Universidade Técnica de Lisboa
Instituto Superior Técnico
Departamento de Engenharia Informática

PhD Dissertation Proposal

Retrieval of 3D Models using Partial Matching

by

Alfredo Manuel dos Santos Ferreira Júnior

Lisboa, October 2007

Contents

List of Figures

List of Tables

Abstract

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. In recent years, some researchers focused their investigation on 3D shape retrieval with partial matching. However, they rely on representing only some parts of the model, such as salient regions or distinctive features, and not the complete set of model subparts.

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 deterministic 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.

Keywords: Content-based retrieval, 3D shapes, partial matching.

1 Introduction

In the last decades the volume of multimedia information, such as images and video, stored into databases and over the internet has been growing. In particular, recent advances on modelling, digitising and visualising techniques led to a clear tendency to increase the number of 3D models both on the internet and in domain-specific databases. Three-dimensional models are used in a wide range of fields, such as engineering (CAD/CAM), virtual reality, medicine (CAT scans), molecular biology, geography (SIG) or not to mention the growing entertainment industry.

As a result, many collections of 3D models are now available for usage on a wide range of areas. Following the evolution of such collections, a considerable effort has been dedicated by researchers all over the world to analysis, classification and retrieval of three-dimensional models. Indeed, the problem of comparing 3D objects arises not only in one specific area or type of model but is a general issue, embracing a wide range of applications. For instance, besides the obvious applications in CAD model retrieval, similarity detection in organ deformation can help diagnosing diseases [72] and structural classification, a basic task in molecular biology, can be approached by 3D similarity search [9].

Thus, there is a growing necessity in classifying and retrieving 3D geometrical models using shape properties instead of (or combined with) text. To that end, several research groups focused their work on classification and retrieval algorithms that rely mainly on shape descriptors (feature vectors). Indeed, finding adequate descriptors that capture global or local shape characteristics while allowing computational viability of the representation has become a main investigation goal in this area. As a result of this research, there are several emerging solutions to search 3D objects based on their geometrical content.

In this context, most current techniques to describe 3D models are based on computable geometrical attributes of the shape. However, these techniques do not use the structure and details of the desired objects to index. Moreover, 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 explains 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 subparts 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.

Aiming an effective solution for 3D shape search, our investigation work will focus on object retrieval mechanisms to analyse and describe the content of models based on techniques for structure-driven partial matching of 3D models. Indeed, we intend to develop techniques that will provide some of the functionality necessary to build a retrieval system with effective support to partial matching but classifying complete models. To that end, we plan to investigate the viability of transposing matching and indexing approaches widely-used in text information retrieval to 3D shape searching.

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.

On this proposal we will firstly present a brief overview of research context underneath the classification and retrieval of three dimensional shapes. To that end we identify the key players in this area and the most relevant events regarding our work. Additionally, we present a short list of existing model databases.

Next, we will focus on the theoretical background and state of the art in 3D shape retrieval by analysing three distinct research areas: shape descriptors; content-based retrieval of 3D models and retrieval of 3D shapes using partial queries. Since shape descriptors are generally used to represent objects in retrieval systems, as referred above, algorithms and techniques to extract such information from models are considered to be among the most important topics in shape retrieval research. Thus, we give special attention to shape descriptor techniques by providing an extensive overview to existing approaches.

In the remaining of this document we will focus in the proposed research problem and present the thesis statement, defining the hypothesis. Next, we will describe the approach we will follow in our research work and present a three-year work plan for the proposed PhD thesis. Finally, we will present some preliminary results obtained in early stages of this work.

2 Background and State-of-the-Art

In this section we intend to introduce the reader into the current context of research on three dimensional shape classification retrieval. To that end, we will start by presenting some people whose work is closely related or has a major relevance to our research. Additionally, we will identify a small set of journals and conferences where our research topics are usually presented and published.

Besides presenting the current research context we focus on important topics in 3D shape retrieval. Regarding our work, these are the shape description, query and matching techniques and benchmarking methodologies. Due to the unquestionable importance of descriptor computation techniques on content-based retrieval system, we dedicate special attention to present later in this section a short but comprehensive survey on 3D shape descriptors. We will also briefly review some existing solutions for 3D shape retrieval. Obviously, this section could not be complete without presenting the state-of-the-art on partial matching of 3D shapes, with which we close this section.

2.1 The key players

Following the increasing importance of content-based retrieval of three-dimensional shapes, several research groups around the world focused their interests on this area, which we will present here. However, producing an exhaustive list with all of them is difficult, with the additional risk of excluding some important researchers. Therefore, we will only present a short list containing those we consider more relevant regarding our work.

The concise list of research groups working on analysis and retrieval of three-dimensional shapes, and corresponding researchers is presented in Table 1. As we mentioned before, the list of research groups referred in this section is far from exhaustive. Indeed, many relevant researchers were not included here since it will result in a quite extensive list of every people of every group working on analysis and retrieval of three-dimensional shapes. Thus, we will only refer the groups whose work had greater impact in our research until the present time and the most prominent researchers of these groups.

Princeton During the last six years, the Princeton Shape Retrieval and Analysis Group, leaded by Thomas Funkhouser, have been addressing key issues in shape-based retrieval and analysis of 3D models. Focusing on effective shape representations and query interfaces, they have developed a search engine for 3D polygonal models and later released a classified set of 3D models that can be used by researchers in this area. The Princeton Shape Benchmark is now widely accepted as a major benchmarking tool within 3D shape retrieval and classification.

Research Group	**Institution**	**People**
Shape Retrieval and Analysis Group	Princeton University (USA)	Thomas Funkhouser David Dobkin Adam Finkelstein Szymon Rusinkiewicz Philip Shilane Josh Podolak
3D Model Similarity Search Project	Konstanz University (Germany)	Daniel A. Keim Dietmar Saupe Benjamin Bustos Tobias Schreck Dejan Vranić
PRECISE	Purdue University (USA)	Karthik Ramani Suyu Hou Y. Kalyanaraman
Shape Modelling Group	CNR IMATI-Ge (Italy)	Bianca Falcidieno Michela Spagnuolo Silvia Biasotti Simone Marini Francesco Robbiano
Multimedia and Geometry Group	Utrecht University (The Netherlands)	Remco Veltkamp Frank Ter Haar Reinier Van Leuken
FOX-MIIRE group	University of Lille (France)	Mohamed Daoudi Jean-Philippe Vandeborre Tarik Filali-Ansary
Media Integration and Communication Center	University of Florence (Italy)	Alberto del Bimbo Pietro Pala Jürgen Assfalg
Image and Video Processing Group	Boğaziçi University (Turkey)	Bülent Sankur Yücel Yemez (Koç University) Ceyhun Burak Akgül Helin Dutağaci

Table 1: Some key players on 3D shape retrieval research

Konstanz The Multimedia Signal Processing Group chaired by Dietmar Saupe and the Databases, Mining and Visualization Group leaded by Daniel Keim, both from Konstanz University, join their efforts in a research project focusing the retrieval of three-dimensional shapes. The 3D Model Similarity Search project is part of the strategic research initiative on Distributed Processing and Delivery of Digital Documents and aims at effective content based model retrieval and efficient indexing and accessing methods.

PRECISE In Purdue University, Karthik Ramani leads the Purdue Research and Education Center for Information Sciences in Engineering (PRECISE). Their research lies at the intersection of design, shape analysis, and information sciences. They focus on developing representations for two and three-dimensional shapes for engineering and proteomics. In the last few years, they developed interfaces for querying, interacting, orienteering, and navigating intelligently in 3D shape databases. Their 3D Sketch-Based System for Conceptual Design, also known as ShapeLab, aims on empower computer-aided design users to retrieve, modify and reuse 3D models through freehand sketches.

CNR IMATI-Ge The people at the Shape Modelling Group, a research team of the Institute of Applied Mathematics and Information Technology at Genova headed by Bianca Falcidieno, have been working on geometric modelling for several years. Their main research goal is to describe the shape of an object through the definition of geometric primitive entities and the classification of the reference context. To that end, they have been working in a variety of topics, ranging from graph comparison to free-form modelling, producing some interesting work on shape retrieval. Moreover, Bianca Falcidieno is coordinator of the AIM@SHAPE network of excellence.

AIM@SHAPE The Advanced and Innovative Models And Tools for the development of Semantic-based systems for Handling, Acquiring, and Processing knowledge Embedded in multidimensional digital objects (AIM@SHAPE) is a sixth framework program project that fosters the development of new methodologies for modelling and processing the knowledge related to digital shapes. This project embraces a multi-disciplinary field, which integrates Computer Graphics and Vision with Knowledge Technologies and builds on using knowledge formalisation mechanisms for linking semantics to shape or shape parts.

Utrecht Also participating on AIM@SHAPE, Multimedia and Geometry group headed by Remco Veltkamp at the Center for Geometry, Imaging and Virtual Environments of Utrecht University has been working on multimedia information retrieval. Along with their research in areas such as music or image retrieval, they have developing some interesting work on 3D shape analysis and retrieval, namely on 3D facial models.

FOX-MIIRE In the Multimedia, Images, Indexation and Recognition (FOX-MIIRE) research group at the University of Sciences and Technologies of Lille, Jean-Philippe Vandeborre, Mohamed Daoudi and their teams are working on three-dimensional model indexing and topological analysis. The recently published FOX-MIIRE 3D-Models Search Engine based on adaptive views clustering algorithm is the first search engine that accepts 3D-Models retrieval from photos [46]. Additionally, they developed a 3D retrieval application for mobile devices, which were presented in the 2007 ACM International Conference on Image and Video Retrieval. The FOX-MIIRE group is a partner of DELOS Network of Excellence.

DELOS Partially funded by the European Commission in the frame of the Information Society Technologies Programme, DELOS is a Network of Excellence on Digital Libraries. The main goal behind DELOS is to provide global access to knowledge contained in the digital collections created by organisations and individuals around the world. To that end, DELOS is conducting a joint program of activities aimed at developing the next generation of Digital Library technologies, based on sound comprehensive theories and frameworks for the life-cycle of Digital Library information. Within the large number of research topics covered by DELOS, the work leaded by Alberto del Bimbo from the University of Florence is of major relevance for our research.

Firenze At the Media Integration and Communication Center of the University of Florence, researchers leaded by Alberto del Bimbo presented a prototype system for content based retrieval of 3D objects, briefly described in Section 2.8. Besides, they have been developing relevant research on curvature maps and spin images as descriptors for 3D shape retrieval.

Boğaziçi At the Image and Video Processing Group of Boğaziçi University, Bülent Sankur's team is developing, together with people from the multimedia, vision and graphics laboratory at Koç University relevant research on 3D shape retrieval. In the last two years, they focused their research on analysis and description of 3D shapes and on 3D face recognition, achieving remarking results.

2.2 The most relevant events and journals

In order to disseminate the results of our research and to allow an easier search of related work, we identified the most relevant events and journals that focus the 3D shape analysis and retrieval area. When gathering this information we decided not to include some major computer graphics publications and events, such as ACM Transaction on Graphics, since their importance is obvious. Instead, we list a short number of selected events and publications with high visibility where researchers on our field have been sharing their work.

SMI The IEEE International Conference on Shape Modeling and Applications, also known as Shape Modelling International (SMI), was launched in 1997 with the goal to join a multi-disciplinary community concerned with computation techniques for modelling and processing digital representations of shapes and their properties. In 2001, SMI has merged with the Implict Surface Workshop and is now run as an annual event alternating between Asia, Europe and America. Today, SMI addresses all aspects of shape acquisition, processing, retrieval and understanding. Since 2006, it includes a related event of great importance for our work, the 3D Shape Retrieval Contest (SHREC).

SHREC Similarly to what happens in other information retrieval research areas, investigators that work on 3D object retrieval established recently an international retrieval contest. Organised by the Network of Excellence AIM@SHAPE, the SHREC aims at evaluating the effectiveness of 3D-shape retrieval algorithms. In its initial version it was designed around the Princeton Shape Benchmark. The success of this contest provided a good feedback of the vitality of the research on 3D shape retrieval. In the following year, the scope of the contest specialised toward problems involving CAD content and partial similarity tasks, among others. Indeed, SHREC is now widely accepted as a reference benchmark in 3D shape retrieval and is expected to become an objective tool for evaluating and comparing 3D retrieval techniques.

SGP The first Eurographics Symposium on Geometry Processing was held in 2003 to allow researchers of the geometry processing area to present and share their work. This emerging research field aims at designing efficient algorithms for acquisition, manipulation, animation and transmission of complex 3D models. The processing techniques are based on recent developments in applied mathematics, computer science, and engineering. Closely related to our work, shape analysis is one of the topics covered by this symposium, making it an interesting event for us. After the success of the first edition, this event are being organised on a yearly base. Papers presented in this event appear in the Eurographics Proceedings Series in cooperation with ACM SIGGRAPH.

MIR Initially launched as Workshop Multimedia Intelligent Storage and Retrieval Management (MISRM) in 1999 and later renamed to ACM SIGMM International Workshop on Multimedia Information Retrieval (ACM MIR), this workshop has the purpose of bringing together researchers, developers, and practitioners from academia and industry, working on multimedia information retrieval. The workshop proceedings are printed by ACM and indexed in the ACM Digital Library, which makes it a good selection when choosing where to publish our work.

SPM Since its inception in 1991, the ACM Symposium on Solid and Physical Modeling Symposium (SPM), then Symposium on Solid Modeling and Applications, has been the primary venue for disseminating research results in the design, representation, analysis, visualisation, and use of digital models. Within the topics of interest covered by this event are geometry processing and shape analysis, two important subjects in our research work. Moreover, the proceedings of this event are published by ACM Press and papers of outstanding quality are selected for publication in the following international journals: Computer Aided Geometric Design, Computer-Aided Design and IEEE Transactions on Automation Science and Engineering.

CAD Published by Elsevier, Computer-Aided Design is an established international journal that provides engineers, designers and computer scientists with key papers on research and developments in the application of computers to the product design process. Although it focus mainly on engineering design, it covers a wide range of topics, such as the management of design databases, which includes the study of techniques for analysis and retrieval of 3D models.

CADG The international journal of Computer Aided Geometric Design, also published by Elsevier, is closely related with CAD journal referred before, but it main focus is on mathematical and computational methods for the description of geometric objects as they arise in areas ranging from CAD/CAM to robotics and scientific visualisation. Indeed, this journal does not have any topic focusing on retrieval of 3D shapes, but is an interesting place to find and publish work on shape description and simplification algorithms.

TOMCCAP The ACM Transactions on Multimedia Computing, Communications, and Applications (TOMCAP) is the flagship publication of the ACM Special Interest Group in Multimedia (SIGMM), published quarterly. It focuses on multimedia computing, multimedia communications and multimedia applications. Since our research field can be considered as part of multimedia information retrieval, this journal assumes major relevance both in our studies and in our dissemination goals.

IJDL The International Journal on Digital Libraries (IJDL), published by Springer, is a quarterly journal aimed at advancing the theory and practice of acquisition, definition, organization, management, and dissemination of digital information. Indeed, databases of 3D models are digital libraries, which justifies the importance of this publication for our work.

IJSM The International Journal of Shape Modeling is aimed at creating a suitable environment for exchanging research results obtained in advanced theories and techniques devised for handling the shape of objects, pointing out main aspects of modeling. Besides the pure shape modeling, this journal also includes topics regarding shape classification, recognition and characterisation.

2.3 Model databases

Besides identifying the key players in three dimensional shape retrieval and the most relevant events for this research area, it is important to identify the existing sources of three dimensional models to which shape retrieval can be useful. Indeed, as a result of recent advances on modelling, digitising and visualising techniques, there are a large number of 3D model collections available for usage both on the internet and in domain-specific databases. Since it will be hard and out of scope of this work to exhaustively enumerate all of this databases, we will refer in the following paragraphs just a few of these databases offering public access.

The Protein Data Bank (PDB) [14, 15], an archival for macro molecular structures, is an early example of such collections now considered the single worldwide archive for biological macro molecules. The PDB stores atomic 3D coordinates and partial bond connectivity for around 29,000 protein molecules, as derived from crystallographic studies.

Another example of a 3D model collection was produced during the Digital Michelangelo Project [77, 76], an archive with digital models of Michelangelo sculptures and architecture containing an aggregate of nearly eight billion polygons.

The National Design Repository [94, 113, 95] is a collection of public domain computer-aided design data from a variety of sources. This data includes tens of thousand solid models and CAD files. Recently, Google^{TM} released the Google^{TM} 3D Warehouse [63], an online service that hosts 3D models of existing objects (mainly buildings) created in Google^{TM} SketchUp [62]. These models can be downloaded into Google^{TM} Earth [61] and placed in their actual location on earth.

Relevant for our work are the databases associated with shape benchmarks, very useful for evaluating retrieval algorithms. From these, the most important is Princeton Shape Database [50] which stores polygonal surface models for more than thirty thousand objects crawled from the web. Additionally, the PRECISE group provides a fully classified database of mechanical parts [60]. This shape benchmark has special interest for people working on analysis and retrieval of CAD models. The objects in this database were collected from various sources including industrial partners of Purdue University and from the web.

2.4 From shapes to descriptors

The application diversity of three-dimensional models in a wide range of fields lead to several distinct forms of model representation [29]. However, these can always be converted or approximated to a more generic one, such as a polygonal mesh, which could be interpreted by classification and retrieval algorithms. These algorithms usually rely on feature vectors to describe both models and queries and then perform the search over the feature vectors instead of comparing the objects directly. Indeed, the usage of feature vectors is the standard approach for multimedia retrieval [44]. A feature vector, also referred as descriptor, is a set of values extracted from a multimedia object that describe it numerically in a high dimensional space.

In 3D models, these values usually describe the shape or certain aspects of the object and form a feature vector of high dimensionality. Thus, a 3D shape descriptor can be considered as a representation of a three dimensional object in a high-dimensional vector space. However, an important goal of any shape description approach is to preserve the maximum shape information on a feature vector with the lower dimensionality possible. Indeed, finding such computational representation of a shape is considered as the primary challenge in building a shape based retrieval system [49]. Therefore, we will present in the next section a survey on most important shape description techniques.

Some relevant recent work had focused on the computation of the shape descriptor, *i.e.* , the extraction of global or local object characteristics. Indeed, several research groups are working in this field, developing mechanisms to support 3D shape retrieval from large

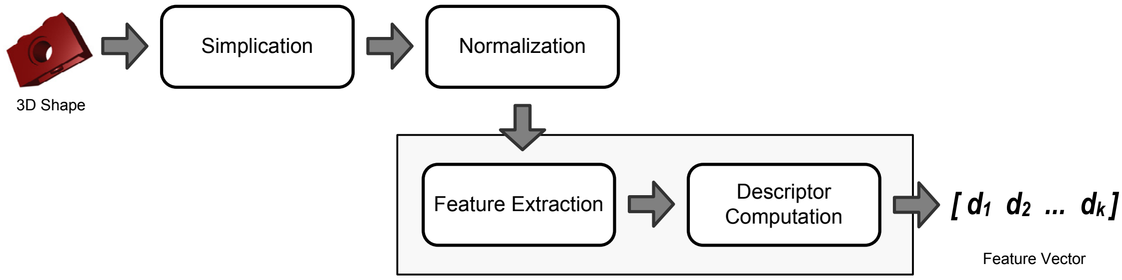

Figure 1: Shape descriptor computation process.

collections using geometrical queries. One of the main goals of their research is to devise descriptors that represent efficiently the shape geometry and that allow computation of ill-defined properties, such as "geometric similarity". A wide variety of extraction algorithms have been proposed, ranging from basic approaches, which use properties of an object bounding box, to more complex ones, like the distribution of normal vectors or curvature, or the Fourier transform of some spherical functions that characterise objects. Depending on the method used, the feature vector describes particular characteristics of an object, capturing different features.

Despite the wide variety of approaches to 3D shape descriptor computation, this process generally follows the basic steps depicted in Figure 1. The first step is the object simplification. In this step is usually performed a noise removal algorithm, especially in digitised models. Additionally, irrelevant geometrical features are removed from the object. In the second step, the resulting shape is pre-processed for rotation, translation and/or scale invariance, when required by the algorithm, followed by pose normalisation. The third step starts by converting the resulting shape to an object abstraction according to descriptor needs. Thus, it is converted to surfaces, volumes or images, depending on the approach. Then, the object abstract representation is transformed into a numerical representation using, for instance, histograms, spherical harmonics or discrete Fourier transform. The last step uses this numerical representation to generate a feature vector. Although this basic steps are similar in most approaches, different authors use different techniques, making descriptor computation a rapidly evolving field.

2.5 Shape descriptors

In this section we will briefly describe existing approaches to 3D shape description. Although we tried, in this document to produce a comprehensive survey, readers are encouraged to consult recent publications entirely dedicated to this topic. Bustos *et al.* extensively surveyed methods for 3D shape descriptor computation [27] and proposed a taxonomy for these methods. However, due to the variety of distinct approaches, there is no universally accepted taxonomy of 3D shape descriptors.

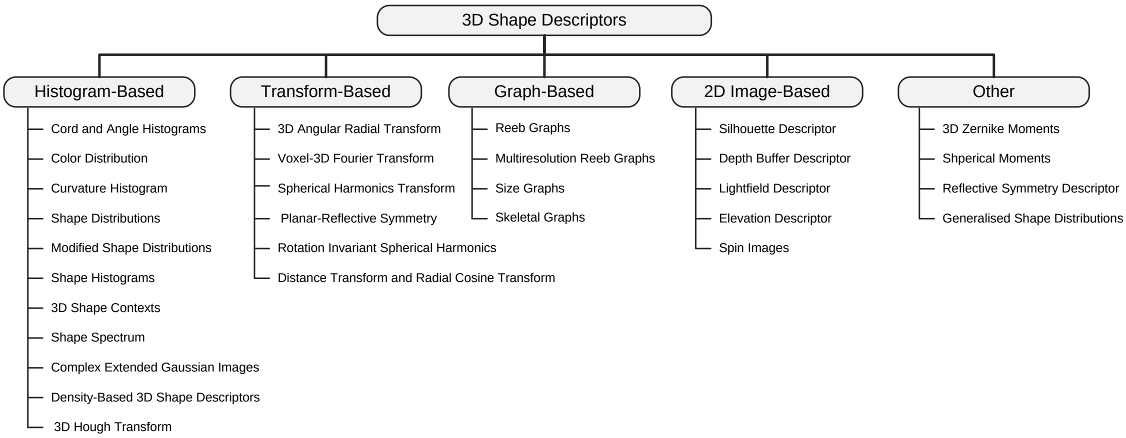

Figure 2: Taxonomy of 3D shape descriptors.

Nevertheless, I will follow a general classification scheme, very similar to the taxonomy proposed earlier by Bustos *et al.*. This scheme, depicted in Figure 2, emphasises the specific way to exploit the shape information contained in a 3D object and was suggested by Akgül [4]. However, the research on 3D shape descriptors is a rapidly evolving field and, besides the recent work referred above, two other relevant surveys were published in recent years. A complete review on content based 3D shape retrieval methods was presented in 2004 by Tangelder and Veltkamp [115]. Additionally, in the following year, Iyer *et al.* published another state-of-the-art review [64] where they classify and compare several 3D shape searching techniques from a CAD/CAM perspective and suggest future trends.

2.5.1 Histogram-based descriptors

In statistics, a histogram is a summary graph showing a count of the data points falling within tabulated frequencies. Basically, it is the graphical version of a table which shows what proportion of cases fall into each of several or many specified categories. Widely used in computer graphics to represent the color distribution in an image, the color histogram is computed by counting the number of pixels for each color. Adopted to 3D shape description, an histogram is often referred as an accumulator that collects numerical values of certain features calculated from the shape to represent. Based on this loose definition, many 3D shape descriptors can be considered as histogram-based methods, although they are not based on histograms in the rigorous statistical sense of the term.

Cord and angle histograms The use of cord and angle histograms for 3D shape descriptors were presented by Paquet *et al.* in [90, 91]. The authors define a cord as a vector that goes from the centre of mass of an object to the centre of mass of a bounded region on the surface of the object. They state that most often in 3-D applications the bounded region is a triangle, which is true due to broad use of triangular meshes. Even in other cases, the object can be transformed into a triangular mesh in the second step of shape descriptor computation process.

This 3D object descriptor are computed based on a collection of three histograms. The first histogram represents the distribution of the angles between the cords and the first reference axis. The second histogram represents the distribution of the angles between the cords and the second reference axis. The third histogram provides the distribution of the radius.

Despite its computational simplicity and efficiency, this approach simplifies triangles to their centres and does not consider the size and shape of the mesh triangles. Therefore, triangles of different sizes have equal weight in the final distribution and centres may not represent adequately the impact of the triangle on the shape distribution because of arbitrary triangle orientations. Moreover, since only global features are used to characterise the overall shape of the objects this method is not very discriminating about objects details, but their implementation is straightforward. It is often used in object retrieval as an active filter, after which more detailed comparisons can be made, or can be used in combination with other methods to improve shape descriptors.

Color distribution Paquet and Rioux [90] proposed, along with the cord and angle histograms, a peculiar color based descriptor for 3D shapes. In their approach, a voxelised representation of the 3D object, where each voxel has a color value associated with it. This value is computed using information from the texture map, material properties and vertex color extracted from the object representation.

Authors suggest three distinct approaches to compute the histogram that describes the object. If color location is relevant, it is only necessary to compute the color histogram of the object, based on the triplets (r, g, b) that represent the model colors. Alternatively, the dominant color is determined for each triplet and then the angle between the normal corresponding to that point and the first eigenvector is calculated. The statistical distribution of these angles is represented by a set of three histograms according to the dominant color. Otherwise, if color location is relevant, authors suggest using a wavelet approach based on a model with six dimensions: x, y, z, R, G, B. The six-dimensional wavelet transform is computed and then used to construct a histogram.

Curvature histogram Koenderink and van Doorn [73] defined the curvature index as a function of the two principal curvatures of the surface. This index gives the possibility to describe the shape of the object at a given point. However, it loses the information about the amplitude of the surface shape and is sensitive to noise. Later, Vandeborre *et al.* used this index to compute a curvature histogram of the shape [119], a local descriptor invariant to geometric transformations.

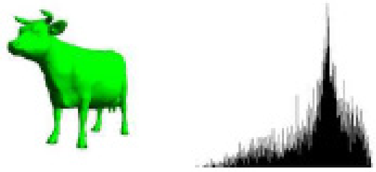

Figure 3: 3D shape and corresponding curvature histogram. (Figure taken from [119])

The surface curvatures are computed at a generic vertex v of a three-dimensional polygonal mesh as expressed in Equation 1, with $k_v^1 \geq k_v^2$, where k_v^1 and k_v^2 are the principal curvatures associated with the point v.

$$I_v = \frac{2}{\pi} \arctan \frac{k_v^1 + k_v^2}{k_v^1 - k_v^2} \tag{1}$$

In curvature histogram construction, the computation of principal curvatures is a key step. Estimating these curvatures can be achieved in several different ways. Authors suggest computing the curvature at each face of the mesh by fitting a quadric to the neighbourhood of this face (*i.e.* the centroid of this face and the centroids of its 1-adjacent faces) using the least-squares method[1]and then determine the principal curvatures k^1 and k^2 using the eigenvalues of a Weingarten endomorphism[2]. An alternative way to compute the curvature histogram descriptor is described in [24]. Figure 3 depicts the curvature histogram with 1024 intervals and the three-dimensional object from where it was extracted.

Shape distribution Osada *et al.* proposed a method for computing shape signatures for arbitrary 3D polygonal models [88]. They use a collection of shape functions computed with random sampling of the surface of the 3D object to describe. This shape function measures global geometric properties of the shape, based on distance, angle, area and volume measurements between random surface points.

Authors suggest five distinct one dimension functions to measure the object properties, which are quick to compute, easy to understand, and produce distributions that are invariant to rigid motions (translations and rotations). The function to compute $D1$ shape distribution measures the distance of a surface point to the centre of mass of the model. Depicted in Figure 4, the $D2$ shape distribution is a function measuring the dis-

[1]Least squares is a method for linear regression that determines the values of unknown quantities in a statistical model by minimizing the sum of the residuals (difference between the predicted and observed values) squared. It is used to find or estimate numerical values of parameters to fit a function to a set of data and to characterise the statistical properties of estimates [1].

[2]Principal curvatures are usually defined as the eigenvalues of the Weingarten map $W = I^{-1}II$, where I and II denote respectively the first and second differential forms.

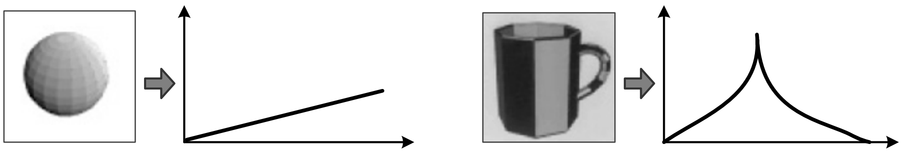

Figure 4: Two distinct objects and corresponding $D2$ shape distributions.

tance between two surface points. The $D3$ shape function measures the square root of the area of the triangle defined by three surface points. The cube root of the volume of the tetrahedron defined by four surface points is measured by $D4$ shape function. Finally, $A3$ distribution function measures the angle formed by three random surface points.

The shape descriptors are constructed from the histograms of a set of the above mentioned shape functions, controlling histogram accuracy through sampling density. The descriptor of shape distributions is fast, simple to implement, and useful for 3D shapes discrimination. However, the proposed shape functions are not adequate to fully describe the 3D shape effectively. Indeed, this approach distinguish models in broad categories very well, but perform poorly when used to discriminate between models with similar gross shape properties but vastly different detailed shape properties.

Modified shape distribution Ohbuchi *et al.* [87] extended the $D2$ shape functions proposed by Osada *et al.*, by devising a set of shape features that are tolerant to topological variations and geometrical degeneration. These are the modified shape $D2$ ($mD2$), the angle and distance histogram (AD) and the absolute angle and distance histogram (AAD).

In the proposed technique, the $mD2$ is similar to original $D2$, but authors used a quasi-random number sequence[3] to select points instead of the pseudo-number sequence[4] suggested by Osada *et al.*. To compute AD, authors measure both distance between a pair of points and angle formed by the surfaces the pair of points are located. Then, a 2D histogram is computed using the angle and distance as two independent variables. The AAD is computed in an identical manner, but while the AD histogram respects the sign of the angle, the AAD feature takes the absolute value of the inner product in order to increase robustness against inconsistently orientated surface. Thus, these two descriptors are not supposed to be used together. Indeed, the choice of using AD or AAD depends only whether the surfaces are properly and consistently orientated.

[3]The quasi random number sequence, also called low-discrepancy sequence, is not really random since a predictable sequence of numbers is generated. However, the numbers generated will provide uniform sampling with a suitable number of samples is required.

[4]A pseudo random number sequence exhibits statistical randomness while being generated by an entirely deterministic causal process. Indeed, a pseudo random generator calculates a number in the range $[0.0, 1.0]$ given an initial seed which is updated each time a number is requested. With enough random numbers sampled, the distribution will be uniform. This, however, could be a large number of samples.

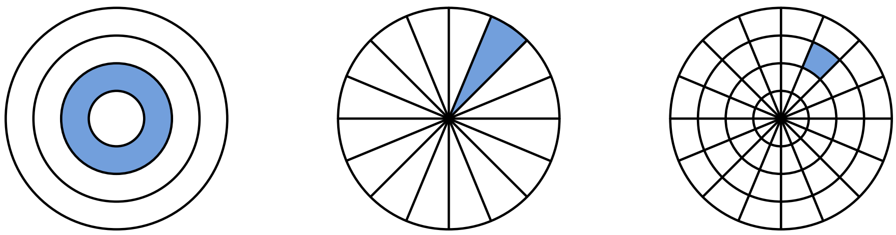

Figure 5: 2D examples of the space decomposition techniques proposed by Akerst *et al.* in [9]. With a single bin marked, are depicted, from left to right, shell, sector and spider web model.

Shape histograms The shape histograms were proposed by Akerst *et al.* as an intuitive approach to describe 3D solid models [9]. In this approach the space where the object resides is partitioned using one of three space decomposition techniques: a shell model, a sector model and a spider web model (depicted in Figure 5). Shell model consists in decomposing the 3D space into concentric shells around the centre point. In the sector model the 3D is decomposed into sectors that emerge from the centre point. Finally, the spider web model (sometimes referred as combined model) is a simple combination of the two decomposition models described above. In any of these techniques, each cell of the decomposed 3D space correspond to a bin in the histogram. Then, the histogram can be constructed by accumulating the surface points in the bins.

The shape histograms method is an intuitive and discrete representation of complex spatial objects. However, authors illustrate the shortcomings of Euclidean distance to compare two shape histograms and make use of a Mahalanobis[5] quadratic distance measure taking into account the distances between histogram bins. On the other hand, this approach needs pose normalisation to be performed in the pre-processing stage. The pose normalisation is necessary because the sector model is only scaling invariant, while the shell model is only rotation invariant.

3D shape contexts The shape contexts was initially introduced by Belongie *et al.* [13] as a descriptor for computing similarity between 2D images. Using reference points, authors assign a shape context to each one, by capturing the distribution of the remaining points relative to it. Körtgen *et al.* extends the 2D shape contexts into a 3D shape description and combines it with the shape histogram, thus proposing a set of descriptors called 3D shape contexts [74].

To compute the 3D shape contexts, N points are sampled from the shape boundaries, as in the shape distribution approach [88]. Then, the vectors originating from one sample point to all other points in the shape are computed, as depicted in Figure 6. Using the distribution over relative positions, is computed for this point a coarse histogram of the

[5]A short explanation of Mahalanobis distance can be found in Section 2.6.3

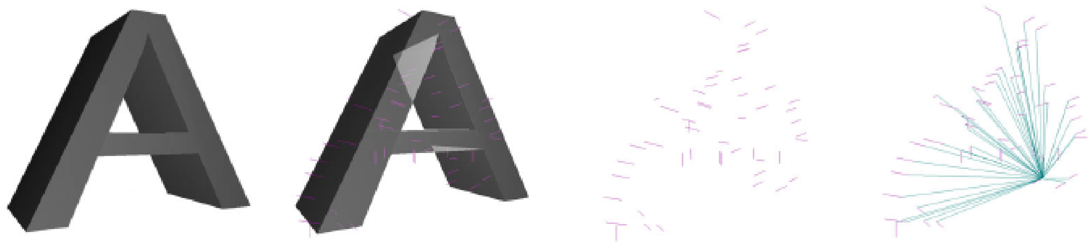

Figure 6: From left to right: original 3D object; mesh with fifty samples; just the fifty samples; vectors originating from one sample point to all others (Figures taken from [74]).

relative coordinates of the remaining $N-1$ points. Indeed, this histogram is an adapted version of one of the Ankerst's [9] shape histograms centred upon the sample point. This method is applied to all N sampled points, producing a 3D shape descriptor that is simply a set of N histograms.

Extended Gaussian images Defined by Horn in [59], the extended Gaussian images (EGI) is a histogram-based technique to represent the shapes of surfaces that define a 3D object. Initially devised for recognition in machine vision systems, this approach was later adapted for pose determination and for computing 3D shape descriptors [68], by means of the complex extended Gaussian images (CEGI) representation.

Basically, the EGI of a 3D object is a histogram that records the variation of surface area with surface orientation. Each bin of this histogram is defined by a pair (θ_i , ϕ_k) and corresponds to some quantum of the spherical azimuth and elevation angles (θ, ϕ) in the range $0 < \theta < 2\pi$ and $0 < \phi < \pi$. These bins accumulate the count of the spherical angles of the face normal per surface, weighted by the scalar that represent the associated visible face area. The CEGI concept, presented by Kang and Ikeuchi, extended the EGI approach by adding the normal distance of each face to the origin. Thus, the weight associated with a particular normal in the CEGI is a complex number whose magnitude and phase are the corresponding visible face area and its signed distance to the origin, respectively. Thus, for a given point in the CEGI associated to normal $\vec{n}_k$, the point weight w_k can be estimated through Equation 2:

$$w_k = |A_{\vec{n}_k} e^{jd_k}|, \tag{2}$$

where $A_{\vec{n}_k}$ is the area of the corresponding face and d_k is the distance of that face to the origin. Concluding, the CEGI method combines the pose determination into shape descriptors and even into the shape matching process. It enables estimation of both the orientation and translation of a given object with respect to a stored model or prototype.

More recently, Zhang *et al.* further extended the EGI representation to capture the volume distribution of an object without canonical alignment, while maintaining the translation, scale and orientation invariance. The volumetric extended Gaussian image (VEGI) shape descriptor [133] is thus able to differentiate between convex and non-convex shapes

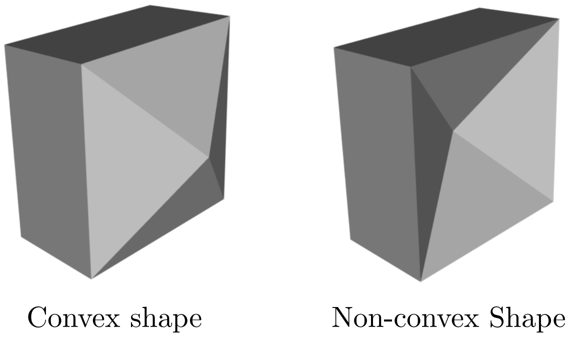

Figure 7: Two objects with the same EGI descriptor, but different volumes, thus different VEGI descriptors.

that share the same EGI, such as the objects depicted in Figure 7. Furthermore, VEGI directly extracts shape features avoiding pose normalization, since it does not depends on canonical alignment of shapes.

3D Hough transform Based on the 2D generalised Hough transform [12], Zaharia and Prêtoux proposed in [132] the 3D Hough transform descriptor (3DHT) to represent a three-dimensional object. In their later work [131], developed canonical 3D Hough transform descriptor (C3DHT). These descriptors are constructed by accumulating points within a set of planes in 3D space, as described below.

Considering that a plane can be defined by a triplet (s, θ, ϕ), where $s > 0$ denotes the distance from the origin of the coordinate system to plan and $0 < \theta < 2\pi$ and $-\frac{\pi}{2} < \phi < \frac{\pi}{2}$ respectively denote the azimuth and elevation associated with the spherical representation of the plane's unit length normal vector. Then, each axis of the parameter space (s, θ, ϕ) can be uniformly sampled and a set of planes with orientation (θ_k, ϕ_j) passing through the p is created. Considering a polygonal model, p stands for the centre of mass of a mesh face g. Finally, for each plane, the quantized normal distance to the origin s_i is calculated. If this value is positive, the bin corresponding to $(s_i,\ \theta_j\ ,\ \phi_k)$ is augmented by a weight factor w^p_{jk} defined in Equation 3, where n^p denotes the unit length normal vector of face g and n_{jk} denotes the normal of the plane defined by pair (θ_i, ϕ_k).

$$w^p_{jk} = A_p |\langle n^p, n_{jk} \rangle| \tag{3}$$

Indeed, 3DHT can be considered as a generalized version of EGI since, for a given s_i, (θ_j, ϕ_k)-bins correspond to an EGI at distance s_i, except for the way the contributions of the faces are assessed. In fact, Agkül *et al.* have experimentally proven [6] that the 3DHT descriptor captures the shape information better than the EGI descriptor.

Shape spectrum The shape spectrum was introduced by Dorai and Jani [39] as a view-based representation of 3D free-form objects. In their work, authors focus on generation of a set of representative views of a 3D object suitable for efficient clustering and retrieval. Indeed, they apply this method to devise a general and powerful technique for organising multiple views of objects of complex shape and geometry into compact and homogeneous clusters. The shape spectrum characterises quantitatively the object shape by summarising the area on the surface of an object at each shape index value. This value is a quantitative measure of the shape of a 3D regular surface at a point p, denoted by I_p and defined as expressed in Equation 4, where k_p^1 and k_p^2 are the principal curvatures of the surface associated with point p, with $k_p^1 > k_p^2$.

$$I(p) = \frac{1}{2} - \frac{1}{\pi} \arctan \frac{k_p^1 + k_p^2}{k_p^1 - k_p^2} \tag{4}$$

Therefore, the shape index is a local geometrical attribute of a 3D surface, expressed as the angular coordinate of a polar representation of the principal curvature vector. It ranges in the interval $[0, 1]$ and is not defined for planar surfaces, since those have $k^1 = k^2 = 0$ which will result in an indeterminate shape index[6]. Figure 8 illustrates the shape index of a free-form 3D object computed by this technique. The shape index provides a scale for representing salient elementary shapes and is invariant with respect to scale and Euclidean transforms. However, an important problem regarding the shape index, in fact all curvature-related quantities, is the estimation unreliability leading to a lack of robustness. Such shortcoming was alleviated by Zaharia and Prêtoux in [131] by augmenting the shape index histogram by two additional attributes named planar surface and singular surface. Then, the proposed 3D shape spectrum descriptor (3D SSD) was applied to 3D retrieval within the MPEG-7 framework for multimedia content description. The 3D SSD of a 3D mesh is defined as the histogram of the shape index values, calculated over the entire mesh.

Therefore, the 3D SDD locally characterises free-form surfaces represented as discrete polygonal 3D meshes. One major advantage of this descriptor is its generality, since 3D meshes may include open surfaces that have not an associated volume. Furthermore, inherited from the shape index properties, the 3D-SSD is invariant with respect to scale, translation, rotation and reflection transforms. On the other hand, this descriptor, as a simple local feature representation, is better to be combined with some global representation schemes to effectively describe 3D object for shape retrieval purposes.

[6]For computational purposes, the shape index of a planar surface is represented by a symbolic label, usually a predefined shape index value to indicate surface planarity. In their implementation, Dorai and Jain [39] used a shape index value of 2.0.

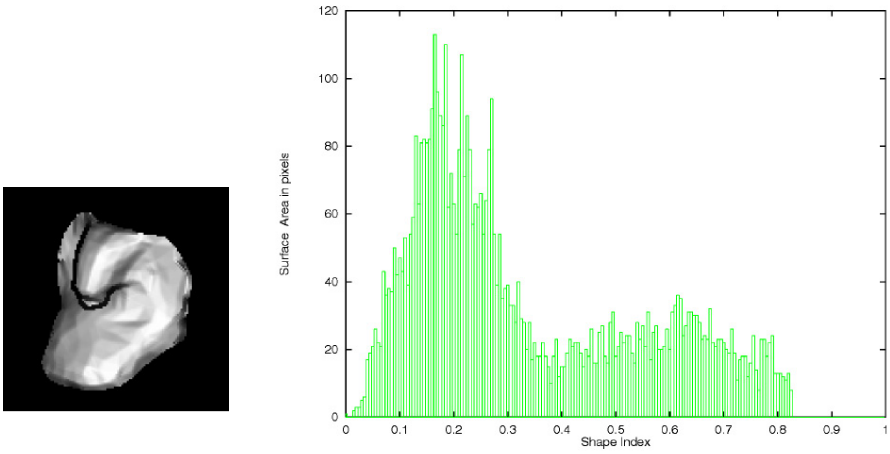

Figure 8: View of a 3D object and corresponding shape index (Figure taken from [39] © 1997 IEEE).

Density-based shape descriptors A density-based descriptor of a 3D object is defined as the sampled probability density function[7] (PDF) of some surface feature. The feature is local to the surface patch and treated as a random variable. This analytical framework was proposed by Akgül and Sankur [6, 7] to extract 3D shape descriptors from local surface features characterising the object geometry. In [5], the authors suggest using as features the radial distance, the radial direction, the normal direction, the radial-normal alignment and the tangent-plane distance. Additionally, authors considered the shape index proposed in [73] as a second order feature that provides categorisation of the shape into primitive forms.

The density-based shape descriptor can be considered as a probability modeling problem. As explained by Akgul *et al.* in [8], the local surface properties are first measured via various features, such as the referred above or a subset of the most discriminating of them. This feature information is then processed with the kernel methodology for density estimation [106, 41] and the PDF of the local feature is estimated at chosen target points. Then, the shape descriptor vector is then simply a sampled version of this probability density function.

As Akgul mentioned in [6], when performing on the Princeton Shape Benchmark, the density-based descriptor shown very good results when compared with other well-known 3D shape descriptors, such as cord and angle histogram, shape distribution, 3D Hough transform, extended Gaussian images, among others. However, this descriptor still suffers a problem that the features are neither scale- nor rotation-invariant and, since the method depends on them, pose normalisation must be accomplished during the prepossessing phase.

[7]The probability density function [3], also called probability function or density function, is a function that represents a probability distribution in terms of integrals.

2.5.2 Transform-based descriptors

Instead of estimating the shape descriptor from the three dimensional space, some approaches rely on mathematical transformations to switch from the spatial domain to a more suitable one and compute from there a shape descriptor. These new spaces are usually the frequency domain, although some recent approaches use different spaces.

Voxel-based 3D Fourier transform Vranić and Saupe suggest switching from the spatial domain to the frequency domain via a 3D Fourier transform [123] (3DFT). Authors start by performing pose normalisation, then voxelise the object using the so-called bounding cube[8] (BC). This voxelisation is achieved by subdividing the BC into N^3 equal sized cubes (cells) and calculating the proportion of the total surface area of the object inside each cell. Then, authors apply a 3D discrete Fourier transform to the voxelised model, *i.e.* calculated values in cells, to compute the descriptor that represents the feature in the frequency domain.

Distance transform and radial cosine transform Following the initial ideas from Vranić and Saupe of applying Fourier transforms to feature extraction, Dutagaci *et al.* proposed estimating a 3D discrete Fourier transform descriptor using two different voxel representations of 3D objects [42], namely, binary and continuous, as depicted in Figure 9. While in the first case the voxel values are simply set to 1 in the surface of the object and 0 elsewhere, in the continuous representation the space is filled with a function of distance transformation. This function, v_d, is called inverse distance function and is calculated as in Equation 5, where $d(x)$ is the minimum L_1 distance from point x to the object surface. Thus, the function $v_d(x)$ has is maximum value of 1 on the object surface and decreases when moving away from it.

$$v_d(x) = \frac{1}{d(x)+1} \tag{5}$$

To compute a quasi rotation invariant descriptor, authors suggest a measure of the spectral energy in a sphere of radius r, because the spectral energy in a sphere centred at the origin of the frequency domain remains constant under rotation. Afterwards, they define the incremental spectral energy (ISE) as the difference of the spectral energies contained within concentric spheres, normalise ISE by r^2 and take its square root to balance out large values accumulated in the low-pass shells. The normalised spectral energy(NSE), which has the property of rotation invariance, is used as the 3DFT-based descriptors of the object. However, slight deviation from rotation invariance can occur due to voxelisation distortion, together with the distance transform values at the corners of the bounding box, which decays rapidly to zero towards there, makes this representation not

[8]Vranić and Saupe consider the bounding cube of a 3D-model as the tightest cube in the canonical coordinate frame that encloses the model, with the centre in the origin and the edges parallel to the coordinate axes.

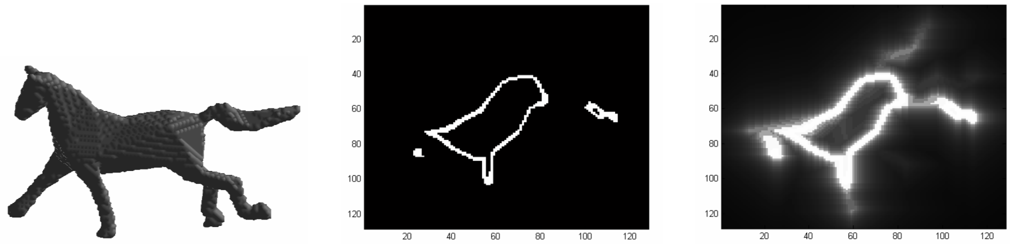

Figure 9: From left to right: voxelised 3D object; cross section of the binary function; cross section of the inverse distance function (Figures taken from [42] © 2005 IEEE).

totally rotation invariant, but instead quasi rotation invariant. Furthermore, authors claim that this descriptor provide a multi-resolution representation because NSE descriptors values at small radii (low-pass region) carry information about the gross shape of the object, while shape details are encoded in the spectral shells at high-frequency radii.

Additionally, Dutagaci *et al.* suggest the use of 3D radial cosine transform (RCT) as an alternative to 3DFT [42]. The RCT coefficients constitute a set of rotation invariant shape descriptors. Such descriptor represents a 3D model with a small number of features, thus being easy and fast to be calculated. However, the retrieval results of RCT are generally worse than 3DFT or some other approaches. Therefore, it is always considered to be mainly suitable to be used together with other descriptors as an preliminary filter.

Spherical harmonics transform A 3D object can be characterized by a function r on the sphere S^2. To that end, rays are cast from the centre of mass the object and, for each ray defined by $u \in S^2$, is estimated the value equal to distance d from origin to the last point of intersection with the object surface. These values yield a sample of function r, called spherical extent function, for a shape I. This function can be defined as

$$\begin{aligned} r: \quad & S^2 \rightarrow \mathbb{R} \\ & u \longmapsto max\{d \geq 0 | du \in I \cup \{0\}\} \end{aligned}$$

where 0 is the origin. In their initial approach, Saupe and Vranić [122] took a coarse number of samples of $r(u)$ to construct a feature vector. However, this simple technique is sensitive to small perturbations of the model. Therefore, to improve the robustness of this approach, they later proposed [97] extracting a dense sample of $r(u)$ and then compute spherical harmonics for this function to describe the shape. Spherical harmonics form a Fourier basis on a sphere, like the sine and cosine do on a line or circle, and allow a spherical function to be decomposed into the sum of its harmonics. Thus, the function r can be represented by the Fourier transform on the sphere defined in Equation 6, using the spherical harmonic function Y_l^m and the Fourier coefficient $\hat{r}(l,m)$.

$$r(\theta, \phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^{m=l} \hat{r}(l,m) Y_l^m(\theta, \phi) \tag{6}$$

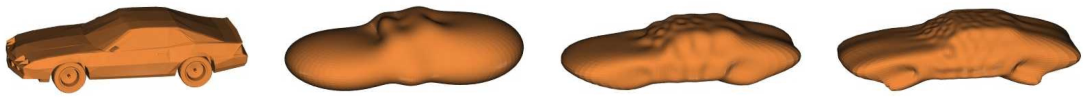

Figure 10: Multi-resolution representation of the function $r(u)$ applied to the original model using k^2 spherical harmonic coefficients (Figures taken from [97]).

To efficiently compute the Fourier coefficient, authors apply a spherical FFT algorithm to N samples taken at points u_{ij}, where $i, j = 0, ..., N$. Thus, the descriptor accuracy can be defined by changing the parameter that defines the sampling size, as well as the number of spherical harmonic coefficients to use. Figure 10 depicts the reconstruction of an object by using different number of coefficients. The shape descriptor is derived from these of coefficients, thus providing an embedded multi-resolution approach for 3D shape description.

Vranić and Saupe [123] improved the robustness of the proposed feature vector by taking samples of the spherical function $r(u)$ at many points, but characterising the map by just a few coefficients in the spectral domain. In [124], they enhanced the spherical harmonics transform approach described above by taking in account the orientation of the surface, along with the extent vector. To that end, they proposed the use of a complex feature vector, in which the real component of the complex function is the extent of the object from the origin and the imaginary part is computed using the normal vector of the surface at sampled points. Thus, in this approach, the complex function $r(u)$ for a shape I was defined as

$$\begin{aligned} r: \quad & S^2 \rightarrow \mathbb{C} \\ & u \longmapsto x(u) + \mathbf{i}y(u) \end{aligned}$$

$$\begin{aligned} x: \quad & S^2 \rightarrow [0, +\infty[\in \mathbb{R} \\ & u \longmapsto max\{x \geq 0 | xu \in I \cup \{0\}\} \end{aligned}$$

$$\begin{aligned} y: \quad & S^2 \rightarrow [0, 1[\in \mathbb{R} \\ & u \longmapsto \begin{cases} 0, & \text{if} x(u) = 0 \\ u \cdot n(u), & \text{otherwise} \end{cases} \end{aligned}$$

where $n(u)$ is the normal vector at the point $ux(u)$. To compute the complex feature vector, authors apply the spherical Fourier transform to the function $r(u)$, as they did in their previous work. Later, Vranić [120] proposed considering a set of concentric spheres with different radii, instead of a single one, thus using a collection of spherical functions to compute the descriptor.

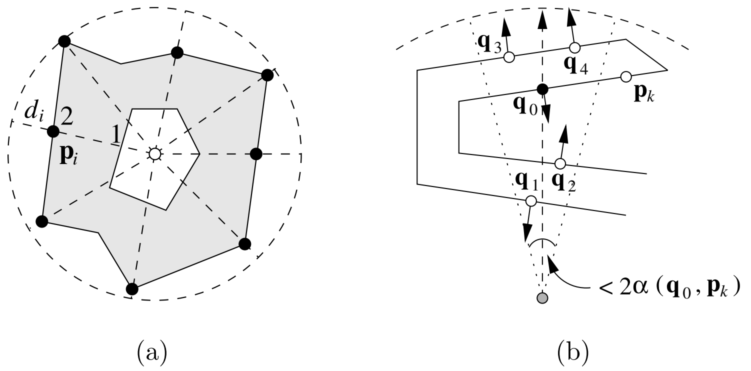

Figure 11: Computing feature maps. Measuring distance from p_i to sphere surface and counting number of object surfaces intersected by a ray (Figures taken from [129] © 2003 IEEE).

On a slightly different approach, Yu *et al.* proposed measuring the amount of effort required to morph[9] a 3D object into a canonical sphere [129] to describe the geometry and topology of the object. To that end, they use a descriptor similar to the spherical extent function together with a descriptor counting the number of intersections from a ray casted from the origin with the object surface. Unlike many other approaches, the authors construct descriptors from feature maps instead of histograms in order to capture spatial information. The distance map is computed by measuring the distance d_i from the normalised object surface to the bounding shape surface, as depicted in Figure 11 (a). Authors consider that the sum of all d_i correspond to the total energy required to deform the object into a sphere. The surface penetration map is constructed by counting the shape surfaces intersected by a ray shot from the centre of the sphere. To efficiently compute the penetration map, is considered a conical volume with the ray as its axis and a small angle α, such as the nearest neighbor point p_k of the ray intersection with the surface q_0 is not contained in the cone, as illustrated in Figure 11 (b). Then, surface normal sign changes are used to calculate the number of surfaces that the ray penetrates, since it is equal to the number of sign changes plus one. The resulting surface penetration map provides information about object topology and concavity.

Papadakis *et al.* presented recently [89] a 3D shape retrieval methodology based on spherical harmonics. The proposed model decomposition and feature extraction is very similar to previous approaches. They compute the spherical functions using not only the intersections of the surface with emanating rays but also points in the direction of each ray which are closer to the origin than the furthest intersection point.

[9]Morphing is a technique that changes one object into another through a seamless transition, generally by producing a sequence of intermediate objects.

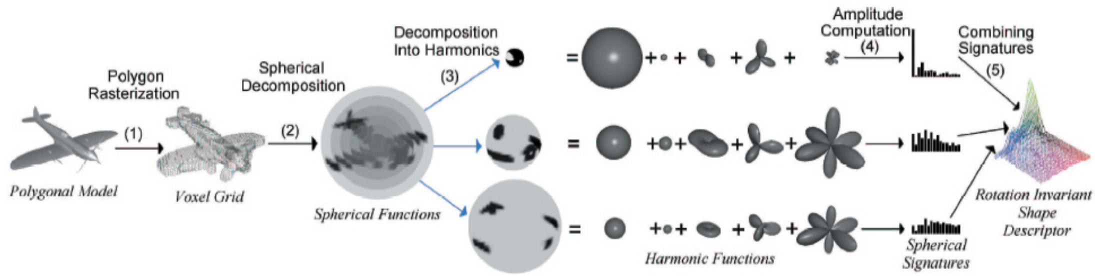

Figure 12: Princeton methodology for computing spherical harmonics shape descriptor (Figure taken from [50] © 2003 ACM).

Rotation invariant spherical harmonics After identifying several limitations of canonical alignment used in other approaches, Kazhdan *et al.*proposed [70] an alternate method to obtain rotation invariant representation of three-dimensional objects based on spherical harmonics. First, the object to describe should be converted into a voxel grid. Then authors intersect the object with a set of concentric spheres and construct a spherical function from voxel values for each sphere. Next, the frequency decomposition of each one of these functions is computed, as well as the norms of each frequency component at each radius. The resulting rotation invariant descriptor is a 2D grid indexed by radius and frequency.

Following this idea, the Princeton group derived a practical methodology, illustrated in Figure 12, to compute rotation invariant descriptor using spherical harmonics [50]. First, they rasterise the object into a $2R \times 2R \times 2R$ voxel grid, with $R \approx 32$ to provide adequate granularity for discriminating shapes while filtering out high-frequency noise in the original data. Then a value of one is assigned to each voxel if it is within one voxel width of object surface, or zero otherwise. The model is translated in order to move its centre of mass to point (R, R, R) and scaled in order to make the average distance from nonzero voxels to the centre of mass $\frac{R}{2}$. By applying these two transforms, translation and scale normalisation are obtained. In next step, a collection of spherical functions $f_r(\theta, \phi)$ defined as

$$f_r(\theta, \phi) = voxel(r\sin(\theta)\cos(\phi) + R, r\cos(\theta) + R, r\sin(\theta)\sin(\phi) + R), \tag{7}$$

where $r \in [0, R]$, $\theta \in [0, \pi]$ and $\phi \in [0, 2\pi]$. Then, using spherical harmonic transform, each f_r function is expressed as a sum of its frequencies and a rotation invariant signature is computed for f_r as a collection of scalars from the L_2 norm of its frequencies f_r^m. The two-dimensional rotation invariant spherical harmonics descriptor is obtained by combining these different signatures over the different radii.

3D angular radial transform Adopted as region-based shape descriptor in MPEG-7 standardisation, the angular radial transform [25] is a moment based description method that expresses pixel distribution within a 2D region. Ricard *et al.* proposed a generalisation of this descriptor to index 3D models [96]. The 3D angular radial transform (3D ART) descriptor preserves the properties of the original 2D descriptor, such as robustness to rotation, translation, noise and scaling. Moreover, the 3D ART produces compact descriptors and allows short retrieval times.

The 3D ART transform is a complex unitary transform defined on a unit sphere. Thus, to compute the 3D ART descriptor, the object must be represented in spherical coordinates, which can be insured during preprocessing stage. Considering ϕ the azimuth angle in the xy-plane from the x-axis, θ the polar angle from the z-axis and ρ the 3D ART coefficient of order n, m_θ and m_ϕ is defined by

$$F_{nm_\theta m_\phi} = \int_0^{2\pi} \int_0^{\pi} \int_0^{1} V_{nm_\theta m_\phi}(\rho, \theta, \phi) \times f(\rho, \theta, \phi) \rho d\rho d\theta d\phi, \tag{8}$$

where $f(\rho, \theta, \phi)$ is a 3D object function in spherical coordinates and $V_{nm_\theta m_\phi}(\rho, \theta, \phi)$ is a basis function composed by one radial function and two angular functions:

$$V_{nm_\theta m_\phi}(\rho, \theta, \phi) = A_{m_\theta}(\theta) A_{m_\phi}(\phi) R_n(\rho) \tag{9}$$

To achieve rotation invariance, the angular basis functions are defined by complex exponential functions as expressed in Equations 10 and 11 for polar and azimuth angles respectively. Without such constraints, the radial basis function is a simple cosine function, as shown in Equation 12. Since the values of parameters n, m_θ and m_ϕ are trade-offs between efficiency and accuracy, these must be chosen upon experimental evaluation, and may change depending on the data set.

$$A_{m_\theta}(\theta) = \frac{1}{2\pi} \exp(2\mathrm{j} m_\theta \theta) \tag{10}$$

$$A_{m_\phi}(\phi) = \frac{1}{2\pi} \exp(\mathrm{j} m_\phi \phi) \tag{11}$$

$$R_n(\rho) = \begin{cases} 0, & \text{if} n = 0 \\ 2\cos(\pi n \rho), & \text{otherwise} \end{cases} \tag{12}$$

In order to determine n, m_θ and m_ϕ, authors measured the recall response for several possible values of these parameters and found a good compromise between the efficiency and accuracy. For instance, in the particular case of a database of technical models, authors have chosen $n = 3$, $m_\theta = 5$ and $m_\phi = 5$. With these values, authors argue that their approach outperforms the spherical harmonics descriptor in speed while keeping a close accuracy.

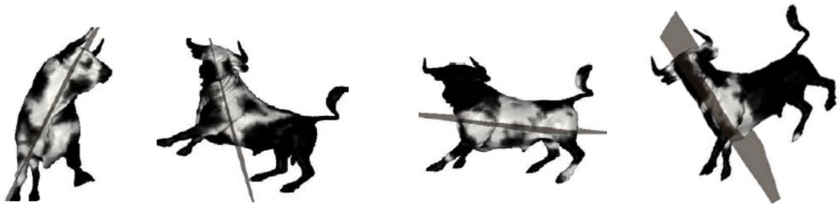

Figure 13: Symmetries with respect to planes representing the four strong local maxima of the PRST. (Figures taken from [92] © 2006 ACM).

Planar-reflective symmetry transform In [92] Podolak *et al.* introduced the planar reflective symmetry transform (PRST). The PRST is a transform from the space of points to the space of planes that captures a continuous measure of the symmetry of a shape with respect to all planes through its bounding volume. This transform combines and extends previous work that has focused on global symmetries with respect to the centre of mass in 3D meshes [69, 71], briefly described in Section 2.5.5, and local symmetries with respect to points in 2D images [130].

Authors also provide an iterative refinement algorithm to find local maxima of the transform precisely. In Figure 13 triangles of the model are colored to show how symmetric they are with respect to the plane of symmetry displayed. The triangles with highest symmetry values are lighter than the ones with less or no reflection in the given plane.

In addition, Podolak *et al.* use the planar reflective symmetry transform to define two geometric properties, the centre of symmetry and the principal symmetry axes. These properties are useful for aligning objects in a canonical coordinate system. Indeed, the planar reflective symmetry transform can be useful for several applications in computer graphics, including segmentation of meshes into parts, and automatic viewpoint selection, besides shape matching.

2.5.3 Graph-based descriptors

The above referred approaches focus on describing the geometry of model to classify, ignoring or giving just few relevance to topological information. At most, these approaches attempt to integrate topological information in the shape descriptor of the object. In contrast, graph-based approaches extract both topology and geometry of 3D objects, focusing on topological relationships between object components and using graphs to represent such relationships. These approaches are generally more complex than the previous ones, but despite their ability to encode geometry and topology, they do not generalise for any type of 3D shape, forcing each approach to restrict its scope to a specific type of object. Therefore, graph-based approaches are not effective in general-purpose retrieval applications.

Furthermore, due to the complexity associated to graph matching, alternative solutions to graph isomorfism are used, such as application of techniques from spectral graph theory

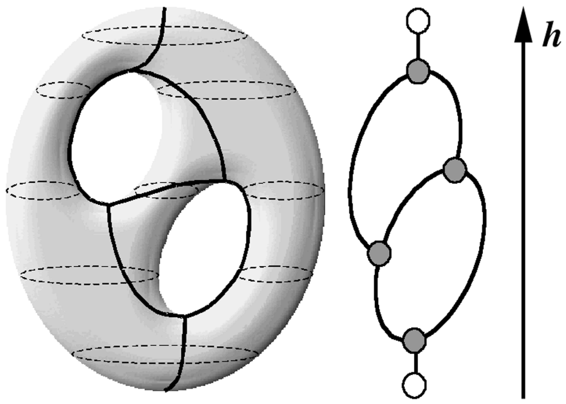

Figure 14: Reeb graph of a bi-torus and some cross sections. (Figure taken from [20])

to convert graphs into numeric descriptors [48, 103]. However, during recent years several researchers focused their attention in graph-based descriptors due to its potential. In the following paragraphs we will present the most commonly used approaches for graph-based descriptors.

Reeb graphs Back in 1946, Georges Reeb proposed considering a topological graph defined as a quotient spaceof a manifold[10] which, under opportune hypotheses defines the skeleton of the manifold itself [93]. Indeed, the Reeb graph is just a topological skeleton determined using a scalar function defined on an 3D object. To automatically construct a Reeb graph, Shinagawa and Kunii [102] proposed defining a scalar function and using a series of cross-sections of the object to determine nodes and arcs of the graph. Considering, for instance, the height function f associated to a manifold $M = M(x, y, z)$, the Reeb graph is the quotient space given by the relationship which identifies the points x_1 and x_2 having same function values and belonging to the same connected component of the inverse image of f. Figure 14 illustrates a Reeb graph of a bi-torus computed using a height function as mapping function.

Several approaches to 3D shape classification and retrieval based on Reeb graphs were proposed in recent years. Biasotti *et al.* obtain graphs by using different quotient functions f and suggest that a good choice of f is necessary to achieve good matching results [22]. Indeed, they conclude that f function must be determined based on the object type, since the same function produces different matching performance for different kinds of models. For instance, the authors proved that using the integral geodetic distance as a quotient function is especially suited for articulated objects.

[10]A manifold is an abstract mathematical space that is locally Euclidean. This means that around every point there is a neighbourhood that is topologically the same as the open unit ball in $\mathbb{R}^n$, *i.e.* the neighbourhood resembles Euclidean space, but the global structure may be more complicated.

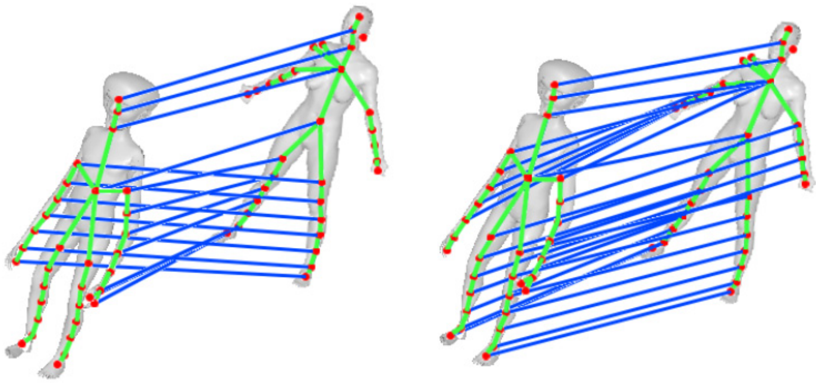

Figure 15: Results of shape matching with MRG *versus* ARG (Figure taken from [117] © 2004 IEEE).

Multiresolution Reeb graphs Hilaga *et al.* introduced the concept of topology matching for 3D object retrieval [57], describing a matching method best suited for articulated objects. Their method constructs Reeb graphs at multiple levels of resolution of a function f, the Multiresolution Reeb Graph (MRG). They proposed using the integral geodetic distance as f, since this function is invariant to rotation and translation and is also robust against changes caused by mesh simplification or subdivision. However, they also defined f as the height of a point on the surface of the object or the curvature value at that point. According to the chosen function, the resulting descriptor has certain properties for a different kinds of models. To quickly determine similarity between polyhedral models they compare the graphs using a coarse-to-fine strategy while preserving the consistency of the graph structures, which results in results in establishing a correspondence between the parts of objects. This graph matching is achieved through sophisticated heuristics proposed by authors and improved later by Tung *et al.* [118].

Tung and Schmitt [117] took further the approach by Hilaga and augmented the Reeb graph by storing geometric attributes in each node, since the original method only takes into account topological information, which is often insufficient for effective shape matching. In this approach authors use, as geometrical information, features such as the cord histograms, local curvature and volume associated with each node. Moreover, they also provided a new topological coherence condition to improve the graph matching.

Using the proposed Augmented Reeb Graph (ARG), Tung and Schmidt could overcome some issues raised during matching with Hilaga's MRG. For instance, graph edges topologically similar might not really be geometrically similar, thus being wrongly matched, as depicted in Figure 15. This figure illustrates the gain obtained by introducing geometrical information in the nodes. While matching without geometrical information (left) legs can be matched with arms, since they are topologically equivalent, by adding geometrical information, arms and legs are well matched.

Aware of MRG drawbacks when models become geometrically and topologically detailed, Bespalov *et al.* studied the application of Hilaga's method to matching of complex machined parts [17]. They stated that such solution produces poor results when directly applied to 3D solid models in engineering databases. Since for this kind of models topological insensitivity is important, they conclude that some improvements should me made to MRG technique. Thus, they present in [19] an alternative to Hilaga's approach. Their method computes the scale-space decomposition of a shape, represented as a rooted tree. Through spectral decomposition the problem is of matching reduced to that of computing a mapping and distance measure between vertex labeled rooted-trees. Indeed, authors claim that their method represents a computationally efficient approach to matching of 3D models, enabling highly accurate matching of solid models of 3D mechanical parts.

Size graphs Following the previous approaches, Biasotti *et al.* use the Reeb graph to construct a centerline skeleton of a 3D model and apply a size function to create a size graph [21]. Their idea is to associate with a 3D object a graph (G^f, ϕ), where G^f is the centerline skeleton computed using the quotient function f and ϕ is a measuring function labelling each node of the graph with local geometrical properties of the model.

In this approach, authors consider four distinct mapping functions f, namely the distance from the barycenter, the distance from the center of the bounding sphere, the integral geodetic distance and the topological distance from curvature extrema [84]. Based on these functions, a centerline skeleton is extracted from the original model. Then, for each node of this skeleton, the value of function ϕ must be calculated to obtain the size graph. Biasotti *et al.* suggest measuring a set of features of the corresponding region on the model, such as the area of the region or the minimum, maximum and average distance of the barycenter of the region to region vertices. To compare models authors use the matching between their size functions, as discussed by d'Amico *et al.* in [36].

Skeletal Graphs In a slightly different approach, skeletons can be derived from solid objects and represented as a direct acyclic graph (DAG). These skeletons capture important information about the object. However, when using shape skeletons in 3D object retrieval, two major challenges arise. Suitable skeleton computation algorithms and similarity functions should be defined. Sundar *et al.* presented a framework that provides both [107]. They propose, as shape descriptor for three-dimensional models, a skeletal graph encoding geometrical and topological information of the object. Then they apply graph matching techniques to match the skeletons and compare them.

To compute the skeleton, Sundar *et al.* first perform a voxelisation of the object. From the voxelised model, the skeletal points are calculated using a distance transform-based algorithm proposed by Gagvani and Silver [53] with a thinness parameter. This method reduces the model voxels to those voxels that are important for reconstruction. These remaining voxels are clustered and the corresponding skeletal points are connected

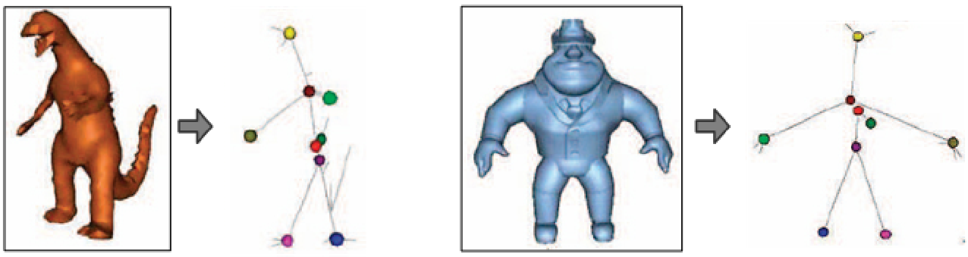

Figure 16: Skeletal graphs for a pair of matching objects (Figure taken from [107] © 2003 IEEE).

in a DAG by applying the minimum spanning tree algorithm. By decreasing the thinness parameter denser graphs can be obtained. Thus, authors suggest varying the thinness this parameter to obtain a hierarchical graph structure. Finally, to each node of the DAG is associated a set of geometric features and a signature vector that encodes topological information of subtrees rooted at this node. This topological signature vector is derived recursively over the sub-graphs of the node using eigenvalues of their adjacency matrices.

The matching procedure proposed by Sundar *et al.* consists of two stages. In the first stage shapes are matched by approximate comparison of their hierarchical skeletal graphs using a greedy algorithm to find the maximum cardinality, minimum weight matching in a bipartite graph. In the second stage is performed geometry matching over the information stored on nodes to refine the results. Figure 16 illustrates skeletal graph matching accomplished using Sundar *et al.* technique, showing the node-to-node correspondence based upon the topology and the radial distance about the edge.

2.5.4 Image-based descriptors

An approach completely different from the previous ones consider representing a three dimensional model in a set of two dimensional spaces. The basic idea behind such approach is that when two 3D models are similar, images captured from the same points of view are also similar. From this idea several researchers were able to reduce the problem of comparing 3D shapes to image matching. Thus, taking advantage of the existing reasonable amount of work in this area capable of producing good retrieval results.

Spin Images Johnson and Herbert proposed a 3D object recognition system [67] based on matching surfaces using the spin image representation. To produce spin images authors use oriented points on the model surface, *i.e.* points associated with the surface normal at that point. Each oriented point (p, n) corresponds to a spin image and defines a local coordinate system using the tangent plane P through p oriented perpendicular to n and the line L through p parallel to n, as illustrated in Figure 17. Then, two coordinates are defined with respect to the oriented point (p, n): the radial coordinate α and the elevation

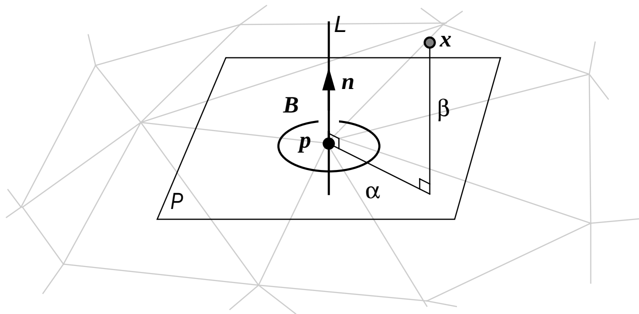

Figure 17: Building a spin image with respect to point p (Figure taken from [66]).

coordinate β. The cylindrical angular coordinate is omitted because it cannot be defined robustly and unambiguously on planar surfaces.

To create the spin image, a 2D accumulator indexed by α and β is created. Next, the coordinates (α, β) are computed for every vertex in the surface mesh that is within the support of the spin image. The bin indexed by (α, β) in the accumulator is then incremented. The resulting accumulator can be thought of as an image where dark areas in the image correspond to bins that contain many projected points. Figure 18 shows the projected (α, β) coordinates and spin images for three oriented points on a model of a valve. For 3D object matching, spin images can be constructed for every vertex in the surface mesh, producing a set of two-dimensional histograms representing the object geometry.

Aware of the high storage necessary for their approach, of the computational overhead when comparing all spin images of two objects and of the existence of redundant information among close or symmetrically related spin images, Johnson and Herbert suggest performing compression on the set of an object spin images.

The spin images approach to 3D shape retrieval was later improved by de Alarcón *et al.* [37]. Instead of compression method proposed by Johnson and Herbert, they suggest data reduction by clustering the spin image set using a self organising map algorithm to group similar spin images, followed by a clustering algorithm. This way, the number of descriptor comparisons during matching is reduced, having only to check the spin image prototypes of each cluster. Moreover, authors introduce a three-level indexing schema based on artificial neural networks, which improves significantly the efficiency in matching query spin images against those stored in the database.

More recently, Assfalg *et al.* suggested a 3D shape retrieval method based on spin images, but using global features [10]. In their approach, spin images are used to derive a view-independent object description. First a set of spin images is created. Then, a descriptor is computed for each spin image in the set. Considering spin images as grey-scale images, these could be efficiently described by a low-dimensional region-based description

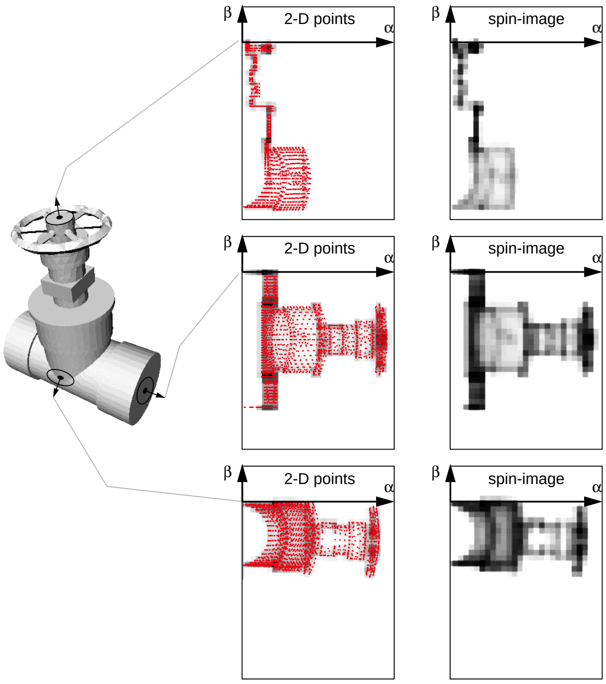

Figure 18: Spin-images for three oriented points on the surface of a model of a valve (Figure taken from [66]).

scheme from the content-based image retrieval domain. Next, authors apply a fuzzy clustering algorithm on the set of resulting image-based descriptors to reduce the number of spin images to a smaller number of prototypes, achieving a compact representation of the initial model, thus allowing efficient indexing and matching.

Silhouette descriptor Vranić presented, in his PhD thesis [121], a 3D shape descriptor based on 2D silhouettes. In this approach, a axis aligned 3D-object is projected on the coordinate hyperplanes, in order to generate three monochrome images as depicted in Figure 19. Next, author founds the outer contour of each image, approximating it by a polygonal line. Then, a commonly used technique on 2D shape description, the discrete Fourier transform, is used to represent the shape features in the spectral domain. The absolute values of the obtained coefficients are used to form the silhouette-based feature vector.

The PCA preprocessing stage necessary in this approach makes the silhouette descriptor pose and scale invariant. Additionally, author stresses that this approach is invariant to rotation due to an interesting property of the discrete Fourier transform, which makes the magnitudes of obtained coefficients (approximately) invariant with respect to rotation of the underlying silhouette image. Moreover, the magnitudes of obtained coefficients are invariant with respect to reflections of a mesh around the coordinate hyper-planes. How-

Figure 19: Silhouette images of an aeroplane model obtained by projecting the model on the coordinate hyper-planes (Figure taken from [121]).

ever, if it is not possible to determine a single unique contour of the silhouette image (for 3D models with holes or disjoint parts), only the longest contour is processed. Therefore, certain parts of 3D-objects might not be described.

Depth buffer descriptor When silhouette images of 3D-objects are created, all the information about shape is contained in contour points, because each interior point of the silhouette has the same attribute. Therefore, in order to capture 3D-shape characteristics Vranić considered other approaches for creating 2D images from 3D-objects [121]. He proposed another feature vector, which is obtained from six depth-buffer images formed using the faces of an appropriate cuboid region.

This approach starts the same way as the silhouette descriptor computation. The model is oriented and scaled into the canonical unit cube. Now, instead of three silhouettes, six grey scale images are used. Each image are rendered on each canonical cube face using a technique similar to the well-known z-buffer algorithm used in computer graphics, but instead of color, the attribute used to fill the interior of the image is the distance to the front clipping pane. After rendering the six images, the three-dimensional discrete Fourier transform is used to represent the image in the spectral domain instead of spacial domain. Figure 20 illustrates the extraction of the depth buffer-based shape descriptor. In the first row the depth-buffer images are formed using the canonical bounding cube. Lighter pixels indicate that distance between view plane and object is smaller than at darker pixels. In the second row are shown the coefficient magnitudes of the two-dimensional Fourier transform of the six depth-buffers.

Elevation descriptor Shih *et al.* proposed [98] a descriptor that shares the basic idea behind the depth buffer descriptor referred above. To compute the elevation descriptor, six different views of the 3D object are captured, corresponding to 2D projections of the object on the faces of the tightest bounding box circumscribing the 3D model. These views encode elevation maps describing the altitude information of the model relative to

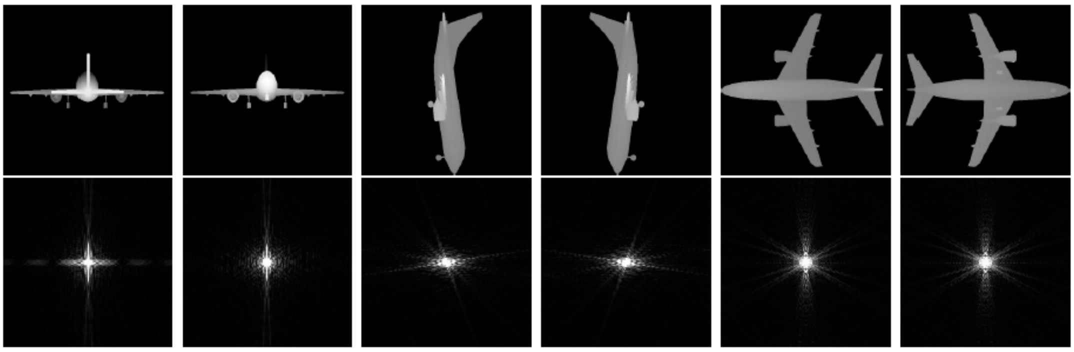

Figure 20: Extraction of the depth buffer-based shape descriptor (Figure taken from [121]).

the corresponding view plane and are represented as gray-scale images. Then, each one of these images is decomposed into a set of concentric circles around the centre point and the elevation descriptor is obtained by taking the difference between the altitude sums of two successive circles.

Unlike depth buffer approach, where authors rasterise the views using a technique similar to z-buffer algorithm, Shih *et al.* decompose the bounding box into a $64 \times 64 \times 64$ voxel grid and identifies the voxels intersected by the object surface as opaque. Then, to compute the elevation maps they simply have to found, for each cell within a 64×64 grid on the box face, the opaque voxel closer to that face and store in that cell the corresponding distance. Figure 21 illustrates the major steps of elevation descriptor computation. The 3D Jeep model is decomposed into a voxel grid and then elevation maps are computed for each k-face of enclosing box. Finally, the elevation descriptor is estimated by finding the distance within two successive concentric circles along the radius j. The complete elevation descriptor is obtained by concatenating the six partial descriptors.

Authors are aware that performing a full matching between two models will require a large number of elevation comparisons. Therefore, to reduce matching time Shih *et al.* provide an efficient similarity computation that finds the best match for a given query model.

Light field descriptor Following the idea that if two 3D models are similar, they also look familiar from all viewing points, Chen *et al.* [34] proposed a descriptor based on silhouettes from many different viewing directions. The light field descriptor encode one hundred orthogonal projections of an object, excluding symmetry, with both Zernike moments and Fourier Descriptors to produce feature vectors that describe the object.

To obtain the silhouettes of the object, authors define a camera system where twenty light field cameras are located on the vertices of a regular dodecahedron[11] centred at the

[11] A general dodecahedron is any polyhedron with twelve faces. A regular dodecahedron is a Platonic solid composed of twelve pentagonal faces, twenty vertices and thirty edges [126].

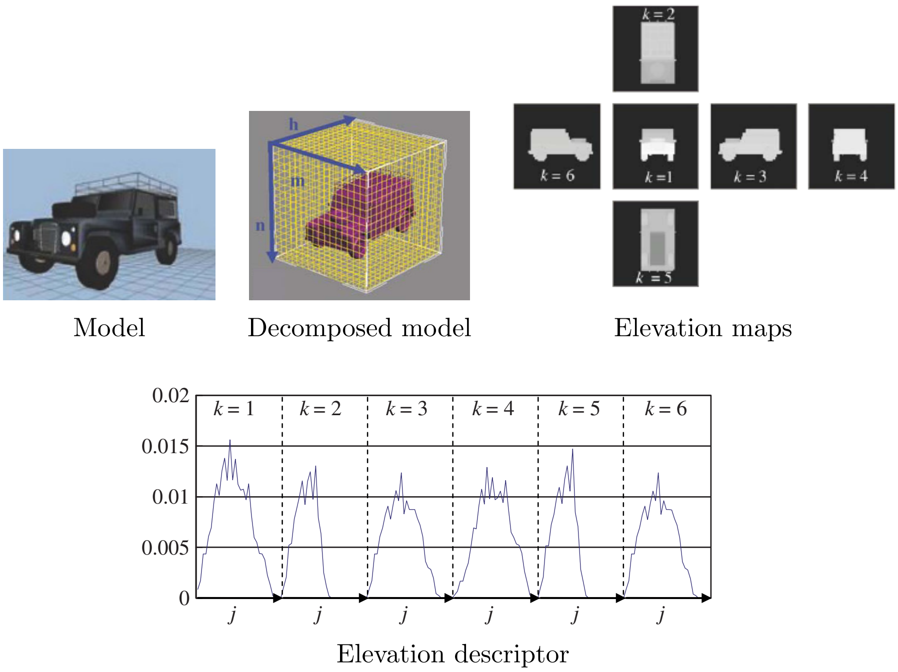

Figure 21: Computing elevation descriptor from a 3D model of a Jeep (Figures taken from [98] © 2006 Pattern Recognition Society).

object centre. The cameras viewing direction (view plane normal) is pointing towards the centre of the object and the camera up-vector is uniquely defined. This means that twenty different views, distributed uniformly over a 3D model, are captured. Therefore, ten different silhouettes are produced for the object, because silhouettes captured by cameras on opposite vertices of the dodecahedron are identical.

However since the cameras are placed on the vertices of the dodecahedron, the camera system must by rotated sixty times[12] in order to switch the cameras onto different vertices. This way the dissimilarity between light field descriptors of two objects are defined as the minimum of the sum of the distances between all corresponding image pairs when rotating one camera system relative to another.

To improve robustness against invariance a set of ten light field descriptors is applied to each 3D model, which lead to the one hundred orthogonal projections. These ten descriptors are created from different camera system orientations. Thus, the dissimilarity between two models is the minimum difference between all combinations light fields. Therefore, the similarity between two 3D models is obtained from the best one of 5,460 different rotations, which determines the final similarity distance.

[12] For a regular dodecahedron, each of the twenty vertices is connected by three edges, which results in sixty different rotations for one camera system.

2.5.5 Other methods

In the previous sections, we described several methods to represent 3D shapes. These were classified according to the technique behind the descriptor computation, may it be histograms, transforms, graphs or images. However there are some approaches that does not fit on these classifications, since the computation of these shape descriptors does not use any of this techniques or combines different methods to achieve a distinct result. In this section we will briefly describe a few of such techniques.

3D Zernike moments Novotni and Klein [86] advocate the usage of a specific kind of shape moment that has the advantage of capturing global information about the 3D shape and not requiring closed boundaries as boundary-based methods. The 3D Zernike descriptors are a projection of the function defining the object onto a set of orthonormal functions within the unit ball. They can be considered as the magnitudes of a set of orthogonal complex moments of the 3D shape and the natural extensions of spherical harmonics based descriptors.

Spherical moments Based on the concepts underneath moment-based method proposed by Saupe [97], Wei and Yuanjun introduced spherical moments as a shape comparison method for 3D model retrieval [125]. This method employs a multi-level spherical moments analysis approach relying on voxelization and spherical mapping of the 3D models. Authors claim that, despite the simplicity of this method, it outperforms in retrieval performance many previously proposed ones.

To compute the shape descriptor of a model, firstly a pose normalization step is done to align it into a canonical coordinate frame. Afterwards, them model is rasterised into a cubical voxel grid, then a series of homocentric spheres centred at the center of the voxel grid are used to produce a series of spherical images, by simply checking the intersection between trigonal pixels on the spheres surface and the object voxels and labelling the pixels accordingly, as illustrated in Figure 22. Finally moments of each sphere are computed and the moments belong to all spheres constitute the descriptor of the model. To estimate the similarity between models, Wei and Yuanjun suggest comparing feature vectors using Euclidean distance.

Reflective symmetry descriptor Kazhdan *et al.* proposed describing 3D models by measuring its amount of symmetry [69]. While such approach in two dimensions is quite simple, since it works by averaging an image against itself reflected along a line of symmetry, with 3D shapes the symmetry computation is more complex. Indeed, the reflective symmetry descriptor [71] of a 3D model is a collection of functions that measure the rotational and reflective symmetry with respect to every axis passing through its centre of mass. Kazhdan *et al.* present an efficient algorithm for computing the reflective

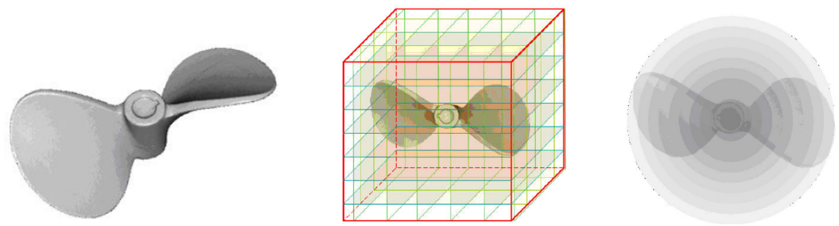

Model　Decomposed model　Syncopation of model

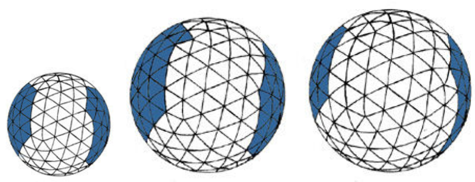

Spherical images with radius 5, 10 and 15.

Figure 22: Computing spherical images from a 3D model (Figures taken from [125]).

symmetry descriptor from a 3D voxel representation of a model, and show that,in addition, planar symmetries can be used for alignment of 3D meshes.

Moreover, authors suggest using this approach to improve existing shape descriptors with symmetry information. In particular, they describe the symmetry augmented descriptor, based on the spherical harmonic representation, described in Section 2.5.2. According to Kazhdan *et al.*, this augmented descriptor provides a highly discriminating representation of the shape.

Generalized shape distributions Liu *et al.* presented last year a combined approach to 3D shape description [78]. The Generalized Shape Distributions (GSD) takes advantage of both local and global shape signatures. The start by generating spin images, on meshes, producing a set of local shape descriptors, which are then clustered in what authors call ”words” in a ”dictionary” of local shapes. This way, they represent a global 3D shape as the spatial configuration of a set of specific local shapes by computing the distributions of the Euclidean distance of pairs of local shape clusters. Then, they store the descriptor in an indexing data structure to reduce the space complexity of the proposed shape descriptor.

Authors claim that their approach is robust to non-trivial shape occlusions and deformations and is more discriminative than a simple collection of local shape signatures since the spatial layouts of a global shape are explicitly computed. The robustness to shape occlusions and deformations comes from the fact that there are statistically a large number of chances that some local shape signatures and their spatial layouts are unchanged and users can easily identify those unchanged parts. Indeed, their preliminary experiments show the effectiveness of the proposed technique for shape comparison and analysis.

In the present PhD research work we plan to research for a solution somehow similar to the GSD proposed by Liu *et al.*. However, we intend to extend this kind of approach, by incorporating topological information for each shape and by using a shape thesaurus, as described below in Section 4.

2.5.6 Discussion on Shape Descriptors

As we have seen above, there are several three dimensional shape descriptors. In a variety of approaches, these algorithms for feature extraction have one thing in common. All aim at producing a feature vector that provides good discriminating power while keeping the time and space consumption relatively low while computing the descriptor. Although some of the techniques we described in this chapter are outdated, the technique proposed in such approaches remain useful. Indeed, many recent algorithms are basically an evolution of older methods. A summary of all analysed methods for shape description is available in Table 2.

To provide a organised view of this research area, we divided the 3D shape description algorithms into five distinct categories. The histogram based approaches simply accounts one or more shape features and constructs histograms with them. These histograms are then used to estimate the corresponding set of feature vectors. The transform-based approaches rely on mathematical transformations to switch from the spatial domain to a more suitable one and compute from there a shape descriptor. For instance by applying the Fourier transform or the spherical harmonics transform. The graph-based approaches computes a graph representing the topology of the model and then uses one or a combination of several histogram-based or transform-based descriptors to code the shape features for each node of the graph. The image-based approaches rely on multiple 2D representations of the shape to compute the descriptor. Finally, there are a few algorithms that do not fit on any of these four categories, which we classify as other.

Depending on the purposes and scope of the retrieval system, some shape descriptors can perform better than other. As a matter of fact, there are no shape descriptor that is clearly better than all the others. Instead, some are best suited for some kind of 3D models or for specific needs of the retrieval system. While, in some cases the major concern is the effectiveness of the shape descriptor, in others the most important factor could be the efficiency of shape description techniques.

The effectiveness of a shape descriptor indicates the amount of shape information it is able to represent. More effective shape descriptors store more information about the shape. On the other hand, the efficiency of a shape description technique regards on the time and space necessary to compute and store the resulting feature vector. Larger and more complex shape descriptors usually led to slower classification and retrieval. In a near future we will evaluate experimentally several shape descriptors, producing a practical comparison that will allow us to select the most appropriate for our work.

	Descriptor	Feature	Refs.
Histogram-based	Cord and Angle Histograms	Distribution of length and angles of cord rays	[90, 91]
	Color Distribution	Voxel colors computed from shape texture	[90]
	Curvature Histogram	Principal curvatures at each face of the mesh	[73, 119]
	Shape Distribution	Collection of shape functions measuring several shape features using randomly selected surface points.	[88]
	Modified Shape Distribution	Shape functions measuring a pair of features using a quasi-random point selection.	[87]
	Shape Histograms	Surface points in cells of decomposed model.	[9]
	3D Shape Contexts	Distribution of sampled points relative to each other.	[74]
	Extended Gaussian Images	Variation of surface area with surface orientation.	[68]
	3D Hough transform	Variation of surface area with surface orientation.	[132]
	Shape Spectrum	Area of surface object with respect to shape curvature.	[39, 131]
	Density-based Shape Descriptor	Variation of surface area with surface orientation.	[6, 7, 5]
Transform-based	Voxel 3D Fourier Transform	Proportion of total surface area inside each cell of voxelised model.	[123]
	Distance and radial cosine transform	Binary and continuous voxel-based distance from a point to the object surface.	[42]
	Spherical Harmonics Transform	Distance from the object surface to the surface of enclosing sphere.	[97]
	Rotation Invariant Spherical Harmonics	Spherical functions of concentric sphere based on voxelised model.	[70]
	Planar-Reflective Symmetry transform	Symmetry of a shape with respect to all planes through its bounding volume.	[92]
Graph-based	Multi-resolution Reeb Graphs	Reeb graphs at multiple levels of resolution of a function over the surface	[57]
	Size Graphs	A centreline skeleton with nodes labelled with local geometric properties.	[21]
	Skeletal Graphs	A directed acyclic graph associated with a set of geometric features and a signature vector.	[107]
Image-based	Spin Images	Surface points projected on planes defined by oriented points on model surface	[67, 10]
	Silhouette Descriptor	Object projections on coordinate hyperplanes.	[121]
	Depth Buffer	Mappping of surface distances into the six faces of the object bounding cube.	[121]
	Lightfield Descriptor	Object silhouettes captured by cameras on the vertices of a dodecahedron.	[34]
	Elevation Descriptor	Decomposition in concentric circles of surface elevation with respect to the six faces of bounding cube.	[98]
Other	3D Zernike Moments	Magnitudes of a set of orthogonal complex moments of the object.	[86]
	Spherical Moments	Moments of spheres mapping voxelised models.	[125]
	Reflective Symmetry Descriptor	Collection of functions measuring rotational and reflective symmetry with respect to every axes passing on barycenter.	[69]
	Generalised Shape Distributions	Clustering of spin images.	[78]

Table 2: Summary of 3D shape descriptors.

2.6 Query and Matching

As we have already mentioned, 3D shape descriptor computation is crucial on a shape based retrieval system. This explains the large amount of work developed in this particular topic, which we surveyed in Section 2.5. However, effective shape representation is not the only challenge to overcome when developing 3D model retrieval systems. Another important part of a 3D shape retrieval solution, as of any content-based retrieval system, is the query and matching processes. Descriptors extracted from objects are usually represented as feature vectors on a multidimensional space. This is truth not only for 3D models, but also for other data types, such as images or music. Effective content-based retrieval of information indexed on multidimensional spaces depends greatly of query and matching techniques.

2.6.1 Query types

Regardless of the data stored on the database or used as a query, several authors agree on a basic set of different types of queries [52, 56, 31] for multidimensional spatial data. Following the enumeration of query types presented by these authors, we will briefly describe the four most common categories of queries.

Exact match query This type of query aims on finding all objects that have exactly the same spatial extent as the spatial query object. Indeed, exact match queries are only of moderate interest in content-based retrieval , and, when applied on content-based retrieval, are usually based on metadata, managed by a traditional database management system. An example of such query is "find all alloy 17 inches-radius wheels".

Range search query Content-based retrieval approaches rely mostly on retrieval-by-similarity queries. One way to accomplish this is by performing a range search query, *i.e.* find all objects that are within a given range, usually a hyperrectangle [13]. Such query can be specified as finding objects in the multidimensional space that have at least one common point with a query volume in that space. For instance, the following sentence is a range search query: "find all 3D MRI models showing a tumor of size between vol_{min} and vol_{max}".

***k*-nearest-neighbor search** Another way to perform retrieval-by-similarity is to search for a given number of objects similar to the query. Measuring this similarity among objects is an important issue on retrieval that we discuss briefly on Section 2.6.3. Usually such

[13] A hyperrectangle, also called orthotope, is parallelotope whose edges are all mutually perpendicular. Indeed, a hyperrectangle is a generalization of the rectangle to higher dimensions [127]. For instance, the cuboid is a 3-orthotope.

similarity is seen as a distance between object representations in multidimensional space. Thus, a k-nearest-neighbor query is specified by finding k objects with the shorter distance to the given query. An example of such query is given by the sentence: "Find the twenty tumors most similar to a specified example".

Within-distance (or α-cut) This third way to query by content is quite similar to the previous one. The main difference is that in this type of query instead of a pre-defined number of results, the search must return all the results within a given distance to the query. Thus, the objective is to find all objects with a similarity score better than α with respect to a query. This means to find all objects whose representation in multidimensional space has a distance smaller than a threshold from the query representation in that space. An α-cut query can be specified, for instance, by the sentence:"Find all the 3D MRI models containing tumours having a given similarity with respect to an example provided".

2.6.2 Matching properties

The performance of a content-based retrieval system depends on a great extent of the quality of the results produced by matching techniques. To characterise the query processing mechanism Vittorio Castelli [31] defined a set of three properties, which we will briefly present in the following paragraphs.

Exhaustiveness A matching algorithm is exhaustive if it retrieves all the database items satisfying the query. A database item that satisfies the query and does not belong to the result set is called a miss. Non-exhaustive range-query processing fails to return objects that lie within the query range. Non-exhaustive α-cut query processing fails to return points that are closer than α to the query template. Non-exhaustive k-nearest-neighbor query processing either returns fewer than k results, or returns results that are not correct.

Correctness Matching is correct if all the returned items satisfy the query.A database item that belongs to the result set and does not satisfy the query is called a false hit. Non-correct range-query processing returns points outside the specified range. Non-correct α cut-query processing returns points that are farther than α from the template. Non-correct k-nearest-neighbor query processing miss some of the desired results, and therefore is also non-exhaustive.

Determinism Matching is deterministic if it returns the same results every time a query is issued, and for every construction of the index. It is common to have non-deterministic range, α-cut and k-nearest-neighbor queries implemented in content-based retrieval systems. This happens because when designing such mechanisms there are another factor to consider. They must be time efficient.

Indeed, to achieve time efficient solutions, is common to relax some of these properties, or even the three of them. For instance, by relaxing exhaustiveness, alone, we are allowing misses, but not false hits, and determinism is retained.

2.6.3 Similarity measuring

To achieve effective matching, resemblance between objects, *i.e.* feature vectors, must be measured. Indeed, nearest-neighbor queries rely on the definition of a similarity function, while α cut queries rely on a scoring function. Usually, in both cases similarity measuring among objects is done through distance functions that estimate the distance between the closest points of their representations in the multidimensional space. Even in range-queries, a distance function is used to sort the objects within the given range by similarity. Due to the importance of such functions, we will briefly explain them in the following paragraphs.

Following the approach presented by Tangelder and Veltkamp [114], we assume that S is a set of all points representing objects in the n-multidimensional space. Then, a metric distance function d on that set S can be formally defined by a non-negative value function $d : S \times S \rightarrow \mathbb{R}^+ \cup 0$ that should satisfy the following properties:

- **self-identity:** $\forall_{x \in S}, d(x, x) = 0$;
- **positivity:** $\forall_{x,y \in S}, x \neq y \Leftrightarrow d(x, y) > 0$;
- **symmetry:** $\forall_{x,y \in S}, d(x, y) = d(y, x)$;
- **triangle inequality:** $\forall_{x,y,z \in S}, d(x, y) + d(y, z) \geq d(x, z)$.

As a matter of fact, an extensive theory lays behind distance functions and there are a large set of possible approaches to measure the distance between two points in a multidimensional space. However, this issue is not within the main scope of our work. Therefore, we will only focus on a short list of the most useful distance functions. These are Euclidean, Manhattan and Chebychev distances, which are particular cases of the more general family of Minkowsky distances, and the Mahalanobis distance.

The Minkowsky distance of degree p, also called p-distance, between two points in a n-dimensional space is given by:

$$d_p = (\sum_{i=0}^{n} (x_i - y_i)^p)^{\frac{1}{p}}. \tag{13}$$

Indeed, this general equation is not applied in practice. Instead, the parameter p is fixed in a few values commonly used. Thus, the Minkowsky distance of degree 1 ($p = 1$) is

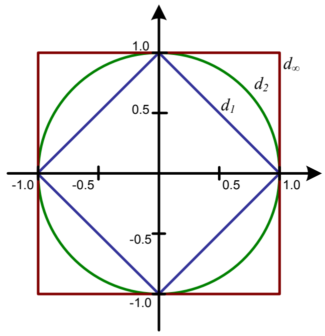

Figure 23: The unit sphere under Manhattan (d_1), Euclidean (d_2) and Chebychev (d_∞) distances.

called Manhattan distance, the usual Euclidean distance is the distance of degree 2 ($p = 2$) and with $p = \infty$ we obtain the Chebychev distance. The difference among these three distance functions is shown in Figure 23, where we depict the unit sphere in each one of these metric spaces. The corresponding distance functions are defined as:

- **Manhattan distance** $d_1(x, y) = \sum_{i=0}^{n}(x_i - y_i)$;
- **Euclidean distance** $d_2(x, y) = \sqrt{\sum_{i=0}^{n}(x_i - y_i)^2}$;
- **Chebychev distance** $d_\infty(x, y) = max_{i=0..n}(x_i - y_i)$.

These three Minkowsky distances are simple, fast to compute and can be generically used. However, in some cases the results obtained by this measurements do not fulfil the needs of retrieval solutions. Thus, to solve problems caused by poorly scaled or highly correlated coefficients of a vector (descriptor), is often used the Mahalanobis [81] distance. It is a computationally expensive generalisation of the Euclidean distance widely used in cluster analysis and other classification techniques to measure the distance between probability distributions. This measurement is based on correlations between variables by which different patterns can be identified and analysed. It is a useful way of determining similarity of an unknown sample set to a known one. It differs from Euclidean distance in that it takes into account the correlations of the data set and is scale-invariant. The Mahalanobis distance is defined in terms of a covariance matrix C, which measures a tendency to vary between two features, as given by the function d_M:

- **Mahalanobis distance** $d_M(x, y) = det|C|^{\frac{1}{d}}(x - y)^T C^{-1}(x - y)$.

There are some other distance functions that are often used in content based retrieval solutions, but we prefer not to mention them all here. Instead, we refer our readers to the comprehensive explanation of similarity measures for retrieval published by Castelli [31] or to a theoretical description of distance function presented by Hervé Abdi [2].

2.7 Benchmarking 3D shape retrieval

Like in other information retrieval research areas, effective testing and comparison of different techniques in multimedia information retrieval requires the existence of widely accepted evaluation frameworks. One of the most complete evaluation projects in this area has been the TRECVid [104]. Focused on information retrieval from digital video, TRECVid provides a large video test collection, uniform scoring procedures, and a forum for organisations interested in comparing their results. Another example is the Music Information Retrieval Evaluation eXchange (MIREX), a formal framework for the scientific evaluation of the many different techniques being employed by researchers in the domains of Music Information Retrieval and Music Digital Libraries [40].

2.7.1 Shape benchmarks

In the three-dimensional models domain, the Princeton Shape Benchmark (PSB), deployed by Thomas Funkhouser team [99], has become the standard and is being widely used for evaluating various representation methods and shape retrieval techniques. The PSB provides a carefully compiled repository of around 1,800 models collected from the web and software tools for evaluating shape-based retrieval and analysis algorithms. The models in this collection are real-world objects such as vehicles, buildings, animals or plants. The popularity of PSB lead to its use in the initial version of the international 3D Shape Retrieval Contest (SHREC), described in Section 2.2.

Another shape repository, concerning evaluation of 3D model retrieval algorithms, was released recently. The AIM@SHAPE network of excellence, introduced in Section 2.1, released their shape repository which makes several 3D models available for researchers to compare shape matching algorithms. It is a shared repository populated with a collection of digital shapes and an integral part of the framework of tools and services for modeling, processing and interpreting digital shapes, developed within the AIM@SHAPE project.

However, despite the existing work on shape repositories, there has been limited work in developing domain dependent benchmark databases for 3D shape searching. Indeed, the most popular benchmarks do not include model classes from specialised application domains, such as molecular biology or CAD engineering. To overcome this, the PRECISE group at Purdue proposed [65] a benchmark database for evaluating shape-based search methods relevant to the mechanical engineering domain. They developed a publicly available Engineering Shape Benchmark (ESB) for comparing various shape-based search algorithms. The ESB includes a set of 867 models along with associated images and a classification schema.

Having also identified the problems referred above, Bespalov *et al.* presented benchmark datasets to assess the relevance of existing 3D shape retrieval techniques for engineering problems [16]. They proposed several distinctive repositories for evaluating techniques for automated classification and retrieval of CAD objects. These collections includes sets

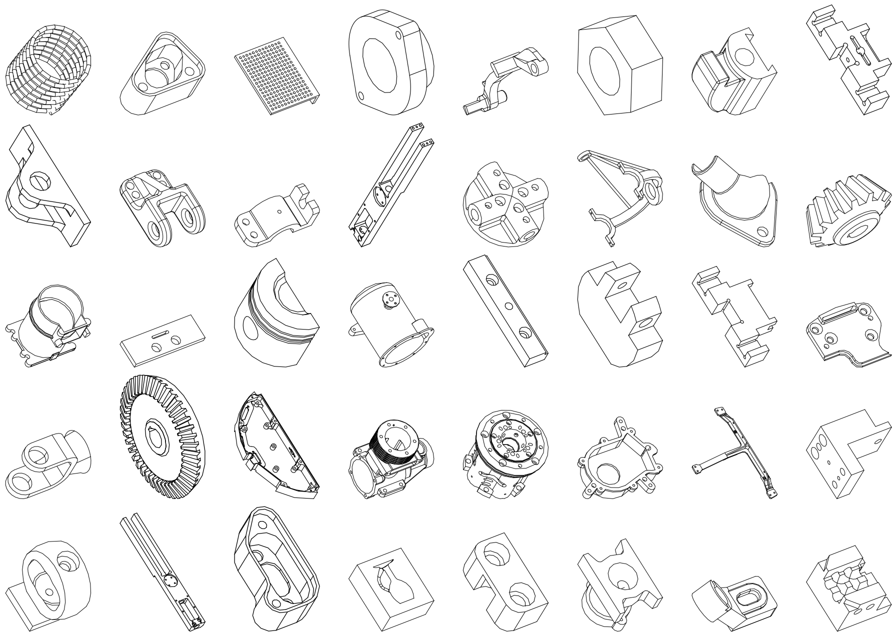

Figure 24: Some models from National Design Repository (Taken from [16] © 2005 ACM).

of CAD primitives, industrial CAD models and LEGO® models. The models on these datasets were extracted from the National Design Repository [94, 113, 95] and manually classified by their appearance, manufacturing process and function, since these are typical classification schemes for CAD models. Figure 24 depicts a few examples of such models.

Clark *et al.* published a novel and interesting work regarding shape signature benchmarking [35]. They state that research on shape retrieval has focused on methods for describing 3D shapes in order to compute properties such as the "geometric similarity". However, authors argue that such properties are ill-defined because little work on assessing how closely these measures match human perceptions has been reported. Thus, they performed an experiment comparing the part families identified by both human subjects and three published shape signatures, using a subset of the National Design Repository.

From this experiment, Clark *et al.* reported two main conclusions. First, they conclude that the results of the human perception test suggest that the used dataset contains a large number of perceptually similar objects. Second, results obtained when benchmarking shape signatures against human perception indicates a low false-positives rate. Moreover, the false-negative rate varied almost linearly with the amount of similarity perceived by human subjects. Such experiment opened a new field of study, where the importance of human perception in the 3D shape retrieval should be seriously taken into account.

2.7.2 Comparative studies

As we have already shown, there are many different techniques to describe three-dimensional shapes, each one with its own strengths and drawbacks. Due to differences between existing approaches to shape description, comparing these is a difficult task. Nevertheless, to assess the effectiveness of shape description methods several comparative studies has been carried out recently. In the following paragraphs, we will briefly describe three of these studies published last year.

Bustos *et al.* presented in 2006 [28] an experimental effectiveness comparison of methods for 3D similarity search. In this study, authors surveyed some approaches to 3D shape retrieval and presented an extensive experimental effectiveness and efficiency evaluation of these techniques, using several 3D collections. Among a total of sixteen shape descriptors, they studied the rotation invariant spherical harmonics descriptor, shape distribution descriptor, shape spectrum descriptor, silhouette descriptor and depth-buffer descriptor.

After comparing the computational complexity of the analysed descriptors, they discuss its retrieval performance. Authors concluded that there is a number of descriptors that have good database-average effectiveness and work well in general, while others work better with some specific model classes but have poorer results on generic models. Finally, authors argue that most descriptors can be considered robust, as they can effectively retrieve similar objects with different level of detail.

In a distinct study, Alberto del Bimbo and Pietro Pala performed a comparative analysis of a few different solutions for description and retrieval by similarity of 3D models [24]. For this study, authors selected descriptors that are representative of the principal classes of approaches. From the class of histogram-based descriptors, authors selected the curvature histograms and the shape functions, while Spin-image signatures and light field descriptors were used to represent the image-based approaches. Authors also included on the comparative study the geometric moments [43] used by Elad *et al.* to describe 3D shapes.

Bimbo and Pala focused their experimental analysis on comparing the four methods referred above according to their robustness to deformations and their ability to capture the structural complexity of 3D objects, as well as the resolution at which models are considered. To that end, authors used two different shape databases. The *Art-Model* database composed by around three hundred high-resolution models from miscellaneous sources was used to test the robustness to geometric deformations. To that end, these models were subjected to special deformations, as illustrated by the examples depicted in Figure 25, generating an extra set of deformed versions of each original model. The other collection used in this comparative study used the well known Princeton shape database.

From the results obtained in this experiment, authors could achieve an extensive set of conclusions. Particularly, they concluded that the light fields descriptor and spin image signatures have superior capability to capture the structural peculiarities of the mod-

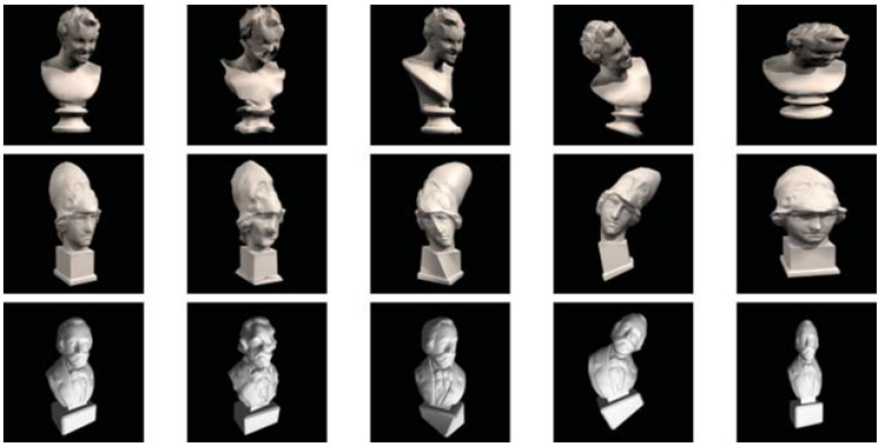

Figure 25: Three busts from *Art-Models* database and corresponding deformed models (Taken from [24] © 2006 ACM).

els, with highest insensitivity noise provided by the light fields representation. Moreover, Bimbo and Pala claim that geometric moments are not able to capture salient and discriminating features of 3D objects and the histogram-based approaches never provide better retrieval performance than the other solutions. However, the computational complexity were not considered in this study, which will eventually uncover additional drawbacks of the image-based descriptors when compared to the histogram-based approaches.

Together with the ESB, Jyanti *et al.* presented a comparative study between twelve different shape descriptors evaluating the effectiveness of these representations on the mechanical engineering domain [65]. Within the set of tested descriptors are spherical harmonics, shape distributions, shape histogram and light field descriptors, among others. In these experiment, authors performed an unusual test. They compared the results retrieved by using every shape descriptor against the random retrieval method. As expected, all shape representation methods outperformed the random retrieval.

Additionally, in a paper published this year, Bustos *et al.* present two recently proposed approaches to shape description and discuss methods for benchmarking the 3D retrieval systems' qualitative performance [26]. Indeed, they suggest as best options for shape retrieval evaluation the Princeton Shape Benchmark and the actual version of the benchmark used in the SHREC 3D retrieval contest.

2.8 Content-based Retrieval of 3D models

During recent years, several 3D shape search engines have been introduced. One of the earliest of such systems was proposed by Paquet and Rioux in 1997. Nefertiti [90] is the first well documented query by content software for three-dimensional model databases. It incorporates a set of retrieval algorithms that allows database searches by scale, shape, color or any combination of these parameters.

Later, in 2001, Thomas Funkhouser and his team released the Princeton 3D model search engine [50]. This system is now the best known solution for shape retrieval, indexing more that thirty six thousand models. Its authors claim that they have developed the search engine to be the "Google^{TM} for 3D models" [49].

Unlike the Princeton team, whose search engines aims on generic 3D models, the PRECISE group at Purdue University developed a search engine for a specific domain [79]. The 3D Engineering Shape Search system integrates a set of existing shape description techniques to compute the feature vectors of a model. This search engine incorporates a 3D interface that allows users to submit a shape as a query, to select the feature vectors that will be used for shape representation and to search the database by browsing.

Starting from the idea that if two 3D models are similar they also look similar from all viewing angles, Chen *et al.* introduced a retrieval system [34] based on the light field descriptor. The 3D Model Retrieval System from National Taiwan University is available on the web and its database contains more than ten thousand publicly available 3D generic models. A simple interface is integrated in this search engine, allowing users to retrieve 3D models by drawing 2D silhouettes.

To serve as a proof-of-concept to methods and tools for content-based search for 3D-mesh models proposed during his PhD research [121], Vranić deployed a web-based retrieval system for 3D models. The Content-based Classification of 3D-models by Capturing spatial Characteristics (CCCC) 3D search engine uses a set of model databases, including the Princeton Shape Benchmark test and training databases, providing around three thousand classified objects.

Based on two distinct approaches to description and matching of 3D objects, Assfalg *et al.* developed a content based retrieval system for 3D shapes [11]. Using a curvature map of the shape surface, authors propose, in one approach subdividing the map into a grid of rectangular tiles and then use these to compute a shape histogram, while in the other approach, the map is segmented into regions of homogeneous curvature, and regions are described with weighted walkthroughs. This search engine allows users to perform queries by example through a web interface on a database of three-dimensional models.

More recently, in 2007, researchers from the FOX-MIIRE group released a on-line search engine for 3D content [47]. Their search engine implements the adaptive views clustering technique, a method proposed by the authors to index 3D models based on two-dimensional views. Besides the good retrieval results offered by the FOX-MIIRE search engine, it has a unique feature when comparing with previous approaches. This search engine is the first that accepts 3D-Models retrieval from photos [46] and can be reached through a mobile device. Figure 26 depicts both the standard and the mobile device interfaces of the FOX-MIIRE search engine. Indeed, the idea of incorporating a 3D model retrieval system in a mobile device was proposed by Suzuki *et al.* [110]. They developed

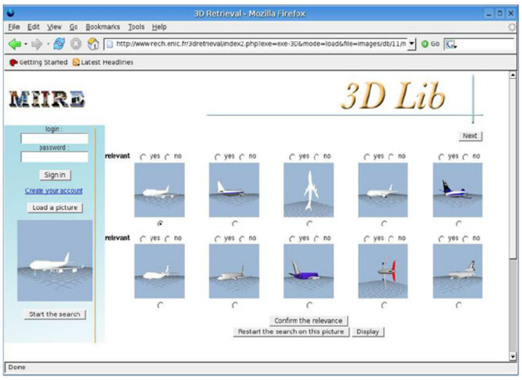

(a)

(b)

Figure 26: MIIRE search engine on PC (a) and PDA (b).

an experimental 3D shape retrieval system for cellular phones where users can search for a model similar to a given example.

Additionally, Qin Lv *et al.* introduced a toolkit to support the construction of content-based similarity search systems [80]. The Ferret toolkit is a content-based similarity search engine for generic, multi-feature object representations. It was designed to solve the similarity search problem in high-dimensional spaces. Indeed, this solution can be used to successfully construct content-based similarity search systems for audio recordings, digital images, 3D shape models and genomic microarray data. Regarding 3D object retrieval, authors collaborated with the PRECISE group to create a search system based on Ferret toolkit. According to Qin Lv *et al.* this system was developed in a few hours by adapting existing solutions for segmentation, feature extraction and similarity measurement into their toolkit.

2.9 Retrieval using Partial Matching

As we have shown in previous sections, there is plentiful work on 3D shape analysis and retrieval. However, 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 paragraphs we will present some recent research work tackling the partial matching of three dimensional models.

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.

2.9.1 Partial matching at NIME

Suzuki *et al.* proposed [111, 109] 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 sub-components and then the 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 [33, 32, 134]. 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 [108]. 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 [112]. One of the decomposition 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 27. 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

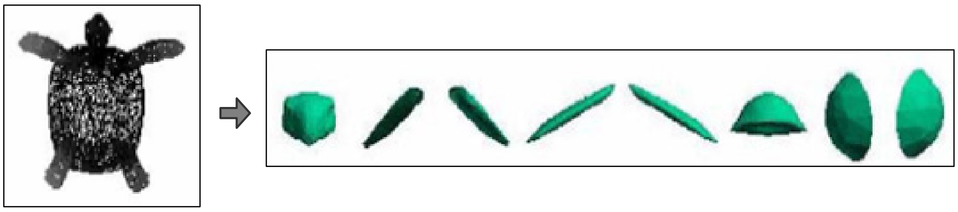

Figure 27: Example of the decomposition of a 3D model of a turtle (Figures taken from [112]) © 2006 IEEE.

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.

2.9.2 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 [23] 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 [82] 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) [105] 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 [57], 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 rotation invariant spherical descriptor used in this approach has been defined by Kazhdan *et al.* in [70] and is briefly described in Section 2.5.2. Addi-

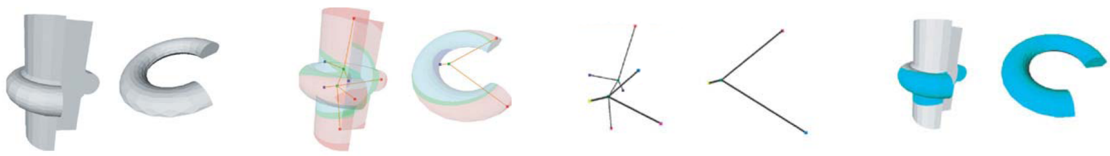

Figure 28: Sub-part correspondence of two mechanical parts (Figures taken from [23] © 2006 Elsevier Ltd.).

tionally, 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 [83] 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 28. 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.

2.9.3 Scale-space feature extraction for partial matching

Focusing on mechanical CAD models, Bespalov *et al.* [18] 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 [14] of a 3D mesh model [85]. The discrete scale-space of a three dimensional model is

[14]Scale-space is widely used by the computer vision and image processing communities for handling image structures at different scales. With this framework, the fine-scale features are iteratively suppressed while the level in the scale-space representation increases. The idea behind this theory is that objects are composed by different structures at different scales. For instance, it is appropriate to represent a dog at

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 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 [55] and the Canny edge detector algorithm [30] respectively. Figure 29 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 [85] 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.

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

---

the scale of meters, but not the hair of its fur or the molecules that compose its skin, which should be represented at much finer scales. Therefore, the scale-space approach consider multiple descriptions for an object at different scales to be able to capture its complete description.

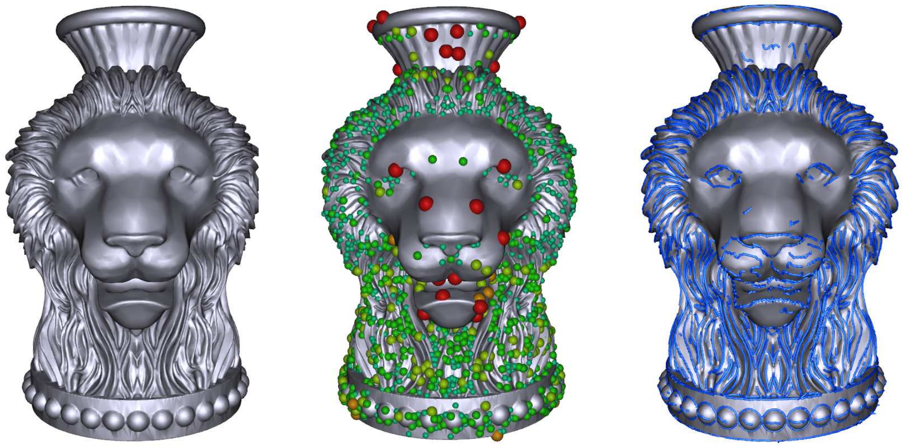

Original model Extracted corners Extracted edges

Figure 29: Combined set of scale-dependent corners and edges extracted from polygonal model (Figures taken from [85] © 2006 IEEE).

function, CAD mechanical models are decomposed and the resulting combined feature set is used for partial matching of 3D models. Figure 30 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.

2.9.4 Salient geometric features for partial matching

In their approach to partial matching, Ran Gal and Daniel Cohen-Or [54] 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,

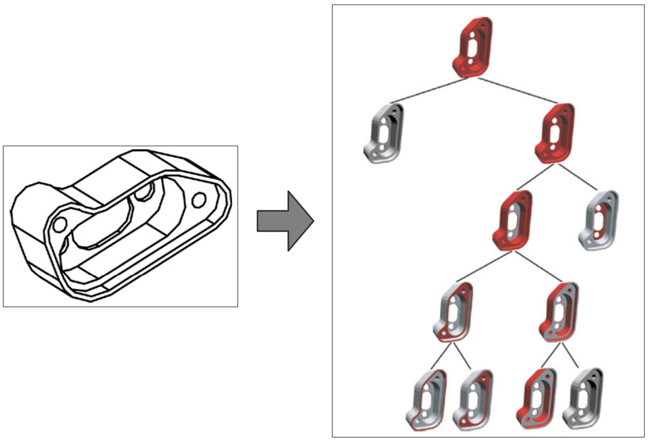

Figure 30: Scale-space decomposition of a mechanical part (Figures taken from [18] © 2006 Elsevier Ltd.).

such option is grounded on previous work by Hoffman and Singh [58]. They have found that human vision defines boundaries along negative 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.

Authors 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. The saliency grade S of a descriptor cluster F is estimated through Equation 14:

$$S = \sum_{d \in F} W_1 Area(d) Curv(d)^3 + W_2 N(F) Var(F). \tag{14}$$

While the term $Area(d)Curv(d)$ expresses the saliency of the region represented by descriptor d, the term $N(F)Var(F)$ expresses the degree of relevance of cluster F, with $N(F)$ representing the number of local minimum and maximum curvatures in the cluster and $Var(F)$ representing the curvature variance of the cluster. The weights W_1 and W_2 can be used to fine-tune the saliency grade function, but authors claim that no manual tuning is required and that using $W_1 = W_2 = 0.5$ should produce good results according to their tests.

After estimating all clusters, authors select from these a set of clusters with higher values of S 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 31 illustrates the result of applying this method to four different

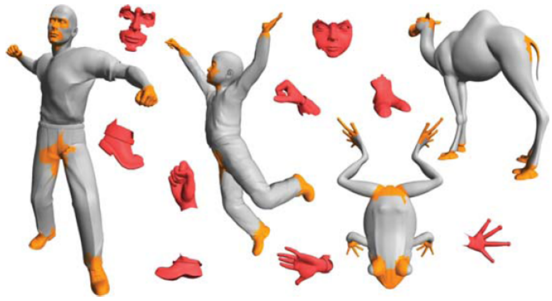

Figure 31: Salient geometric features from four models and corresponding individual subparts (Figure taken from [54] © 2006 ACM).

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[15]. 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.

2.9.5 Partial matching at Princeton 3D shape retrieval and analysis group

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 [54], Shilane and Funkhouser [100] 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.

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.

[15]Geometric hashing is an highly efficient technique with low polynomial complexity developed for matching geometric features against a database of such features [75]. 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 [128].

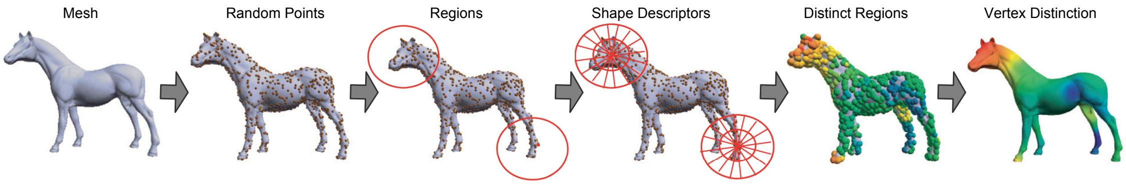

Figure 32: Selecting distinctive regions of an object (Figures taken from [101] © 2007 ACM).

In their method, Shilane and Funkhouser [101] 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[16] [70]. Figure 32 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 [51]. 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 [38]. 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.

[16]Rotation invariant spherical harmonics were briefly described in Section 2.5.2.

3 Thesis Statement

Now that we have briefly analysed the background and state-of-the-art on 3D model retrieval with partial matching, we can focus on the topic we plan to explore. As we have shown in previous sections, during the last decade most research on three dimensional shape retrieval has focused mainly on global matching, sometimes using local features, but always by measuring similarity between complete models [115]. Comparing subparts of models is a harder challenge than global matching, since these are not predefined and can be any subshape of a larger object with any orientation and scale.

Indeed, to accomplish partial matching it is necessary to identify and isolate sub-parts in models before 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 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 The ultimate goal

Ignoring for a while the problems ahead, we present an overall description of what we think a content based model retrieval system should be. In our opinion, an ideal three-dimensional retrieval system should be able to decompose models into its components, eventually in multiple scales, and classify all of them correctly. After an effective object decomposition and subpart classification, it should be possible to submit a query to the system, representing a subpart of existing objects in the database, and it will return a list of models containing similar parts.

Despite no such fully functional system had been devised yet, we illustrate in Figure 33 an example of a 3D model search with a partial query. In this case a screw was submitted as a query to a retrieval system with partial matching. As a result of this query, models that has screws are returned, including, of course, objects that are screws. Achieving such solution is indeed the ultimate goal of most researchers working on 3D shape retrieval. It will constitute the equivalent of the Google^{TM} search engine but for 3D model collections.

During our research work we intend to develop novel techniques that will provide some of the functionality necessary to build a retrieval system with effective support to partial matching. Assuming that models in the database are not previously decomposed by human operators, devising a system that behaves like the example given above is a complex task. Probably we would not be able to devise the complete solution for this problem, but we hope to provide some relevant contributions to it.

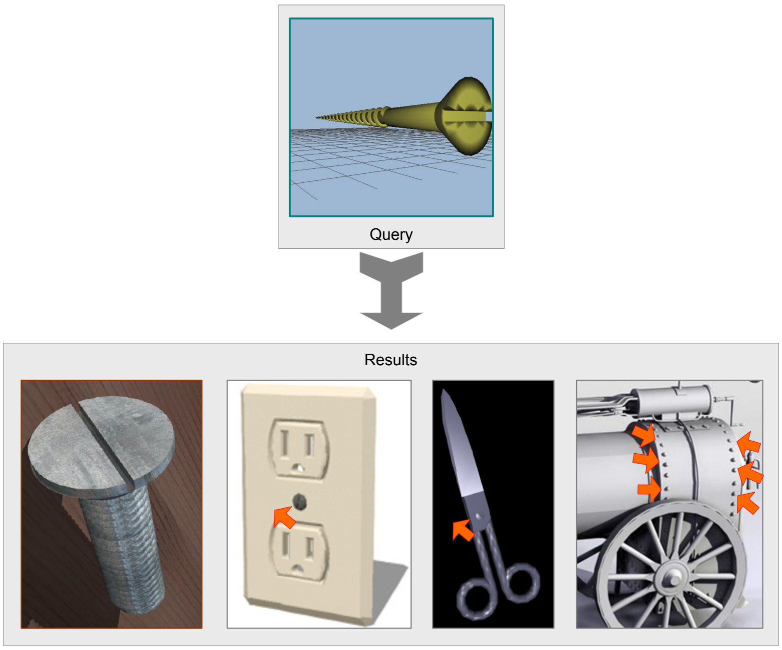

Figure 33: Example of partial matching. Users provide an element and the system returns objects containing it.

3.2 Problem statement

Indeed, a currently open challenge is how to perform retrieval with partial matching on large collections of three dimensional models using all of its features. Although 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 2.9. 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.

3.3 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. 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 3D shapes.

Besides, we plan to adapt from text information retrieval some other well-known techniques because these have already prove their efficiency handling large collections of documents, such as **tf-idf** [17] weight measure or data compression techniques. Since we will index all subparts of 3D models, it is necessary to rank them according to its importance for retrieval. Indeed, some authors have already proposed some geometry-based methods to select distinctive regions of 3D surfaces regarding the whole collection [101]. However, we will extend their approach and combine it with concepts form information retrieval. We will use the **tf-idf** to classify the relevance of distinctive components of three-dimensional models. Moreover, since we will use inverted files for indexing the shapes we will take advantages of existing inverted file compression techniques to optimise our solution.

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 devise 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.

[17] The term frequency and inverse document frequency (**tf-idf**) is a quantitative measure commonly used in information retrieval and text mining to determine the relevance of a word to a document in a collection, regarding to the contents of the whole collection.

3.4 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 for model segmentation that allows proper description based on a thesaurus of shapes;

2. Identify a method or combination of methods for shape matching and indexing that supports a thesaurus-based approach on 3d shape retrieval;

3. Develop and evaluate a thesaurus-based framework for 3D shape retrieval with partial matching.

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 identify clearly 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 segmentation and topology extraction 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 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.

After overcoming the challenge of geometrical shape description we will focus on three-dimensional shape segmentation. First, we must devise a suitable solution to decompose 3D models unambiguously and efficiently. Eventually, this can be done through extending existing segmentation techniques. As we are currently doing with shape descriptor computation, we will compare several segmentation approaches and study carefully the results to determine the best solution for our purposes. Anyhow, during preliminary studies on 3D model decomposition we identified the work developed by Tierny *et al.* [116] as a promising approach to skeleton driven segmentation.

Additionally to shape geometry description and model segmentation, we will focus on finding a method to represent 3D model topology that suits our needs of efficient indexing and matching. To that end, we will probably use a graph-based representation. Such representation usually encode both topology and geometry of the whole object and its components, which is exactly what we need. Within this subject, the research group lead by Bianca Falcidieno at IMATI-Ge have been investigating techniques for structural model geometry description [105]. We plan to follow closely their work and eventually integrate some of their findings in our solution.

Finally, we will focus on the core investigation challenge of this work. Having already devised mechanisms to segment three-dimensional models into a set of subshapes and having found shape description algorithms that allow efficient shape matching, we will move on to the ground-breaking technique we propose to study in this thesis. We will use a thesaurus of shapes to describe a collection of three-dimensional objects. To that end we must transpose indexing techniques commonly-used in textual information retrieval to be used in 3D shape retrieval.

In addition to geometric information we will also use topology for classification, matching and retrieval purposes. However, the topology will be used as an auxiliary strategy to distinguish models that share geometrically similar components. To that end, for indexing and matching purposes we might take advantage of the NB-Tree structure proposed by Fonsecain his PhD thesis [48].

4.1 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 should be develop research work that will produce a general 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. Indeed, all published results in this area were achieved through research work focused in a specific type of models. For instance, an approach developed for retrieval on protein databases might not even function on an engineering database.

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. We select the CAD model due to our partnerships with the mould industry, described with some detail in Section 5.3.

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 issues. However, such solution dos 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 issue 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.

4.2 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 [99] 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.

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 34. Thus, we expect to finish this doctoral program on August 2009, by submitting the preliminary version of the PhD thesis dissertation.

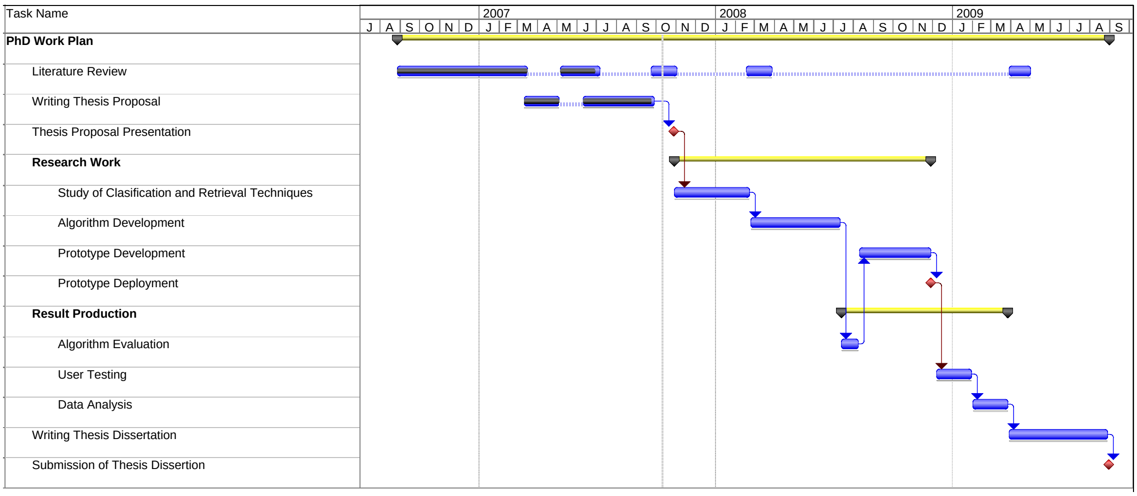

Figure 34: Gantt chart representing PhD work plan.

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 solution for 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, besides this PhD proposal, 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 indexing and matching mechanisms. Within this study we will implement a preliminary solution for 3D shape retrieval. Such system shall be able to perform classification, indexing and retrieval of 3D shapes combining existing algorithms and techniques. This solution might be considered as an adaptable framework where we will integrate and test our algorithms devised during the rest of the research work.

After we have finished the 3D shape retrieval solution described above, we will start to develop the algorithms proposed to support retrieval of 3D shapes with 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.

When we have finally achieved a feasible partial matching solution for 3D shape retrieval, we will 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 during this period 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.

5.3 Third year

We plan to start the third year of this work with a user evaluation of the CAD models retrieval prototype developed during the second year. We would like to involve in these tests real users of such tool. Since it is supposed to be used by designers, the main idea is to test it with professionals from the mould industry, as we did in the past during our research on retrieval of 2D technical drawings [45]. For that end, we intend to continue our seven-year long (at present date) partnership with CENTINFE[18]. Results obtained in these tests will allow us to fine-tune our techniques and to validate our approach as a practical solution, not only as a proof-of-concept prototype, but as an effective tool used by real users.

The second half of this year will 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 work a thesis dissertation.

6 Preliminary Results

As part of our initial research we compared the behaviour of existing shape descriptors to identify the most suitable for the solution we propose in this document. To that end, we developed a prototype that computes 3D shape descriptors and performs shape matching, based on these descriptors and on combinations of several of these descriptors. Then, to evaluate the shape descriptors, we used a subset of the Princeton Shape Benchmark (PSB) database.

6.1 Prototype

The prototype is divided in three main modules, as depicted in Figure 35. The *file reader* module reads the source files containing the three dimensional models and stores the corresponding polygonal mesh. We implemented two distinct versions of this module. While one reads the widely used VRML files, the other reads the OFF files provided by the PSB database. In the *feature extractor* module we coded the algorithms to compute the shape descriptors, constructing the corresponding feature vectors. Finally, the *shape matching* module compares a given feature vector with the existing in the model dataset and produce a similarity list enumerating the best matches.

Currently the *feature extractor* module implements a selection of three shape descriptors and all the combinations of compound descriptors. However, this module was developed in order to support a much larger set of shape descriptors. Indeed, in short term we plan to include a few more, which will allow a more complete comparison of shape descriptors.

6.2 Shape Descriptors

In this preliminary stage of our work we selected three histogram-based shape descriptors, briefly described in Section 2.5.1. These descriptors are the cord and angle his-

[18]CENTIMFE is a technological training centre for the Portuguese Mould Industry.

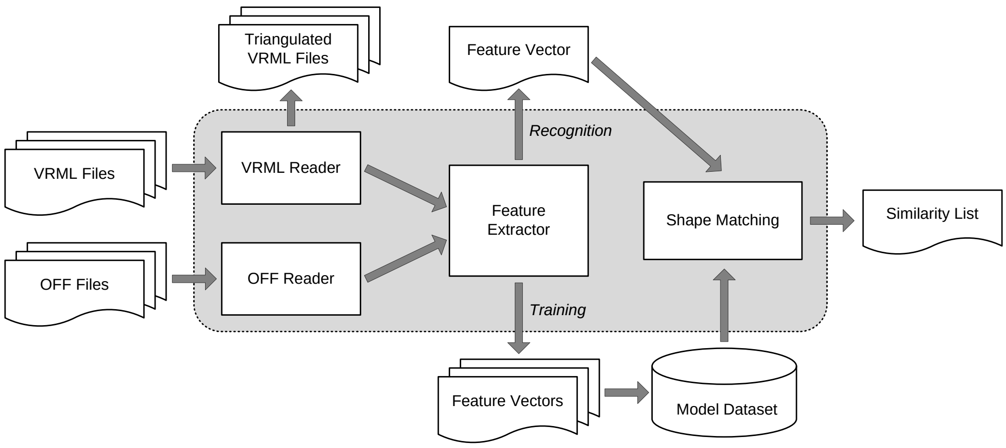

Figure 35: Architecture of preliminary prototype.

togram (CAH), shape distribution (SD) and complex extended Gaussian image (CEGI). For our prototype we implemented versions according to corresponding authors instructions, in order to obtain the more precise results possible.

Therefore, we followed the original method proposed by Paquet *et al.* in [90, 91] to implement the CAH descriptor. We considered a cord as a ray segment which joins the barycentre of the mesh with a triangle centre. We then conctructed the histograms using the length of the rays and the angles of the rays to the three coordinate planes.

Similarly, to estimate the SD descriptor we computed five distinct histograms as suggested by Osada *et al.* in [88] To that end we measured the angle formed by three random surface points, the distance of a random surface point to the centre of mass of the model, the distance between two random surface points, the area of the triangle defined by three random surface points and the volume of the tetrahedron defined by four random surface points.

Finally, to create the CEGI shape descriptors we constructed two histograms based on a complex function, as suggested by Kang and Ikeuchi [68]. One of these histograms measures the visible face area, composing the magnitude of complex function. The other accounts the normal distance of the face from the designated origin in the direction of the normal, composing the phase of complex function.

An example of the results obtained by the algorithms we implemented to compute the shape histogram is illustrated in Figure 36. In this case, a three dimensional model of a dog, from the PSB database, was used and corresponding histograms are shown. The feature vectors that describe the model are then constructed from these histograms and used for representing the object in the model dataset. The shape matching is obtained simply by comparing these feature vectors in order to find similar ones, which should correspond to similar 3D models.

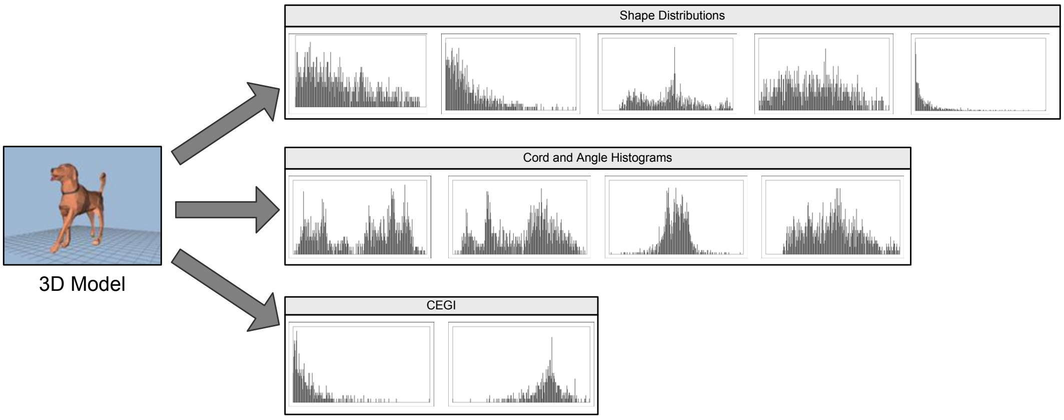

Figure 36: Shape descriptors computed by preliminary prototype.

6.3 Model Database

As referred above we just have implemented three feature extraction algorithms. This only allow us to perform a small comparative study, which will be more complete as other shape descriptor computation algorithms are implemented. Thus, since we are on a initial stage of our shape descriptor evaluation process, we decided to use a reduced model database instead of a more complete one, to be used when we implement the remaining shape description techniques .

Therefore, for the preliminary shape descriptor performance evaluation experiments we used a subset extracted from PSB database, which contains a total of 1914 models. To construct this subset, we selected twelve classes of creatures and eight classes of furniture, summing a total of four hundred models.

6.4 Evaluation results

Using the reduced model database we performed evaluation experiments and tested the three single 3D shape descriptors, as well as all of their combinations. In this experiment, we choose a 3D model which belongs to the class "Table" as the query shape. The three single shape descriptors and all their combinations were employed to retrieve this model from the database of four hundred models and then we observed and compared the retrieval results.

To quantify the performance of shape descriptors, we used a $11 - Point$ (11P) average precision. This measures the average of precision values taken at eleven equally spaced recall values $(0.0, 0.1, 0.2, ..., 0.9, 1.0)$. We use this measurement because 11P average precision value is commonly considered as a summary of the recall-precision plot, which emphasises overall performance. Additionally we also measured the $R - precision$,

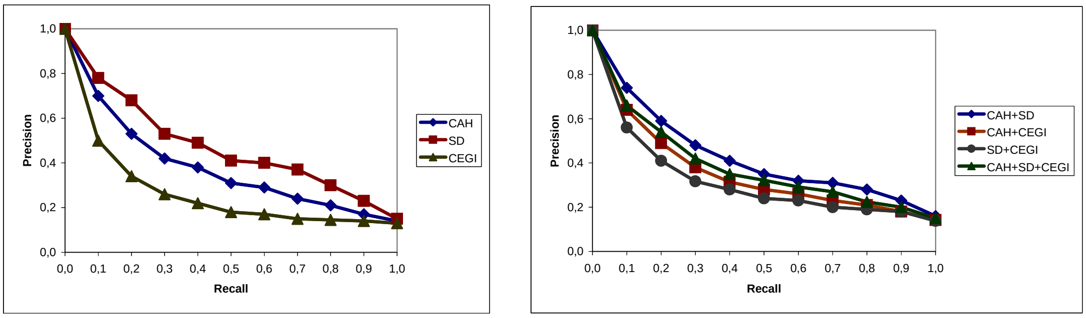

Single descriptors

Combined descriptors

Figure 37: Precision-recall plots of single and combined shape descriptors

the ratio of the models retrieved from the same class as the query in the top R retrievals, and the $2R - precision$, similar to $R - precision$ but computed using the top $2R$ retrievals. However, will discuss here only the results that allow a direct analysis of the overall performance.

Our experiments shown that the most approving results we obtained are produced by SD. This is illustrated in the precision-recall plot depicted in Figure 37, where the plot of the SD shape descriptor is on the top of the others. Furthermore, from all the shape descriptor combinations, the one that performed better was the combination of CAH and SD. On the contrary, is shown in Figure 37 that CEGI descriptor produced a very poor classification result a single descriptors. Moreover, combined descriptors containing CEGI has lower precision than those without CEGI.

6.5 Next steps

Despite the results we obtained in our preliminary evaluation experiment, we recognise that these are not enough for our purposes. Therefore, to obtain more accurate and informative results we need to implement more shape descriptors and use a more complete database.

In a near future we will continue the development of the prototype, integrating more shape descriptors, in order to repeat the experiment, now using the complete PSB database. From the results gathered from this experiment we expect formulate a strongly-based opinion on shape descriptors for 3D models and, eventually, publish it on an international journal.

After identifying the best descriptors for our work, we will focus on shape decomposition and topological description of three dimensional objects. To that end, we plan to work in cooperation with the Shape Modelling group from IMATI-Ge (see Section 2.1) who have been developing some interesting research in this area. Taking advantage of their and our knowledge, we intend to achieve, as a result of our joint research, a solution for efficient 3D shape retrieval with partial matching.

References

[1] Hervé Abdi. *Least-squares*, volume 1 of *Encyclopedia for research methods for the social sciences.* Sage, Thousand Oaks (CA), 2003.

[2] Hervé Abdi. *Distance*, volume 2 of *Encyclopedia of Measurements and Statistics.* Sage, Thousand Oaks (CA), 2007.

[3] Milton Abramowitz and Irene A. Stegun. Probability functions. pages 925–964, 1974.

[4] Ceyhun Burak Akgül. 3d shape descriptors and similarity learning. Phd thesis progress report, Boḡaziçi University, Istanbul, June 2006.

[5] Ceyhun Burak Akgul, Bulent Sankur, Francis Schmitt, and Yucel Yemez. Multivariate density-based 3d shape descriptors. In *SMI '07: Proceedings of the IEEE International Conference on Shape Modeling and Applications 2007*, pages 3–12, Washington, DC, USA, 2007. IEEE Computer Society.

[6] Ceyhun Burak Akgül, Bülent Sankur, Yücel Yemez, and Francis Schmitt. A framework for histogram-induced 3d descriptors. In *Proceedings of the 14th. European Signal Processing Conference (EUSIPCO'06)*, Florence, Italy, September 2006.

[7] Ceyhun Burak Akgül, Bülent Sankur, Yücel Yemez, and Francis Schmitt. Density-based 3d shape descriptors. *EURASIP Journal on Advances in Signal Processing*, 2007:Article ID 32503, 16 pages, 2007. doi:10.1155/2007/32503.

[8] Ceyhun Burak Akgül, Bülent Sankur, Yücel Yemez, and Francis Schmitt. Improving efficiency of density-based shape descriptors for 3d object retrieval. In *Computer Vision/Computer Graphics Collaboration Techniques, Third International Conference (MIRAGE'07)*, pages 330–340, Rocquencourt, France, March 2007.

[9] Mihael Ankerst, Gabi Kastenmller, Hans-Peter Kriegel, and Thomas Seidl. 3d shape histograms for similarity search and classification in spatial databases. In *SSD99: Proceedings of the 6th International Symposium on Advances in Spatial Databases*, pages 207–226, London, UK, 1999. Springer-Verlag.

[10] Jürgen Assfalg, Alberto Del Bimbo, and Pietro Pala. Retrieval of 3d objects by visual similarity. In *MIR '04: Proceedings of the 6th ACM SIGMM international workshop on Multimedia information retrieval*, pages 77–83, New York, NY, USA, 2004. ACM Press.

[11] Jürgen Assfalg, Alberto Del Bimbo, and Pietro Pala. Content-based retrieval of 3d models through curvature maps: a cbr approach exploiting media conversion. *Multimedia Tools Appl.*, 31(1):29–50, 2006.

[12] D. H. Ballard. Generalizing the Hough transform to detect arbitrary shapes. *Readings in computer vision: issues, problems, principles, and paradigms*, pages 714–725, 1987.

[13] S. Belongie, J. Malik, and J. Puzicha. Shape matching and object recognition using shape contexts. *IEEE Trans. Pattern Anal. Mach. Intell.*, 24(4):509–522, 2002.

[14] Helen M. Berman, John Westbrook, Zukang Feng, Gary Gilliland, T. N. Bhat, Helge Weissig, Ilya N. Shindyalov, and Philip E. Bourne. The protein data bank. *Nucleic Acids Research*, 28(1):235–242, 2000.

[15] F.C. Bernstein, T.F. Koetzle, G.J. Williams, E.F. Meyer Jr, M.D. Brice, J.R. Rodgers, O. Kennard, T. Shimanouchi, and M. Tasumi. The protein data bank. a computer-based archival file for macromolecular structures. *European Journal of Biochemistry*, 80:319–324, 1977.

[16] Dmitriy Bespalov, Cheuk Yiu Ip, William C. Regli, and Joshua Shaffer. Benchmarking cad search techniques. In *Proceedings of the 2005 ACM symposium on Solid and physical modeling (SPM '05)*, pages 275–286, New York, NY, USA, 2005. ACM Press.

[17] Dmitriy Bespalov, William C. Regli, and Ali Shokoufandeh. Reeb graph based shape retrieval for cad. *Proceedings of DETC03 2003 ASME Design Engineering Technical Conferences*, September 2003.

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

[19] Dmitriy Bespalov, Ali Shokoufandeh, William C. Regli, and Wei Sun. Scale-space representation of 3d models and topological matching. In *SM '03: Proceedings of the eighth ACM symposium on Solid modeling and applications*, pages 208–215, New York, NY, USA, 2003. ACM Press.

[20] Silvia Biasotti. Topological techniques for shape understanding. Central European Seminar on Computer Graphics, Bratislava, Slovakia, 2001.

[21] Silvia Biasotti, Ennio De Giorgi, Michela Spagnuolo, and Bianca Falcidieno. Size functions for 3d shape retrieval. In Konrad Polthier and Alla Sheffer, editors, *Proceedings of the Fourth Eurographics Symposium on Geometry Processing*, pages 239–242, Cagliari, Sardinia, Italy, June 2006. Eurographics Association.

[22] Silvia Biasotti, Simone Marini, Michela Mortara, Giuseppe Patan, Michela Spagnuolo, and Bianca Falcidieno. 3d shape matching through topological structures. In Ingela Nystrm, Gabriella Sanniti di Baja, and Stina Svensson, editors, *Discrete Geometry for Computer Imagery*, volume 2886 of *Lecture Notes in Computer Science*, pages 194–203. Springer, 2003.

[23] 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.

[24] Alberto Del Bimbo and Pietro Pala. Content-based retrieval of 3d models. *ACM Trans. Multimedia Comput. Commun. Appl.*, 2(1):20–43, 2006.

[25] Miroslaw Bober. Mpeg-7 visual shape descriptors. *IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY*, 11(6):716–719, 2001.

[26] Benjamin Bustos, Daniel Keim, Dietmar Saupe, and Tobias Schreck. Content-based 3d object retrieval. *IEEE Computer Graphics and Applications*, 27(4):22–27, 2007.

[27] Benjamin Bustos, Daniel A. Keim, Dietmar Saupe, Tobias Schreck, and Dejan V. Vranić. Feature-based similarity search in 3d object databases. *ACM Comput. Surv.*, 37(4):345–387, 2005.

[28] Benjamin Bustos, Daniel A. Keim, Dietmar Saupe, Tobias Schreck, and Dejan V. Vranić. An experimental effectiveness comparison of methods for 3d similarity search. *Int. J. on Digital Libraries*, 6(1):39–54, 2006.

[29] Richard J. Campbell and Patrick J. Flynn. A survey of free-form object representation and recognition techniques. *Comput. Vis. Image Underst.*, 81(2):166–210, 2001.

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

[31] Vittorio Castelli. Multidimensional indexing structures for content-based retrieval. *Image Databases*, 2002.

[32] 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.

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

[34] Ding-Yun Chen, Xiao-Pei Tian, Yu-Te Shen, and Ming Ouhyoung. On visual similarity based 3d model retrieval. *Computer Graphics Forum*, 22(3):223–232, 2003.

[35] Doug E.R. Clark, Jonathan R. Corney, Frank Mill, Heather J. Rea, Andrew Sherlock, and Nick K. Taylor. Benchmarking shape signatures against human perceptions of geometric similarity. *Computer-Aided Design*, 38(9):1038–1051, September 2006.

[36] Michele d'Amico, Patrizio Frosini, and Claudia Landi. Using matching distance in size theory: A survey. *International Journal of Imaging Systems and Technology*, 16(5):154–161, 2006.

[37] Pedro A. de Alarcón, Alberto D. Pascual-Montano, and José M. Carazo. Spin images and neural networks for efficient content-based retrieval in 3d object databases. In *CIVR '02: Proceedings of the International Conference on Image and Video Retrieval*, pages 225–234, London, UK, 2002. Springer-Verlag.

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

[39] Chitra Dorai and Anil K. Jain. Shape spectrum based view grouping and matching of 3d free-form objects. *IEEE Trans. Pattern Anal. Mach. Intell.*, 19(10):1139–1146, 1997.

[40] J. Stephen Downie. The music information retrieval evaluation exchange (mirex). *D-Lib Magazine*, 12(12), December 2006.

[41] Richard O. Duda, Peter E. Hart, and David G. Stork. *Pattern Classification (2nd Edition)*. Wiley-Interscience, November 2000.

[42] Helin Dutagaci, Bulent Sankur, and Yucel Yemez. Transform-based methods for indexing and retrieval of 3d objects. In *3DIM '05: Proceedings of the Fifth International Conference on 3-D Digital Imaging and Modeling*, pages 188–195, Washington, DC, USA, 2005. IEEE Computer Society.

[43] Michael Elad, Ayellet Tal, and Sigal Ar. Content based retrieval of vrml objects: an iterative and interactive approach. In *Proceedings of the sixth Eurographics workshop on Multimedia*, pages 107–118, New York, NY, USA, 2002. Springer-Verlag New York, Inc.

[44] Christos Faloutsos. *Searching Multimedia Databases by Content*. Kluwer Academic Publishers, Norwell, MA, USA, 1996.

[45] Alfredo Ferreira. Sketch and image based retrieval of technical drawings. Master's thesis, Universidade Técnica de Lisboa, Instituto Superior Técnico, Lisboa, Portugal, June 2005.

[46] Tarik Filali Ansary, Mohamed Daoudi, and Jean-Philippe Vandeborre. 3d-models search engine from photos. In *ACM International Conference on Image and Video Retrieval (CIVR 2007)*, Amsterdam, The Netherlands, July 2007.

[47] Tarik Filali Ansary, Mohamed Daoudi, and Jean-Philippe Vandeborre. A bayesian 3D search engine using adaptive views clustering. *IEEE Transactions on Multimedia*, 9(1):78–88, January 2007.

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

[49] Thomas Funkhouser, Michael Kazhdan, Patrick Min, and Philip Shilane. Shape-based retrieval and analysis of 3d models. *Communications of the ACM*, 48(6):58–64, 2005.

[50] Thomas Funkhouser, Patrick Min, Michael Kazhdan, Joyce Chen, Alex Halderman, David Dobkin, and David Jacobs. A search engine for 3d models. *ACM Trans. Graph.*, 22(1):83–105, 2003.

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

[52] Volker Gaede and Oliver Günther. Multidimensional access methods. *ACM Computing Surveys (CSUR)*, 30(2):170–231, 1998.

[53] Nikhil Gagvani and Deborah Silver. Parameter-controlled volume thinning. *CVGIP: Graph. Models Image Process.*, 61(3):149–164, 1999.

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

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

[56] Ho Min Wong Hee Kap Ahn, Nikos Mamoulis. A survey on multidimensional access methods. Technical Report UU-CS-2001-14, Institute of Information and Computing Sciences, Utrecht University, 2001.

[57] 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.

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

[59] B. K. P. Horn. Extended gaussian images. *Proceedings of the IEEE*, 72(12):1671–1686, 1984.

[60] Suyu Hou and Karthik Ramani. Sketch-based 3d engineering part class browsing and retrieval. In Thomas Stahovich and Mario Costa Sousa, editors, *Proceedings of the Eurographics Workshop on Sketch-Based Interfaces and Modeling*, pages 131–138, Vienna, Austria, 09 2006. Eurographics.

[61] Google Inc. 'google earth'. http://earth.google.com, 2005.

[62] Google Inc. 'google sketchup'. http://sketchup.google.com, 2006.

[63] Google Inc. '3d warehouse'. http://sketchup.google.com/3dwarehouse, 2007.

[64] Natraj Iyer, Subramaniam Jayanti, Kuiyang Lou, Yagnanarayanan Kalyanaraman, and Karthik Ramani. Three-dimensional shape searching: state-of-the-art review and future trends. *Computer-Aided Design*, 37:509–530, April 2005.

[65] Subramaniam Jayanti, Yagnanarayanan Kalyanaraman, Natraj Iyer, and Karthik Ramani. Developing an engineering shape benchmark for cad models. *Computer-Aided Design*, 39(9):939–953, September 2006.

[66] Andrew Johnson. *Spin-Images: A Representation for 3-D Surface Matching.* PhD thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, August 1997.

[67] Andrew E. Johnson and Martial Hebert. Using spin images for efficient object recognition in cluttered 3d scenes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(5):433–449, 1999.

[68] S. B. Kang and K. Ikeuchi. The complex egi: A new representation for 3-d pose determination. *IEEE Trans. Pattern Anal. Mach. Intell.*, 15(7):707–721, 1993.

[69] Michael Kazhdan, Bernard Chazelle, David Dobkin, Thomas Funkhouser, and Szymon Rusinkiewicz. A reflective symmetry descriptor for 3d models. *Algorithmica*, 38(1):201–225, 2003.

[70] Michael Kazhdan, Thomas Funkhouser, and Szymon Rusinkiewicz. Rotation invariant spherical harmonic representation of 3d shape descriptors. In Leif Kobbelt, Peter Schrder, 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.

[71] Michael Kazhdan, Thomas Funkhouser, and Szymon Rusinkiewicz. Symmetry descriptors and 3d shape matching. In *SGP '04: Proceedings of the 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing*, pages 115–123, New York, NY, USA, 2004. ACM Press.

[72] Daniel A. Keim. Efficient geometry-based similarity search of 3d spatial databases. In *Proceedings of the 1999 ACM SIGMOD international conference on Management of Data*, pages 419–430, New York, NY, USA, 1999. ACM Press.

[73] Jan J. Koenderink and Andrea J. van Doorn. Surface shape and curvature scales. *Image Vision Comput.*, 10(8):557–565, 1992.

[74] M. Körtgen, G.-J. Park, M. Novotni, and R. Klein. 3d shape matching with 3d shape contexts. In *The 7th Central European Seminar on Computer Graphics*, Budmerice, Slovakia, April 2003.

[75] 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.

[76] Marc Levoy. The digital michelangelo project. In *Proceedings of the Second International Conference on 3-D Digital Imaging and Modeling*, pages 2–11, Ottawa, Ont., Canada, October 1999.

[77] Marc Levoy, Kari Pulli, Brian Curless, Szymon Rusinkiewicz, David Koller, Lucas Pereira, Matt Ginzton, Sean Anderson, James Davis, Jeremy Ginsberg, Jonathan Shade, and Duane Fulk. The digital michelangelo project: 3d scanning of large statues. In *SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 131–144, New York, NY, USA, 2000. ACM Press/Addison-Wesley Publishing Co.

[78] Yi Liu, Hongbin Zha, and Hong Qin. The generalized shape distributions for shape matching and analysis. In *Proceedings of the IEEE International Conference on Shape Modeling and Applications 2006 (SMI'06)*, page 16, Washington, DC, USA, 2006. IEEE Computer Society.

[79] K. Lou, S. Prabhakar, and K. Ramani. Content-based three-dimensional engineering shape search. *icde*, 00:754, 2004.

[80] Qin Lv, William Josephson, Zhe Wang, Moses Charikar, and Kai Li. Ferret: a toolkit for content-based similarity search of feature-rich data. *SIGOPS Oper. Syst. Rev.*, 40(4):317–330, 2006.

[81] P.C. Mahalanobis. On the generalized distance in statistics. *National Institute of Sciences of India*, 2(1):49–55, 1936.

[82] 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.

[83] 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.

[84] Michela Mortara and Giuseppe Patanè. Shape-covering for skeleton extraction. *International Journal of Shape Modeling*, 8(2):139–158, 2002.

[85] 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.

[86] Marcin Novotni and Reinhard Klein. 3d zernike descriptors for content based shape retrieval. In *SM '03: Proceedings of the eighth ACM symposium on Solid modeling and applications*, pages 216–225, New York, NY, USA, 2003. ACM Press.

[87] Ryutarou Ohbuchi, Takahiro Minamitani, and Tsuyoshi Takei. Shape-similarity search of 3d models by using enhanced shape functions. In *TPCG '03: Proceedings of the Theory and Practice of Computer Graphics 2003*, page 97, Washington, DC, USA, 2003. IEEE Computer Society.

[88] Robert Osada, Thomas Funkhouser, Bernard Chazelle, and David Dobkin. Shape distributions. *ACM Trans. Graph.*, 21(4):807–832, 2002.

[89] Panagiotis Papadakis, Ioannis Pratikakis, Stavros Perantonis, and Theoharis Theoharis. Efficient 3d shape matching and retrieval using a concrete radialized spherical projection representation. *Pattern Recogn.*, 40(9):2437–2452, 2007.

[90] E. Paquet and M. Rioux. Nefertiti: A query by content software for three-dimensional models databases management. *3dim*, 00:345, 1997.

[91] Eric Paquet, Marc Rioux, Anil Murching, Thumpudi Naveen, and Ali Tabatabai. Description of shape information for 2-d and 3-d objects. *Signal Processing: Image Communication*, 16:103–122, September 2000.

[92] Joshua Podolak, Philip Shilane, Aleksey Golovinskiy, Szymon Rusinkiewicz, and Thomas Funkhouser. A planar-reflective symmetry transform for 3d shapes. In *Proceedings of International Conference on Computer Graphics and Interactive Techniques: ACM SIGGRAPH*, pages 549–559, New York, NY, USA, 2006. ACM Press.

[93] Georges Reeb. Sur les points singuliers dúne forme de pfaff completement integrable ou dúne fonction numrique. *Comptes Rendus des séances de lÁcadémie des sciences*, 222:847–849, 1946.

[94] William C. Regli, Cheryl Foster, Erik Hayes, Cheuk Yiu Ip, David McWherter, Mitchell Peabody, Yuriy Shapirsteyn, and Vera Zaychik. National design repository project: A status report. In *International Joint Conferences on Artificial Intelligence (IJCAI) and AAAI/SIGMAN Workshop on AI in Manufacturing Systems*, Seattle, Washington, USA, August 2001.

[95] William C Regli and Daniel M Gaines. A repository for design, process planning and assembly. *Computer-Aided Design*, 29(12):895–905, December 1997.

[96] Julien Ricard, David Coeurjolly, and Atilla Baskurt. Generalizations of angular radial transform for 2d and 3d shape retrieval. *Pattern Recogn. Lett.*, 26(14):2174–2186, 2005.

[97] Dietmar Saupe and Dejan V. Vranić. 3d model retrieval with spherical harmonics and moments. In *Proceedings of the 23rd DAGM-Symposium on Pattern Recognition*, pages 392–397, London, UK, 2001. Springer-Verlag.

[98] Jau-Ling Shih, Chang-Hsing Lee, and Jian Tang Wang. A new 3d model retrieval approach based on the elevation descriptor. *Pattern Recogn.*, 40(1):283–295, 2007.

[99] P. Shilane, M. Kazhdan, P. Min, and T. Funkhouser. The princeton shape benchmark, 2004.

[100] 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.

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

[102] Yoshihisa Shinagawa and Tosiyasu L. Kunii. Constructing a reeb graph automatically from cross sections. *IEEE Computer Graphics and Applications*, 11(6):44–51, 1991.

[103] A. Shokoufandeh, S. Dickson, K. Siddiqi, and S. Zucker. Indexing Using a Spectral Encoding of Topological Structure. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR'99)*, pages 2491–2497. IEEE Computer Society, 1999.

[104] Alan F. Smeaton, Paul Over, and Wessel Kraaij. Trecvid: evaluating the effectiveness of information retrieval tasks on digital video. In *Proceedings of the 12th annual ACM international conference on Multimedia (MULTIMEDIA '04)*, pages 652–655, New York, NY, USA, 2004. ACM Press.

[105] 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.

[106] Paul Speckman. Nonparametric and semiparametric models. *Journal of the American Statistical Association*, 101:1308–1309, September 2006.

[107] H. Sundar, D. Silver, N. Gagvani, and S. Dickinson. Skeleton based shape matching and retrieval. *Proceedings of the Shape Modeling International 2003*, page 290, 2003.

[108] 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.

[109] 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.

[110] Motofumi T. Suzuki, Yoshitomo Yaginuma, and Yuji Y. Sugimoto. A 3d model retrieval system for cellular phones. In *Proceedings of IEEE International Conference on Systems, Man and Cybernetics*, volume 4, pages 3846–3851. IEEE Computer Society, October 2003.

[111] 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.

[112] 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.

[113] S. Szykman, R.D. Sriram, C. Bochenek, J.W. Racz, and J. Senfaute. Design repositories: engineering design's new knowledge base. *Intelligent Systems and Their Applications, IEEE*, 15(3):48–55, 2000.

[114] Johan W. H. Tangelder and Remco C. Veltkamp. Polyhedral model retrieval using weighted point sets. In *Proceedings of the Shape Modeling International 2003 (SMI'03)*, page 119, Los Alamitos, CA, USA, 2003. IEEE Computer Society.

[115] Johan W.H. Tangelder and Remco C. Veltkamp. A survey of content based 3d shape retrieval methods. In *SMI '04: Proceedings of the Shape Modeling International 2004 (SMI'04)*, pages 145–156, Los Alamitos, CA, USA, 2004. IEEE Computer Society.

[116] Julien Tierny, Jean-Philippe Vandeborre, and Mohamed Daoudi. Topology driven 3d mesh hierarchical segmentation. *IEEE International Conference on Shape Modeling and Applications (Shape Modeling International 2007)*, pages 215–220, 2007.

[117] Tony Tung and Francis Schmitt. Augmented reeb graphs for content-based retrieval of 3d mesh models. In *Proceedings of the Shape Modeling International 2004 (SMIÓ4)*, pages 157–166, Los Alamitos, CA, USA, 2004. IEEE Computer Society.

[118] Tony Tung and Francis Schmitt. The augmented multiresolution reeb graph approach for content-based retrieval of 3d shapes. *International Journal of Shape Modeling*, 11(1):91–120, 2005.

[119] Jean-Philippe Vandeborre, Vincent Couillet, and Mohamed Daoudi. A practical approach for 3D model indexing by combining local and global invariants. In *1st IEEE International Symposium on 3D Data Processing Visualization Transmission (3DPVT'02)*, Padova, Italy, June, 19-21 2002.

[120] Dejan V. Vranić. An improvement of rotation invariant 3d-shape based on functions on concentric spheres. In *Proceedings of the 2003 IEEE International Conference on Image Processing (ICIP 2003)*, volume 3, pages 757–760, Barcelona, Spain, September 2003.

[121] Dejan V. Vranić. *3D Model Retrieval*. PhD thesis, University of Leipzig, Germany, June 2004.

[122] Dejan V. Vranić and Dietmar Saupe. 3d model retrieval. In *Proceeedings of Spring Conference on Computer Graphics and its Applications (SCCG2000)*, pages 89–93, Budmerice, Slovakia, May 2000.

[123] Dejan V. Vranić and Dietmar Saupe. 3d shape descriptor based on 3d fourier transform. In *Proceedings of the EURASIP Conference on Digital Signal Processing for Multimedia Communications and Services (ECMCS 2001)*, pages 271–274, Budapest, Hungary, 2001.

[124] Dejan V. Vranić and Dietmar Saupe. Description of 3d-shape using a complex function on the sphere. In *Proceeedings of 2002 IEEE International Conference on Multimedia and Expo (ICME '02)*, volume 1, pages 177–180, 2002.

[125] Liu Wei and He Yuanjun. 3d shape similarity comparison using multi-level spherical moments. *Computer-Aided Design Applications*, 3(1–4):307–314, 2006.

[126] Eric W. Weisstein. Dodecahedron. From MathWorld–A Wolfram Web Resource.

[127] Eric W. Weisstein. Orthotope. From MathWorld–A Wolfram Web Resource.

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

[129] Meng Yu, Indriyati Atmosukarto, Wee Kheng Leow, Zhiyong Huang, and Rong Xu. 3d model retrieval with morphing-based geometric and topological feature maps. *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR '03)*, 02:656, 2003.

[130] Hagit Zabrodsky, Shmuel Peleg, and David Avnir. Symmetry as a continuous feature. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(12):1154–1166, 1995.

[131] T. Zaharia and F. J. Prêteux. 3D shape-based retrieval within the MPEG-7 framework. In Edward R. Dougherty and Jaakko T. Astola, editors, *Proceedings of the SPIE conference on Nonlinear Image Processing and Pattern Analysis XII*, volume 4304, pages 133–145, May 2001.

[132] T. Zaharia and F. J. Prêteux. Hough transform-based 3D mesh retrieval. In L. J. Latecki, D. M. Mount, A. Y. Wu, and R. A. Melter, editors, *Proceedings of the SPIE conference on Vision Geometry X*, volume 4476, pages 175–185, November 2001.

[133] Jiqi Zhang, Hau-San Wong, and Zhiwen Yu. 3d model metrieval based on volumetric extended gaussian image and hierarchical self organizing map. In *MULTIMEDIA '06: Proceedings of the 14th annual ACM international conference on Multimedia*, pages 121–124, New York, NY, USA, 2006. ACM Press.

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