

Salient Geometric Features for Partial Shape Matching and Similarity

RAN GAL and DANIEL COHEN-OR
Tel-Aviv University

This article introduces a method for partial matching of surfaces represented by triangular meshes. Our method matches surface regions that are numerically and topologically dissimilar, but approximately similar regions. We introduce novel local surface descriptors which efficiently represent the geometry of local regions of the surface. The descriptors are defined independently of the underlying triangulation, and form a compatible representation that allows matching of surfaces with different triangulations. To cope with the combinatorial complexity of partial matching of large meshes, we introduce the abstraction of salient geometric features and present a method to construct them. A salient geometric feature is a compound high-level feature of nontrivial local shapes. We show that a relatively small number of such salient geometric features characterizes the surface well for various similarity applications. Matching salient geometric features is based on indexing rotation-invariant features and a voting scheme accelerated by geometric hashing. We demonstrate the effectiveness of our method with a number of applications, such as computing self-similarity, alignments, and subparts similarity.

Categories and Subject Descriptors: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling

General Terms: Algorithms

Additional Key Words and Phrases: Partial matching, shape retrieval, salient features, similarity, geometric transformations

1. INTRODUCTION

Matching is a fundamental task in a vast number of geometric applications in numerous fields such as computer-vision, robotics, molecular-biology, and others [Veltkamp and Hagedoorn 2001]. Recently, as a result of the extensive availability of three-dimensional (3D) models, interest in 3D shape-retrieval techniques has increased. Until now, research has focused on global matching, where similarity is measured between entire models [Tangelder and Veltkamp 2004]. Partial matching is the task of matching subparts or regions. The parts that are matched are not predefined, can be any subshape of a larger shape, and can be in any orientation or scale. Partial matching is a much harder problem than global matching, since it needs to search for and define the subparts prior to measuring similarities.

In this article we focus on partial shape matching of surfaces represented by triangular meshes. Generally, it is difficult to define a metric that agrees with the human perception of resemblance. Even with a given similarity measure, it is a challenge to match shapes that have different representations. In addition, there is a huge combinatorial complexity barrier to cross to make partial matching practical. For example, the four lotus flowers shown in Figure 1 are part of the surface of the Buddha model. Matching them requires an extensive search of the over 1M triangles of the Buddha model. Note that

Authors' address: School of Computer Science, Tel-Aviv University, Tel-Aviv 69978 Israel; email: {galran,dcor}@tau.ac.il.
Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 1515 Broadway, New York, NY 10036 USA, fax: +1 (212) 869-0481, or permissions@acm.org.
© 2006 ACM 0730-0301/06/0100-0130 \$5.00

ACM Transactions on Graphics, Vol. 25, No. 1, January 2006, Pages 130–150.



Fig. 1. Left: curvature analysis of the Buddha model and the local surface descriptors. Blue is low curvature and red is high. Right: the self-similarity of the four lotus flowers, which are detected by our partial matching technique.

the shapes of the flowers are incomplete and nonidentical; also, the flowers are positioned at various orientations, and their extent is not known a priori.

To alleviate the problem, we define a sparse set of *local surface descriptors* across the surface. These descriptors represent local regions of the surface, defined independently of the underlying triangulation, and thus form a compatible representation that allows measuring the similarity between regions with possibly dissimilar triangulations. Since the number of descriptors is significantly smaller than the number of vertices, they reduce the combinatorial complexity of the surface representation. By carefully defining the descriptors, their descriptive power is sufficiently high and effective. Nevertheless, the number of descriptors is still too large to allow efficient partial matching of large and complex surfaces. Thus, we define *salient geometric features* which form compound higher-level descriptors. A salient geometric feature, or in short, a salient feature, consists of a cluster of descriptors that locally describe a nontrivial region of the surface.

The salient features of a surface typically characterize the surface well and form a basis for a nonglobal similarity measure among subparts of shapes. Salient feature matching permits various notoriously hard applications, such as finding similar parts across a single surface or a number of surfaces. For example, in Figure 1, the four lotus flowers of the Buddha model are detected by searching for self-similarity, that is, nontrivial subparts which match other subparts over the same surface.

In the paper we make the following contributions:

- Partial matching*. We address the problem of partial-matching for triangular meshes. The method extends recent results on global matching of 3D models.
- Local surface descriptors*. We introduce new local surface descriptors which efficiently encode regions of the surface. These are defined independently of the underlying triangulation, and form a sparse compatible representation that allows matching of surfaces with dissimilar triangulations.

- Salient geometric features.* We introduce the abstraction of salient geometric features and present a method to construct them. A salient geometric feature is a compound high-level feature that characterizes a local partial shape. With each salient feature we associate a number of rotation-and-scale invariant indices to accelerate matching operations and similarity measures.
- Applications.* We show a number of geometric applications in which partial matching is vital. We apply our method on various large and complex meshes to demonstrate its efficiency.

2. BACKGROUND

The problem of similarity and matching of shapes has been extensively studied in numerous fields such as computer-vision, robotics, molecular-biology, and others [Veltkamp and Hagedoorn 2001]. Most of the work has focused primarily on matching shapes in 2D images. Matching 3D models seems an easier problem since the geometry is given, and there is no occlusion or disrupting external effects such as lighting and reflections. On the other hand, 3D models typically lack a simple parameterization domain, and thus registration and feature correspondence are more difficult tasks.

In the computer graphics field, matching of 3D shapes was developed mainly for shape retrieval. Recently, new methods were developed for the retrieval of 3D models in the context of a Web search engine, based on geometric properties rather than textual ones [Elad et al. 2001; Vranic et al. 2001; Iyer et al. 2003; Shilane et al. 2004; Min 2004]. However, partial matching methods, which are vital in 2D matching with occlusions, have not received much attention in 3D shape retrieval [Tangelder and Veltkamp 2004].

Most of the techniques aim at retrieving similar models based on their overall geometries. In other words, the similarity is based on the global properties of the 3D models such as moment invariants [Elad et al. 2001; Cybenko et al. 1997], Fourier descriptors [Vranic et al. 2001; Saupe and Vranic 2001; Ohbuchi et al. 2003b], histograms and shape distributions [Osada et al. 2001; Ohbuchi et al. 2003a], and harmonic-based representations [Funkhouser et al. 2003; Kazhdan et al. 2003].

These methods rely on the ability to first align and normalize models by some global similarity transformation (rotation + uniform scale) that normalizes the models and establishes some correspondence between them. Such global alignment methods do not discriminate between the details, and can easily cause similar local features to be misaligned and, consequently, result in an improper global similarity measure. This has motivated Kazhdan et al. [2004] to control the anisotropy influence as a means to compensate for the misalignments of the global methods. In this respect, the partial matching technique that we introduce here can be used to align two models that have little global similarity. With our partial matching method, we can explicitly match corresponding prominent features and establish a global alignment derived from the partial matching (see Figure 2).

Graph-based methods rely on the fact that 3D model topology is an important shape characteristic. Topology is typically represented as a relational data structure such as graphs [Hilaga et al. 2001; Chen and Ouhyoung 2002; Sundar et al. 2003]. A graph-based representation facilitates the partial matching of subparts that have the same topology.

In Funkhouser et al. [2004] the authors presented a method for 3D matching based on the sum of squared distances between two model's shape descriptors. The method can also allow varying weights (importance) of different parts of the model using the original model shape descriptor, thus enabling a part-in-a-whole matching. The method relies on the assumption that the models are aligned to bring the matching parts closer to one another in 3D space.

A part-in-a-whole matching method can be classified as a partial matching under an identity transformation. For a general partial matching, one would like to match parts under any rigid motion transformation and sometimes under a scale transformation as well. One can also consider a broader spectrum of transformations, for example, affine or perspective. In this article we develop a partial shape matching




Fig. 2. Matching a model of a man and a model of a woman in different poses. In spite of the obvious difference between them, the algorithm matches the two (bottom left) better than a global PCA alignment (bottom right).

method that supports matching under rigid transformations and uniform scale, known as a *similarity transformation*. Our method does not globally align the shapes, a fact that negates the possibility of matching, for example, the hand of a statue with a large base, to the hand of the same statue without the base.

Partial matching is a fundamental building block in the recognition and the registration of 3D objects from depth images generated by range scanners [Frome et al. 2004]. In such applications the query is usually view-dependent and resolution-dependent. Usually, these techniques are based on matching local descriptors, such as Spin images and shape context [Johnson and Hebert 1999; Huber et al. 2004]. Unlike those descriptors, the surface descriptors we develop in our work are scale-independent and adaptive in the sense that they are carefully built to efficiently represent mesh regions as large as possible. We further elaborate on these methods in Section 9.

The extraction of “interesting” points or salient features in image space has been studied extensively [Marr 1982]. The definition of interesting surface features, or salient features, is a central issue in our work. As we elaborate below, our definition of interesting parts is heavily based on the surface curvature. Shum et al. [1996] recognized the importance of the surface curvature as a means to match 3D shapes (see also Zaharia and Preteux [2001]).

2.1 Salient Geometric Features

Human perception and the ability to recognize and interpret shapes is a broad and extensively studied subject. Indeed, it is not one task, but many: color, shading, shape, motion, texture, and context are all typically used in the process, not to mention cultural background and personal associations. A central problem in computer graphics and modeling is the understanding of the concept of “shape.” For our purpose we use the formal definition of shape defined in Dryden and Mardia [1998]: “all the geometrical information that remains when location, scale, and rotational effects (Euclidean transformations) are filtered out from an object.”

There is a consensus that representing shapes in terms of their parts may aid the recognition process in human vision [Hebb 1949]. In 2D, it can help in the handling of occlusions and in 3D, in the handling of nonrigid bodies. A critical question is: “which parts?” In our research we define, by means of general computational rules, the salient parts of an object that encapsulate enough of its characteristics to be

used in a “first-index” search [Hoffman and Singh 1997]. By “salience of a part” we mean that we aim at a small number of parts that capture the object’s shape. Their salience determines, in part, their efficacy as an index.

Our approach is built on the theory of salience of visual parts proposed by Hoffman and Singh [1997]. According to their theory, the salience of a part is a function of its size, relative to the whole object, the degree to which it protrudes, and the strength of its boundaries. In this work, we aim at developing quantitative definitions for these measures that can be applied to meshes. The emphasis is on the computation of the “strength of the boundaries”. Our approach is that this is a function of the curvature values. We propose that the salience of a part depends on (at least) two factors: its size relative to the whole object, and the number of curvature changes and strength. In Section 5 we describe the computation method for defining salient geometric features.

3. OVERVIEW

Given a 3D mesh, we analyze its geometric properties and define a sparse representation of the mesh with local shape descriptors. Then we extract interesting and salient geometric features from the mesh. These features are indexed and stored in a model database that can be queried efficiently for partial matching. Matching subparts requires transforming them into a common coordinate system. A naive implementation would be to store all the models in all the possible coordinate systems defined by each vertex and its neighborhood. Then, for a given query, we would test all possible orientations and scalings defined by the vertices of the query. This approach would cause a “combinatorial explosion” of tests, far beyond our computational abilities, especially with models containing thousands of vertices and more.

To overcome these combinatorial barriers, we present three methods:

- Represent the given mesh with a sparse set of shape descriptors. Each descriptor represents a local region that has a good quadric fit.
- Define a small set of salient geometric features and index them. The salient features represent interesting parts of a given shape.
- Precompute a geometric hash table that enables fast partial matching under rotation and scale transformation.

In Sections 4–6 we describe the methods which

- analyze the model’s surface;
- define local shape descriptors;
- define salient geometric features and store a vector of indices;
- query the model database for partial matches using indexing and geometric hashing.

Then in Section 7 we discuss performance issues of our method. In Section 8 we show a number of applications and results based on our partial matching algorithm. In Section 9 we elaborate on the relation between our geometric hashing method and the histogram based methods, and we conclude in Section 10 and discuss future work.

Before we move on, we refer to Figure 3, where the essence of partial matching is illustrated. The figure shows three partial matches of the models in (a) and (b). Note that only small portions of the two models actually match. This example emphasizes the fact that partial matching is more than matching subparts of the given models, and that the matched portions are not predefined.

4. LOCAL SURFACE DESCRIPTORS

A local surface descriptor is a point p on a surface and its associated quadric patch that approximate the surface in a local neighborhood of p . The key idea is that a small set of descriptors can represent




Fig. 3. Partial matching. Three portions of the eagle model (b) are superimposed on a frame over which elevation deformation is applied (a). The three copies have different scales and orientations. The three portions of the eagle are matched with the full eagle model by applying partial matching (c)–(e). Note that the overlap between each matched pair is only partial. The matching is performed by transforming pairs of triplets of local shape descriptors into a common coordinate system. The local shape descriptors are illustrated with small yellow spheres.

a shape, provided that each descriptor effectively represents the local surface region around it. This representation is adaptive to the geometry of the shape. Smooth regions are represented by a relatively small number of descriptors, since each quadric patch can locally approximate a large region. Less smooth regions of high frequency require a relatively large number of descriptors. Note in Figure 5(b) the density of local shape descriptors around the suction cups and at the smoother regions. As a means to analyze the local surface, we associate curvature values with each descriptor. In the following, we describe the quadric fitting scheme and the algorithm to construct an efficient set of descriptors.

4.1 Quadric Fitting and Curvature Estimation

Using the method presented in Douros and Buxton [2002], the surface is locally approximated at each vertex by an analytic quadric patch that approximates the surface well. For each quadric patch, a representative point is selected and the differential properties of the patch at that point are then calculated analytically.

For the implicit surface $F(x, y, z) = 0$, we use a quadric which has nine coefficients. The fitting of such a surface is easy to formulate as a least squares minimization that can be solved as an eigenproblem. This allows the method to work on any type of point cloud in 3D, regardless of the orientation of the underlying surface. Note that such a representation does not single out any coordinate axis.

For each vertex v of the mesh, we fit an implicit surface, and its projection onto the implicit surface is denoted by \tilde{v} . The curvature tensor and curvature derivatives at v are estimated by those computed analytically at \tilde{v} [Ohtake et al. 2004]. To increase the robustness, we test a number of neighborhood sizes around each vertex, and select the quadric that has a median curvature. There are faster techniques to estimate discrete curvatures [Alliez et al. 2003; Rusinkiewicz 2004], but since these curvature computations are offline, we implemented a quadric fitting technique which we found to be generally more robust. Figure 4 illustrates the local quadric fittings and the estimated curvature map.

4.2 Compact Representation

To generate a rather small and yet effective set of local surface descriptors, we define a local surface patch and associate it with a descriptor that represents the patch. Our method uses a region-growing




Fig. 4. Examples of the quadric fitting. Red points are samples on a grid defined by the principal axes (in blue). The mesh is colored by the Gaussian curvature.




Fig. 5. Example of suction cups on a tentacle of the octopus model. (a) original tessellation and the quadric patches in bold. (b) The local surface descriptors are located at the yellow points. Note that their density is adaptive with respect to the surface shape.

technique to iteratively define the local patches. We sort the mesh vertices according to their absolute Gaussian curvature in descending order, and for each vertex, we grow the largest possible “quadric patch” that approximates its neighborhood. All the vertices included in the local patch are then excluded from the sorted list, and the iterative process continues to define the next patch. Our method is greedy in the sense that in each step we define the largest possible region around the selected vertex that meets the prescribed error threshold. We use 10^{-4} of the model bounding box diagonal length for all the examples shown in this article. We measure the quality of a quadric fitting using the sum of squared algebraic distances of the fitted points from the fitting surface. An algebraic distance is much easier to evaluate than a *Euclidean* distance and although it is less robust it is an effective measure of quality.

Once a patch is defined, we select a representative point at the center of its mass, and associate with it the highest curvature across the patch. Figure 5 illustrates the result on a tentacle of the octopus model. Recall that one of the main objectives of the local surface descriptors is to define a representation independent of the given model tessellation. Thus, we randomly (not necessarily uniformly) sample the surface [Osada et al. 2001; Elad et al. 2001], and use these random samples for the quadric fitting. See Figure 6 for an example of the curvature computed over two different tessellations. It can be seen that the curvature maps and distribution of the local surface descriptors are similar despite the different tessellations.

5. COMPUTING SALIENT GEOMETRIC FEATURES

Loosely speaking, a salient geometric feature is a region of the surface which has a nontrivial shape. We construct a salient geometric feature by clustering together a set of descriptors that are interesting




Fig. 6. The curvature map computed over the same model with different tessellations. Blue is low curvature and red is high.

enough in the sense that they have a high curvature relative to their surroundings, and a high variance of curvature values. The salient geometric feature is a function of one free parameter h , the *scale*, which is defined by the radius of a sphere that bounds the salient feature.

For each descriptor, we grow a cluster of descriptors around it by incrementally adding a descriptor from its neighborhood that maximizes the current saliency grade, which will be shortly defined. The process stops when the contribution of an additional descriptor is insignificant. We perform this process on all the descriptors while allowing some degree of overlapping between clusters (in our implementation we allowed up to 20% overlapping).

The saliency grade of a cluster F consisting of $d \in F$ descriptors is a function of the following four terms:

- $\sum_{d \in F} Area(d)$, where $Area(d)$ is the area of the patch associated with d relative to the sphere size;
- $\sum_{d \in F} Curv(d)$, where $Curv(d)$ is the curvature associated with d ;
- $N(F)$ —the number of local minimum(s) or maximum(s) curvatures in the cluster; and
- $Var(F)$ —the curvature variance in the cluster.

A saliency grade S can be defined as a linear combination of these four terms. However, empirically we found the following expression more effective:

$$S = \sum_{d \in F} W_1 Area(d)Curv(d)^3 + W_2 N(F)Var(F). \quad (1)$$

The first term $Area(d)Curv(d)^3$ expresses the saliency of the region, by combining its relative size and curvature. $Curv(d)$ can be either the Gaussian curvature or the maximal curvature component. Figure 7 shows the two curvature maps of the women from the Thai statue, showing no significant differences. Note that a maximal curvature can be a better choice in case of CAD models that have large areas with zero Gaussian curvature. The second term $N(F)Var(F)$ expresses the degree of interestingness of the cluster, and adds up to the saliency of the cluster, as shown in Figure 8. We used 0.5 for the weights W_1 and W_2 . The definition of salient parts is not too sensitive to the heuristics used to define them, since we are not interested in a particular part, but rather, in generating a number of salient features that characterize the model well (see Figure 9). It should be emphasized that the constants employed in Equation (1) require no manual tuning, and all the examples and experiments used in this article use the same expression.

In other words, we define a salient geometric feature as a region of the model which is salient and interesting compared to other parts of the model. The top graded clusters define the salient geometric




Fig. 7. A curvature analysis over the three women from the Thai statue. The images are results of using Gaussian curvature (left) versus Max curvature (right). The minor difference indicates that definition of salient features is insensitive to the analysis.



Fig. 8. The left image shows the local surface descriptors colored by the curvatures. In the image on the right, the saliency grades are visualized on a scale from blue to red. The patches are colored according to the saliency grade of the salient feature that contains them.

feature of a given shape. Some results of salient geometric feature are illustrated in Figure 9. We extract only the most salient features, either those with grades in the top 10% or those that meet some prescribed threshold.




Fig. 9. Examples of salient geometric features (in orange) defined automatically by analyzing the surfaces. The individual parts are in red.

6. INDEXING AND GEOMETRIC HASHING

The representative salient features of the model are defined and extracted offline and stored in a rotation-and-scale-invariant database. Each salient feature is associated with a vector index (a signature) and inserted into a geometric hash table. The vector of indices is used as a geometric signature that allows fast access to parts that have general similar shapes. To be effective, we use rotation-and-scale-invariant geometric signatures. However, the signature alone is not discriminative enough and does not provide the actual transformation that matches them. By employing geometric hashing [Lamdan and Wolfson 1988] the explicit transformation that best matches the shapes is determined.

In our implementation we use the four terms that were used to define the saliency grade, as expressed by Eq. (1), as a similarity-invariant vector index. Note that other elaborated indices, like normalized moments could be used as well. However, implementing the indexing using the saliency terms reinforces our claims for the efficiency of salient features.

We use geometric hashing to accelerate the retrieval of partial matches from a model database. Briefly, geometric hashing is a mechanism that accelerates partial matching under rigid transformation and possibly a uniform scale. The method determines the best transformation from among a large set of transformations by means of a voting scheme. The method quantizes the transformation space into a six-dimensional table. Given two objects, A and B , for each triplet of points on A and each triplet on B , the transformation between the triplets is computed and a vote is recorded in the corresponding cells of the table. The entry with the most votes defines the aligning transformation. To avoid an exhaustive search of all the $\binom{|A|}{3} \cdot \binom{|B|}{3}$ transformations, geometric hashing encodes offline all the candidate transformations in a large hash table, so that, given a query model of size $|Q|$, it needs to query the hash table only for $\binom{|Q|}{3}$ transformations, regardless of how many models are encoded in the hash table. The precomputed hash table obviates a search over each of the model triplets.

In our implementation, we use salient geometric features as object and local surface descriptors as points that are used to define the triplets and votes in the hash table. The hash table contains all the salient geometric features from all models in the database. Typically, each salient feature consists of around 20–30 local surface descriptors. We define candidate transformations between triplets of points

Table I. Some Statistics for Large and Small Models (LDSs are *local surface descriptors* and SGFs are *salient geometric features*.)

	Buddha	Octopus	Enterprise	Fort
# of Vertices	543,652	149,666	26,270	2,328
# of Faces	1,087,716	299,332	51,355	4,460
# of LDSs	5,330	1,912	647	132
# of SGFs	113	47	19	12
Building Time	19 min	6.5 min	65 s	4 s
Storage Size	3.5 GB	930 MB	97 MB	9.4 MB
Query Time	36 s	13 s	3 s	1 s

taken from a small neighborhood (up to two rings), so that the number of alignment tests typically required is only in the order of $|Q|$, rather than $\binom{|Q|}{3}$.

By combining indexing and geometric hashing, we successfully find partial matches between the query and the salient features stored in the hash table. The index is used to quickly reject salient features during the voting phase. The transformation that matches two given salient features is the building block for various matching and similarity applications as described below.

7. TIME AND STORAGE CONSIDERATIONS

The proposed method enables performing most of the work necessary for partial matching in offline mode, allowing fast query time. Table I shows examples of the time and storage space needed for our preprocess step.

In its core, a partial matching problem require intensive computations for searching and comparisons. We use a rather time consuming analysis of the surface to create a sparse and adaptive set of descriptors. To index our constructed salient geometric features, we use geometric hashing. The main advantage of geometric hashing for indexing is that it allows a scale independent mechanism which is important for our applications. For simple cases, where no scale is needed, one can consider implementing other methods for indexing the salient geometric features, such as spin images and shape context [Johnson and Hebert 1999; Huber et al. 2004].

Some applications or scenarios may demand a more efficient preprocess step. As we mentioned above, there are fast methods to estimate discrete curvatures and quadric coefficients [Petitjean 2002; Alliez et al. 2003; Rusinkiewicz 2004]. These methods can speed up the discrete curvature estimation for good-quality models, but are less robust for models with irregular tessellations.

In addition, there is a large body of work on how to improve the efficiency of geometric hashing, both in time and space requirements [Wolfson and Rigoutsos 1997].

8. APPLICATIONS AND RESULTS

Partial matching, in general, is a fundamental tool for various geometric applications. In particular, there are numerous applications for surface processing. We describe here three applications that we have experimented with so far: self-similarity, shape alignment, and subpart retrieval.

8.1 Self-Similarity

Figure 10 shows the results of a search for self-similarity. The surface is analyzed by identifying a popular feature that characterizes the surface. The idea is to automatically identify salient features that have a multitude of similar occurrences across the given surface. In the figure, the suction cups on the tentacle of the octopus model are identified (see Figure 5 for a close view of these cups). Note that the suction cups have a relatively large variety of sizes. Figure 1 is another result of a self-similarity application, by which we learned about the existence of the four lotus flowers on the Buddha model.



Fig. 10. Self-similarity. The suction cups on the tentacle of the octopus model have a similar shape. Note that their size varies across the model.

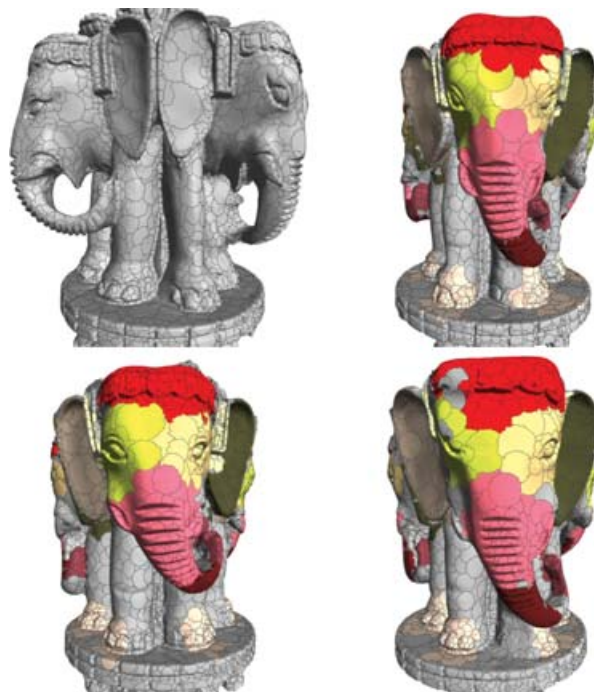


Fig. 11. Symmetry analysis of the elephants (up left). Similar regions on the three elephants are detected and colored.

Table I shows some statistics for applying a query to the Buddha and Octopus models which are large, and to two other smaller models.

In the above examples, copies of all the salient features are searched under rotation and uniform scale transformations. The algorithm first sorts the salient features of the given surface, and then, for each of the k salient features with largest saliency grades, we apply a search for partial matching. Figure 11



Fig. 12. The Thai statue from the Stanford 3D scanning repository. The model, which consists of 10M polygons, has a rough three-way symmetry.

shows a self-similarity analysis applied to part of the huge model of the Thai statue from the Stanford 3D scanning repository. This model, shown in Figure 12, consists of 10M polygons, which is beyond the capacity of our system. In our experiments, we processed only parts of it: the three women shown in Figure 7, and the three elephants shown in Figure 12. The statue has a rough three-way symmetry, but the three parts of the handmade statue do not match since they are not perfect copies. We applied our self-similarity analysis and for each match we grow larger regions that form compatible parts. The regions are grown as long as they match up to some error tolerance. The results of this three-way symmetry analysis are shown in Figure 11.

8.2 Shape Alignment

Most global matching methods first determine a rigid-body transformation and a uniform scale, which align two models together as closely as possible, before measuring the distance between them. This is typically achieved by a principal component analysis (PCA) applied to the whole model. Such a global analysis is prone to errors when the shapes disagree even in an apparently minor region. Figure 2 shows an alignment between models of a man and a woman. The example shows that, although the two models have obvious differences in their geometric shapes and poses, by partial matching of their salient geometric features, it is possible to align the two where they match.

Another example is shown in Figure 13 where the same models appear in two different poses. The latter example emphasizes that a global PCA alignment fails to align them correctly, whereas the partial match alignment correctly aligns wherever possible. Our alignment is based on a majority scheme that finds the transformation which satisfies most of the salient features.



Fig. 13. Model alignment based on global principle component analysis (PCA) (right) and partial matching (left).

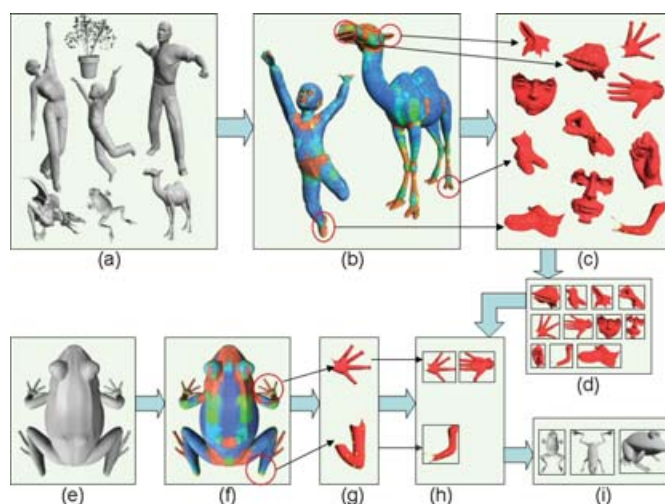


Fig. 14. Overview of our framework for shape retrieval. Given a database of models (a), we analyze each model's geometric properties and define a sparse representation of its mesh with local shape descriptors (b). Then we extract interesting and salient geometric features from each model (c). These features are indexed and stored in a model database that can be queried efficiently for partial matching (d). The steps described above are performed offline. When given a query model online (e), we analyze it in the same fashion, defining local surface descriptors (f) and salient geometric features (g). For each salient geometric feature, we query the model database for a list of matching features (h). Using these matches, we return the models that have larger number of matches (i).

Our alignment algorithm first searches for matches between pairs of salient features one from each model. Then for each such match, the associated transformation gets a grade that reflects the number of salient features it successfully align. The most successful transformation is voted, and applied to bring the two given models close wherever possible. The voted transformation defines a correspondence between the two models. Once this correspondence is defined, the corresponding features can be brought closer in the least-squares sense.

8.3 Partial Shape Retrieval

Applying partial shape retrieval allows subparts of a larger shape which might be dissimilar to the query to be retrieved. An overview of our framework for partial shape retrieval is shown in Figure 14.




Fig. 15. Two examples ((b) and (c)) of successful matches of the coarse human hand (a), and two false positive matches ((d) and (e)).

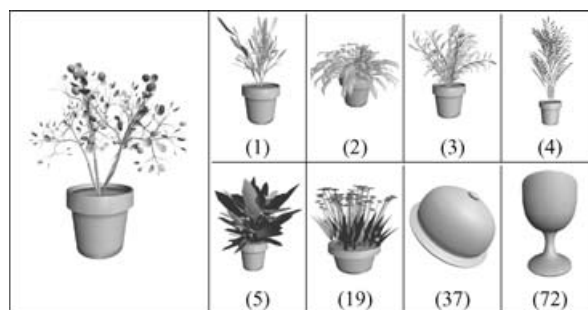


Fig. 16. Similarity retrieval results for a query with a plant model. The models retrieved are not necessarily globally similar, but contain partial similarities. Note that the figure does not show the eight best matches, but the first five and some samples from the first 72 matches.

Figure 15 shows an example of retrieving human hands. Such a mechanism is vital for an efficient modeling by example [Funkhouser et al. 2004], since it finds similar shapes under rotation and uniform scale transformations which are part of a larger dissimilar shape. See, for example, Figure 16 where the query is a plant. The models retrieved are not necessarily globally similar (e.g., some plants are tall and some short and wide), but they have enough partial similarities. Note that Figure 16 does not show the best matches, but the first five and some samples from the first 72 matches.

Interesting results, which cannot be achieved by global matching techniques, are shown in Figures 3 and 17. The two examples in Figure 17 show cases where the semantics of the retrieved shapes are different from the query object; nevertheless, the subparts are geometrically similar to the geometry of the query object.

To evaluate the performance of our method for partial matching, we report on the precision-recall experiments conducted over the Princeton Shape Benchmark (PSB) model database [Shilane et al. 2004]. We used only the natural models where the classifications included classes such as human hands, human feet, human faces, animal limbs, etc. Table II shows a summary of our results using the evaluation methods from [Shilane et al. 2004]:

- nearest neighbor*: the percentage of closest matches that belong to the same class as the query;
- first tier and second tier*: the percentage of models in the query’s class that appear within the top K matches, where K depends on the size of the query class; and
- e-measure*: a composite measure of the precision and recall for a fixed number of retrieved results.

For more details on the discrimination methods, see Shilane et al. [2004].

Our experiments were conducted on the P4-3.0 GHz with 2 GB of memory. We used the “saliency grade” (Equation (1)) as an index. However, it should be stressed that most models in the PSB database




Fig. 17. Partial matching. The submarine (in orange) is matched to two parts of the Enterprise (left) and four towers of the castle (right).

Table II. Summary of Our Experiments on the Princeton Shape Benchmark Database and the Precision-Recall Measures

General	# of salient features	1843
	Average # of surface descriptors	18
	Storage size	1850 MB
Timing	Generate time	55 min
	Average compare time	23 s
Discrimination (all models)	Nearest neighbor	87.3%
	First tier	78.2%
	Second tier	84.9%
	E-measure	42.5%
Nearest-neighbor per class examples	Human hand	93.4%
	Human feet	86.2%
	Animal limb	81.6%
	Human head	76.1%

have a coarse triangulation (see, for example, Figure 15(a)). In such cases, the saliency grade is too crude and can be quite different for similar coarse shapes. To avoid this, for coarse models, we directly used the local surface descriptors in the geometric hash table. Nevertheless, our precision-recall scores were quite good with respect to the typical results achieved in global matching [Shilane et al. 2004]. It should be noted that much of the success depended on the definition of the classes. For example, human hands are well defined, while other classes, like animal limbs or human faces, are harder to define.

9. HISTOGRAM-BASED DESCRIPTORS

Partial matching plays an important role in the recognition and registration of 3D objects from depth images or point clouds produced by range scanners [Johnson and Hebert 1999; Huber and Hebert 2003; Frome et al. 2004]. In these applications, the query scene is usually view-dependent and resolution-dependent, and the data is often noisy and incomplete.

Treating partial matching with purely local descriptors, such as surface curvature, is less stable in the presence of noise. Thus, stable point descriptors rely on local neighborhood support. For example, spin images [Johnson and Hebert 1999] and shape contexts [Frome et al. 2004] are well-known quantized representations of the local neighborhood of a given point.

These descriptors take as input a point cloud P and a basis point p , and capture the regional shape of the scene at p using the distribution of points in a small support surrounding p . The support region




Fig. 18. Spin images computed at similar shapes of different scales on the stars model. The figure illustrates four spin images at locations for which humans perceive as similar shapes of different scales. It can be seen that, while the spin images are similar when computed over shapes with similar scale (left images), they are dissimilar for similar shapes of different scales (right images).

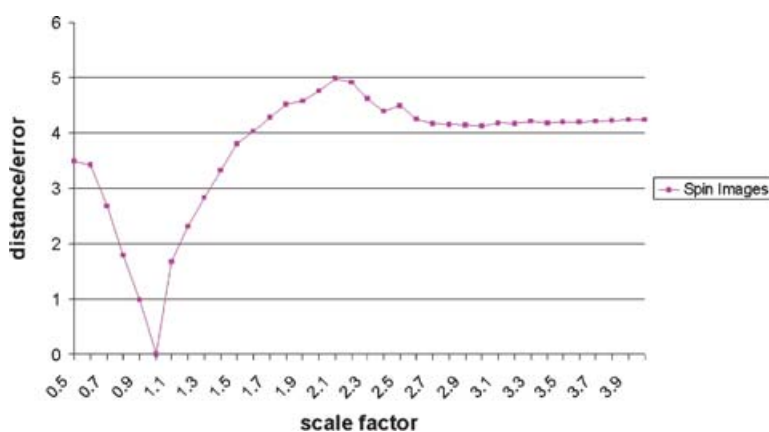


Fig. 19. The L_2 distance between spin images computed around the same points at different scales and the spin image at scale 1.0. The graph shows the sensitivity of the spin-image signature to scale.

is discretized into bins, and a histogram is defined by counting the number of points contained within each bin. These histograms can be used directly as the local descriptor or can be further enhanced by additional transformations such as spherical-harmonic representation.

These local descriptors can serve as an alternative for our local surface descriptors. However, these descriptors are not scale-invariant and they are not adaptive to the geometry, as we shall explain next.

The histogram-based descriptors [Johnson and Hebert 1999; Huber and Hebert 2003; Frome et al. 2004] are not scale-invariant. For example, in Figure 18, we illustrate four spin images at locations for which humans perceive as similar shapes of different scales. It can be seen that, while the spin images are similar when computed over shapes with similar scale (left images), they are dissimilar for similar shapes of different scales (right images). Spin images' sensitivity to scale is shown in Figure 19. Applying some scale normalization is not applicable since the matched instances are not known a priori.




Fig. 20. The stars model (a) consists of 18 stars of three different scales. A medium-scale star is used as a query. The result of using spin images and our method are shown in (b) and (c), respectively. This example illustrates the fact that our method is scale-independent.

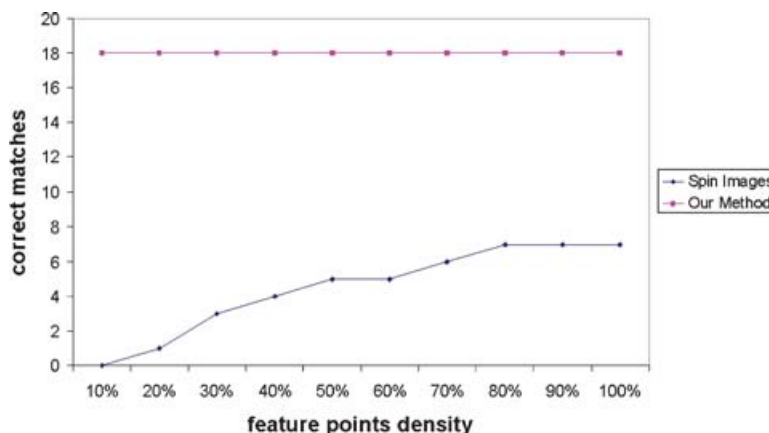


Fig. 21. This figure shows the tradeoff between the number of descriptors and the number of correct matches. It can be seen that our method requires only a small number of descriptors to successfully find all the partial matches.

In our method, we match salient features using geometric hashing, which is inherently a scale-invariant method, as can be seen in Figures 3 and 20. To illustrate the difference between spin image descriptors, and our descriptor, we employ the spin-image descriptor in a different setting. We distribute a number of spin images over the surfaces, and for each spin-image descriptor we compute its best matches rather than a single match. Then we employ a voting scheme by computing the transformation that gets the most of matchings between corresponding descriptors (see Figure 20).

Another interesting issue is the tradeoff between the efficiency and the fidelity of the matching. A dense set of spin images is likely not to miss the best match. However, an exhaustive approach would be overly costly. In our method a sparse set of descriptors are carefully defined by with respect to underlined geometry by analyzing the surface as described in Section 4. Figure 21 shows the tradeoff between the number of descriptors and the number of correct matches. It can be seen that our method requires only a small number of descriptors to successfully find all the partial matches.

It should be noted that in our method the selection of the descriptors which requires the analysis of the surface is time-consuming and computed offline. Unlike the histogram-based techniques, we invest computation time to create a sparse and adaptive set of descriptors with a good discriminative power.



Fig. 22. Registration of four synthetic scans. The surface of each scan has normals enclosed in a normal cone. Matches of overlapping areas were found using our method.

10. CONCLUSIONS AND FUTURE WORK

In this article, we have described a method for partial matching of triangular meshes. The method generalizes the recent global geometric matching techniques developed for shape retrieval. The methodology we used is based on an analysis of the surface to define a set of sparse local surface descriptors that captures the essence of the geometry of the original mesh. The sparseness of the descriptors allows one to deal with complex meshes and with large regions that have no significant details. Our method is designed to handle fine meshes since it analyzes the surface curvature over discrete representation. When the mesh is coarse, the local curvature measure is too crude to convey significant information for the analysis. The fidelity of the local surface descriptors, and consequently the efficiency of the salient features, directly depends on the quality of the mesh and the curvature analysis. We believe that both local surface descriptors and salient features are of considerable importance for mesh processing, and more research is definitely needed to investigate other approaches.

The local surface descriptors are a set of points with no connectivity. This permits one to deal with nonmanifold objects or with point sets. However, more work is required to investigate the extension of partial matching to points.

We have presented a number of applications based on partial matching of surfaces. However, these are applications we have already experimented with; there are other possible directions that we would like to pursue:

- Surface registration.* Partial matching of two pieces of surface is an essential step in the registration and stitching of two surfaces into a coherent larger surface. This may be applicable to range scanners and to the reconstruction of a complete shape from a number of scans. Figure 22 shows a synthetic example of four scans of the dragon model registered into one coherent shape using our method. Note that surface registration scenario usually does not require scale invariant matching, so a more simple indexing method can be considered to reduce the preprocessing time of geometric hashing.
- Surface reconstruction.* Another application is the reconstruction of a surface from a sparse and noisy samples. Figure 23 shows a piece of a CAD model reconstructed by tessellating the sampled region by copying matching parts from the given model itself.
- Surface style.* Two surfaces, which may be globally dissimilar, can share similar closely related geometric features. The style of a shape may be defined by the self-similarity of its geometric features. This can lead to a classification process based on common features.




Fig. 23. Reconstructing part of a CAD model from sparse samples by copying matching parts from the model itself.

—*Object correspondence.* An important extension of the alignment application is, as we have shown, the definition of cross-mesh partial correspondence. The correspondence of features can also consider the geometric location of the features to define a more global similarity measure. Correspondence, in general, can be applicable to numerous applications, for example, object-space morphing.

ACKNOWLEDGMENTS

We are grateful to Patrick Min for his helpful comments. The *Happy Buddha* and *Thai Statue* are courtesy of of Stanford University; the *Octopus* is courtesy of Mark Pauly.

REFERENCES

- ALLIEZ, P., COHEN-STEINER, D., DEVILLERS, O., LEVY, B., AND DESBRUN, M. 2003. Anisotropic polygonal remeshing. *ACM Trans. Graph.* 22, 3, 485–493.
- CHEN, D.-Y. AND OUHYOUNG, M. 2002. A 3D object retrieval system based on multi-resolution REEB graph. In *Proceedings of the Computer Graphics Workshop* (Taiwan). 16–20.
- CYBENKO, G., BHASIN, A., AND COHEN, K. D. 1997. Pattern recognition of 3D CAD objects. *Smart Eng. Syst. Des.* 1, 1–13.
- DOUROS, I. AND BUXTON, B. F. 2002. Three-dimensional surface curvature estimation using quadric surface patches. *Proceedings of Scanning 2002*.
- DRYDEN, I. AND MARDIA, K. 1998. *Statistical Shape Analysis*. John Wiley and Sons.
- ELAD, M., TAL, A., AND AR, S. 2001. Content based retrieval of VRML objects: An iterative and interactive approach. In *Proceedings of the Sixth Eurographics Workshop on Multimedia 2001*. Springer-Verlag New York, NY, 107–118.
- FROME, A., HUBER, D., KOLLURI, R., BULOW, T., AND MALIK, J. 2004. Recognizing objects in range data using regional point descriptors. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- FUNKHOUSER, T., KAZHDAN, M., SHILANE, P., MIN, P., KIEFER, W., TAL, A., RUSINKIEWICZ, S., AND DOBKIN, D. 2004. Modeling by example. *ACM Trans. Graph.* 23, 3, 652–663.
- FUNKHOUSER, T., MIN, P., KAZHDAN, M., CHEN, J., HALDERMAN, A., DOBKIN, D., AND JACOBS, D. 2003. A search engine for 3D models. *ACM Trans. Graph.* 22, 1, 83–105.
- HEBB, D. O. 1949. *The Organization of Behavior*. John Wiley, New York, NY.
- HILAGA, M., SHINAGAWA, Y., KOHMURA, T., AND KUNII, T. L. 2001. Topology matching for fully automatic similarity estimation of 3D shapes. In *Proceedings of SIGGRAPH 2001* (Los Angeles, CA). 203–212.
- HOFFMAN, D. D. AND SINGH, M. 1997. Saliency of visual parts. In *Cognition* 63, 1, 29–78.
- HUBER, D. AND HEBERT, M. 2003. Fully automatic registration of multiple 3D data sets. *Image Vis. Comput.* 21, 7 (July), 637–650.
- HUBER, D., KAPURIA, A., DONAMUKKALA, R. R., AND HEBERT, M. 2004. Parts-based 3-D object classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR 04)*.
- IYER, N., KALYANARAMAN, Y., LOU, K., JAYANTI, S., AND RAMANI, K. 2003. A reconfigurable 3D engineering shape search system part i: Shape representation. In *Proceedings of ASME DETC 03: Computers and Information in Engineering (CIE) Conference* (Chicago, IL).
- JOHNSON, A. E. AND HEBERT, M. 1999. Using spin images for efficient object recognition in cluttered 3D scenes. *IEEE Trans. Pattern Anal. Mach. Intell.* 21, 5, 433–449.

- KAZHDAN, M., FUNKHOUSER, T., AND RUSINKIEWICZ, S. 2003. Rotation invariant spherical harmonic representation of 3D shape descriptors. In *Proceedings of the Symposium on Geometry Processing*.
- KAZHDAN, M., FUNKHOUSER, T., AND RUSINKIEWICZ, S. 2004. Shape matching and anisotropy. *ACM Trans. Graph.* 23, 3, 623–629.
- LAMDAN, Y. AND WOLFSON, H. J. 1988. Geometric hashing: A general and efficient model-based recognition scheme. In *Proceedings of the International Conference on Computer Vision*. 238–249.
- MARR, D. 1982. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. W. H. Freeman, San Francisco, CA.
- MIN, P. 2004. A 3D model search engine. dissertation. Department of Computer Science, Princeton University, Princeton, NJ.
- OHBUCHI, R., MINAMITANI, T., AND TAKEI, T. 2003a. Shape-similarity search of 3D models by using enhanced shape functions. In *Proceedings of the Theory and Practice of Computer Graphics 2003*. IEEE Computer Society Press, Los Alamitos, CA, 97.
- OHBUCHI, R., NAKAZAWA, M., AND TAKEI, T. 2003b. Retrieving 3D shapes based on their appearance. In *Proceedings of the 5th ACM SIGMM International Workshop on Multimedia Information Retrieval*. ACM Press, New York, NY, 39–45.
- OHTAKE, Y., BELYAEV, A., AND SEIDEL, H.-P. 2004. Ridge-valley lines on meshes via implicit surface fitting. *ACM Trans. Graph.* 23, 3, 609–612.
- OSADA, R., FUNKHOUSER, T., CHAZELLE, B., AND DOBKIN, D. 2001. Matching 3D models with shape distributions. In *Proceedings of the International Conference on Shape Modeling & Applications*. IEEE Computer Society Press, Los Alamitos, CA, 154.
- PETTITJEAN, S. 2002. A survey of methods for recovering quadrics in triangle meshes. *ACM Comput. Surv.* 34, 2, 211–262.
- RUSINKIEWICZ, S. 2004. Estimating curvatures and their derivatives on triangle meshes. In *Proceedings of the Symposium on 3D Data Processing, Visualization, and Transmission*.
- SAUPE, D. AND VRANIC, D. V. 2001. 3D model retrieval with spherical harmonics and moments. In *Proceedings of the 23rd DAGM-Symposium on Pattern Recognition*. Springer-Verlag, Berlin, Germany, 392–397.
- SHILANE, P., MIN, P., KAZHDAN, M., AND FUNKHOUSER, T. 2004. The Princeton Shape Benchmark. In *Proceedings of the Conference on Shape Modeling International*. 167–178.
- SHUM, H.-Y., HEBERT, M., AND IKEUCHI, K. 1996. On 3D shape similarity. In *Proceedings of the 1996 Conference on Computer Vision and Pattern Recognition (CVPR '96)*. IEEE Computer Society Press, Los Alamitos, CA, 526.
- SUNDAR, H., SILVER, D., GAGVANI, N., AND DICKINSON, S. 2003. Skeleton based shape matching and retrieval. In *Proceedings of the Conference on Shape modeling International (Seoul, Korea)*. 130–139.
- TANGELDER, J. W. AND VELTKAMP, R. C. 2004. A survey of content based 3D shape retrieval methods. In *Proceedings of the International Conference on Shape Modeling and Applications*, 145–156.
- VELTKAMP, R. C. AND HAGEDOORN, M. 2001. State of the art in shape matching. In *Proceedings of the Conference on Principles of Visual Information Retrieval*. 87–119.
- VRANIC, D. V., SAUPE, D., AND RICHTER, J. 2001. Tools for 3D object retrieval: Karhunen-Loeve transform and spherical harmonics. In *Proceedings of the IEEE 2001 Workshop on Multimedia Signal Processing*. 293–298.
- WOLFSON, H. J. AND RIGOUTSOS, I. 1997. Geometric hashing: An overview. *IEEE Comput. Sci. Eng.* 4, 4, 10–21.
- ZAHARIA, T. AND PRETEUX, F. 2001. Three-dimensional shape-based retrieval within the MPEG-7 framework. In *Proceedings of the SPIE Conference on Nonlinear Image Processing and Pattern Analysis XII*. 133–145.

Received June 2005; revised October 2005; accepted December 2005