. Experimental results show that the proposed method finds the rotation that best aligns a 3D mesh.

1. Introduction

Normalization of 3D models is a common pre-processing stage in many applications in computer graphics, such as, visualization, 3D object recognition, 3D shape matching and retrieval [2,19,22,26]. 3D models are generally given in arbitrary scale, position and orientation in 3D-space. Most of the methods do not satisfy geometrical invariance; hence it is important to normalize the models into a canonical coordinate frame before any processing. The normalization consists of two steps: the alignment to determine the pose-invariant and the scaling to make the scale-invariant. The alignment is the most difficult point in the normalization process.

To perform an alignment, a concatenation of isometries in 3D-space (translation, rotation and reflection) must be selected to determine the canonical coordinate system. In most of the methods, the center of gravity of the model is chosen as the origin to secure the translation invariance. However, the choice of a suitable rotation is still a well discussed topic [2,5,11,16,19,24,26]. Note that the alignment problem addressed in this paper is different from the alignment approaches of [5,11], where the purpose is to find the best alignment between two given 3D models. Here, we want to compute an intrinsic global coordinate system for each 3D object.

When looking at a 3D model, one can say whether it is well aligned or not and knows, in most of the cases, how to find its good alignment. When the 3D model has symmetries, it is aligned with particular axes or symmetry planes. This is confirmed by Ferguson [9] who noticed that symmetry detection is a key part of human perception and this fact has guided Podolak et al. [17] when introducing principal symmetry axes. Our goal is to find a method that best aligns any 3D model (an alignment similar to what a human would select – see left part of Fig. 1) and will consequently align two similar 3D models in the same way. In this paper, we show that by detecting *the planar reflection symmetries* we can select a set of good alignment axes. However, this method is guaranteed to give the correct alignment for only some cases and keeping only this type of symmetry is insufficient for computing the best alignment for any 3D model. An alternative method is to detect also *the local translational symmetry* that has an interesting semantic meaning: the object has the same geometrical properties in different parts along a given direction.

* Corresponding author. Fax: +33 139635674.
E-mail addresses: mohamed.chaouch@inria.fr (M. Chaouch), anne.verroust@inria.fr (A. Verroust-Blondet).

1524-0703/$ - see front matter © 2009 Elsevier Inc. All rights reserved.
doi:10.1016/j.gmod.2008.12.006

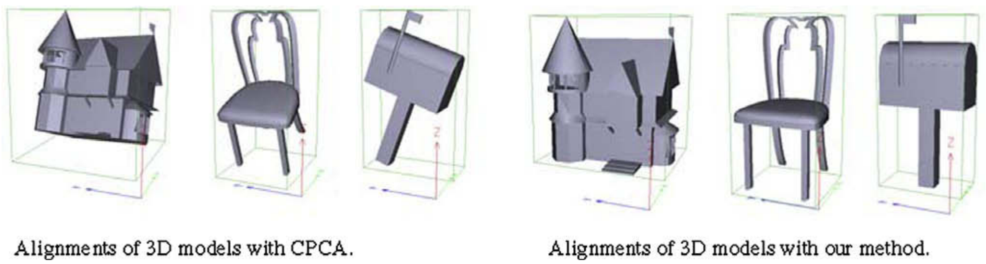

Fig. 1. Comparing CPCA based alignment and our alignment. In each case, green, blue and red arrows represent the principal axes (CPCA) and the suitable axes (our method). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

To build our general alignment algorithm, we first focus on discrete detection of plane reflection symmetries and classify a model in terms of its symmetry group and the number of its mirror planes. This classification is used to select the good alignment axes among those found by the principal components analysis (PCA). Then we introduce local translational invariance cost (*LTIC*) that measures the invariance of a model with respect to local translation along a given direction. This measure is used to compute the remaining alignment axes when the model has at most one good alignment axis given by the PCA. This paper is an extension of [4]: it gives more details on our alignment algorithm and discuss on methods computing reference frames from our alignment axes.

We first review related work on alignment and symmetry detection for 3D models in Section 2. Then we present our selection of the best alignment axes within the PCA-eigenvectors by analyzing the plane reflection symmetries (Section 3), and we describe our alignment method (Section 4). Experimental results evaluating our method are presented in Section 5. Finally, we discuss on the ordering and the orientation of the alignment axes in Section 6 and we conclude in Section 7.

2. Related work

The most well-known approach computing the alignment of 3D objects is the principle component analysis method (PCA) [2,16,19,24,26], which is based on the computation of moments of 3D models. After a translation of the center of mass to the origin of the coordinate system, three principal axes computed with PCA are used to determine the orientation. Experiences show that PCA-alignment has two disadvantages: (i) it is often imprecise and can produce poor alignments and (ii) the principal axes are not always good at aligning orientations of different models within the same semantic class (as noticed by Chen et al. [6] on the mug example).

Podolak et al. [17] introduce a planar reflective symmetry transform (PRST) that computes a measure of the reflectional symmetry of a 3D shape with respect to all possible planes. They use it to define two new concepts for the global coordinate system, *the center of symmetry* and *the principal symmetry axes*. The principal symmetry axes are the normals of the orthogonal set of planes with maximal symmetry, and the center of symmetry is the intersection of those three planes. This approach has been improved by Rustamov with the augmented symmetry transform [18].

Other methods finding symmetries in 3D models have been presented. These include Minovic et al. [14], who compute symmetries of a 3D object represented by an octree. Their method is based on the computation of a *principal octree* aligned with the principal axes. Then they compute a measure of symmetry, the *symmetry degree*, reasoning with the number of distinct eigenvalues associated to the principal axes. Furthermore, Sun and Sherrah [21] convert the symmetry detection problem to the correlation of the Gaussian image. Then rotational and reflectional symmetry directions are determined using the statistics of the orientation histogram. Finally, Martinet et al. [13] use generalized moments to detect perfect symmetries in 3D shapes and Mitra et al. [15] and Simari et al. [20] compute partial and approximate symmetries in 3D objects.

Our goal is to align 3D models using their planar symmetry properties. Our method must be such that similar objects (i.e., objects belonging to the same semantic class) have similar alignments. As noticed in [14], any plane of symmetry of a body is perpendicular to a principal axis. As a result, for models that have plane reflection symmetries, some PCA-coordinate planes coincide with some mirror planes. Therefore, we have chosen to use the PCA, not for global alignment, but for selection of robust partial alignment features of a model (i.e., only the principal axes that we consider good for a perfect alignment).

Given a 3D model, the first key idea is to test the reflection symmetry of the PCA-coordinate planes. According to the result of this test, we select a set of principal axes and use them in our alignment method. When the model has at least two orthogonal mirror symmetries, the PCA gives the good alignment. In the other cases we use the local translational invariance cost along a direction to compute the good alignment axes.

Before describing our alignment procedure, let us classify the 3D polygonal models with respect to their plane reflection symmetry and select classes of objects where PCA gives a good alignment.

3. Symmetry and 3D objects

In what follows $\mathscr{M}$ denotes a 3D polygonal model represented by its surface $\mathfrak{S}$ composed of a set of triangular facets $\mathfrak{T} = \{T_1, \ldots, T_{N_T}\}$. We study the reflection planes in the symmetry groups [8], and use them to discriminate differ-