Retrieval of 3D Models using Partial Matching

Concise version of PhD Dissertation Proposal

Alfredo Ferreira
Department of Information Systems and Computer Science
INESC-ID/IST/Technical University of Lisbon
alfredo.ferreira@inesc-id.pt

January 2008

Abstract

The growth of 3D object collections and their dissemination into a wide range of application domains originated the necessity of effective retrieval solutions, similar to those that exist for textual information. However, despite the efforts of researchers during the last years, such solution was not found yet. Most of the existing retrieval solutions only support queries by the complete object. Those that provide some way to search for a model using a partial query are either limited to small databases or index just a small portion of each model, discarding the so-called irrelevant parts. Indeed, the efficiency and success of such solutions are far from those obtained by their textual counterparts

In this document we study the current state-of-the-art in retrieval of 3D shapes with partial matching and propose a research path that will focus on a novel approach to this topic. We think that transposing some ideas traditionally used in text information retrieval we will get closer to an effective solution for 3D shape retrieval. To that end and to overcome some problems faced by previous attempts to shape retrieval with partial matching, we suggest decomposing each object in the collection into its subparts considering for that matter all the other objects and then index these subparts using a shape thesaurus.

1 Introduction

The growing number of three dimensional objects in digital libraries led to a problematic situation. Searching and browsing collections of models is no longer a trivial task. Today, a regular domain-specific database can contain thousands of items, and the number of generic 3D models available on the internet is much larger. Indeed, unless some meta-data have been assigned to models, finding the desired model in large collections is an hard task.

To ease this task, researchers proposed, during the last decade, several approaches to retrieve 3D models based on shape similarity. Some of these content-based retrieval systems are able to find a model in a database from a sketched query or using query-by-example. However, results produced by such systems are far from the successful query results obtained by their textual counterparts.

A major handicap of most retrieval systems is the fact that they only support queries of the complete object. Even those who use local features to represent a model in a database usually do not allow matching of object subparts. Indeed existing approaches work mostly by comparing the complete models and do not allow partial queries to be formulated, which greatly hinders their usefulness. This is similar to a text-based system requiring detailed specifications of pages or complete documents to find a given document, in which only example documents would serve as query initiators instead of typing a few words to a search engine. This might explain why 3D model retrieval systems enjoy limited usefulness and there is no equivalent of a Google™ or Yahoo® search engine for three-dimensional geometric shapes.

Nevertheless, in recent years a few 3D shape retrieval approaches with partial matching has been proposed. These approaches allow searching for a model supplying as a query only part of the desired model . However, such solutions rely on representing some sub-parts of the model and not the complete model. Indeed, considering only a small set of distinctive features of an object to classify it proved to be an efficient shortcut, but some eventually relevant object information is discarded in this process. Informally, a 3D search engine using these approaches can be compared to a text retrieval system that uses automatically extracted keywords to classify document instead of the whole content.

In this thesis proposal we will focus on 3D shape retrieval with partial matching. Unlike other approaches, we will represent complete models, in a way similar to text retrieval systems that classify all words from the entire document and not only some select words. We believe that through the use of a context-aware shape decomposition and a shape thesaurus with inverted indexes we will be able to describe and retrieve 3D models partially. This approach will allow us to take advantage of some well known techniques from text information retrieval, such as the term frequency and inverse document frequency to rank the relevance of every subpart in the database. We hope that, at the end, our research contribute to devising an effective solution for 3D shape retrieval with partial matching.

We believe that through the combination of effective shape decomposition techniques with a shape thesaurus we will be able to describe and retrieve 3D models with partial queries. However, despite the success of word thesaurus in text documents, while words can be easily extracted from documents, shape identification in a three-dimensional object is a much harder task and the success of such approach is not guaranteed. Indeed, the difficulties of this task came not only from its computational complexity but also from the ambiguity of such identification.

Therefore, our research will focus on devising a technique for model segmentation suitable to be used in conjunction with a shape thesaurus. Moreover, we will investigate shape retrieval, indexing and matching methods that support the proposed thesaurus-based approach. Additionally, we will develop a framework for 3D shape classification and retrieval with partial matching based on the proposed approach.

In the remaining of this document we will present a brief review of the most recent related work in partial matching of 3D shapes and the we focus in the proposed research problem and present the thesis statement, defining the hypothesis. Finally, we will describe the approach we will follow in our research work and present a three-year work plan for the proposed PhD thesis.

2 Related Work

Despite the plentiful of work on 3D shape analysis and retrieval developed in last decade, most research focus on description and matching of complete shapes. These approaches are usually based on global or local features and sometimes a mix of both to improve results. But even when taking advantage of local features, most common 3D shape retrieval approaches do not support partial matching. In the following sections we will present some recent research work tackling the partial matching of three dimensional models.

2.1 Partial matching through shape decomposition

Suzuki *et al.* proposed [23, 24] a solution that follows a process commonly used for partial matching using the information of the entire object. The 3D model is initially decomposed into its subcomponents and then shape descriptors for these shapes are computed. The shape descriptor computation and matching techniques used for partial matching are identical to the techniques used in typical 3D shape retrieval.

There are a multiplicity of different ways to decompose a 3D object. Indeed, besides the impractical user-assisted 3D model decomposition, several automatic techniques have been proposed. These usually rely on object attributes such as color, texture or shape. Detailed explanations of these techniques can be found in several papers [4, 5, 27]. As a matter of fact, object decomposition is so relevant for our research, that further details will be presented in a future survey.

In their work, Suzuki *et al.* apply a simple and automatic decomposition technique. They decompose 3D models into several parts by comparing angles created by normal vectors of each polygonal face, and the technique finds sharp angles and cuts polygonal faces into parts based on a typical clustering approach. To tune the decomposition granularity is used a threshold for the angle size. A wide angle size produces a large number of shape parts while a sharp angle size produces a small number of components.

To compute the shape descriptors for extracted components, Suzuki *et al.* used a rotation invariant shape descriptors they proposed earlier for their similarity retrieval system [22]. In this method the object part is initially normalised for scale and then for orientation by using principal component analysis pose normalisation. Next it is voxelised and inserted into a cube divided in a three dimensional grid. The number of voxels contained in each cell are computed and then a clustering technique is applied. Finally, the descriptor are constructed from a voxel distribution function.

Authors acknowledge that, although their decomposition technique is fast and fully automatic, occasionally the algorithm can not efficiently handle highly complex 3D models. However, they suggest using a powerful algorithm from the several algorithms available to decompose 3D models to produce better object decomposition. Additionally, time complexity is also a problem of the proposed method, since the decomposition process is a time consuming task and shape matching requires a considerable amount of time due to the high number of shape descriptors for each model.

More recently, Suzuki *et al.* improved their decomposition method and partial shape descriptors construction algorithm to attain better similarity retrieval results [25]. One of the decomposition

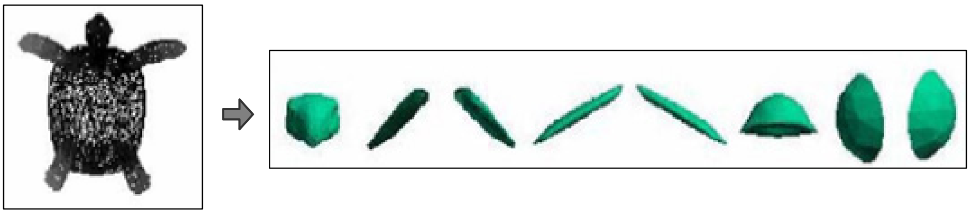

Figure 1: Example of the decomposition of a 3D model of a turtle (Figures taken from [25]) © 2006 IEEE.

enhancements was the use of the area proportion to identify irrelevant parts that should be merged into other. An example of model decomposition obtained with the enhanced algorithm is depicted in Figure 1. Other improvement in this approach was the use of multiple bounding boxes in descriptor computation. Authors use a bounding box for each decomposed part, instead of only one for the entire object used in their previous solution. However, despite for most models this approach proved better, the time complexity problems were not solved and when 3D models does not have visually irrelevant parts the previous technique works better.

2.2 Partial matching by structural descriptors

It is widely accepted that humans recognise and code mentally shapes in terms of relevant parts and their spatial configuration. Therefore, geometric features are insufficient to fully describe a three dimensional model for retrieval. It is necessary to combine geometric data with structural information.

Biasotti *et al.* described [2] an interesting method for partial shape matching that couples geometry and structure in a single descriptor. Based on the theory of Reeb graphs, as an alternative to commonly used skeletal graphs, authors compute the so-called structural descriptor. They suggest [15] encoding the shape and all its relevant sub-parts in a graph which represents the structure of the object and its geometry at the same time.

The proposed extended Reeb graph (ERG) [21] generalises the original Reeb graph definition to a surface on which a finite set of contour levels given by a mapping function f is defined. In their work, authors compare two distinct mapping functions, since choosing this function is an important aspect of the proposed method. One option is using the distance from the centre of mass of the object as a mapping function, which makes f rotation invariant, but sensitive to pose changes. The other option is estimating f as suggested by Hilaga *et al.* in [11], using the integral geodesic distance to the surface centre, which is also pose invariant. Biasotti *et al.* conclude that the latter is best suited for retrieving articulated objects disregarding its pose, while the first option distinguishes articulated models in different poses.

Using the selected mapping function, the ERG is constructed and represents the topology of the model. Then, the corresponding value of f and a geometric descriptor is assigned to each node of the graph, which represents a sub-part of the model. To compute the geometric descriptor assigned to each node, authors use spherical harmonic analysis of the corresponding sub-part. The

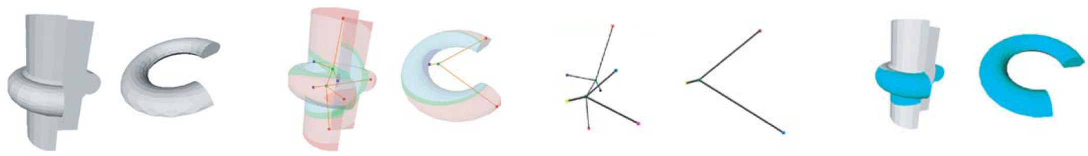

Figure 2: Sub-part correspondence of two models (Figures taken from [2] © 2006 Elsevier Ltd.).

rotation invariant spherical descriptor used in this approach has been defined by Kazhdan *et al.* in [13]. Additionally, each sub-part is uniformly scaled separately before computing the descriptor to guarantee that retrieval is scale invariant. Indeed, due to the necessity of finding similar sub-parts with different sizes, scale invariance is an important feature in retrieval with partial matching approaches.

Since the structural descriptor is coded as a directed attributed graph, the sub-part correspondence between models is obtained by matching its descriptors, *i.e.* matching its graphs. Using inexact graph matching, the authors adapted the algorithm proposed by Marini [16] for the computation of the maximum common sub-graph between two directed, acyclic graphs with attributes. The specialised version of this algorithm produces a set of all common sub-graphs between two extended Reeb graphs, considering not only the topological structure but also node attributes such as the geometric descriptor. The similarity estimation between models is obtained by considering the size of the common sub-graphs with respect to the size of the corresponding graphs and the similarity distance between the nodes belonging to the common sub-graphs.

An example of the above described technique is shown in Figure 2. To obtain partial matching between two models the ERG are extracted from each object and the structural descriptor are computed based on it. Then, a graph matching technique is applied to compare the structural descriptors, identifying the common sub-graphs. Finally, the similar subparts are identified in both objects by comparing the common sub-graphs.

2.3 Scale-space feature extraction and object decomposition for partial matching

Focusing on mechanical CAD models, Bespalov *et al.* [1] proposed a partial matching technique for finding similarities across part models constructed from data acquired in 3D scanners. For that end they propose a feature extraction technique based on recursive decomposition of polyhedral surfaces into patches which applies the method introduced by Novatnack *et al.* for extracting and integrating shape features in the discrete scale-space of a 3D mesh model [17]. The discrete scale-space of a three dimensional model is constructed by unwrapping the shape surface onto a planar domain, as a two dimensional image of surface normals. After this initial step, the scale-space operator used in image processing can be applied to the 3D shape.

However, the parametrization of original mesh to the planar domain that produces the surface unwrapping is not isometric, introducing distortion in the image. As a result of this distortion, relative geodesic distances between points on the original 3D model are not equivalent to relative

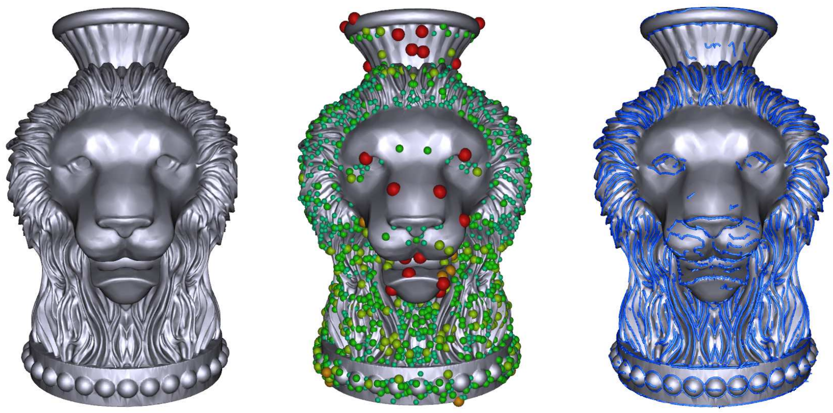

Figure 3: Combined set of scale-dependent corners and edges extracted from polygonal model (Figures taken from [17] © 2006 IEEE).

distances between corresponding points on the 2D normal map. Therefore, to correct this distortion, authors compute the distortion for each point in the 2D image and then construct a dense distortion map with these values. Then, this map is used to approximate the geodesic distances between two points in the two dimensional image representing the unwrapped model surface. Finally, the discrete scale-space of the original model is constructed from finer to coarse by iteratively convolving the normal map with a distortion adapted Gaussian kernels, as commonly done when computing the scale-space of a two dimensional image.

After the discrete scale-space of the model has been constructed, scale-dependent shape features can be extracted in a similar manner to image feature detection. To that end, a gradient of the normal map that correctly accounts for the distortion is defined. This gradient is then used to detect edges and corner of the original shape in the normal map. Since a 3D corner is a point with geometric changes in more than one direction, these points can be detected in the normal map by identifying large local changes in the normal directions. On the other hand, an edge in the 3D model corresponds to a line of points with significant changes in the surface geometry. Therefore, edges are detected by finding maxima along gradients previously computed. Indeed, to detect corners and edges authors suggest methodologies analogous to the Harris corner detection algorithm [10] and the Canny edge detector algorithm [3] respectively. Figure 3 depicts the results of scale-dependent corners and edges extraction from a 3D mesh model.

Once the features have been extracted at individual scales these are combined into a unified feature set which encodes the scale-dependent geometric structure of the shape, providing a concise representation of the original model. Authors argue that, with the appropriate parameters, the method can be tuned to extract local features of engineering relevance from CAD mechanical models. Thus, they adapted feature extraction in scale-space proposed by Novatnack [17] discussed above by replacing the geodesic distance function by a new distance function computed with respect to triangular faces of the model. This function measures the maximum angle between adjacent faces on the shortest path between two surface polygons.

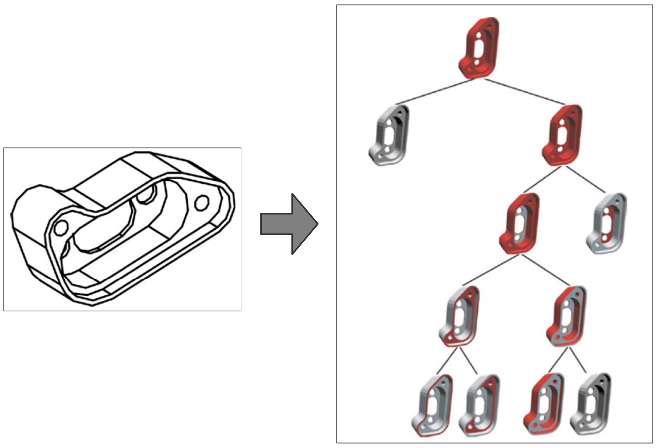

Figure 4: Scale-space decomposition of a model (Figures taken from [1] © 2006 Elsevier Ltd.).

In practice, the maximum angle function introduced by Bespalov *et al.* quantifies the smoothness of the surface, since smaller angles correspond to smoother surfaces. Using this function, CAD mechanical models are decomposed and the resulting combined feature set is used for partial matching of 3D models. Figure 4 illustrates a scale-space decomposition of a CAD model. In this example the presented tree are not full, since it will be hard to understand the results if the whole tree was depicted.

2.4 Salient geometric features for partial matching

In their approach to partial matching, Ran Gal and Daniel Cohen-Or [9] shown that a relatively small number of salient geometric features can describe a three-dimensional model with sufficient detail for various applications of content-based shape retrieval. Based on this idea they introduced the abstraction of salient geometric features and presented a method to extract these features from polygonal meshes.

The first step of this method is computing a sparse set of local surface descriptors across the surface and use these to measure similarity between regions of the model, even if they have dissimilar polygonal meshes. Then, these descriptors are clustered in order to locally describe a nontrivial region of the surface. Each one of these clusters form a compound higher-level descriptor that represent a salient geometric feature characterising a local partial shape. In this approach trivial regions of the model are considered irrelevant and discarded.

A major challenge facing the Gal and Cohen-Or was correctly identifying the salient features. To that end, they start by making a loose definition of salient geometric feature. In this definition, a salient geometric feature is a region of the object surface with a non-trivial shape. Based on this definition, they select regions with high curvature relative to their surroundings and high variance of curvature values as geometrically salient. Indeed, such option is grounded on previous work by Hoffman and Singh [12]. They have found that human vision defines boundaries along negative

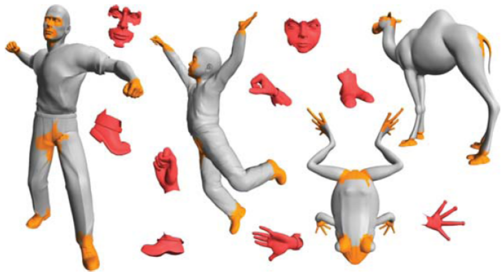

Figure 5: Salient geometric features from four models and corresponding individual sub-parts (Figure taken from [9] © 2006 ACM).

minima of the principal curvatures on surfaces. From this, Hoffman and Singh suggest that salience of a region depends on its size relative to the whole object, the degree to which it protrudes, and the strength of its boundaries.

Ran Gal and Daniel Cohen-Or identify salient regions by growing, for each descriptor from the sparse set, a cluster of descriptors. Such cluster is constructed by incrementally adding descriptors from its neighbourhood that maximise the saliency of the cluster until the contribution of neighbour cluster become insignificant. After estimating all clusters, authors select from these a set of clusters with higher values of saliency grade and use them to identify the set of salient geometric features of the model. This set should include model regions that are salient and interesting compared with other parts of the model. Figure 5 illustrates the result of applying this method to four different models and selecting as salient the top 10% cluster ordered according to saliency grade.

In this approach each model is represented by a set of descriptor clusters corresponding to the salient geometric features of the object. Ran Gal and Cohen-Or associate each one of these features with a vector index (a signature) and insert it in a geometric hash table[1]. Authors recognise that elaborate indices, such as normalised moments can be used to describe the geometric features. However, they simply use the terms employed for defining the saliency grade to construct the vector index, reinforcing their claims for the efficiency of salient features in shape retrieval.

2.4.1 Partial matching using distinctive regions

The researchers at the Princeton 3D shape retrieval group follows a slightly different path. Instead of identifying the salient regions of an object, as proposed by Ran Gal and Cohen-Or [9], Shilane and Funkhouser [18] suggest selecting the distinctive regions of a 3D surface. The basic idea behind their approach is to focus the shape matching process on local features of shapes that are consistent among objects of the same class and distinctive relative to object of other classes.

[1] Geometric hashing is an highly efficient technique with low polynomial complexity developed for matching geometric features against a database of such features [14]. This technique uses a grid-based hash table to store every feature of every object but only a limited number of features is used to determine a mapping into the hash. During a query, the remaining features are used when hash collisions exist. With this technique matching is possible even when the recognisable database objects have undergone transformations or when only partial information is present [26].

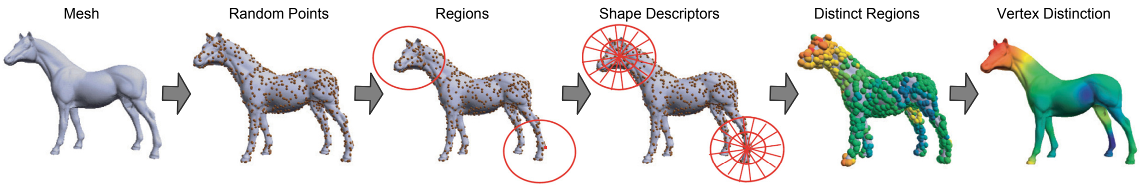

Figure 6: Selecting distinctive regions of an object (Figures taken from [19] © 2007 ACM).

Instead of using global descriptors, which represent global features of the model and fail when local properties of an object distinguishes it from others, in their approach authors use local shape descriptors. However, computing and storing local shape descriptors for the whole shape is time consuming and space expensive. To overcome this, they proposed a method for finding distinctive features of an object that are more relevant for shape retrieval.

In their method, Shilane and Funkhouser [19] define a distinctive region as a region with features that are only found on objects of a single class, while a not distinctive region is a region common to many objects of different classes. Therefore, in this approach to find the distinctive regions of an object the complete model database should be initially classified into object types. Otherwise it will not be possible to establish which are the objects of the same class. And such relationship is necessary to identify common features.

The distinctive region identification process starts by randomly sample each mesh on the database in order to obtain a set of spherical regions, covering the object at different scales. For every region, authors compute the corresponding shape descriptor that represents the distribution of surface area within that region. Next, by comparing all the descriptors of the database, they produce a ranked list of matches for each descriptor and use it to produce measures of region distinctiveness, thus identifying the most distinctive regions of each model.

Identifying distinctive regions is, therefore, a pipeline of relatively simple steps. Although other sampling methods could be used, authors propose selecting points randomly with uniform distribution with respect to surface area. Likewise, several shape descriptors can be used, but authors suggest describing the shape of every spherical region using rotation invariant spherical harmonics [13]. Figure 6 illustrates the different stages of the process of partitioning a model into distinctive regions with respect to a set of object classes in a given database. In the final result, regions in red are the most distinctive while regions in blue are least distinctive.

To perform partial matching retrieval on large model databases, Funkhouser and Shilane proposed a priority-driven search algorithm [8]. This kind of backtracking search algorithm considers only partial matches that can possibly lead to the lowest cost matching, as in the widely known shortest path algorithm by Dijkstra [6]. Therefore, authors use a cost function that accounts for both feature dissimilarity and geometric deformation to order the list of pairwise matches between features of query and of objects in database. The proposed algorithm produces a list of best target objects sorted by the similarity of a subset of matching features between the object and the query.

2.5 Discussion on approaches to partial matching

Existing techniques to partial 3D shape matching can be roughly divided into two distinct approaches. One uses only relevant features of the model while the other uses the entire object, usually decomposed into sub-parts. The research work developed by Suzuki *et al.* at the National Institute of Multimedia Education (NIME), in Japan, fits in this last category, as well as the work by Ran Gal and Daniel Cohen-Or or the approach followed in Princeton by Shilane and Funkhouser.

The main disadvantage of using only relevant features to classify models in a collection is that most features of that object are simply ignored. Worst, it is impossible to guarantee that the so-called relevant features are indeed the ones that best describe the model. Moreover, even a clearly relevant feature of a model might be shared by all objects in the collection, making it useless for shape retrieval. Additionally, it is not even trivial to define which features are relevant. Indeed, different researchers might identify different relevant regions in the same object.

On the other hand, current approaches to partial matching that does not discard any subpart of the object, *i.e.* index the whole object, such as the solution proposed by the Shape Modelling Group lead by Bianca Falcidieno in the CNR IMATI-Ge [2], faces the problem of time and space complexity. While this approach work well for small collections, since retrieval is based on pairwise matching, it is not feasible to be used in large datasets. Indeed, besides the time necessary to make all the comparisons, there is also the necessity of large storage space to keep the shape descriptors for all shapes and all its subparts, especially if one considers that each model can be decomposed into hundreds of subparts.

3 Thesis Statement

Now that we have briefly analysed the current research work on 3D model retrieval with partial matching, we can focus on the topic we plan to explore. As we have seen in previous section, to accomplish partial matching it is necessary to identify and isolate subparts in models before classifying and measuring similarity.

Therefore, devising a retrieval solution with partial matching faces two major challenges, besides the ones shared with the global matching approaches. The first is the correct and efficient decomposition of models into its subparts, identifying the relevant or, better, all of them. The second is to devise an effective way to index the extracted information that allows fast and accurate search. In the present research work we will dedicate a special attention to these two problems.

3.1 Problem statement

A currently open challenge in this area is how to perform retrieval with partial matching on large collections of three dimensional models using all of its features. As discussed above, although

[2]The Shape Modelling Group is a research team of the Institute of Applied Mathematics and Information Technology, branch of Genova (IMATI-Ge), of the Italian National Council of Research (CNR).

some techniques seems practical for indexing large models, and even large collections of complex models, these only consider a small set of relevant local features of each object. Such approaches do not fulfil the main goal of the research on shape retrieval with partial matching: the ability to find models with different global shape properties but having just some characteristics in common, which might not even be the most relevant geometric features. In the proposed research work we intend to achieve this goal.

To that end, we plan to follow a slightly different approach from those proposed by other researchers and described in Section 2. Like them, we aim to provide a solution that will allow successful searches on three-dimensional databases with partial queries. However, we will consider not only relevant parts of models, but the whole set of parts that compose objects. Thus, we need to overcome one major problem. The time and space complexity associated to indexing all components of each model, even the irrelevant ones. This is even worst when considering a multi-scale approach to shape decomposition. Nevertheless, indexing contents of large collections have been already addressed with success in text information retrieval. We suggest adapting techniques from this area to 3D shape retrieval.

3.2 Research hypothesis

In the proposed research work we intend to research the viability of transposing the matching and indexing approaches widely-used in text information retrieval. It is well known that these approaches produce successful practical results, such as the obtained within Google^{TM} search engine. Therefore, we propose using a shape thesaurus for model classification and indexing, similarly to what happens with words in text documents. Unfortunately, while words can be easily extracted from documents, shape subpart identification in a three-dimensional object is a much harder task. The difficulties of this task came not only from its computational complexity but also from the ambiguity of such decomposition.

We believe that it is possible to achieve a practical solution, overcoming the major challenges of computational complexity and ambiguous shape decomposition. Moreover, it will be necessary to devise effective methods for shape matching and indexing suited to a thesaurus-based approach. In our work we plan to create the above referred shape thesaurus and then transpose well known inverted indexing techniques used in text data to support retrieval of 3D shapes.

Having identified the major challenges we plan to address, we can now formulate the following research hypothesis:

> **Partial 3D object matching can be achieved by decomposing a model into a set of subshapes that describe the whole model, combined with a shape thesaurus for indexing.**

This hypothesis summarises our main ideas and clearly indicates our final objective: to contribute to the development of a 3D shape retrieval solution that supports partial matching. But such goal should be decomposed in a set of research objectives presented in the next section.

3.3 Research objectives

Following the hypothesis stated above, we already identified the overall objective of our work. However, we can decompose it in a set three major research goals:

1. Devise a solution to decompose a model according to the context were it lies, *i.e.* other objects in the collection;

2. Identify methods for shape matching and indexing that supports a thesaurus-based approach on 3D shape retrieval;

3. Develop a thesaurus-based framework for 3D shape retrieval with partial matching combining the methodologies referred above.

To achieve these objectives we have already defined an overall plan identifying the path we must pursue. Hence, in the next section we will describe the approach we intend to follow.

4 Proposed Approach

At the current stage of our research we are not yet able to clearly determine the methods and techniques we will use during the remaining of our research. However, we have already selected the shape decomposition approach we plan to use in our work. We intend to extend existing graph-based segmentation techniques to not only construct the model topology, but also identify the major shape components. Then, we must devise a suitable composition of several shape descriptors to extract the geometrical information from the previously identified shape components.

The success of the proposed work depends in great extents from two major issues: the accuracy of the topology extraction and segmentation process and the robustness and efficiency of shape descriptors. Consequently, our efforts in the initial stages of this research work were focused on studying and comparing existing alternatives for shape description. In a short term, we plan to develop an appropriate feature extraction algorithms to be used for shape description in our approach. Indeed, besides simply combine and extend existing algorithms into our solution we might found necessary to develop completely new methods from scratch. Additionally to obvious shape matching, the efficiency of these methods will be of great importance in our approach to shape decomposition.

4.1 Context-aware 3D shape decomposition

During the last years, several approaches to 3D shape decomposition have been published, with recognised success in some domain-specific models, such as articulated characters. However, independent of the methodology used, these approaches only consider the model that should be decomposed, ignoring the context were it lies. Although such decomposition techniques could be very effective, they does not completely fulfil our needs to decompose models in a collection in order to classify them using a shape thesaurus.

In this thesis work we will investigate a completely novel approach to 3D shape decomposition. In this approach, during object segmentation we will consider not only the object that is being decomposed, but also other objects in the collection. Such context-aware 3D shape decomposition will start by identifying distinctive regions of objects in the collection and then perform a multi-level shape decomposition in each model based on those distinctive regions.

4.2 Shape thesaurus

After devising an effective methodology to decompose objects in the collection into their subparts as described above, these subparts will be clustered according to their geometric properties. Based on the results of this clustering we will construct a thesaurus of 3D shapes. Such thesaurus includes shapes extracted from objects on the collection, *i.e.* its subparts, besides the whole objects.

The so-constructed shape thesaurus can then be used to index the objects in the collection. On the proposed approach, we plan to directly transpose the ideas and techniques used in text information retrieval. Therefore, we will use an inverted index which maps, for each shape in the thesaurus, objects that are or contain a subpart geometrically similar to that shape.

From this shape thesaurus and corresponding inverted index, we believe we will be able to efficiently retrieve models in a collection based on a partial query. Indeed, in this approach we reduce the necessity of pairwise matching between the query and all subparts of all models to a simpler and faster matching between the query and the shapes in the thesaurus. Moreover, to improve this matching we plan take advantage of the NB-Tree developed by Fonseca in his PhD thesis [7].

4.3 Framework for 3D shape retrieval with partial matching

In order to integrate the above referred techniques according to the proposed approach, we must develop a framework for 3D shape retrieval with partial matching. Such framework should, first of all, provide the functionality necessary to process the 3D models and the query, which will be a 3D shape. Then, it should implement shape classification, indexing and matching using the techniques we develop in this research work, together with methods widely used in traditional information retrieval.

Indeed, despite the importance of this framework to implement and evaluate the proposed solution to shape retrieval, it is not the core of our research work. Instead, it can be simply seen as the "skeleton" were the different methods and techniques we propose (which constitute the core of our work) will be put together. Basically, such framework will provide us a way to achieve our overall goal, *i.e.* prove that a a context aware decomposition technique in conjunction with a shape thesaurus can efficiently retrieve 3D models from partial queries.

4.4 Research focus

From medicine to forensics, from archeology to biology, from architecture to mechanical CAD, there are a broad range of application for 3D model retrieval. The ultimate goal of any research in this topic should be develop an approach that will provide a generic solution, effective in all these areas. However, such goal is not feasible. Indeed, from all research groups working on 3D retrieval none express such intention.

As a matter of fact, most published results in 3D model retrieval were achieved through research work focused in specific domains. For instance, a successful approach developed for retrieval on protein databases might not even function on engineering databases. Therefore, in the proposed research work, we will focus on a single type of three-dimensional models. Nonetheless, we intend to achieve a solution as generic as possible, within the selected data type. Thus, we will exclude from the proposed study some shape collections, such as medical and chemical databases.

In a first stage we will use existing data sets of 3D models of everyday objects to help us devising a good solution. However, to achieve effective practical results, we plan to constrain our research to CAD model collections, which will allow us to fine-tune our algorithms, taking advantage of some particularities of this type of models. More precisely, CAD models of mould parts due to our partnerships with the mould industry, described with some detail in Section 5.3.

On the other hand, a major challenge to overcome on 3D shape retrieval is the definition of an effective interface that will allow users to define queries. The most simple approach rely only in query-by-example, where users submit an existing model as a query. In this particular case, devising the interface raises no relevant issues. However, such solution does not fulfil the real users needs. Indeed, there are a myriad of different approaches to 3D query specification interfaces, from more or less traditional modeling tools to image-based queries. Anyhow, this topic are not within the scope of the proposed research work. Thus, initially we will employ a query-by-example interface and later we might integrate other solutions, such as sketch-based queries.

4.5 The 3D test data set

A subset of Princeton Shape database is now widely use for evaluating shape-based retrieval and analysis algorithms. The Princeton Shape Benchmark [20] provides a repository of 3D models to promote the use of standardized data sets and evaluation methods for research in matching, classification, clustering, and recognition of 3D models. Therefore, we plan to use mainly this collection for evaluation of algorithms and techniques developed during our research work.

However, it is expected that during the evolution of our research, novel 3D model databases appear and become accepted by the research community. In that case, we will adapt our work in order to take into account the trends in the 3D shape retrieval area. For instance, the SHREC 2007 involved multiple tracks, one of these is exactly partial matching, the primary goal of our work. If this model remains in the future, it is quite obvious that we will use extensively the datasets recommended in the partial matching task on SHREC.

Additionally, we plan to use a collection of mechanical CAD models of manufactured parts

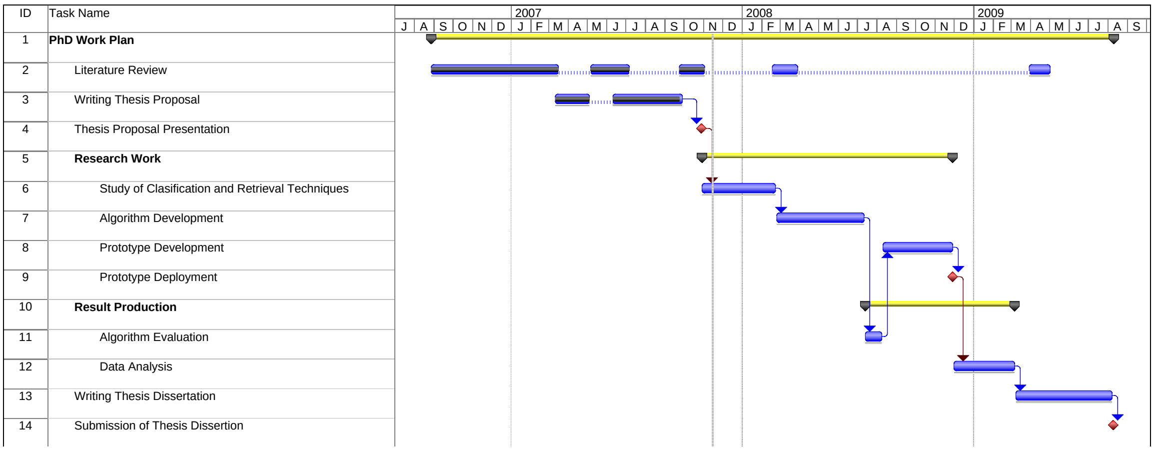

Figure 7: Gantt chart representing PhD work plan.

supplied by the mould industry. With such collection we hope to fine-tune our method for a particular domain-specific application. Moreover, implementing the proposed methods in the industry will provide us a good opportunity to prove the efficiency of our approach among real users outside laboratory.

5 Work Plan

In this section we will present the research work plan, based on a three-year doctoral program. Indeed, it started last year and some work has already been done while other is currently ongoing, as depicted in the Gantt chart at Figure 7. Thus, we expect to finish this doctoral program on August 2009, by submitting the preliminary version of the PhD thesis dissertation.

5.1 First year

In the early stages of this work we made a brief literature review on multimedia information retrieval, focusing mostly on three dimensional shape analysis and retrieval. The main goal of this study was to identify the main research topic and corresponding open issues that we must address in this PhD research. As a result of this preliminary review we determined the topic of the present thesis. We will aim on devising a novel approach to 3D shape retrieval with partial matching. Then, after establishing the scope of our research, we restarted the bibliographic research. In this second part of the literature review, we aimed on finding related work on areas of relevance to this investigation and identifying within each area who are the key authors, who are the most prominent research groups and what are the main approaches followed by them with success in recent years.

From the initial research we concluded that computation of shape descriptors is a key topic in any approach to 3D shape retrieval. Therefore, we are currently developing preliminary prototypes which implement some existing techniques on 3D shape descriptor computation. These prototypes receive as input models from the most relevant shape databases, identified during literature review, and produce several distinct feature vectors. Furthermore, the prototypes will provide the most common similarity measures methodologies.

Information produced by these preliminary prototypes will allow us to collect and study comparative data about existing approaches to help identifying techniques that best suits the planned research. As a result of the work developed during the first year we will produce technical reports on the state-of-the-art and we intend to submit a survey of the state-of-the-art on 3D shape description to an international journal.

5.2 Second year

Based on the conclusions from the work developed during the first year, we will start the second year by performing a theoretical and experimental study on selected classification and retrieval techniques. In this study we plan to broad the focus from shape description techniques to their integration with segmentation and matching mechanisms. Then we will implement algorithms for context-aware 3D shape decomposition and use them with thesaurus-based indexing techniques in a 3D shape retrieval solution with support to partial matching.

Since we will be working on a novel approach, the techniques we propose must be extensively experimented with the most relevant 3D shape databases, using pier recognised benchmarks for that purpose. In this stage we will perform a formal evaluation of our algorithms and check their efficiency, which will eventually led to discarding some less successful methods and improving others.

5.3 Third year

At the end of the second year we should have finally achieved a feasible partial matching solution for 3D shape retrieval. Thus, in the beginning of the third year, we plan to develop a fully functional prototype for classification and retrieval of 3D shapes with partial matching. Such prototype shall incorporate the feature extraction, indexing and matching techniques devised during this work. Moreover, in order to attain a practical application for this prototype we will focus our efforts on a specific class of 3D models: CAD technical drawings.

The work developed until this date will be published in technical reports and submitted to international conferences and journals. Roughly, two different groups of publications are expected to be produced. In one group we will present the results obtained during the study on classification and retrieval techniques. On the other group of publications we will present our approach to classification and retrieval of 3D shapes. Namely, we will focus on presenting our shape descriptors, classification, indexing and matching techniques, as well as the whole solution.

The second half of this year be dedicated to analyse and reflect on the work developed during this doctoral program. Furthermore, we shall identify unsolved issues while revealing paths for future research. Finally, we will compile, organise and synthesise all the information, producing, as an outcome of all the work described above, a thesis dissertation.

References

[1] Dmitriy Bespalov, William C. Regli, and Ali Shokoufandeh. Local feature extraction and matching partial objects. *Computer-Aided Design*, 38(9):1020–2037, September 2006.

[2] Silvia Biasotti, Simone Marini, Michela Spagnuolo, and Bianca Falcidieno. Sub-part correspondence by structural descriptors of 3d shapes. *Computer-Aided Design*, 38(9):1002–1019, 2006.

[3] J. F. Canny. A computational approach to edge detection. *IEEE Trans. Pattern Anal. Mach. Intell.*, 8(6):679–698, 1986.

[4] Bernard Chazelle, David P. Dobkin, Nadia Shouraboura, and Ayellet Tal. Strategies for polyhedral surface decomposition: an experimental study. In *Proceedings of the eleventh annual symposium on Computational geometry (SCG '95)*, pages 297–305, New York, NY, USA, 1995. ACM Press.

[5] Bernard Chazelle and Leonidas Palios. Decomposition algorithms in geometry. *Algebraic Geometry and its Applications*, pages 419–447, 1994.

[6] Edsger W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1(1):269–271, 1959.

[7] Manuel J. Fonseca. *Sketch-Based Retrieval in Large Sets of Drawings*. PhD thesis, Instituto Superior Técnico / Universidade Técnica de Lisboa, 07 2004.

[8] Thomas Funkhouser and Philip Shilane. Partial matching of 3D shapes with priority-driven search. In *Symposium on Geometry Processing*. Eurographics, June 2006.

[9] Ran Gal and Daniel Cohen-Or. Salient geometric features for partial shape matching and similarity. *ACM Transactions on Graphics*, 25(1):130–150, 2006.

[10] C. Harris and M.J.Stephens. A combined corner and edge detector. In *Fourth Alvey Vision Conference*, pages 147–152, 1988.

[11] Masaki Hilaga, Yoshihisa Shinagawa, Taku Kohmura, and Tosiyasu L. Kunii. Topology matching for fully automatic similarity estimation of 3d shapes. In *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 203–212, New York, NY, USA, 2001. ACM Press.

[12] Donald D. Hoffman and Manish Singh. Salience of visual parts. *Cognition*, 63(1):29–78, April 1997.

[13] Michael Kazhdan, Thomas Funkhouser, and Szymon Rusinkiewicz. Rotation invariant spherical harmonic representation of 3d shape descriptors. In Leif Kobbelt, Peter Schröder, and Hugues Hoppe, editors, *Proceedings of the 2003 Eurographics/ACM SIGGRAPH symposium on Geometry processing*, pages 156–164, Aire-la-Ville, Switzerland, 2003. Eurographics Association.

[14] Yehezkel Lamdan and Haim J. Wolfson. Geometric hashing: A general and efficient model-based recognition scheme. In *Proceeedings of the Second International Conference on Computer Vision*, pages 238–249, December 1988.

[15] Simone Marini, Biasotti Silvia, and Falcidieno Bianca. Partial matching by structural descriptors. In Tim Crawford and Remco C. Veltkamp, editors, *Content-Based Retrieval*, number 06171 in Dagstuhl Seminar Proceedings. Internationales Begegnungs- und Forschungszentrum fuer Informatik (IBFI), Schloss Dagstuhl, Germany, 2006.

[16] Simone Marini, Michela Spagnuolo, and Bianca Falcidieno. From exact to approximate maximum common subgraph. In Luc Brun and Mario Vento, editors, *Graph-Based Representations in Pattern Recognition*, volume 3434 of *Lecture Notes in Computer Science*, pages 263–272. Springer, 2005.

[17] John Novatnack, Ko Nishino, and Ali Shokoufandeh. Extracting 3d shape features in discrete scale-space. In *Proceeding of the Third International Symposium on 3D Data Processing, Visualization, and Transmission*, pages 946–953, Los Alamitos, CA, USA, 2006. IEEE Computer Society.

[18] Philip Shilane and Thomas Funkhouser. Selecting distinctive 3d shape descriptors for similarity retrieval. In *Proceedings of the IEEE International Conference on Shape Modeling and Applications 2006 (SMI'06)*, page 18, Washington, DC, USA, 2006. IEEE Computer Society.

[19] Philip Shilane and Thomas Funkhouser. Distinctive regions of 3d surfaces. *ACM Transactions on Graphics*, 26(2):7, 2007.

[20] Philip Shilane, Patrick Min, Michael Kazhdan, and Thomas Funkhouser. The princeton shape benchmark. In *Shape Modeling International*, June 2004.

[21] Michela Spagnuolo, Silvia Biasotti, Bianca Falcidieno, and Simone Marini. Structural descriptors for 3d shapes. In Tim Crawford and Remco C. Veltkamp, editors, *Content-Based Retrieval*, number 06171 in Dagstuhl Seminar Proceedings. Internationales Begegnungs- und Forschungszentrum fuer Informatik (IBFI), Schloss Dagstuhl, Germany, 2006.

[22] Motofumi T. Suzuki, Toshikazu Kato, and N. Otsu. A similarity retrieval of 3d polygonal models using rotation invariant shape descriptors. In *Proceedings of IEEE International Conference on Systems, Man and Cybernetics*, volume 4, pages 2946–2952, Nashville, TN, USA, 2000. IEEE Computer Society.

[23] Motofumi T. Suzuki, Yoshitomo Yaginuma, and Yasutaka Shimizu. A partial shape matching technique for 3d model retrieval systems. In *ACM SIGGRAPH 2005 Posters*, page 128, New York, NY, USA, 2005. ACM Press.

[24] Motofumi T. Suzuki, Yoshitomo Yaginuma, Tsuneo Yamada, and Yasutaka Shimizu. A partial shape matching method for 3d model databases. In *Proceedings of the Ninth IASTED International Conference on Software Engineering and Applications (SEA2005)*, pages 389–394, Phoenix, USA, November 2005. ACTA Press.

[25] Motofumi T. Suzuki, Yoshitomo Yaginuma, Tsuneo Yamada, and Yasutaka Shimizu. A 3d model retrieval based on combinations of partial shape descriptors. In *Proceedings of IEEE North American Fuzzy Information Processing Society Annual Conference (NAFIPS 2006)*, pages 273–278, Montreal, Canada, June 2006. IEEE Computer Society.

[26] Haim J. Wolfson and Isidore Rigoutsos. Geometric hashing: An overview. *IEEE Comput. Sci. Eng.*, 4(4):10–21, 1997.

[27] Emanuel Zuckerberger, Ayellet Tal, and Shymon Shlafman. Polyhedral surface decomposition with applications. *Computers & Graphics*, 26(5):733–743, 2002.