THE UNIVERSITY OF OSLO

DEPARTMENT FOR INFORMATICS

Improving the performance of shape similarity retrieval systems

Mats Carlin

Doctoral dissertation

Submitted to the Faculty of Mathematics and Natural Sciences at the University of Oslo in partial fulfillment of the requirements for the degree Dr.Scient in Informatics

May 2000

Abstract

The main aim of this thesis is to develop methods and strategies to improve the performance of feature-based shape similarity retrieval systems.

We have developed a framework for measuring the performance of shape similarity retrieval methods and systems based on human perception, application-specific information and the mathematical shape representations. Several performance measures have been applied to test the performance of commonly used groups of shape features.

We have introduced the concepts of relevance and redundancy with respect to shape similarity retrieval and have applied these to select optimal feature subsets. On a large feature set including many redundant and some irrelevant features, we have showed that it is possible to reduce the number of features by 80 % and at the same time improve performance significantly. Redundant and irrelevant features are frequently encountered for many applications since it is easy to compute many features.

We have introduced a hierarchy of necessary, objective and wanted requirements to the individual shape features. The hierarchy has allowed us to focus on relevance and redundancy as the objective requirements to assess which features are most important to shape similarity retrieval for our specific application. The most important features seem to be features with a high-level abstract interpretation.

The issues above are main issues to enable improvement of shape similarity retrieval systems. The methods have been implemented and tested in a system for retrieval of Computer Aided Design (CAD) drawings of aluminium sections called the DieFinder, yielding significant improvement.

Dedication

This thesis is dedicated in memorian to my maternal grandfather Lars Joel Larsson (1894-1986) who never had the opportunity to attend university despite his obvious interest and talent.

"For the next half-hour Clara was hard at work, putting in marks and rubbing them out again, and hunting up and down for a suitable picture. This she found the hardest part of all."

Lewis Carroll, Tangled tales, knot 5, Oughts and crosses, 1885

Contents

Contents **1**

List of Figures **3**

Preface **5**

1 Introduction **9**
1.1 Justification . . . 9
1.2 Objectives . . . 11
1.2.1 Scientific objectives . . . 11
1.2.2 Application-specific objectives . . . 12
1.3 Contribution . . . 12
1.4 Thesis organisation . . . 13

2 Applications **15**
2.1 The aluminium extrusion process . . . 15
2.1.1 Aluminium extrusion . . . 15
2.1.2 Parts of a die . . . 15
2.1.3 The die life cycle . . . 16
2.1.4 This thesis and aluminium extrusion . . . 17
2.2 The DieFinder . . . 18
2.2.1 Data representation . . . 18
2.2.2 Modes of operation . . . 19
2.2.3 Feature based shape similarity retrieval . . . 20
2.3 The SpeedPredictor . . . 21

3 Shape similarity retrieval **23**
3.1 What is shape similarity? . . . 23
3.1.1 The shape space . . . 23
3.1.2 Exact and approximate shape similarity . . . 24
3.1.3 Shape and causal history . . . 25
3.1.4 Shape and functionality . . . 25
3.1.5 Shape and perception . . . 26
3.1.6 Invariant shape features . . . 26
3.1.7 The shape similarity retrieval problem . . . 27
3.2 Shape similarity retrieval . . . 27

3.2.1 System requirements 28
3.2.2 Information refinement blocks 29
3.2.3 Matching blocks 35
3.2.4 Performance measures 36
3.3 Image retrieval systems 36
3.3.1 The difference between image retrieval and vector graphics retrieval . . 37
3.3.2 Shape, colour, texture and composition 37
3.3.3 State-of-the-art within image retrieval 39
3.3.4 Applications of image retrieval systems 40

4 Summary of papers **43**
4.1 First paper: Measuring the performance of shape similarity retrieval methods . 43
4.2 Second paper: Selecting feature subsets and distance measures for shape similarity retrieval 45
4.3 Third paper: Which shape features are most important for shape similarity retrieval? 46
4.4 Fourth paper: Computing geometric moments for objects with an exact polygon representation 47
4.5 Fifth paper: Measuring the complexity of non-fractal shapes by fractal methods 48

5 Discussion and future work **49**
5.1 Discussion of the main results 49
5.1.1 Performance evaluation of shape similarity retrieval 49
5.1.2 Selecting the optimal set of features 50
5.1.3 Assessing individual shape features 50
5.1.4 Theoretical treatment of the shape space 51
5.1.5 Computability 51

Bibliography **52**

I Measuring the performance of shape similarity retrieval methods **59**

II Selecting feature subsets and distance measures for shape similarity retrieval **81**

III Which shape features are most important for shape similarity retrieval? **107**

IV Computing geometric moments for objects with an exact polygon representation **137**

V Measuring the complexity of non-fractal shapes by fractal methods **145**

List of Figures

1 The beauty of a Voronoi diagram 8
1.1 The fi

Preface

"And what is the use of a book", thought Alice, "without pictures and conversations".
Lewis Carroll, Alice in Wonderland, 1865

ing courses at the University (which was not included in my study). My study includes the following courses; Theory of Science (MNVIT401), Multivariate Statistical Modelling (ST341), Inverse Problems and Parameter Estimation (INVPAR), Estimation in practise (IN370), Combinatorial optimisation (IN330B) and Spline modelling (IN329).

PROSMAT is a large research programme for the process and materials industries in Norway. It spans materials like plastic, wood and metal and technology from development of new processes and products to investigation and understanding of existing production processes. The PROSMAT project ”Empirical modelling techniques for the future extrusion technology” is part of a cluster of projects aimed at the Norwegian aluminium extrusion industry with Hydro Aluminium as a major partner. The cluster of projects cover all stages of production of aluminium extruded sections and include Finite-Element Methods (FEM) modelling, development of new die designs, investigation of material properties and empirical modelling of the production process. The long-term goal of these projects is to improve the production process to be able to increase productivity and reduce scrap percentage. The project cluster is organised such that researchers from different fields regularly meet and discuss project improvements and may build on each others experience. This is a very fruitful model for organising research work.

It has been three exciting years with a continuous development of new thoughts and ideas. Many of these ideas remain to be explored, but some have been explored in this thesis. The five papers in this thesis are:

1. Mats Carlin. Measuring the performance of shape similarity retrieval methods. Submitted to *Computer Vision and Image Understanding*, special issue on empirical evaluation of computer vision algorithms.

2. Mats Carlin. Selecting feature subsets and distance measures for shape similarity retrieval. Submitted to *Pattern Recognition*.

3. Mats Carlin. Which shape features are most important to shape similarity retrieval? Submitted to *Pattern Recognition*.

4. Mats Carlin. Computing geometric moments for objects with an exact polygon representation. *Proceedings of VI-98, Vision Interface*, pp.319-324, Vancouver, Canada, June 1998.

5. Mats Carlin. Measuring the complexity of non-fractal shapes by fractal methods. *Pattern Recognition Letters* 21:1013-1017, 2000.

Early versions and abstracts of some of the ideas presented in the above papers have been or will be published in *Proceedings of NOBIM-98, Norwegian Image Processing and Pattern Recognition conference*, pp.28-38, Oslo, Norway, June 1998, *Proceedings of the 2nd Workshop on Industrial Mathematics for Nordic PhD students*, September 1998, Orkdal, Norway, *Proceedings of the 10th Norwegian seminar on statistics*, Larvik, Norway, June 1999 and *Proceedings of NOBIM-2000, Norwegian Image Processing and Pattern Recognition conference*, Trondheim, Norway, June 2000.

Acknowledgements

During my dr.scient thesis I have encountered a number of inspiring and encouraging people

Chapter 1

Introduction

The philosophical concept of shape has intrigued mankind since the dawn of civilisation. The Greek philosopher Aristotle (384-322 BC) was one of the first to divide the properties of any object into shape and material. To Aristotle shape was an expression of the objects' nature. Material was an expression of the content of the object. Shape was reality and material possibility. One philosophical view at that time was that shape is the comprehensible or conceivable by an object as opposed to material which is perceivable. Since then physics has almost altered our view of shape and material. Material can be described in detail using the laws of physics and is almost comprehensible in scientific terms, while the shape of things are described by the same basic geometry as two thousand years ago and we do not fully understand how we perceive shape or shape similarity.

The problem of identifying exact shape similarity between polygonal objects under a Euclidean transformation was solved by Euclid (ca. 300 BC) in his famous geometry primer *The elements* [14]. However approximate shape similarity between objects with an arbitrary representation is not yet fully understood and cannot be achieved with reasonable performance using automata today, despite the many technological and scientific breakthroughs since the Greek civilisation. It is still a surprisingly difficult problem.

Figure 1.1: The first two objects are exactly similar under rotation and translation, while the last two objects are only approximately similar under the same transformation

1.1 Justification

The current state-of-the-art of digital image retrieval systems can be divided in three major types based on their basic indexing method.

- **Textual descriptions** where a human has indexed the images by textual annotations describing the contents of each image.

- **Feature descriptions** where a number of features are extracted from the images and are used during retrieval.

- **Template matching** where a template image or shape is deformed and fit to the individual images in the database.

None of these techniques seems superior since all of them have shortcomings. Humans are able to create textual descriptions of images, but the cost is high and many problems still remain, as the descriptions must be universally consistent and complete using a common and fully comprehended dictionary. Cultural, perceptual and lingual barriers are difficult to cross in such systems.

Template matching is difficult and time consuming as the template must be compared with all the other images in the database prohibiting fast retrieval. To the author's knowledge, there exists no complete investigation of which template deformations model our perception of shape similarity. Combined with hierarchical techniques, this approach is however very promising. When computers become increasingly fast retrieval times will decrease and larger image databases can be accessed using template matching approaches.

Feature descriptions can be computed automatically for large image databases provided that we have robust and accurate implementations of the feature extractors. Retrieval is fast even for huge image databases, but we still lack an understanding of which features are most important for retrieval performance. This thesis try to address this last issue and create a methodology for improving feature-based shape similarity retrieval systems.

Images contain four important groups of features; Colour, texture, shape and composition. Colour, texture and composition are suitable descriptors for natural objects and is extensively used in state-of-the-art commercial and scientific image retrieval system. Shape is the most important descriptor of many man-made objects as these often may have arbitrary colour and texture. Our database of Computer Aided Design (CAD) drawings of extruded aluminium sections contain only man-made objects and only shape information, no colour, texture or composition. Shape is also the least explored type of descriptor in image retrieval systems. Shape is however a difficult issue as objects do have a 3-dimensional shape and an image only contain one single 2-dimensional projection of that 3-dimensional shape. In our case we need not to worry about the third dimension. We work with purely 2-dimensional objects.

In this thesis I have concentrated the work on solving some general problems of feature-based shape similarity retrieval system for retrieval of aluminium extruded sections. The work has some important practical applications, but at the same time it raises a number of universal scientific issues:

- What is the essence of shape?

- How do we perceive shape similarity?

- Which aspects of the concept shape are important for shape similarity?

- How do we measure the performance of a shape similarity retrieval system?

While designing and implementing a working feature-based shape similarity retrieval system many features had to be computed and new problems appeared which had to be solved. The prerequisite of computing features robustly and accurately has also been addressed in this thesis. These features should contain the most important information for retrieval.

1.2 Objectives

The main part of this thesis and the three first papers mainly discuss some universal issues within shape similarity retrieval.

1.2.1 Scientific objectives

The key scientific objectives are

- Creating a methodology for evaluation of shape similarity retrieval methods and systems
- Creating a methodology for feature selection for optimal shape similarity retrieval
- Gaining insight about which features are important for shape similarity retrieval
- Comparisons of different feature extraction methods with respect to shape similarity retrieval

To be able to evaluate the performance is a crucial factor when implementing a retrieval system of any kind. While reviewing the state-of-the-art of shape similarity retrieval, it became apparent that many researchers today compare shape similarity retrieval methods based on either visual presentation of the retrieval results or based on simple evaluation measures such as relevance feedback. Relevance feedback is a subjective measure of which of the retrieved objects is relevant or not. Relevance feedback is not very suitable for large-scale comparisons of shape similarity retrieval systems or methods as the measure is not repeatable. We therefor wished to go beyond the current approaches and create several independent measures of retrieval performance based on the available sources of knowledge about the problem. The first paper in this thesis concerns performance evaluation.

It was also clear that in a feature-based shape similarity retrieval system, the major problem is not to compute many features. There is a wealth of literature describing different features and their computation. Two main issues still remain:

1. *The computational methods must be robust and accurate.* It soon became apparent that some commonly used computational methods are neither robust nor accurate for several types of features. The fourth paper in this thesis is concerned with robust computation of geometric moments from polygons.
2. *We must also be able to identify which features are important or not for retrieval.* Adding too many features curbs performance while indexing and retrieval time increase. Finding an **optimal** feature subset is therefor of huge importance. This is the issue discussed in the second and third paper of this thesis.

When we do have many features and methods for measuring the performance of shape similarity retrieval, it is of some interest to compare specific groups of features. If one single group of features is superior, we may prefer to implement this group of features first to boost early performance. It is however expected that a mixture of features will attain the best possible performance. A comparison of different groups of features is part of the first paper in this thesis.

1.2.2 Application-specific objectives

The key application-specific objectives are

- Retrieving similar objects from a Computer Aided Design (CAD) drawing database
- Identification of the most important application-specific shape features
- Presentation of process and geometry data in an informative and orderly manner

Except for the last application objective, there is a one-to-one relationship between the scientific and application-specific objectives.

1.3 Contribution

The major contributions in this thesis are:

- A framework for **measuring the performance** of shape similarity retrieval methods based on human perception of similarity, application-specific information and the mathematical representation has been established. Four new performance measures have been used to test the performance of several groups of shape features providing some interesting results.
- We have defined the two concepts of **relevance** and **redundancy** with respect to shape similarity retrieval. This approach is new within shape similarity retrieval. We have used these two concepts to select **optimal** feature subsets containing those features which are most important to shape similarity retrieval. We have showed empirically that we can reduce a large feature set with many redundant and a few irrelevant features by 80 % and at the same time increase retrieval performance significantly.
- We have assessed the **importance** of individual shape features with respect to shape similarity retrieval by the relevance and redundancy of the individual features.

In addition some minor contributions have been achieved:

- A robust method for computing geometric moments from polygons have been described and we have proved that some commonly used methods are sensitive to the slope of the individual line segments in a polygon.
- We have presented a new method for measuring complexity of non-fractal objects based on the divider-step method which is commonly used to estimate the fractal dimension of fractal objects.

These contributions are described in detail in chapter 4 and in the accompanying papers.

1.4 Thesis organisation

This thesis is organised as follows. Chapter two briefly describes the aluminium extrusion production process and the two applications, the DieFinder and the SpeedPredictor. We have focused on functional requirements of the applications. This chapter forms background information for this thesis. Chapter three introduces the scientific and technological challenges of image retrieval in general and shape similarity retrieval specifically, focusing on the problems sought to be solved in this thesis. Chapter four provides a brief introduction and overview of the individual papers in the papers. Chapter five contains a discussion of the primary results and suggestions for further work.

A full description of many of the shape features used throughout this thesis is given in the appendices, as well as a complete listing of all the Computer Aided Design (CAD) drawings in the test database.

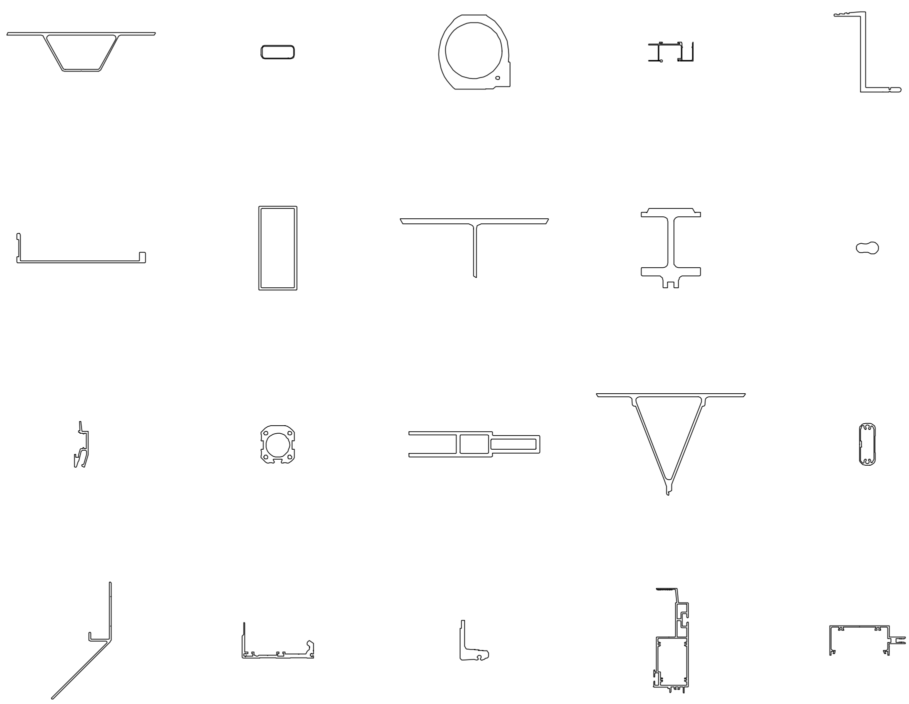

Figure 1.2: The essence of shape similarity retrieval: Which shapes are similar and which are not?

Chapter 2

Applications

This thesis is application-oriented in the sense that all problems sought to be solved are spurred by the interest in a specific application. A scientific work like this thesis has the aim of reaching beyond the application issues and focus on issues of general academic interest. To make the reader familiar with our application is however of significant interest, since it provides the context for the thesis.

This chapter first describes the aluminium extrusion process and then the two current applications, the DieFinder and the SpeedPredictor. We have focused on the functional requirements of the applications. In the next chapter we will discuss the scientific issues of a shape similarity retrieval system.

2.1 The aluminium extrusion process

Below follows a brief description of the aluminium extrusion process intended for the reader not familiar with this production process.

2.1.1 Aluminium extrusion

Aluminium extrusion is a process where cylindrical billets of aluminium are transformed to aluminium sections with a fixed 2-dimensional cross-section by pushing the aluminium through a 3-dimensional die, much like toothpaste through the mouth of the tooth paste tube. The billets are heated in induction ovens to a temperature of about 450°C. The heated billet of aluminium is then placed in the press container and a large ram pushes the relatively soft aluminium billet through the extrusion die. The 3 dimensional extrusion die forms the actual aluminium section with a fixed 2 dimensional cross-section.

2.1.2 Parts of a die

Figure 2.1 sketches a typical longitudinal cross-section of a die package containing a feeder, die, backer and support ring.

Flat dies are usually made in one piece and are used to produce sections without hollows. A hollow die consists of minimum two pieces, a die plate and a mandrel plate. The die plate

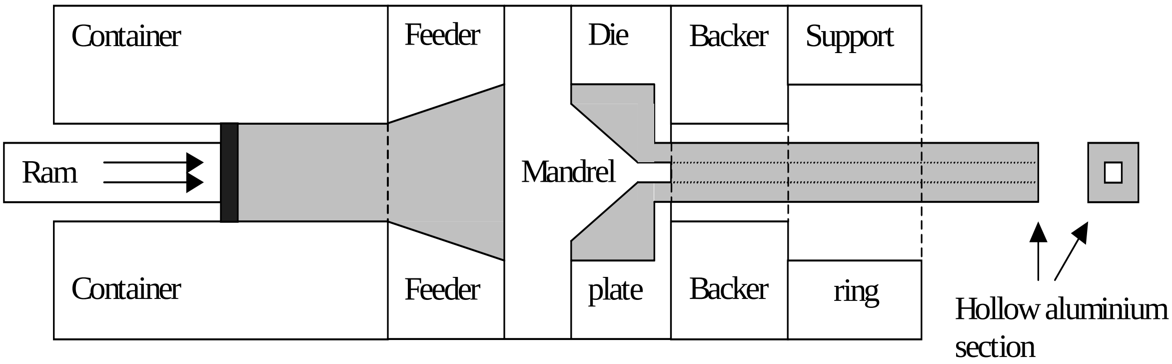

Figure 2.1: Longitudinal cross section of a die

specifies the outer surfaces of the aluminium section, while the mandrels defines the hollows. The mandrel is supported by webs. The metal flows around and weld together behind these webs. Both flat dies and hollow dies may have several exits allowing several smaller sections being produced at the same time. At the exits the die is designed with bearing surfaces. The length and angle of the bearing surface are used to guide and enhance the material flow. The main aim of the bearing surfaces is to decrease or increase friction between the die and the aluminium. The die is a fairly complex high-tech product made of high quality steel.

The die is the heart of the extrusion process and the shape of the section is important for the productivity.

2.1.3 The die life cycle

The die life cycle is described in figure 2.3. When a new die is designed, the process often starts with the idea that an aluminium section could be used for a certain application or in a certain product. The application requirements decide the overall section shape, alloy and characteristics like surface quality and shear strength of the section. The application design results in a draft section shape design. Based on the production limitations and opportunities a final section design is worked out in close collaboration between the customer, the die maker and the extruder. Based on the section shape the die maker designs a new die based on his experience and feedback from the extruder on earlier designs.

During this design process, there is an immense need for a software tool that identifies similar sections from which the section and die designers can draw experience.

A first version of a die is then produced and tested with some trial runs on the press. Corrections are made during these test runs until the extruder is satisfied. We now have a production-ready die. During normal production the die may be corrected again if needed. The die is preheated in a die oven before use. Each die can be used for a production lot of anywhere from one single billet to over 100 billets. The die is used in production until it is worn out, some part of it breaks or the customer does not order more sections.

If there is still a need for producing the sections, a second die is designed and produced by the die maker based on the experience made in production with the first die. For sections with a large production volume this loop may continue for a long time. Each new version

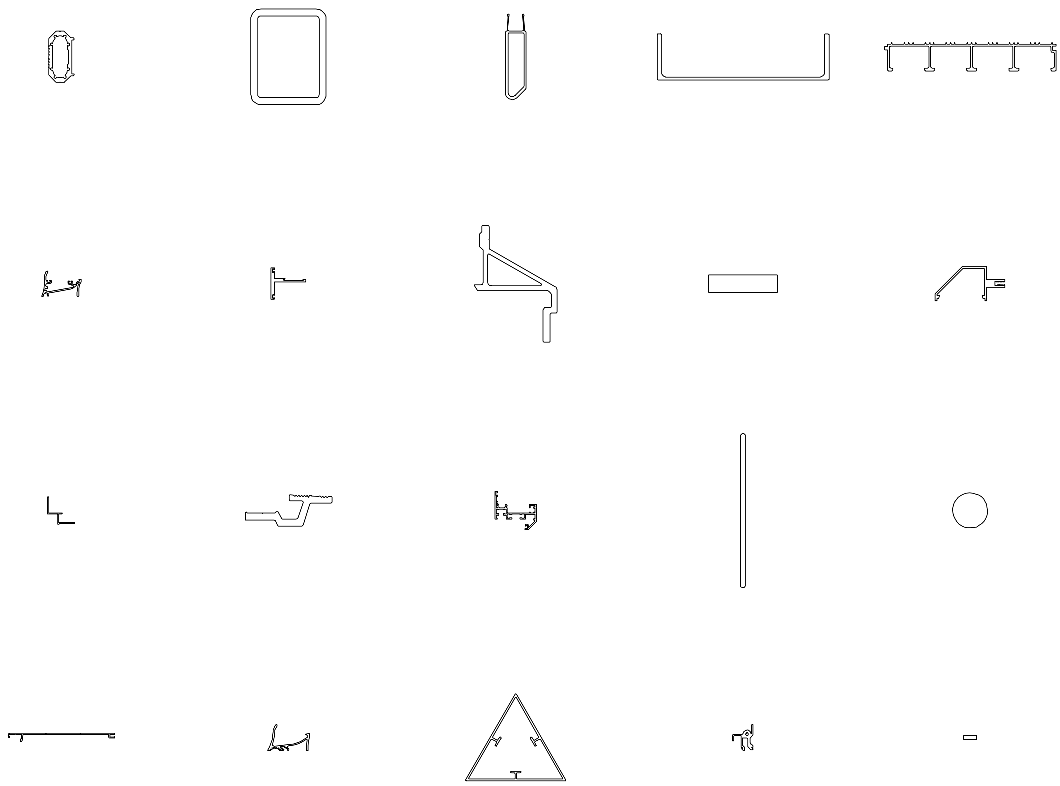

Figure 2.2: A small selection of extruded sections. Which shapes are similar and which are not?

may differ somewhat from the earlier versions, but the section shape remains the same. Over time the productivity will increase as new improvements of the die design are suggested. Most production parameters are measured and logged in a large production database. At an average extrusion plant about 4 000 different sections are produced each year, including about 1 000 completely new section shapes and dies each year.

A shape similarity retrieval software tool allows the designers and managers natural access to this large knowledge database by inspecting similar sections and dies.

2.1.4 This thesis and aluminium extrusion

The purpose of the applications of this thesis is mainly to enhance the die design process by enable feedback from the production process, minimise the number of trial runs and corrections needed before and during production. Another goal is to provide good initial estimates for the production speed and productivity. The thesis itself will only in small degree be concerned with aluminium extrusion specific problems, but will concentrate on the key task of shape similarity retrieval of Computer Aided Design (CAD) drawings.

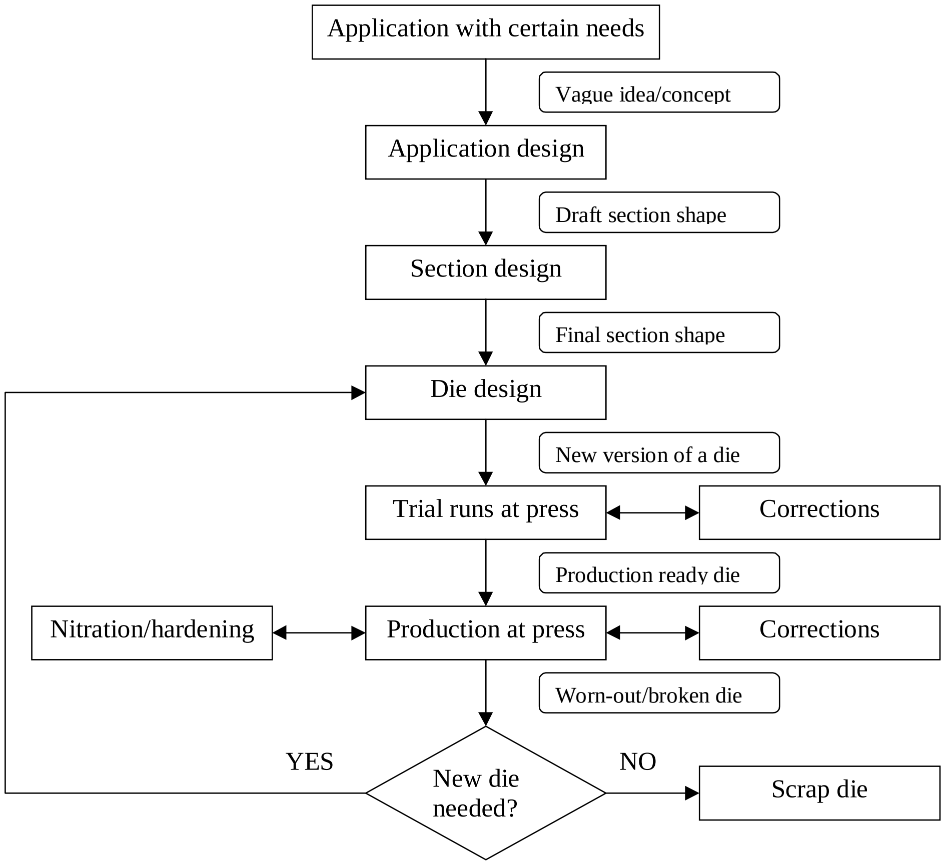

Figure 2.3: The die life cycle

2.2 The DieFinder

The DieFinder is a software tool that identifies similar dies for aluminium extrusion to enable the die designer to learn from previous designs. Most of the production in the industrial world is today computer aided using Computer Aided Design (CAD) and Computer Aided Manufacturing (CAM). *The next giant step within industry is to create working information feedback loops along the production chain.* The DieFinder is a tool primarily created to allow such feedback from production to the designer of the actual dies.

The main goal of the DieFinder is to search for similar die designs in a large database of CAD-drawings. In our case the dies are aluminium extrusion dies and in the current system these 3-dimensional dies are represented by a 2-dimensional lateral cross-section of the actual aluminium extrusion. The DieFinder has several applications including identifying good designs for new dies or learning from previous designs.

2.2.1 Data representation

The DieFinder as a database tool consists of two elements; a **production database** and a **collection of section CAD-drawings**.

The production database will not be described in detail in this thesis, but it contains all important production data and we may access these data for each section in the database. The collection of CAD-drawings is used to identify similar sections in the database. As most CAD-systems represent the shape information based on surfaces and curves in some CAD-system specific format, we have chosen an intermediate format called DXF (AutoCAD Data eXchange Format) which most CAD-systems support. From the DXF file we extract only the shape information as lines, curves and circle arcs and process these further from this level by approximating the shape contours by polygons. This is our basic representation for shapes.

2.2.2 Modes of operation

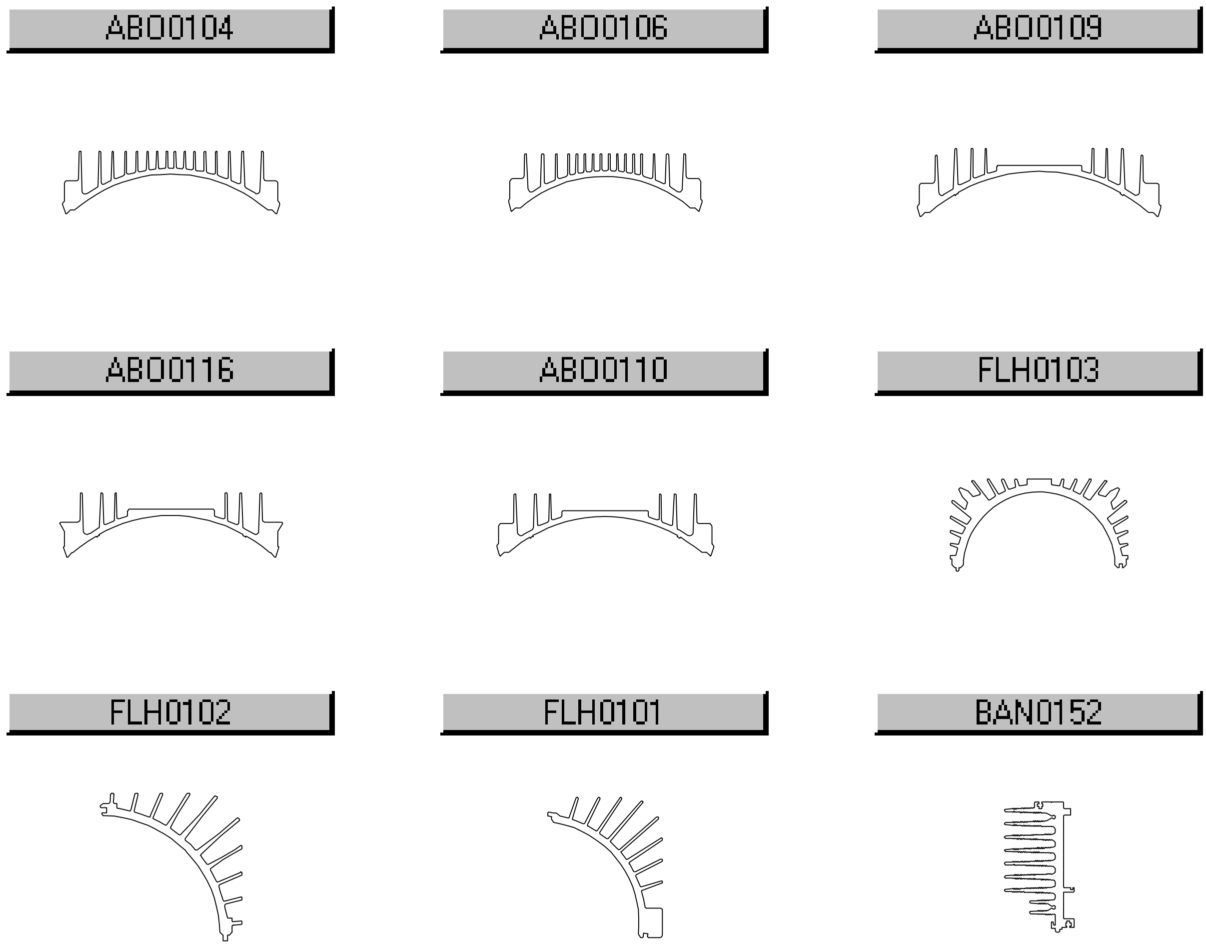

Figure 2.4: Query result for one specific section (ABO0104)

The query interface of the DieFinder can be operated in three different modes:

- **Manual search mode**, where the user may give each individual feature either none, low, medium or high weight. The manual search mode is suitable for searching for special details or sections as the user has full control of the weighting of the features.

- **Group search mode**, where the weights are predefined for searching for a special type of sections which appears often. One example is tubes, another example is rectangular bars, a third example is cooling elements with many fins.

- **Automatic search mode**, where the best possible general feature subset is used to identify similar sections.

The retrieval results are presented as the p-nearest matches with a measure of similarity ranking each object. By pushing a button above each image the essential data of the production database can be accessed. Figure 2.4 shows a typical retrieval result for the section ABO0104 which is a cooling element with 16 cooling fins. The 8 retrieved sections are all relevant since they are all cooling elements. It is obvious that ABO0106 is very similar to ABO0104, while the next three sections ABO01xx have the same basic shape, but with fewer cooling fins.

2.2.3 Feature based shape similarity retrieval

We have chosen a feature based approach to shape similarity retrieval since:

- Feature based shape similarity retrieval offers relatively fast indexing and retrieval in linear time

- Fast implementation due to the extensive literature on feature extraction

- Features can be interpreted as parameters of physical phenomena and hence provide us with deeper understanding of the process

We have focused on shape features with a physical or statistical interpretation:

- **General** shape features such as area, perimeter and diameter

- **Mass distribution** features based on geometric moments

- **Frequency distribution** features such as elliptic Fourier descriptors

- **Structural features** based on a skeleton representation such as tongues, fins and walls

- **Symmetry** as either mirror, rotation or radial symmetry and degree of symmetry

- **Fuzzy measures** describing specific basic shapes

- **Complexity measures**

These features are computed since we believe that most of them influence the aluminium section production process and hence are relevant to retrieval. As examples the mass distribution and the symmetry of a section are important for the metal flow in the die, while the perimeter length is roughly proportional to the friction between the aluminium and the steel die during extrusion. Sections with long and thin fins or walls are difficult to produce. The shape features will be described with greater detail later in this thesis.

2.3 The SpeedPredictor

The SpeedPredictor is a software tool for estimation or prediction of the extrusion ram speed for new aluminium extrusion dies based on some key parameters describing the geometry of the die and the production settings. The ram speed is a key production parameter as it is proportional to productivity assuming no production bottlenecks and constant accelleration and cycle times. The application could be altered to model any other parameter in the process through a process of learning the new dependencies.

There is an immense need in the aluminium industry for working prediction models for a number of key parameters of the press. The production on each die is influenced by the 3-dimensional shape of the die, the temperature distribution of the die and the aluminium, metal flow including friction between the aluminium and the die, material states at different temperatures and several other factors in a fairly complex model. The state-of-the-art today is to run large Finite-Element-Method (FEM) models for a single billet with a set of initial values of important parameters. These FEM models may contain several tens of thousands of elements even for relatively simple dies which have to be iterated through small time steps for each set of constraints. Even simple FEM models may take hours to run, must be fine-tuned to produce valid results and they still represent a simplification of the actual extrusion process. The FEM models are mainly used as a research tool to gain new basic knowledge about the extrusion process, not for estimating the key parameters of the press.

The SpeedPredictor provides two estimates for the ram speed. The first estimate is based on a regression model of the historical production values combined with some key geometric parameters describing the shape of each individual section.

The second estimate is based purely on historical production variables of resembling sections. Similarity retrieval is used to identify a certain number of sections sufficiently similar to the new one. The historical production data of these resembling sections are then used to estimate the ram speed for a new section. *The key assumptions are that similar sections are expected to produce in a similar manner and have a similar productivity.*

Chapter 3

Shape similarity retrieval

In this chapter we discuss important issues within shape similarity retrieval related to the contributions of this thesis. We also review the current state-of-the-art of commercial and scientific image retrieval systems and discuss alternative approaches.

3.1 What is shape similarity?

We start our discussion on shape and shape similarity by two different definitions of shape.

Definition 1 *The **shape** of a physical object is the external form or contour, the geometry of its external surfaces or contours, the boundary between the objects interior and exterior. Shape is the outline or characteristic surface configuration of the object.*

Definition 2 ***Shape** is all the geometrical information that remains when location, scale and rotational effects are filtered out from an object.*

The first definition is a lexicographical definition based on several dictionaries, while the second is a mathematical definition [10]. We will use the first definition as the starting point for our discussion on shape and shape similarity in relation to mathematical models, physical models and human perception. The discussion on the shape space is inspired by the work on shape spaces for landmark data [10], curvature-scale space [11] and reaction-diffusion space [12], but the formulation below differs from [10] as we use shape spaces for continuous contours with a finite basis, not shape spaces for landmark data.

3.1.1 The shape space

Definition 3 *The **shape space** is the space spanning all possible configurations of shape.*

The shape space is here a purely abstract concept, it is no well-defined mathematical vector space. Any mathematical shape representation forms a basis for a well-defined vector shape space. Mathematical shape representations include polygons, splines, skeletons or polynomial curves. The vector shape spaces are restricted compared to the abstract shape space since they only span those configurations of shape which can be represented by the specific representation.

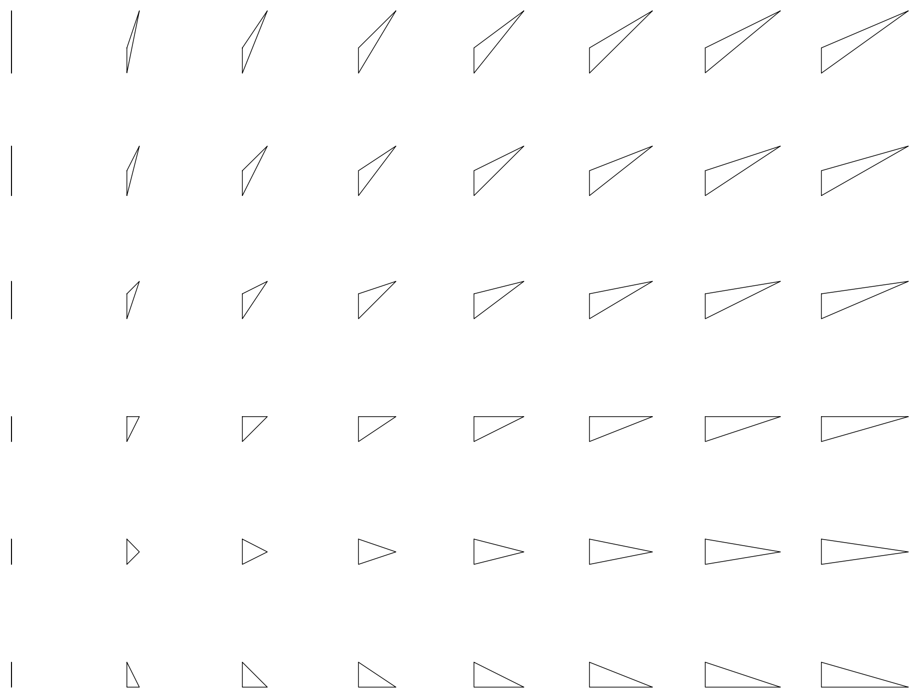

Figure 3.1: Sample instances from a plane in the planar triangle shape space.

We may illustrate these ideas in the **planar triangle shape space** which spans all possible triangles. The ideas can easily be transferred to a more general shape space. Each point in the triangle shape space may be represented by the three vertices of each triangle $((x_1, y_1), (x_2, y_2), (x_3, y_3))$. Fig. 3.1 [10] shows some samples from a plane in the triangle shape space with two fixed vertices (v_1, v_2) and a varying vertex (v_3). Notice that row 1 and row 3 from the bottom are reflections of each other and hence similar under a similarity transformation.

In our case we are working in a **planar polygon shape space** using a polygon basis. It is obvious that the planar polygon shape space is a subspace of the more general shape space. Working with a finite basis, there are many shapes which cannot be represented exactly in the planar polygon shape space, such as circles, circle arcs, ellipses, polynomial arcs etc. On the other hand we are able to represent all types of polygons in the polygon shape space. We are only interested in simple polygons where no non-adjacent edges intersect, since no physical objects will have contours that intersect.

3.1.2 Exact and approximate shape similarity

Exact mathematical similarity is always defined with respect to a set of geometric transformations. In our case *we consider similarity with respect to a similarity transform*, which includes translation, rotation, reflection and scale, *or a isometric transform*, which includes only translation, rotation and reflection. Other common sets of transformations are the Euclidean transformation (only translation and rotation), the affine transformation (translation, rotation, reflection, scale, skew and shear) and the projective transformation (including a viewpoint and a horizon).

Mathematical similarity and invariance are closely related. Similarity is a measure on the objects under geometric transformation, while invariance is a measure of a single property of the object under geometric transformation [13].

The Greek mathematician Euclid assembled most of that period's knowledge of mathematics in his astonishing 13 volume work *the Elements*. The Elements was reassembled from the different surviving pieces and fragments a century ago. Euclid covers exact triangle and polygon similarity in his sixth book, proposition 19 and 20 [14]. Euclid only considered exact similarity under a Euclidean transformation, including translation and rotation. Since then, exact similarity under a range of geometric transformations has been explored and fairly well understood.

Definition 4 *Exact shape similarity under a similarity transform creates a* ***similarity shape subspace*** *to the shape space upon which all shapes can be transformed. All similar shapes are mapped to the same point in this subspace.*

Formally this subspace is an orbit space of shape configurations in R^n under the action of a similarity transformation, see [10]. Using different geometric transform, we may create the Euclidean shape subspace, the affine shape subspace or the projective shape subspace.

When we discuss shape similarity, *it is implicit that we are discussing approximate shape similarity*. Approximate similarity is poorly understood and this fact directly influences the shape similarity retrieval problem and makes it a fairly difficult one. When we try to identify approximate shape similarity, we may think of it as creating approximate measures in the shape similarity subspace above.

3.1.3 Shape and causal history

One way of regarding shape is that shape is the result of the causal history of the object, the result of how the material was shaped [15]. We may actually identify the minimal deformation between two objects and use the size of the deformation to judge the similarity of the two objects. In addition the process that shaped the objects influences the distribution of shapes in shape space. *Two objects that have been exposed to a similar set of deformations are often similar in shape*, see fig. 3.2.

3.1.4 Shape and functionality

Not only the shaping of the object is important for the concept of shape, but also the functionality of the object. A fish has fins to be able to swim, not only because the fish was shaped with fins. A knife has a sharp edge to allow us to cut with it, not only because it was shaped as it is. Fig. 2.4 shows that cross-sections of aluminium cooling elements have a large perimeter (large surface in 3 dimensions) and many fins to allow effective heat transfer from the element. Fig. 3.3 shows that cross-sections of aluminium building sections have a similar structure.

The functionality of the object does to some extent governs the shape of the object. The functionality of the objects influences the distribution of shapes in shape space. *Two objects that have a common use and functionality are often similar in shape.*

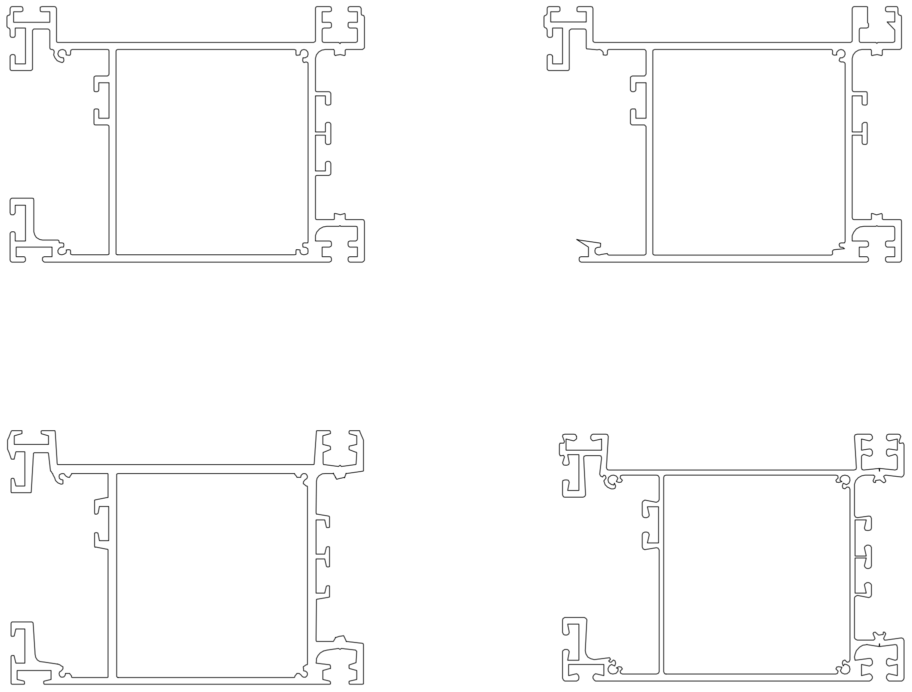

Figure 3.2: These objects are perceived as similar despite the deformations applied;

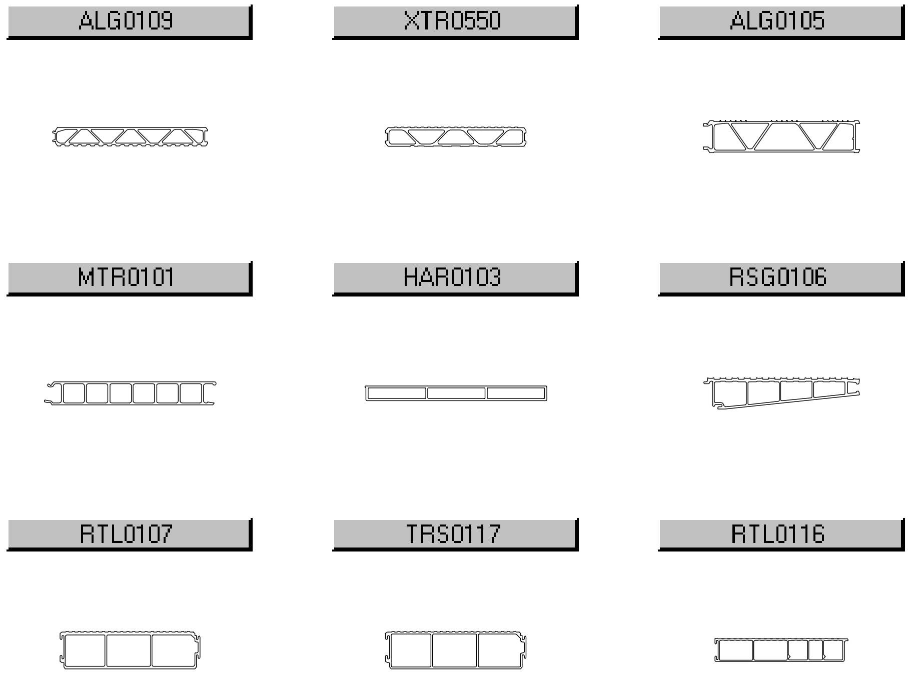

Figure 3.3: Building sections are similar due to their functionality.

pretation. It is easy to apply a distance measure in the feature space which measures similarity. *Two objects that are represented by the same set of invariant features are similar with respect to this feature set.*

The distance between two objects in the feature space indicates how similar they are in terms of those features spanning the feature space. **Shape similarity has become measurable.**

We believe that the properties of the mapping between the shape space and feature space is very important to be able to understand shape similarity in a feature-base shape similarity retrieval system.

3.1.7 The shape similarity retrieval problem

The basic problem that we try to solve in this thesis can be described as

Problem 5 (Shape similarity retrieval) *For each query template we apply a distance measure d_k to each object k which ranks the objects with decreasing shape similarity. Retrieve those objects with smallest distance to the query template.*

What shape similarity retrieval requires in practise, is further discussed in the first paper of this thesis in conjunction with our quest to establish valid performance measures for shape similarity retrieval.

3.2 Shape similarity retrieval

This section provides a general description of a general-purpose shape similarity retrieval system with its specific building blocks. The system is designed for the task of retrieving similar

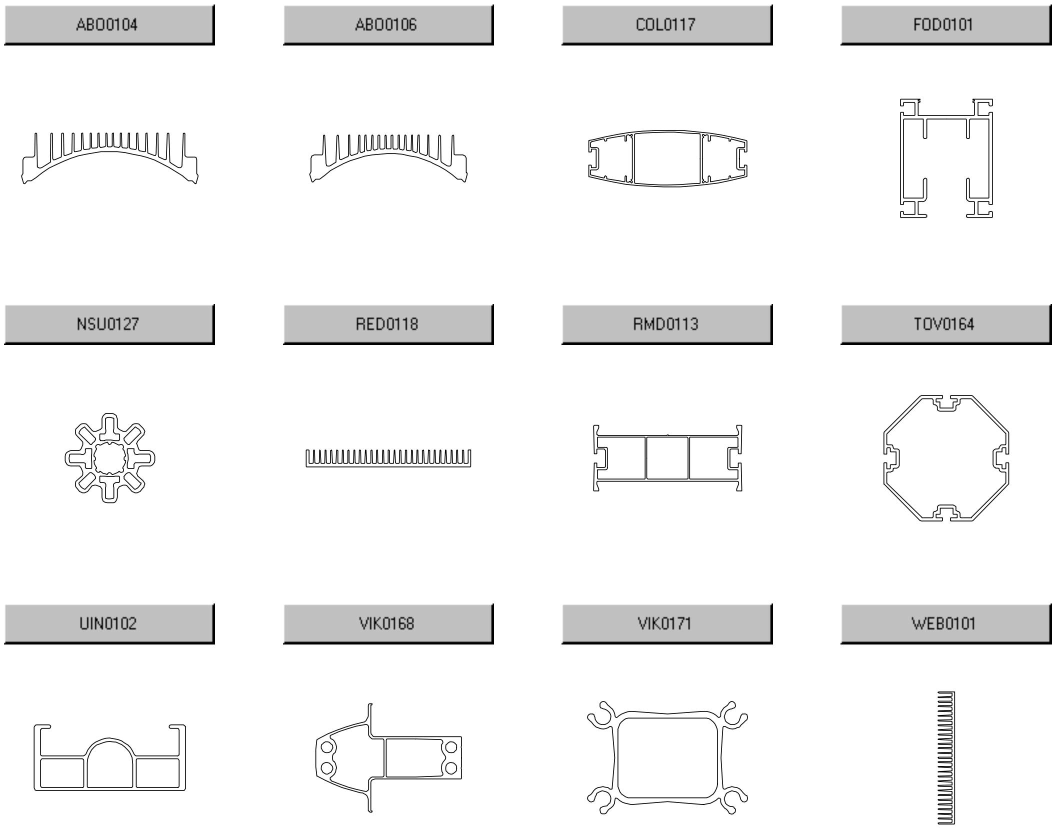

Figure 3.4: These objects are similar with respect to the two features **mirror symmetry** and **perimeter length**.

shapes when presented with a new 2-dimensional object. We only consider shapes represented by parametric contours, not objects represented by images. The system requirements are described from a methodological viewpoint discussing important technological and scientific issues.

We have already described the application and the intentions of the retrieval system in section 2.2. Here we describe the system requirements of each individual subsystem, such as contour extraction, representation transformation, feature extraction, feature selection, feature clustering, different matching techniques and evaluation of cluster or similarity measures.

The purpose of this section is to give the reader an introduction to shape similarity retrieval in general and which alternatives exist to our implementation.

3.2.1 System requirements

The available input data to our shape similarity retrieval system is a set of Computer Aided Design (CAD) drawings. The geometric shape information in the CAD drawings must be extracted and processed to build a shape similarity retrieval system. The important building blocks of the shape similarity retrieval system can be divided in information refinement blocks and matching blocks (see fig. 3.5).

The major information refinement blocks are contour extraction, representation transforma-


Cluster space

Contour extraction

First we must automatically extract the contour from the original format. The contours must be represented in a uniform representation to enable comparison. The basic shape representation should abide these criteria:

- The representation should be accurate enough for our purpose
- The representation should not introduce any large systematic or random errors when subject to later processing
- The representation should preserve important measures of the original shape such as symmetry and area or perimeter length.
- The representation should easily transform into other suitable representations.

In our case the CAD drawings contain parametric descriptions of the contours of the individual sections as a list of geometric elements such as line segments, circle arcs and different kinds of smooth splines (polynomial curves). These parametric segments are linked together and joined into chains of curve segments for each contour. We approximate the original geometric elements by line segments producing a single polygon for each contour with a maximum approximation error for circle arcs and splines. There exist formal techniques for assessing polygonal approximations to curves [18], but these do not take into account the further processing of the contours, putting weight on approximation error, not on preservation properties.

There are mainly four sources of errors:

- There may be **man- or machine-made errors** in the drawing compared to the actual produced object. In such cases the map does not match the real world.
- There may be **inconsistencies** in the contour representations, in the form of gaps, loops or superfluous line segments. CAD systems today do not always check the consistency of the data on this level.
- The representation will have small random errors due to **rounding** and the **finite precision** of computer based systems.
- The **approximation** in itself is producing **small systematic errors** due to the difference between approximation and the original parametric contour.

Representation transformation

The polygon contour representation is a boundary representation in the spatial domain. The polygons can relatively easily be transformed to structural representations such as skeletons, mass distribution representations, frequency representations, pixel representations (see fig. 3.6) or deformation representations like the curvature-scale space.

The basic requirements to the representation transformations are given below.

- The transformation must be **reversible** within some prespecified accuracy

- The **geometry information must be preserved**

- The new representation should make it **easier to extract important information** about the object

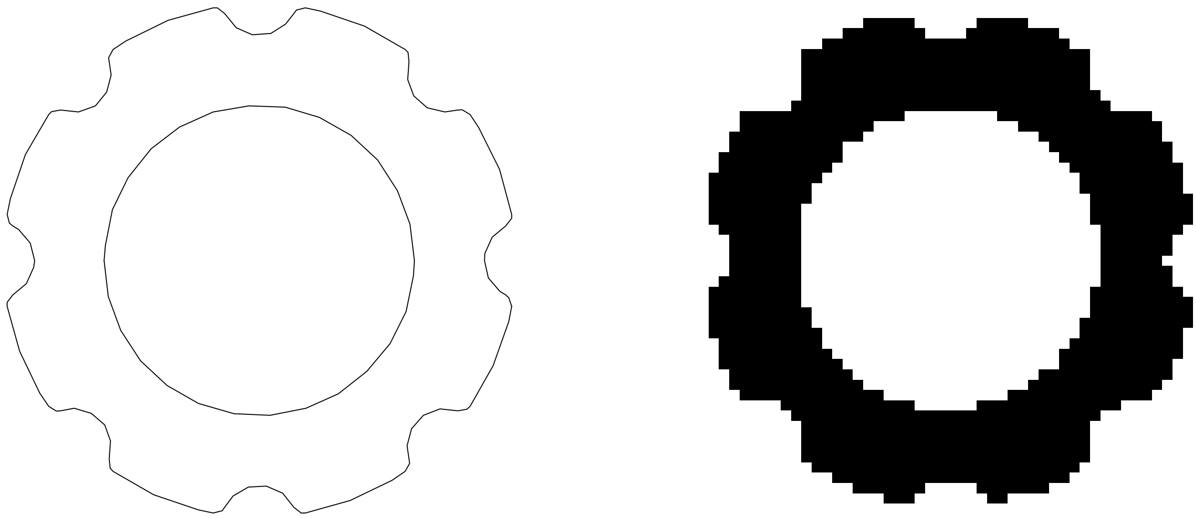

Figure 3.6: Transformation from a polygon to a pixel representation

The medial axis or the skeleton of an object is an interior representation of the original object [22]. We have implemented routines for transformation of a polygon to a skeleton utilising the generalised Voronoi diagram [23]. Each point in the plane can be said to be uniquely closest to a set of points on the contour. A generalised Voronoi region consists of those points in the plane, which are uniquely closest to any point on a generating element. In our case the generating elements of the polygon is its line segments and its vertices. The generalised Voronoi diagram consists of the borderlines between the generalised Voronoi regions. The points on the generalised Voronoi diagram are equally close to at least two generating elements [24]. The generalised Voronoi diagram can easily be reduced to a skeleton by removing all edges touching the reflex vertices of the original polygon, see figure 3.7. Every point on the skeleton is represented by a co-ordinate position and a radius. The radius is the distance to the closest points on the contours. The radius information can be used to reconstruct the original contours from the skeleton points, as we can observe in the figure. The skeleton representation allows us to extract structural information in an ordered manner.

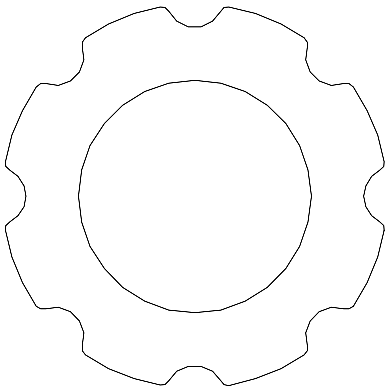
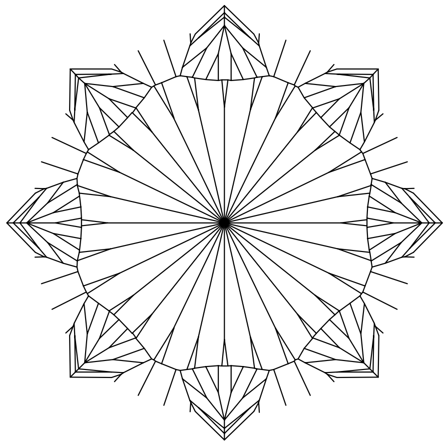
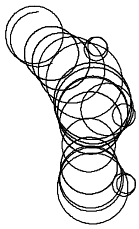
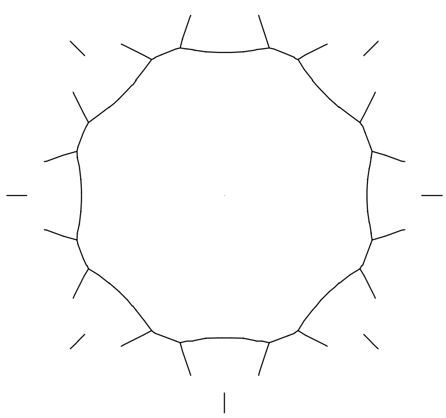

- The feature should have a **global scope**, i.e. cover all possible shapes [82].
- The feature should be computed by a **robust and stable algorithm** that produce accurate results.
- The feature should contribute to the **indexability** of the objects.

The third paper of this thesis provides a further discussions of the requirements to the features and how these requirements can be used to assess the features. The requirements are discussed with a focus on the mapping from the shape space to the feature space.

Feature transformation

Feature transformation covers two important tasks, normalisation and construction of new features. Each feature should be normalised, either to a unit interval or by zero mean and unit variance. In some cases the feature will also have an unwanted distribution. A common approach is to perform a non-linear feature transformation to a well-known and wanted distribution [25], such as the uniform distribution, the normal distribution or a Gaussian mixture distribution. We should however be careful to apply a non-linear transformation to a feature as we may obscure the possible interpretation of the feature. It is also possible to construct new features as linear combinations of the existing features or non-linear combinations. The last option is not explored in this thesis though.

Feature selection

It is easy to compute many invariant features. Some of these features will in general contain the same information and hence be redundant, while others will be irrelevant for the shape similarity retrieval problem. Some features may even contain no information at all, only noise and perturbations, although this is not very likely. We need methods for evaluating each feature independently and assessing them with respect to other features. Our goal is to select the most representative and minimal subset needed for performing our search for similar 2-dimensional objects, to create a reduced feature space with certain properties.

Some possible requirements to the feature subsets are listed below.

- The **relevance** of each feature [78].
- The linear and non-linear **dependencies or correlation** of each feature with respect to the others (**redundancy**).
- Each feature should be **robust to noise** and certain types of **degradations** of the original shape.
- The **information content and granulation** of each individual feature.
- The **spanning and distribution** of each individual feature.

In the second paper we select feature sets based on relevance and redundancy. In the third paper we discuss the issue of assessing individual features further using many of the above requirements.

Feature clustering

We may perform a clustering of the data in the original or the reduced feature space. The aim of the clustering is to produce classes or groups of objects that are similar to each other.

There are some requirements to the clustering algorithm [26].

- The clustering algorithm must handle **outliers** in the shape space. Outliers are here considered to be shapes that are dissimilar from most of the other shapes in the database.
- The clustering algorithm should allow **graceful degradation when subject to noise and perturbations** in the shape space or the feature space.
- The clustering algorithm should **identify the number of clusters** in the data set and the members of each cluster. The real number of clusters is generally unknown and also the number of samples in each cluster.
- The clustering algorithm should be able to **integrate different data types** such as Boolean, integer and real valued features.
- The clustering algorithm should **handle irrelevant features** with respect to the clustering. The clusters may exist only in a subspace of the reduced feature space.
- The clustering algorithm should be able to **identify the structure** of the clusters. Some assumptions about the structure of the clusters must however be made in advance.

Two common approaches to clustering are **template-based clustering** which identifies templates for the cluster centers and **hierarchical clustering** which organises the feature space in a hierarchy of partitions.

A clustering method is the correct method of choice if all shapes can be classified in distinct classes. We may easily evaluate the performance of clustering algorithms [27], since there is a straightforward answer to whether a shape belongs to a cluster or not.

The main disadvantage of clustering methods is the imposed sharp boundaries between the clusters. There is not necessarily a good correspondance between automatic clustering and human perception of shape similarity [28]. By applying fuzzy techniques, a degree of membership to each cluster may be specified, creating softer boundaries between the clusters [29].

Clustering is in general very difficult to do in high-dimensional spaces due to **the curse of dimensionality** [30]. A high-dimensional space is much more sparsely populated than a low-dimensional space containing the same number of observation points. To obtain the same density, the number of observations must increase exponentially as a function of the number of dimensions.

There is also the philosophical question whether shape similarity retrieval is a clustering problem. In general, we may argue that the shape space is inhabited by a continuous spectrum of shapes. It is impossible to cluster a continuous spectrum of shapes in a meaningful way. In some applications, the shapes will be unevenly distributed in the shape space, resulting in an uneven density concentrating many of the shapes as clusters in shape space. We believe that clustering is application-specific and that the clustering results will be invalid for other populations of shapes in shape space. We have therefor not examined clustering in this thesis.

3.2.3 Matching blocks

The matching blocks perform the similarity search when presented with a new drawing. These blocks create a similarity measure in one of the possible representation spaces, either the shape space, the feature space or some relevant subspace of these spaces (e.g. a cluster space or a feature subset space).

Template shape matching

By a template shape matching technique we mean a method that measures the similarity of an object to a template directly in the shape space. Template shape matching can be performed using a number of alternative methods:

- **Set theoretic measures** by aligning the template with each object in the database and apply a set theoretic measure, such as the Hausdorff distance [31], to measure the distance between two shapes.

- **Deformable templates** by aligning the template with each object in the database and identify the necessary deformation transforming the template shape to equal the sample object shape [19, 20].

- Different **scale spaces**, such as the curvature-scale space [11], the entropy scale-space [21] or the reaction-diffusion space [12], which allow pre-computation and a hierarchical approach to the matching.

- **Reduced resolution approaches** allow a hierarchy of coarse to fine approximations to the original objects [32].

We have already stated that it is not preferable to perform a template shape matching due to the following facts:

- We have insufficient knowledge of similarity in a shape space

- It may become very time consuming to perform a template shape match.

Hierarchical matching techniques can speed the search process, but the remaining open question is how most of these techniques are related to our perception of similarity. Which deformation or set theoretic measure corresponds to our perception of similarity?

Graph-matching

By disregarding the geometric information, the skeleton can be seen as a graph describing the basic structure of the object [33]. We may match these graphs applying graph-matching techniques. Graph-matching can be done fast using reduced resolution hierarchy with the major structure on top and minor details at the leaves or the hierarchy. This approach is conceptually different from most of the above approaches since we remove the geometry and focus on the structure of the object.

Feature matching

Feature matching is performed in an n-dimensional feature space. Each object is represented as a point in the feature space. Feature matching is performed by applying a distance measure to the points representing the objects in the database and measuring the distance to the point representing the query template. Those objects with a small distance in the feature space are considered to be similar in terms of those features spanning the feature space [74]. In the second paper we discuss the issue of selecting features and a distance measure based on redundancy and relevance. An alternative to feature selection is to use an individual performance measure to weight each individual feature.

3.2.4 Performance measures

We need a performance measure for shape similarity retrieval to enable comparison and evaluation of the results of the retrieval. In our first paper we established seven performance measures for shape similarity retrieval.

1. **Subjective ratings** based on our perception of shape similarity between a reference set of objects and the other objects in the database.
2. **Relevance feedback** based on the perceived relevance of the retrieved objects [75].
3. A **contextual classification** of the objects by a common dictionary.
4. Non-geometric **contextual variables** which are related to shape similarity.
5. An expert may use his **interpretation** of the shape features to establish a rating.
6. The **consensus** of several independent representations.
7. **Deformation properties** of the objects.

Only relevance feedback has been thoroughly explored in the image retrieval literature [76]. We have applied some of the above performance measures to compare the performance of different groups of features in the first paper, to select feature subsets in the second paper and to assess individual features in the third paper.

3.3 Image retrieval systems

In this section we provide an overview of the state-of-the-art of image retrieval and list some of the current commercial and scientific image retrieval systems. The aim is to provide the context for our research in a broader sense.

3.3.1 The difference between image retrieval and vector graphics retrieval

Shape similarity retrieval is related to the much wider problem of image retrieval. Shape is combined with colour, texture and composition in most image retrieval systems. There are some fundamental differences between images and vector graphics, like Computer Aided Design (CAD) drawings.

- Most digital images are planar perspective projections of a 3-dimensional world, while vector graphics exactly represent 3-dimensional objects or cross-sections. Since much of the shape information is lost during the projection due to occlusion and single views, vector graphics contain a much more complete shape description of the objects.
- Digital images are composed of a matrix of pixel values, while vector graphics are composed of a parametric description of the contours or surfaces of the object. The sampling resolution and the number of quantisation levels of each pixel limit the accuracy and introduce noise in images, while vector graphics contain more precise shape information.
- The shape information is mixed with colour, texture and intensity information in an image and must be extracted by segmenting the different objects from other objects and background in images, while there is no need for segmentation in vector graphics since the information is object-oriented and non-intersecting.

These 3 factors make it easier to do shape similarity retrieval in a vector graphics representation than in an image. *Vector graphics contain exact shape information about every object.* The literature is very scarce on shape similarity retrieval using vector graphics [34]. We must rely on the much wider literature on image retrieval in our presentation of state-of-the-art in this field.

A few image retrieval applications are closer to vector graphics retrieval than general image retrieval. Scanned images of planar objects like paper documents [41], technical drawings or company trademarks [42] are examples of images without projection and occlusion. The image resolution and segmentation problems still remain, but are easier overcome than in other images.

Our shape similarity problem is different from most other image retrieval problems in another aspect as well since *our objects cannot be grouped consistently in valid classes*. Text retrieval or recognition is an example of a retrieval problem where the letters are easily grouped in a fixed number of classes. Much of the work on text retrieval is therefor directed towards features relevant for letter classification and discrimination. Our shape similarity retrieval problem contain no classes since almost any shape could be formed in aluminium. Much of the literature is therefor partly irrelevant for our problem.

3.3.2 Shape, colour, texture and composition

Most image retrieval systems are based on extracting four distinct types of features; **shape**, **colour**, **texture** and **composition**.

Definition 6 *The **shape** of a physical object is the external form or contour, the geometry of its external surfaces or contours, the boundary between the objects interior and exterior. Shape is the outline or characteristic surface configuration of the object.*

Shape of an object must be segmented from the background and other objects and then represented by a smooth shape representation. Shape is extensively used in retrieval of documents, technical drawings, trademarks and pictograms.

Definition 7 *The **colour** of a physical object is our perception of the spectral characteristics of the light that is emitted by, reflected by or transmitted through the object in the visible electromagnetic spectrum.*

Colour is dependent on the spectral characteristics of not only the material of the object itself, but also the light source, the incidence angle, the optical components and the sensor when recording the image. Three components, red, green and blue (RGB) commonly represent colour. It is possible to transform to an invariant colour representation such as hue, saturation and value (HSV). Hue describes the distinct colour, saturation the level of colour saturation and value describes the total intensity of the three components together. It is common to use the dominant colour, the colour histogram and the colour composition as retrieval criteria [35, 36].

Definition 8 *The **texture** of a physical object is its appearance, considered as the quality and structure of the material or substance it is made of. In an image, texture is the homogeneous spatial variation of intensity over an area on the surface of an object.*

The word texture is formed by the Latin word for web or loom. Texture may appear as a result of the interaction of light with either a variation of the spatial material properties of the object or a spatial surfaces variation, the microstructure of the shape of the object. Texture is often described by terms like smoothness, roughness, coarseness etc. Texture can be characterised by the frequency content of the intensity of an image [37, 38, 39].

Definition 9 *The **composition** of an image is the way in which the parts or objects are arranged, especially how much of each part there is.*

Composition is used in combination with colour [40, 43] or texture.

Definition 10 ***Spatial information** is information about the relative position and orientation of the object with respect to the other objects and the viewpoint of the image.*

The most common features used today for general image retrieval are colour, texture and composition. These features are fairly well understood. Colour and texture are material properties of the objects in the images. Natural objects will often have very characteristic material properties. The sky has a limited set of colours. Grass has a limited set of textures. It is reasonable to believe that images of natural objects can be retrieved using material filters together with information about the composition. The sky tend to be at the top of the image, grass at the bottom.

Artificial objects are however often characterised by their shape and structure, not solely by their material properties. Cars and houses come in many different colours, clothing may have many different textures. In our application with a large database of Computer Aided Design drawings of aluminium sections, only shape information is available. The colour and texture does not matter at all.

Natural objects tend to be self-similar in the sense that fractal and randomised variations in the patterns repeat, while man-made objects in general tend to be more structured. Artificial objects can be described efficiently by classical Euclidean geometry, while many natural objects are too fragmented and irregular to be described by the same geometry. Natural objects are often best described by Mandelbrot fractal geometry [88].

3.3.3 State-of-the-art within image retrieval

Below follows a brief description of some of the existing image retrieval systems.

- The **Querying-By-Image-Content (QBIC)** system was developed at IBM Almaden Research Center during the 1990s. It has been tested on a large number of varying data sets and is probably the most extensive system in the world today. Test data includes stamps, bitmap trademarks, stock photos and art collections. The simple version of QBIC relies heavily on the colour and intensity of the retrieved images, while the advanced version also includes texture and topics. Colour information is retrieved as main colours, colour histograms, colour composition and layout. Much of the work on QBIC is public knowledge due to extensive publishing from the research group [44, 45, 46, 47, 48, 49, 73]. The research group have a licence agreement with Magnifi and Virage.

- **Visual Image Retrieval Engine (VIRAGE)** delivers 3 different products within content-based image retrieval; a video cataloger, a media manager and browser and an **Image Read/Write (IRW) toolkit**. Virage is the only truly commercially available product of high standard known to the author. It relies on colour, composition, texture and structure information. Colour information is identified as hue, saturation, tint, dominant colours,

- **Just Another COntent Based retrieval system (JACOB)** is aimed at video databases. JACOB combines colour and texture in combination with frame-to-frame motion to retrieve video sequences [52].

- **Comparison Algorithm for Navigating Digital Image Databases (CANDID)** is another system aimed at digital image databases. CANDID combines colour, texture and shape in feature signatures [55].

- **Multimedia Analysis and Retrieval System (MARS)** uses colour layout, texture and segmented shape to retrieve images from a database [65].

- **Chabot** is a system for annotated image retrieval and use manual indexing to obtain high-level descriptive texts [58].

- **Shape Queries Using Image Databases (SQUID)** is a system for retrieval of fish images based on silhouette images of fish. SQUID use the curvature-scale space to identify similarities between the different fish types. Contour data of the different fishes is available.

- **WebSeek** and **VisualSeek** at the Columbia University offer content-based image and video search and catalog tools for the world-wide web. They identify video, colour and greyscale images as well as graphics and use colour set, histograms and composition as their main features. WebSeek organises the images by a subject description (e.g. sports, news, art, astronomy) [53, 54].

- **SurfImage** is a content-based image retrieval system with relevance feedback developed at the French research institute INRIA at Rocquencourt. It is tested on a nice collection of greyscale images and texture images, including town and facial images.

- **Excalibur** uses a combination of colour, texture and shape features. Compositional information about brightness and colour is used extensiely.

- **PicHunter** is developed at the NEC Research Institute. PicHunter is based on fairly simple feature extraction techniques, including global statistical and frequency properties and colour percentages. PicHunter use Bayesian probability theory combined with relevance feedback to optimise performance.

3.3.4 Applications of image retrieval systems

General-purpose image retrieval is not yet state-of-the-art, most working systems have been adjusted or designed for specific applications. In this sections we list a number of possible applications and group them.

- In **architecture** there are many possible applications. Building architecture differs from interior architecture. Some applications are due to the need for city planning or landscape modelling. Many architectural items are available as images, some as technical drawings. Architecture museums may want to make historical images, drawings and paintings available.

- **Art museums and galleries** want to make their material accessible to a larger audience than those available to attend the museum or gallery. Basic datatypes are drawings, paintings, sculptures and photographs, but any type of visual art may be presented.
- In **archeology** most items are today photographed and to make archeological image databases searchable will speed the research in this field.
- Another field of interest is **biometry** used for personal identification, either to restrict or allow access, as a signature or to search for criminals. Finger prints, palm prints, iris recognition, voice recognition and face recognition are fields which require a high degree of specialised algorithms and methods.
- **Paper documents** still dominate the archive systems of the world. In many cases there is a need to transfer the data in the documents to electronic format. Image retrieval system can be used to identify hand-written or type-written documents, including as different data types as letters, formal documents, money, checks, forms, logos, trademarks, illustrations and maps, to mention only a few. Paper documents are purely man-made, but may contain images with natural content.
- **Electronic documents** are becoming increasingly common. The technology for searching through text is reasonably mature, but there is lack of algorithms for retrieval of digital images and technical drawings. Our application is to technical drawings of mechanical parts.
- **Geographic information systems** include airborne and satelite images with military, environmental, communications and economic applications. A special type of images is astronomical images. Many satelite and astronomical images are multispectral images.
- In **medicine** large image database are built with pasient data as X-ray images, nuclear magnetic resonance images and other images.These images mainly display biological tissue with a few man-made elements like protheses.
- **Nature** images are common in many scientific databases of species and documenting research. Flowers, plants, animals, insects, birds and fish species are documented by images. Landscape and scenery are documented by panorama images.
- Large databases of **portraits of people** exist both with historical and criminal individuals.

Chapter 4

Summary of papers

The first three papers are concerned with the main issues of this thesis, establishing performance measures for shape similarity retrieval methods, using these performance measures to establish relevant and non-redundant feature subsets with improved performance and finally assessing the individual features importance with respect to shape similarity retrieval.

The last two papers are minor contributions that were formed when I encountered some unsolved problems during the study. The first paper is concerned with the accuracy of the computational methods for geometric moments, while the second outlines a new complexity measure based on a fractal dimension estimation technique.

4.1 First paper: Measuring the performance of shape similarity retrieval methods

This paper has been submitted to the journal *Computer Vision and Image Understanding*, special issue on empirical evaluation of computer vision algorithms.

Motivation

To enable systematic improvement of some computational method one must be able to measure the performance of that computational method. In shape similarity retrieval the common measure is relevance feedback [75, 76]. Relevance feedback is based on performing shape similarity retrieval and counting the number of hits and misses of each new search. The results of relevance feedback can be used to adjust the retrieval algorithm. Our main idea was to use several sources of information to measure the performance of the shape similarity retrieval methods, not only our perceived relevance. We also wanted to rate the similarity of different objects on a continuous scale, not only as hits and misses. When we have a reliable performance measure, we are able to measure the performance of shape similarity retrieval systems.

Methodology

We have created a framework for measuring the performance of shape similarity retrieval methods based on three different knowledge sources; Human perception of similarity, the application-specific information and the mathematical representation of shape. The measures are based on measuring the difference between a number of reference objects and all the other objects in the database. Seven different measures are presented and four of them are tested on a database of Computer Aided Design drawings of aluminium extruded sections.

We apply the different performance measures to different specific groups of features, including skeleton-based features, three different kinds of moments, elliptic Fourier descriptors, fuzzy/symmetry features and a mixed set of features. To each specific group of features, a common group of five important general shape features were added providing basic information about each shape.

Results

A performance measure based on the productivity of the aluminium sections provides completely different result than the three other performance measures, based on perception and customer identification. This result indicates that production similarity and subjective similarity captures different information. The five common features explain much of the variation in productivity, but we may gain somewhat by adding moments. This is not surprising as the moments can be interpreted as physical properties of the material flow during production of the aluminium sections.

The mixed set of features performs better on both the subjective ratings, but is surprisingly bad on the productivity measure. The mixed set of features is obtained by selecting features that are important for subjective similarity. This result confirms that productive similarity is different from subjective shape similarity.

The fuzzy/symmetry, the skeleton and elliptic Fourier descriptors perform well on the subjective ratings and customer identification performance measures.

The set of Hu and Li invariant moments has the worst overall performance of the different feature sets compared. The invariant moments are fairly popular within object recognition, but have little impact on the quality of shape similarity retrieval beyond the five common shape features. These results have been obtained on almost noise-free data, but there is no reason to believe that the invariant moments are more robust to noise than many of the other features. We therefor question the use of these moment invariants for shape similarity retrieval.

Contribution

We have developed a framework for measuring the performance of shape similarity retrieval methods and systems based on three different sources of knowledge; Human perception, application-specific knowledge and the mathematical representation of shape. These performance measures have been used to evaluate different specific groups of features. The comparison shows that moments are important for productivity, but not for human perception, while skeleton features, elliptic Fourier descriptors and fuzzy/symmetry features all improve performance if the performance is measured by a perceptual performance measure. In the next two papers the performance measures of this first paper are used systematically to select optimal subsets of features and to assess which individual features are most important for shape similarity retrieval.

4.2 Second paper: Selecting feature subsets and distance measures for shape similarity retrieval

This paper has been submitted to the journal *Pattern Recognition*.

Motivation

After establishing performance measures for shape similarity retrieval, the next important task becomes to optimise the shape similarity retrieval with respect to the performance measure. In a feature-based shape similarity retrieval system, we wish to select a relatively small subset of features with optimal performance. Since the distance metric used in the feature space influences performance, we thus also want to select an optimal metric. These are difficult tasks since the combinatorial problem of selecting feature subsets is increasing exponentially with the number of available features. Recent cover papers have stated that there is lack of papers testing feature subset selection methods on large real-world problems [77]. This paper addresses these issues.

Methodology

We first define relevance with respect to shape similarity retrieval. Our definition of relevance differs from recent papers on relevance [78, 79]. Then we define redundancy with respect to subsets of features based on covers. Our covers are inspired by the use of blankets to select optimal feature subsets for classification tasks [80], but are defined for continuous variables. These two definitions are used to select feature subsets and distance measures for shape similarity retrieval. We use forward selection, backward elimination, hybrid methods and stochastic methods such as genetic algorithms and simulated annealing. The methods are tested on a database consisting of 113 different features computed for 1686 Computer Aided Design drawings.

Results

The results in the paper are restricted to a subjective performance measure for clarity of presentation. Applying a robust metric instead of a Euclidean metric result in a 4.4 % performance improvement using all 113 features and even higher performance improvement for the selected feature subsets.

The genetic algorithm produce the best feature subsets for retrieval reducing the size of the feature subset by 80 % and increasing performance 12.5 % compared to using the full feature set. Forward selection yields significantly better results than backward elimination in all our tests, while hybrid selection from a small random feature subset excels compared to forward selection and backward elimination.

The selection methods based on relevance are superior to those based on redundancy, but by combining the two measures we are able to identify relevant and non-redundant feature subsets with close to equal performance as those identified using only relevance.

Contribution

The two concepts of relevance and redundancy are introduced with respect to shape similarity retrieval. We have proved that we can reduce the number of features in a feature-based shape similarity retrieval system by 80 % and still improve performance significantly. We have compared a number of different selection algorithms for this task and have shown that we are able to achieve better performance than the expected target performance by random drawn feature subsets.

4.3 Third paper: Which shape features are most important

comparison of such a large set of features. We have managed to identify the most important features with respect to shape similarity retrieval using measures of relevance and redundancy.

4.4 Fourth paper: Computing geometric moments for objects with an exact polygon representation

This paper has been published in the *Proceedings of VI-98, Vision Interface conference,* pp.319-324, Vancouver, Canada, June 1998.

Motivation

It is very important that all shape features are computed according to some robust and precise method. During our work on computing moments from polygons, we revealed that some of the most common methods for computing geometric moments from polygons [83, 84, 85] are sensitive to high slope values. The error increases with increasing order of the moment. For us the high-order moments must be very accurate, since we use them for computation of rotational symmetry.

Methodology

A new method for computing the exact geometric moments based on binomial expansion was developed and proved to be slope insensitive. Four different methods were thoroughly tested by tilting a simple square shape and by applying each of the methods to all the objects in our database and compared to the true analytic moments.

Results

The paper proves that two of the methods for computing moments from polygons are sensitive to high slope values and that up to 26 % of the drawings in our database were affected by this sensitivity. About 3 % of the drawings had a relative error of more than 10 % for the fourth order moments. Our method based on binomial expansion and a method developed by Singer [86] proved to be slope insensitive with no significant error.

Contribution

After the publication of this paper, the same computational solution has been found in an early technical report cited in some of the literature [87]. However, the method had not previously been published in any scientific journal or at any major conference, as far as the author is aware of. The remaining contribution in the paper is still the part concerned with proving the slope sensitivity of some of the existing methods.

4.5 Fifth paper: Measuring the complexity of non-fractal shapes by fractal methods

This paper has been accepted for publication in the journal *Pattern Recognition Letters*.

Motivation

Complexity is one intuitive descriptor of shape. When we researched the scientific literature for complexity measures, none of the complexity measure we found was thoroughly and suitably described and tested. We therefor constructed our own complexity measure. I have always been intrigued by fractal theory [88], but I was also aware that artificial objects like the ones we had in our database are not fractal objects. When investigating a commonly used method for measuring the fractal dimension of fractal objects, some observations led us to the development of the new complexity measure.

Methodology

We propose a new method for computing the complexity of non-fractal man-made objects using the common divider step method originally used to estimate the fractal dimension of fractal objects [89]. The method is based on measuring the perimeter of an object using rulers of different length. The slope of the resulting formula seems to indicate how complex the shape is.

Results

The new method for measuring the complexity of non-fractal shapes seems to work fairly well for the objects in our database, as can be seen in the figures in the paper. The complexity measure seems to correspond to our understanding of complexity. We also showed that the man-made objects of the database are not fractal or self-similar. The complexity measure may have implications for fractal objects as well, but this aspect has not been thoroughly explored in this short letter.

Contribution

A new method for estimating the complexity of non-fractal shapes has been proposed. The method has been tested on the 1686 shapes in our database and seems to correspond with our understanding of the term complexity.

Chapter 5

Discussion and future work

The main theme of this thesis is how to improve feature-based shape similarity retrieval. To be able to improve retrieval we have concentrated the work on three important tasks:

- Establishing performance measures for shape similarity retrieval.
- Optimal feature selection with respect to shape similarity retrieval.
- Assessing the redundancy and relevance of individual features to identify the most important features for shape similarity retrieval.

5.1 Discussion of the main results

In this section we discuss the remaining questions of the thesis with the aim of suggesting future work.

5.1.1 Performance evaluation of shape similarity retrieval

Evaluating shape similarity retrieval systems is not a trivial task due to the subjective character of shape similarity itself. We have assumed that shape similarity is a continuous measure in the space of possible shapes. The obvious alternative is to assume that the shapes are clustered with respect to shape similarity, transforming the problem to a classification problem. This is probably true for some applications where the shapes are naturally classified in distinct classes. In our case it is simply not the case. In a general sense shape similarity is also continuous. Nevertheless it would be of interest to evaluate shape similarity from a cluster identification viewpoint and compare the results with ours.

When we have compared different groups of features, we have tried to make the comparison as fair as possible, using the same number of features in each group and adding some general features which contains additional information that is necessary for the shape similarity retrieval. It is always a question of which features should be included in each group.

We have included up to fourth order moments and Fourier descriptors, since higher order features are susceptible to noise on the representation. The low order features only capture the overall shape of an object. The details are lost. Since our representation is fairly accurate, it would have been interesting to check how much retrieval is improved by including higher order features. We must distinguish between performance detoriation which is due to imprecise representation and detoriation due to the addition of the new higher order features.

We have a much higher degree of freedom when we select which skeleton features should be included. Most of the scientific literature is focusing on how to compute skeletons, only a fraction of the papers discuss which features should be extracted and often very specific features for a specific application. We would like to see a paper on how to systematically compute skeleton features and which of these features are important to shape similarity retrieval. In our group of skeleton features we have included those skeleton features which performed best of about 20 features.

It is relevant to ask if other feature groups should have been compared with the existing, but the comparison contains six common groups of features from the literature. We consider that to be a fairly broad comparison of feature groups.

The evaluation is performed using only three of the seven performance measures presented. It is obvious that we would like to test all the described methods. This remains to be done.

5.1.2 Selecting the optimal set of features

Feature selection is performed using a number of techniques using redundancy and relevance with respect to shape similarity using a perceptual shape similarity measure. Relevance is unfortunately not a monotonous measure in the number of features. Otherwise we could have applied branch-and-bound techniques to find the optimal set of features. A number of techniques have been introduced recently, including beam search and floating search techniques. None of these have been tested in our case.

One main issue that we have avoided is how to validate that the selected feature subsets are optimal. Validation is obscured since many of the features are redundant, i.e. completely described by one or several other features. There is probably not a single optimal feature subset, but many equivalent optimal feature subsets. Repeated runs will converge towards different feature subsets with equal performance. One limitation today have been the available computer power, which limits the number of runs we have been able to run, since we have pushed the limits by evaluating over 100 features simultaneously, the largest study to date on these kind of data.

An alternative to feature selection is feature weighting. This option will probably increase performance further since some features obviously are more important than others. We believe weighting should be performed after selection. Another option that has not been explored is the concept of non-linear redundancy.

It is also clear that the results are application specific. We would like to see the same methods tested on other data sets.

5.1.3 Assessing individual shape features

We have assessed individual shape features for the task of shape similarity retrieval using redundancy and relevance as our key objective requirements. Our theoretical framework of necessary, objective and wanted requirements allows us to focus on the objective requirements. It is obvious that other requirements than redundancy and relevance can be chosen as objective requirements. We would like to see other objective requirements tested. The features robustness to deformations and noise is important for many applications and is a good candidate that should be tested within the same framework. Then we could be able to answer if robust features are more relevant or less redundant than other features. We would also like to investigate the implications of the feature distribution on shape similarity retrieval.

5.1.4 Theoretical treatment of the shape space

In this thesis we have briefly discussed the concept of a shape space in rather vague terms. The shape space theory was originally developed for landmark data [10], data with a point correspondence between different shapes. By elaborating this idea further we believe that many new interesting theoretical results will occur. It is however difficult to develop the theory for continuous shape representations were there is no point correspondence.

5.1.5 Computability

In the literature we often encounter new methods for computing different features. The robustness of the computational methods to slope, divisions, scale, precision, colinearity, cocircularity and rotation symmetry is seldom treated thoroughly. We believe that the computational methods influence the performance of some features in many studies, since the implementation often is not robust to these factors.

Bibliography

[1] Mats Carlin & Tom Kavli. Similarity analysis of extruded aluminium sections. *Proceedings of NOBIM-98, Norwegian Image Processing and Pattern Recognition conference*, 28-38, Oslo, Norway, 1998.

[2] Mats Carlin. Computing geometric moments for objects with an exact polygon representation. *Proceedings of VI-98, Vision Interface*, 319-324, Vancouver, Canada, 1998.

[3] Mats Carlin. Modelling the geometric resemblance of objects in CAD drawings (extended abstract). *Proceedings of the 2nd Workshop on Industrial Mathematics for Nordic PhD students*, Orkdal, Norway, 1998.

[4] Mats Carlin. Evaluering og utvelgelse av relevante egenskaper som beskriver objekters geometriske form. Extended abstract (in Norwegian). *Proceedings of the 10th Norwegian seminar on statistics*, Larvik, Norway, 1999.

[5] Mats Carlin. Measuring the performance of shape similarity retrieval methods. Submitted to *Computer Vision and Image Understanding*, special issue on Empirical Evaluation of Computer Vision Algorithms, 2000.

[6] Mats Carlin. Selecting feature subsets and distance measures for shape similarity retrieval. Submitted to *Pattern Recognition*, 2000.

[7] Mats Carlin. Which shape features are most important for shape similarity retrieval? Submitted to *Pattern Recognition*, 2000.

[8] Mats Carlin. Measuring the complexity of non-fractal shapes by fractal methods. *Pattern Recognition Letters* 21:1013-1017, 2000.

[9] Mats Carlin. Improving the performance of shape similarity retrieval systems. *Proceedings of NOBIM-2000, Norwegian Image Processing and Pattern Recognition conference*, Trondheim, Norway, 2000.

[10] Ian L.Dryden & Kanti V.Mardia. *Statistical shape analysis*. Jon Wiley & Sons, 1998.

[11] Farzin Mokhtarian & Alan K.Mackworth. A theory of multiscale, curvature-based shape representation for planar curves. *IEEE Transactions on Pattern Analysis and Machine Intelligence* PAMI 14(8):789-805, 1992.

[12] Benjamin B.Kimia, Allen R.Tannenbaum & Steven W.Zucker. Shapes, Shocks and Deformations I: The components of two-dimensional shape and the reaction-diffusion space. *International Journal of Computer Vision*

[28] David McG. Squire & Thierry Pun. Assessing agreement between human and machine clusterings of image databases. *Pattern Recognition* 31(12):1905-1919, 1998.

[29] Raghu Krishnapuram, Hichem Frigui & Olfa Nasraoui. Fuzzy and possibilistic shell clustering algorithms and their application to boundary detection and surface approximation. *IEEE Transactions on Fuzzy Systems* 3(1):29-60, 1995.

[30] R. Bellman. *Adaptive control processes: A guided tour*. Princeton, University Press, 1961.

[31] Daniel Huttenlocher, Gregory Klanderman & William Rucklidge. Comparing images using the Hausdorff distance. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI 15(9):850–863, 1993.

[32] Hideyuki Tamura & Naokazu Yokoya. Image database systems: A survey. *Pattern Recognition* 17(1):29-43, 1984.

[33] R. L. Ogniewicz. Skeleton-space. A multiscale shape description combining region and boundary information. *Proceedings of Conference on Computer Vision and Pattern Recognition (CVPR)*, pp.746–751, 1994.

[34] William I. Grosky, Peter Neo & Rajiv Mehrothra. A pictorial index mechanism for model-based matching. *Data & Knowledge Engineering* 8:309-327, 1992.

[35] Michael J. Swain & Dana H. Ballard. Color indexing. *International Journal of Computer Vision* 7(1):11-32, 1991.

[36] Markus Stricker & Markus Orengo. Similarity of color images. *SPIE vol.2420 Storage and retrieval for image and video databases III* 381-392, 1995.

[37] B.S. Manjunath & W.Y. Ma. Texture features for browsing and retrieval of image data. *IEEE Transactions on Pattern Analysis and Machine Intelligence* PAMI 18(8):837-842, 1996.

[38] Fang Liu & Rosalind W. Picard. Periodicity, directionality and randomness: Wold features for image modelling and retrieval. *IEEE Transactions on Pattern Analysis and Machine Intelligence* PAMI 18(7):722-733, 1996.

[39] A. Jain, D. Zongker. Feature selection: Evaluation, application and small sample performance. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-19 (2) (1997) 153-158.

[40] Anil K. Jain & Aditya Vailaya. Image retrieval using color and shape. *Pattern Reocgnition* 29(8):1233-1244, 1996.

[41] Gerd Maderlechner, Peter Suda & Thomas Brückner. Classification of documents by form and content. *Pattern Recognition Letters* 18:1225-1231, 1997.

[42] Anil K. Jain & Aditya Vailaya. Shape-based retrieval: A case study with trademark image databases. *Pattern Reocgnition* 31(9):1369-1390, 1998.

[43] Babu M. Mehtre, Mohan S. Kankanhalli & Wing Foon Lee. Content-based image retrieval using a composite color-shape approach. *Information Processing & Management* 34(1):109-120, 1998.

[44] Wayne Niblack et al. The QBIC project: Querying images by content using color, texture and shape. *SPIE* 1908:173-187, 1993.

[45] C. Falutsos et al. Efficient and effective querying by image content. *Journal of Intelligent Information Systems* 3:231-262, 1994

[46] Ashley et al. Automatic and semi-automatic methods for image annotation and retrievalin QBIC. *IBM Research Report*, RJ9951(87910), April 1995.

[47] Chris Okon. IBM's image recognition tech for databases at work: QBIC or not QBIC. *Advanced Imaging*, 63-65, May 1995.

[48] Myron Flickner et al. Query by image and video content: The QBIC system. *IEEE Computer*, 49-56, September 1995.

[49] Robert Finn. QBIC. *IBM Research*, 3:22-25, 1996.

[50] Jeffrey R. Bach et.al. The Virage image search engine: An open framework for image management. *SPIE, vol.2670 Storage and retrieval for image and video databases IV*, 1996.

[51] A. Pentland, R.W. Picard & S. Sclaroff. Photobook: Content-based manipulation of image databases. *International Journal of Computer Vision* 18(3):233-254, 1996.

[52] Marco L. Cascia & Edoardo Ardizzone. JACOB: Just a content-based query system for video databases. *ICASSP*, pp.1216-1219, 1996.

[53] John R. Smith & Shih-Fu Chang. VisualSEEk: A fully automated content-based image query system. *ACM Multimedia*, pp.87-98, 1996.

[54] John R. Smith & Shih-Fu Chang. An image and video search engine for the world-wide web. *SPIE vol.3022 Storage and retrieval for image and video databases V*, 84-95, 1997.

[55] Patrick M. Kelly & Michael Cannon. Query by image example: The CANDID approach. *SPIE vol.2420 Storage and retrieval for image and video databases III* 238-248, 1995.

[56] P. G. B. Enser. Progress in pictorial information retrieval. *Journal of Documentation* 51(2):126-170, June 1995.

[57] Venkat N. Gudivada & Vijay V. Raghavan. Content-based image retrieval systems. *IEEE Computer*, 18-22, September 1995.

[58] Virginia E. Ogle & Michael Stonebraker. Chabot: Retrieval from a relational database of images. *IEEE Computer*, 23-32, September 1995.

[59] Rohini Srihari. Automatic indexing and content-based retrieval of captioned images. *IEEE Computer*, 49-56, September 1995.

[60] Philippe Aigrain, Hongjiang Zhang & Dragutin Petkovic. Content-based representation and retrieval of visual media: A state-of-the-art review. *Multimedia Tools and Applications*, 3:179-202, 1996.

[61] Graeme Baxter & Douglas Anderson. Image indexing and retrieval: Some problems and proposed solutions *Internet Research: Electronic Networking Applications and Policy*, 6(4):67-76, 1996.

[62] Kersti Aas & Line Eikvil. Survey on content-based access to image and video databases. *NR Research Report*, 1997.

[63] Lois F. Lunin. The big picture: Selection and design issues for image information systems. *In: Heidorn & Sandore(eds): Digital image access and retrieval*, 42-56, 1997.

[64] Rajiv Mehrotra. Content-based image modelling and retrieval. *In: Heidorn & Sandore(eds): Digital image access and retrieval*, 56-67, 1997.

[65] Tom Huang, Sharad Mehrotra & Kannan Ramchandran. Multimedia analysis and retrieval system (MARS) project. *In: Heidorn & Sandore(eds): Digital image access and retrieval*, 100-117, 1997.

[66] David A. Forsyth et al. Finding pictures of objects in large collections of images. *In: Heidorn & Sandore(eds): Digital image access and retrieval*, 118-139, 1997.

[67] Babu M. Mehtre, Mohan S. Kankanhalli & Wing Foon Lee. Shape measures for content-based image retrieval - a comparison. *Information Processing and Management*, 33(3):319-337, 1997.

[68] Venkat N. Gudivada & Vijay V. Raghavan. Modelling and retrieving images by content. *Information Processing and Management*, 33(4):427-452, 1997.

[69] F. Idris & S. Panchanathan. Review of image and video indexing techniques. *Journal of Visual Communication and Image Representation*, 8(2):146-166, June 1997.

[70] David Forsyth, Jitendra Malik & Robert Wilensky. Searching for digital pictures. *Scientific American*, 72-77, June 1997.

[71] Suresh K. Choubey & Vijay V. Raghavan. Generic and fully automatic content-based image retrieval using color. *Pattern Recognition Letters*, 18:1233-1240, 1997.

[72] Chung-Lin Huang & Dai-Hwa Huang. A content-based image retrieval system. *Image and Vision Computing*, 16:149-163, 1998.

[73] Wayne Niblack et al. Updates to the QBIC system. *SPIE vol.3312 Storage & retrieval for image and video databases VI* 150-161, 1998.

[74] Michael Kliot & Ehud Rivlin. Invariant-based shape retrieval in pictorial databases. *Computer Vision and Image Understanding*, 71(2):182-197, 1998.

[75] R. Price, T.-S. Chua & S. Al-Hawamdeh. Applying relevance feedback to a photo archival system. *Journal of Information Science* 18:203-215, 1992.

[76] Jing Peng, Bir Bhanu & Shan Qing. Probabilistic feature relevance learning for content-based image retrieval. *Computer Vision and Image Understanding* 75(1-2):150-164, 1999.

[77] H. Liu & H. Motoda. Feature transformation and subset selection. *IEEE Intelligent Systems* 26-28, March-April 1998.

[78] A. L. Blum & P. Langley. Selection of relevant features and examples in machine learning. *Artificial Intelligence* 97:245-271, 1997.

[79] R. Kohavi & G. H. John. Wrappers for feature subset selection. *Artificial Intelligence* 97:273-324, 1997.

[80] D. Koller & M. Sahami. Toward optimal feature selection. *Proceedings of the 13th International Conference on Machine Learning (ML)*, Bari, Italy, pp.284-292, July 1996.

[81] Øivind Due Trier, Anil K. Jain & Torfinn Taxt. Feature extraction methods for character recognition - a survey. *Pattern Recognition* 29(4):641-662, 1996.

[82] Sven Loncaric. A survey of shape analysis techniques. *Pattern Recognition* 31(8):983-1001, 1998.

[83] N. J. C. Strachan, P. Nevadba & A. R. Allen. A method for working out the moments of a polygon using an integration technique. *Pattern Recognition Letters* 11:351–354, May 1990.

[84] X.Y. Jiang & H. Bunke. Simple and fast computation of moments. *Pattern Recognition* 24(8):801–806, 1991.

[85] Jia-Guu Leu. Computing a shape's moment from its boundary. *Pattern Recognition* 24(10):949–957, 1991.

[86] Mark H. Singer. A general approach to moment calculation for polygons and line segments. *Pattern Recognition* 26(7):1019–1028, 1993.

[87] J.M. Wilf & R.T. Cunningham. Computing region moments from boundary representations. Publication 79-49, Jet Propulsion Lab, California Institute of Technology, Pasadena, California, November 1979.

[88] B. Mandelbrot. Chapter 3-6 in *The fractal geometry of nature*, Freeman, 14–57, 1983.

[89] D. Stoyan and H. Stoyan. Methods for the empirical determination of fractal dimension. Chapter 5 in: D. Stoyan and H. Stoyan, *Fractals, random shapes and point fields*, John Wiley & Sons, 39–45, 1998.

Paper I

Measuring the performance of shape similarity retrieval methods

This paper has been submitted to the journal *Computer Vision and Image Understanding*, special issue on empirical evaluation of computer vision algorithms.

Paper II

Selecting feature subsets and distance measures for shape similarity retrieval

This paper has been submitted to the journal *Pattern Recognition*.

Paper III

Which shape features are most important for shape similarity retrieval?

This paper has been submitted to the journal *Pattern Recognition*.

Paper IV

Computing geometric moments for objects with an exact polygon representation

This paper has been published in the *Proceedings of VI-98, Vision Interface conference*, pp.319-324, Vancouver, Canada, June 1998.

Paper V

Measuring the complexity of non-fractal shapes by fractal methods

This paper has been published in *Pattern Recognition Letters* 21:1013-1017, 2000.