

THE SIXTH FRAMEWORK PROGRAMME
Call identifier FP6/2002/IST/C



Information Society
Technologies

FET Open Scheme

SPECIFIC TARGETED RESEARCH PROJECT

Perceptually-relevant Retrieval Of Figurative Images

Profi



Project Description

Participants:

Utrecht University, Department of Information and Computing Sciences
(Netherlands)

University of York, Advanced Computer Architecture Group (UK)

Free University Berlin, Theoretical Computer Science (Germany)

Aktor Knowledge Technology (Belgium)

Coordinator: Remco Veltkamp, Utrecht University, Remco.Veltkamp@cs.uu.nl

1. Project summary

In this project we aim to invent and develop new techniques for the retrieval of figurative images (such as clip art, logos, signs) from large databases. Our techniques will be based on the extraction and matching of perceptually relevant shape features, thereby overcoming many of the limitations of existing methods. This project will develop and evaluate new algorithms for:

- Perceptual segmentation of raw images, and grouping of shape elements.
- Matching of geometrical patterns representing shape features.
- Partial matching: fitting part of one shape with part of another.
- Indexing shape features in huge databases of figurative images.
- Indexing the relative spatial layout of shape features within these images.

The project meets the objectives of the FET programme through a highly-innovative (and hence high-risk) programme to develop novel techniques for shape matching aimed at tackling one of the key problems limiting the effectiveness of current image retrieval techniques. Our project offers the possibility of significant advances at both the scientific and economic level. New in our approach is the primary role of perceptually relevant shape features, the emphasis on the unsolved problem of partial matching, and indexing over lay-out and shapes rather than over feature vectors.

The newly developed algorithms will be experimentally verified in a prototype system, and subjected to rigorous evaluation on databases with independently-validated ground truth.

2. Project objectives

2.1. Motivation

Despite over a decade of research into content-based image retrieval (CBIR), the task of finding a desired image in a large collection remains problematic. Even in application areas where there is a clear need for effective image retrieval, such as medical diagnosis and trademark registration, current technology fails to meet user needs. Much existing research has concentrated on retrieval techniques for natural images (typically photographs of natural scenes or objects), using various combinations of extracted colour, texture and layout feature (Veltkamp and Tanase, 2001). Techniques for the retrieval of figurative images, a class of artificially-produced images which includes icons, trademarks, coats of arms, and clip-art images, have received less attention, even though there is evidence (e.g. Wu et al, 1996), that these images require different techniques for effective retrieval.

For the purposes of our project, figurative images can be defined as artificially-produced images designed to have visual impact, and consisting of multiple homogeneous elements, which may be closed regions, lines, or areas of texture. They may represent a given type of object (such as a dog or car) in stylised form, or consist purely of abstract patterns. They may be coloured or monochrome. Some monochrome examples of figurative images appear in Figs 1-5 below. A comprehensive investigation of retrieval techniques for such images is in our view long overdue, for the following reasons:

- Current techniques for the retrieval of such images are demonstrably inadequate (Eakins, 2001).
- Figurative images such as trademark images, logos, clip art, coats of arms, and icons do not readily lend themselves to retrieval on the basis of name.
- Accurate retrieval and management of such images is of major economic importance.
- Figurative images provide an ideal vehicle for the development of improved shape retrieval techniques, which could be applicable to a much wider domain of images.

Shape is probably the single most important feature used by human observers to characterize an image - psychological studies (e.g. Biederman, 1987) show that a whole range of familiar objects can be recognized as readily from stylised line drawings as from full-colour natural images. However, the process of automatically extracting image features that characterize these elements has proved remarkably difficult, as illustrated in Fig 1. Professional trademark examiners judge all of the following four images to be similar, because all can be perceived as a triangle enclosing a circle - even though they differ in such basic physical characteristics such as the number of components they contain, and not all of them explicitly contain a triangle and a circle.




Fig. 1. Example of four figurative images judged by professional trademark examiners to be perceptually similar.

Other aspects can also be important when judging similarity, including image **structure**, the layout of individual image elements (Fig 2). Here, the triangular layout of image (b) makes it appear more similar to query image (a) than does (c), despite the similarity in the shape of individual components. For images that can be interpreted as natural or man-made objects, such

as trees or ships (in contrast to abstract shapes illustrated here), there is a further complication: their **semantic** interpretation needs to be considered as well. As discussed below, this is a particularly intractable problem, with no easy solution in sight (Eakins, 2002).

The decision on what constitutes an image element can often be quite subjective (see Fig 1(d)), and is frequently subject to significant individual variation (Ren et al, 2000). The task of devising techniques that can accurately retrieve such images from a database of hundreds of thousands of images is extremely challenging. This is particularly true of trademark image retrieval, where the nature of the application demands virtually 100% recall.

Several further problems are holding back the development of successful retrieval techniques in this area. Partial matching of shapes (see Fig 3(a) and (b) below) is problematic because commonly used feature-based approaches, which generate global feature vectors, do not apply. Instead, geometric algorithms must be devised that compute the similarity between two shapes directly. Developing efficient indexing techniques is crucial when databases can contain literally millions of shapes. However, this is difficult because the ordinary ‘point access methods’ for feature vectors lose efficiency in high-dimensional search space, and hence cannot be used. Therefore, indexing data structures that work on shapes directly must be developed, together with new techniques for indexing their relative spatial layout.




Fig. 2. A typical trademark image (a), together with an image judged to have perceptually similar aspects (b), and one judged to have little perceptual similarity (c).



Fig 3. Examples of inadequacy of whole image based measures. Trademark examiners judge that image (a) should retrieve (b), though its global shape is very different. In contrast, (c) should not retrieve (d), even though their edge direction histograms are virtually identical.

2.2. Specific Objectives

Our project aims to focus on three aspects of figurative image retrieval which we believe to be crucial for further progress - improved perceptual shape modelling, (partial) shape matching, and indexing of shape and layout. We recognize that other issues, particularly semantic image matching, are also important, but feel that semantic retrieval is a sufficiently major problem that it needs to be investigated in a separate project. The Profi project will therefore develop new algorithms for shape modelling, matching, and indexing, and combine them into an experimental system, meeting the following prime objectives:

- 1 To develop shape representations for matching and retrieval based on what the eye actually sees. This in turn breaks down into three more specific objectives:

- 1.1 To establish the principles that account for variability in human image segmentation and shape perception, and to explore the implications for automated image matching and retrieval.
- 1.2 To develop both region-based and line-based representations of shape elements in figurative images based on these findings.
- 1.3 To develop composite shape representations capable of representing a variety of individual ways in which that shape might be perceived.
- 2 To develop improved shape matching methods that accurately reflect human similarity judgements, which can operate in the presence of noise, and which are computationally efficient. Specifically:
 - 2.1 To develop shape matching algorithms to determine image similarity by comparing the similarity of sets of image components, whether region- or line-based. We aim to devise robust similarity measures from which perceptually-salient image similarity can be computed.
 - 2.2 To develop partial matching algorithms, capable of fitting part of one shape with part of another, and test their effectiveness in the presence of occlusion.
- 3 To develop and test indexing methods which permit efficient matching of images in large collections.
 - 3.1 To develop novel shape indexing structures that avoid relatively time-consuming comparisons of shape similarity as much as possible, e.g. by using distance values rather than shapes directly.
 - 3.2 To develop techniques for indexing the relative spatial layout of shape features within images.
- 4 To draw our three strands of work together into a proof-of-concept prototype and evaluate its retrieval effectiveness in a work setting.

It should be emphasised that our objective is not to build a fully-fledged retrieval system for figurative images. This would need to include analysis of semantic content, text identification, and other aspects that fall outside the scope of the current proposal.

Rather, we aim to show proof of concept. We consider these concepts are proven successful when our prototype system performs better than other systems reported in the literature, and also performs better than commercial systems as used in current practice at Aktor's. The performance will be measured in terms of a number of appropriate ways, such as speed, precision and recall, Bull's Eye Percentage (used for MPEG7 evaluations), etc. The performance will be considered 'better' when such measures if the improvement is statistically significant.

3. Relevance to the objectives of the specific programme and/or thematic priority

We consider that the Profi project meets the objectives of FET Open. Specifically:

The proposed research is inherently innovative, high-risk and long-term:

New and innovative techniques will be developed in three separate areas - perceptual image segmentation, matching of complete shapes and shape fragments, and indexing shape and structural features within multi-component images. In each case the proposed new ideas are motivated by results developed through the partners' own previous research, but represent significant advances on their previous work. The integration of these techniques, a crucial aspect of the project, represents a further step into the unknown. Hence, the current proposal is both highly innovative and of high risk. The project is inherently long-term as it involves fundamental investigations into human perceptual behaviour, and the design and integration of novel perceptually-based shape representations and novel shape matching and indexing techniques, all involving extensive research.

It is embryonic research, showing proof-of-concept:

Our aim in this project is to investigate some key problems which limit the effectiveness of current retrieval technology - new segmentation and grouping techniques based on knowledge of human image perception, new geometric algorithms for complete and partial shape matching, and new and more efficient indexing techniques for shape and shape layout. Our research is embryonic in the sense that we cannot define the best way to achieve this in advance. We expect our ideas to grow and develop as the project proceeds. The ultimate success of our approach will be judged by the effectiveness and efficiency of the proof-of-concept prototype retrieval engine for figurative images which we intend to develop in the later stages of the project. These objectives form the first steps on the way to meeting one of the biggest challenges in content-based image retrieval: the development of effective and efficient retrieval systems capable of matching human perceptual similarity judgements.

It holds out the promise of major advances at a foundational level:

If our project is successful, it will pave the way for the development of improved shape matching techniques of potential use in a wide range of application areas, including intellectual property management, design archive management and multimedia asset management (see section 5, especially 5.2). The results of our project also have the potential to influence some of the basic science underpinning developments in CBIR (see section 5.3) - our perceptual modelling experiments could shed light on mechanisms of human perceptual image grouping, our new techniques for matching groups of image components could have implications for reducing the computational complexity of a whole range of processor-intensive algorithms, and our new indexing techniques could be applicable to managing all kinds of complex objects in a database. Finally, our work could also influence future standards development, particularly the recent ISO MPEG-7 standard for image and video metadata description, (see section 5.1).

4. Potential impact

With the growth of image collections in design, medicine, journalism and a host of other areas, and the known limitations of human indexing, better techniques for providing content-based access are urgently needed. Our project could potentially have a major impact on the future development of CBIR systems in a wide range of application areas making use of shape matching. These include trademark image registration, management of engineering and architectural design archives, handling databases of hallmarks in precious metals, multimedia asset management, commercial trademark and patent diagram retrieval, and checking for brand counterfeiting and breach of copyright on the Web. The results of our project will benefit European users and developers of image data management systems in areas as diverse as engineering, architectural and fashion design, multimedia content and cultural heritage management, reinforcing competitiveness.

The individual techniques will have a wider impact still. The results of our human shape perception experiments should have implications for shape matching well outside the domain of figurative images. Our new techniques for shape matching will again be widely applicable, particularly with partly-occluded images. Our novel indexing techniques will not only be relevant to image retrieval systems - they could potentially be applied to similarity retrieval of any kind of object, and could have widespread applications in the data mining area. The market for such applications is already large and is likely to grow further in the coming years.

4.1. Contribution to standards

There are two main standard bodies dealing with shapes: the MPEG group, targeting mainly entertainment applications; and the standards related to the exchange of product model data for industrial applications, such as CAD (STEP, STL, IGES). To date, standardization is extending into many different areas, especially within the Motion Picture Experts Group (MPEG). Currently MPEG-4, although ongoing, has already defined many definitive description formats not only for video and audio, but also for graphics and graphical objects. MPEG-7 has already started to define metadata for objects described in MPEG-4, and MPEG-21 is now defining how to control this content from many different perspectives.

Graphical elements in MPEG-4 have been derived, at least in the most part, from the basis of VRML, and there are still inherent limitations in the descriptions of these graphical elements: the geometric representation is specified, but it could be extended to consider also the intended perceptual level.

With results from Profi, this could be taken a step further by defining smart objects that have their perceptual interaction with other shapes built in. This could be proposed as an extension to the current MPEG standard. This would result in a higher level of shape description, providing for a much more extensive functionality. These elements would be very new and potentially create a separate part to the standard due to their uniqueness among other MPEG sections, complementing the whole standard. The complement to MPEG-4 would be on the definition level, to MPEG-7 on the level of metadata descriptions required for syntactic associations with the objects, providing the semantics and functionality of shape.

If the project is successful, i.e., if the performance proofs favourable compared to current methods and systems, we will explore if there are possibilities to make a formal link to MPEG standards, even though both MPEG-4 and MPEG-7 have already reached the status of an ISO standard

4.2. Scientific and technological impact

This project has the potential to make a major contribution to the further development of the field of content-based image retrieval. Shape matching is a crucial aspect of CBIR, but has been used far less often and less effectively than colour or texture matching (Veltkamp and Tanase, 2001). By providing a fresh look at the process of shape matching, and generating multi-view perceptual shape representations which can be efficiently matched, we plan to remove some major barriers to effective shape matching. This could bring significant benefits for future research in the area. Perhaps the most important of these in the long-term is the potential for improved semantic retrieval systems. Shape is one of the most crucial elements in human object recognition (Biederman, 1987), but has so far been used only to a very limited extent in semantic image retrieval, because of the difficulty of representing object shape in a perceptually-meaningful way.

The project also addresses some of the most fundamental issues emerging from discussions on Content-Based Image Retrieval at the International Conference on Pattern Recognition, December 2001, and the Schloss Dagstuhl Seminar on Content Based Image and Video Retrieval, January 2004:

1. *What are the 'killer applications' for content-based image retrieval?*
Retrieval of figurative images - particularly trademarks - is a prime candidate for this title: it is an application where users already routinely search large databases of images, where they cannot readily search on the basis of textual input, and where companies are keenly interested in the development of better techniques.
2. *How can we make better use of techniques based on human image perception and cognition?*
Our project is grounded in human perception, making use both of experimental findings and theoretical concepts from the field of cognitive psychology.
3. *How can we best assess the effectiveness of CBIR techniques?*
Our evaluation will be rigorously conducted, using real application databases, queries and independent relevance judgements. Our results will encourage other researchers to develop more user-centred (as opposed to technology-centred) CBIR systems.

In addition, the project has the potential to contribute to related areas within both computer science and psychology through the results of individual workpackages. While no-one can predict possible contributions with any certainty, these could include:

- cognitive psychology: a greater understanding of the mechanisms by which humans perceive multi-component images, and how human variability arises;
- algorithm design: a better understanding of the process of similarity matching of all kinds of multi-component objects and how to reduce its computational complexity;
- database: how to improve the indexing efficiency of databases containing complex objects.

4.3. Economic impact

Given a period of further development, the project deliverables should have a ready market for exploitation in the areas of image retrieval and multimedia asset management. With the growth of image collections in a host of areas, and the known limitations of human indexing, better techniques for providing content-based access are urgently needed. A good indication of the size of the problem comes from the recently-released technical White Paper "A new vision for internet search" from the USA-based ICT company eVision (<http://www.evisionglobal.com>). Drawing on a number of government and industry studies, they estimate that the number of images used for commercial purposes worldwide on the internet now exceeds 1 billion, a number set to rise to 10 billion by 2005. Seven million images are added to the Web each day; only around 5% of these are indexed. The inadequacy of current search services is highlighted by the following statistics:

- On average, professionals spend 20 hours per week searching for media on the Web.
- Over 35% of these attempts end in failure.
- Global corporations lose an estimated \$25 billion each year through brand counterfeiting, and misuse of brand names and logos.

Overall, eVision estimates the world market for "visual search solutions" to be worth \$11 billion per year.

If this project is successful, it should be capable of making a major contribution to meeting the market needs outlined above. Firstly, it will pave the way for the development of trademark search systems which (in contrast to current experimental systems) are sufficiently reliable for operational use. The protection of trademarks is an area of considerable economic importance - both trademark registration agencies and commercial search services could benefit enormously from dependable software, increasing the reliability of results, improving the efficiency of the search process, and providing new services which are currently not technically feasible. In this way, these companies can create new market opportunities, and increase the competitiveness of a whole range of European organizations. Large organizations employing their own staff for trademark searching would be able to divert their staff to more productive tasks; other organizations relying on commercial search services would be able to rely on a faster, more reliable and cheaper service. Trademark search firms themselves would see a major increase in productivity through the use of automated methods of trademark searching. The size of the potential market is quite considerable. For example, if just 1% of brand counterfeiting can be prevented by improved search services, this represents a potential market of \$250 million each year for such services.

Secondly, the improved image segmentation, shape matching and indexing techniques developed within the project could be incorporated within a much wider range of CBIR systems, covering application areas such as engineering, fashion, and multimedia content and cultural heritage management. The potential economic impact of the wider applications of our shape matching techniques is quite considerable. For example, the ability to call up similar designs or to identify substitute parts capable of meeting production deadlines can give an engineering firm a significant competitive edge. Improvements in competitiveness of a range of businesses in such areas should lead in the long term to improved employment opportunities. Two further applications of economic importance which do not currently seem to have been exploited are searching databases of hallmarks in gold and silver jewellery (an important part of determining their authenticity) and managing collections of on-screen icons used by software developers. Finally, as noted above, any one of the individual techniques developed within the project could have applications outside the area of CBIR. Examples of this could include the use of our novel shape matching techniques in machine vision, and our indexing techniques in database management and data mining.

The potential economic advantages to European industry of such improved search services are considerable. Preventing even 1% of global brand counterfeiting could generate \$250 million/year. Even if the European market represents only 10% of this (a very conservative estimate), an investment of just under €1 million in the current project could bring potential benefits of €20 million/year or more from a single application area! Like any truly innovative project, this is an inherently risky project because no guarantee of success can be given. However, our consortium has an impressive track record in the area, and has made considerable efforts to minimize the risks inherent in the project. For the main discussion on risk management, see section 6.4. We therefore feel that the potential economic benefits of the project far outweigh the remaining risks.

4.4. European dimension

The main reason why this project needs to be conducted at a European level is that there is no single country in Europe that could provide the depth of expertise in the essential disciplines of the project. One of the reasons for the limited success achieved so far in developing techniques for figurative image retrieval is that the problem has proved to be much more complicated and multi-faceted than any of the research teams in the area originally suspected. Further progress in this area requires a major initiative, drawing on a range of different disciplines. It is unlikely that any single European country would be able to fund this initiative at the level required to meet the critical mass that this consortium achieves. Each of the consortium members has an international reputation in one or more of the fields relevant to this project.

5. Workplan – for the whole duration of the project

5.1. Introduction – general description and milestones

Introduction

Our proposed project aims to develop new techniques that permit markedly improved retrieval of figurative images on the basis of their perceived appearance. To achieve this, we propose separate but linked investigations in three areas which are crucial to further progress in modelling shape and structure within figurative images based on known findings about human perception, and the retrieval of such images:

- Perceptual image modelling and segmentation.
- Matching of complete and partial shapes.
- Shape and spatial indexing.

Details of our overall methodology are presented in subsection Approach below. During the first 30 months of the project, we plan to develop and test novel and improved techniques in each of these three areas. In each of these areas, one partner will be the responsible investigator, while the others also contribute. Regular exchange of information between partners will ensure that research in one area is guided by results from the other two. All three aspects will be brought together during the project, by integrating the algorithms developed in each of these areas into a proof-of-concept prototype retrieval engine. The overall validity of our approach will then be determined through a detailed evaluation of both the retrieval effectiveness and computational efficiency of our prototype. The individual workpackage are described below, section 7.2 presents a Gantt chart of project efforts, and section 7.3 shows the structure of the project and workpackage linkages. Forms listing workpackages, deliverables, and details of workpackages appear in appendices at the end of this document.

Approach

To solve the problems listed above, we plan to use a combination of segmentation and grouping techniques based on knowledge of the subjective human perceptual system, geometric algorithms for (partial) shape matching pattern, and metric space approaches for efficient indexing.




Fig 4. Inadequacy of region-based measures. Simple segmentation of image (a) into regions (b) fails to show crucial elements of its perceived shape (a circle within an ellipse). Similarly, segmentation of (c) into regions (d) conceals its similarity to (e).

Comparative evaluation experiments carried out in collaboration with staff from the UK Patent Office using the most recent version of our ARTISAN system as a test-bed have confirmed that component-based matching of trademark images is significantly more effective than whole-image matching (Eakins et al, 2003). However, they have also indicated that failure to segment multi-component images in a perceptually significant way (as illustrated in Fig 4) is the largest single cause of retrieval failure. Hence further significant improvements in retrieval effectiveness appear to be crucially dependent on better techniques for modelling the way in which humans segment an image into perceptually significant parts.

We plan to achieve objective 1 (developing perceptually-significant shape representations) in the following way. Firstly, we plan to extend our previous experiments on human shape perception (Ren et al, 2000), building up a further body of ground truth in the form of quantitative data on the different ways in which multi-component images can be segmented. We plan to use a wider range of images than in our previous experiments, including those in which colour and texture play a major part. We shall also select images specifically to investigate the relative importance of specific features such as symmetry and singularity (Goldmeier, 1972), with the aim of illuminating the mechanisms behind observed variations in human behaviour. Secondly, we plan to derive and validate a set of shape representations for each image, which are consistent with the findings of our experiments in human perception. Two initial lines of approach appear promising. Firstly, we plan to generate multiple region-based segmentations of figurative images, using perceptually-based techniques such as edge-flow detection (Ma and Manjunath, 1997) and salient contour closure (Mahamud et al, 2003). We then aim to aggregate these regions into perceptually-significant groupings by developing algorithms based on Gestalt principles, and validate our results against the ground truth from our perceptual experiments. A second approach, building on the work of Alwis & Austin (1998), will aim to derive a set of plausible Gestalt-based line and curve representations of figurative images, again validating these against the ground truth from our human experiments. Finally, we aim to use suitable AI techniques to derive rules for selecting the most perceptually-salient combination of alternative region and line-based image representations on which to base feature extraction and shape matching and indexing algorithms. This step is necessary to avoid the risk that many of the models generated by the above technique will represent a nonsense interpretation of the image, or will simply repeat information from other views. Methods we aim to use here include the induction of a rule base from our ground truth examples using algorithms such as C4.5 (Quinlan, 1993), and the use of genetic algorithms (Michalewicz, 1996) to optimize the fit between algorithmic methods for region splitting and grouping and observed human behaviour.

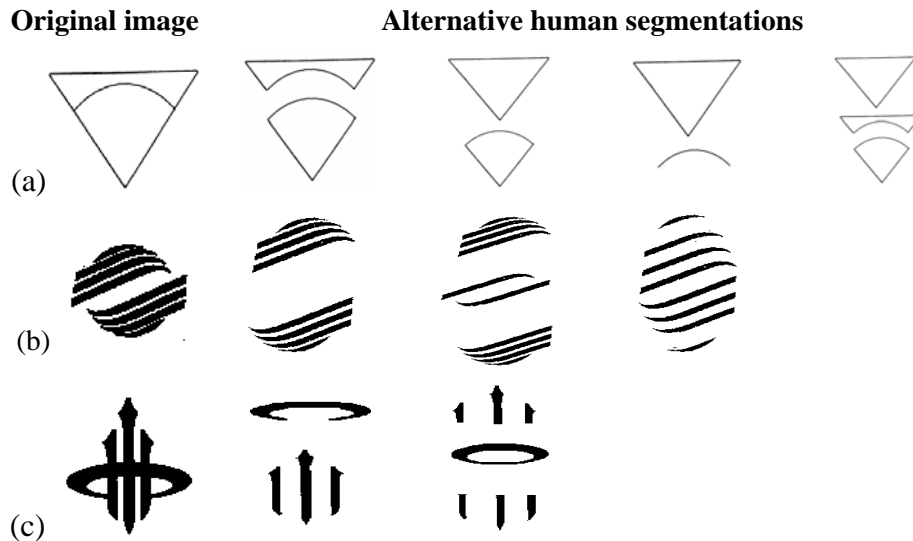


Fig 5. Some of the variation observed in our experiments on human trademark segmentation behaviour.

Our shape matching research aims to devise new algorithms for whole and partial shape matching which both accurately reflect human similarity judgements, and overcome the computational limitations of current techniques. A framework for so-called ‘parametric search’ recently developed by us (Oostrum and Veltkamp 2002) makes geometric optimization much easier. Particular attention will be given to devising novel techniques for directly matching *sets* of

shape elements rather than individual components, as this should allow us to tackle both problems together. We shall also investigate the effectiveness of matching based on boundary regions (such as the neighbourhood of corners) known to be of particular perceptual significance. Our initial approach to partial shape matching will be to develop techniques for curve matching using Hausdorff and Fréchet distances restricted to only part of the curve, and we expect to develop new measures of our own later in the project. As a first filtering step, we will use feature-based matching and matching using shock-graphs (Kimia et al. 1995), so that the number of more complex computations is reduced.

Because the geometric shape matching algorithms are slow, they cannot be used directly during a query on a large database. Hence, standard database indexing techniques cannot be used. Rather than using heuristic feature vector methods, which are not very discriminating, we will use indexing of the pattern metric space, which might result in a few false positives, but guarantees avoiding false negatives. Our initial approach will be to extend our vantage object indexing technique (Vleugels and Veltkamp, 2002). Because of its genericity, this paper has won the Pattern Recognition Society Award for outstanding contribution. We will identify improved ways of selecting vantage objects, and find ways of modifying the technique for partial matching, where the triangle inequality does not hold. A similar metric space approach will be developed for indexing the spatial layout of components in an image. We will also develop a simulation tool for generating synthetic database populations and inter-object distances. Under the condition that the distance measure on the shapes obeys the triangle inequality, it is guaranteed that there are no false negatives. The basic idea is to compute distances between all shapes in the database and a fixed number of vantage (reference) objects. The same is done for the query shape. Now, if the triangle inequality holds, the set of all shapes that are as similar to the vantage objects as the query shape includes all those shapes that are actually similar to the query shape. Not only do we apply this new approach to shape-based retrieval of figurative images, we will investigate a number of important research issues. It is yet an unsolved question how to choose the vantage objects, and how to decide how many are needed and sufficient.

We also plan to devise new structural features of images, based on the relative positions of region centroids, patterns formed by the centroids of regions with similar shape, and topological information such as which regions are contiguous, and which are included within larger regions. What is novel in our approach is that we plan to combine layout descriptors (such as 2D strings) with an indexing scheme, to cope with large numbers of images. Consider a query of multiple shape elements in a certain layout. A naïve approach is to first retrieve all images that contain the individual shape elements, and then to check all these images for containing the specific layout. This will typically still involve many images to check. In this project we will develop an embedding of layouts (object constellations) into a metric space along the same lines as the vantage space approach as used for single shapes. Research issues are the construction of mappings of layout into metric space. Again partial matching, this time of layouts rather than single objects, is a problem that will be investigated. When such issues can be solved, our shape retrieval engine is expected to have unprecedented power in searching for similar figurative images.

Finally, we aim to test both the efficiency and effectiveness of our new algorithms and data structures through an experimentation platform, an integrated retrieval engine capable of supporting whole- and part-image queries on a large database of figurative images. Initial system testing will be performed on image collections such as the MPEG-7 test set (about 7,000 images), publicly-available clip art images (200,000), and images from the UK Trade Marks Registry (160,000), and benchmarks from the International Association for Pattern Recognition, Technical Committee 5, 'Benchmarking and Software' (<http://algoval.essex.ac.uk/tc5>). Our final evaluation studies will be performed on a large trademark database with established ground truth, in cooperation with the commercial trademark search company CompuMark. Results will be disseminated through workshops and seminars.

Introduction of workpackages

Two workpackages will investigate improved methods for figurative image modelling, the first concerned with developing alternative perceptually-based segmentation techniques, and the second with combining and validating the outputs from these techniques. Together, these will provide the basis for generating the set of alternative representations for each image needed to model known human segmentation behaviour.

Workpackage WP2 will develop a range of region-based and line-based perceptual segmentation techniques. After initial segmentation of each image into both region- and line-based components, we plan to use further splitting and merging techniques to generate a wide range of candidate image component sets for subsequent validation in workpackage WP3. This workpackage will aim to develop techniques for selecting the most perceptually-significant views of each image, validating these against ground truth from human perceptual studies. The outcome of these workpackages will be a set of algorithms to identify and apply the most appropriate segmentation techniques to each incoming image, to generate a set of candidate segmentations, and automatically select from this the final set of segmentations to represent each image in the shape database. Image representations will be made available to the other project partners in a standard XML-based exchange format, specifying key shape descriptors as well as the geometry of each component. Note that provision of test data from WP2 to other workpackages will begin in month 3, using results from the ARTISAN system. (This is of lesser quality than will be developed in the course of this project, but is suitable as a starting-point). Further image representations will be distributed as the relevant algorithms are developed and implemented, up to month 30.

Two workpackages will aim to develop novel and improved techniques for shape matching, based on methods from the domain of computational geometry, where partners FUB and UU already have extensive experience. Workpackage WP4 will develop algorithms for matching sets of curves and sets of regions, an extension to existing methods that is essential if we are to compute similarity measures between images consisting of variable numbers of line- or region-based components. These will be assessed both for their ability to generate similarity measures which match human judgements, and for their computational efficiency. It is worth noting that it is not at all trivial to make good implementations of geometric algorithms. Such algorithms are usually designed with the implicit assumption that we can perform real arithmetic. Implementations have to face the problem of limited precision of floating point representations, which can lead to dramatically wrong outcomes, unless special precautions have been taken. We will focus on robust implementations, for example by using exact arithmetic and using robust libraries such as CGAL, the Computational Geometry Algorithms Library (resulted from an Esprit IV project). Shape features and similarity measures developed in this project will be compared to existing features which have been used in image retrieval in the past (such as Fourier descriptors), and to features recommended by the MPEG-7 standards committee (such as the angular radial transform ART and curvature scale space). Workpackage WP5 will extend this investigation to the important area of partial shape matching.

Two further workpackages are concerned with developing efficient methods of indexing both shape and structural aspects of images to permit efficient matching in large collections. Workpackage WP6 aims to use the promising underlying idea of the indexing method previously developed by UU, based on computing distances between all shapes in the database and a fixed number of vantage (reference) objects. The most effective strategy for choosing vantage objects will be investigated, requiring detailed analysis of distributions of inter-shape distances in object space. Finally, the technique will be adapted for use with partial shape matching and other situations where the triangle inequality may not hold. Workpackage WP7 aims to develop improved methods for matching and indexing shapes on the basis of the spatial layout of components, such as the relative positions of region centroids, patterns formed by the centroids of

regions with similar shape, and topological information such as which regions are contiguous, and which included within larger regions. These will be compared with existing solutions described in the literature.

Workpackage 8 will combine output from several other workpackages and perform similarity assessment in a graph-matching framework that is based on Bayesian Statistics. Components identified using the deliverables from WP2 and WP3 will be used. To describe those components, the proposed matching framework will use features including the geometric features used in WP4 and WP5. It aims to exploit the advantages offered by the graph matching framework by combining those features with topological and directional features used in WP7.

Workpackage WP9 will aim to integrate all the improved techniques developed within WP2-WP8 into a single proof-of-concept shape retrieval prototype, which will then be extensively evaluated. While not intended to be a complete system (as indicated above, it will make no attempt to model aspects such as image semantics), the prototype will be capable of generating perceptually-significant representations of stored images using the algorithms developed in WP2 and WP3, using these as the basis for the feature and shape matching techniques developed in WP4 and WP5, and generating shape and layout indexes using techniques developed in WP6 and WP7. Searchers will be able to define whole or part-image queries for matching against the image database. An extensive evaluation will be conducted in WP10, enabling us to draw conclusions about the effectiveness of our techniques in combination as well as in isolation.

Finally, workpackage WP11 will explore the possibilities for commercial exploitation of the project results, and will disseminate the scientific results.

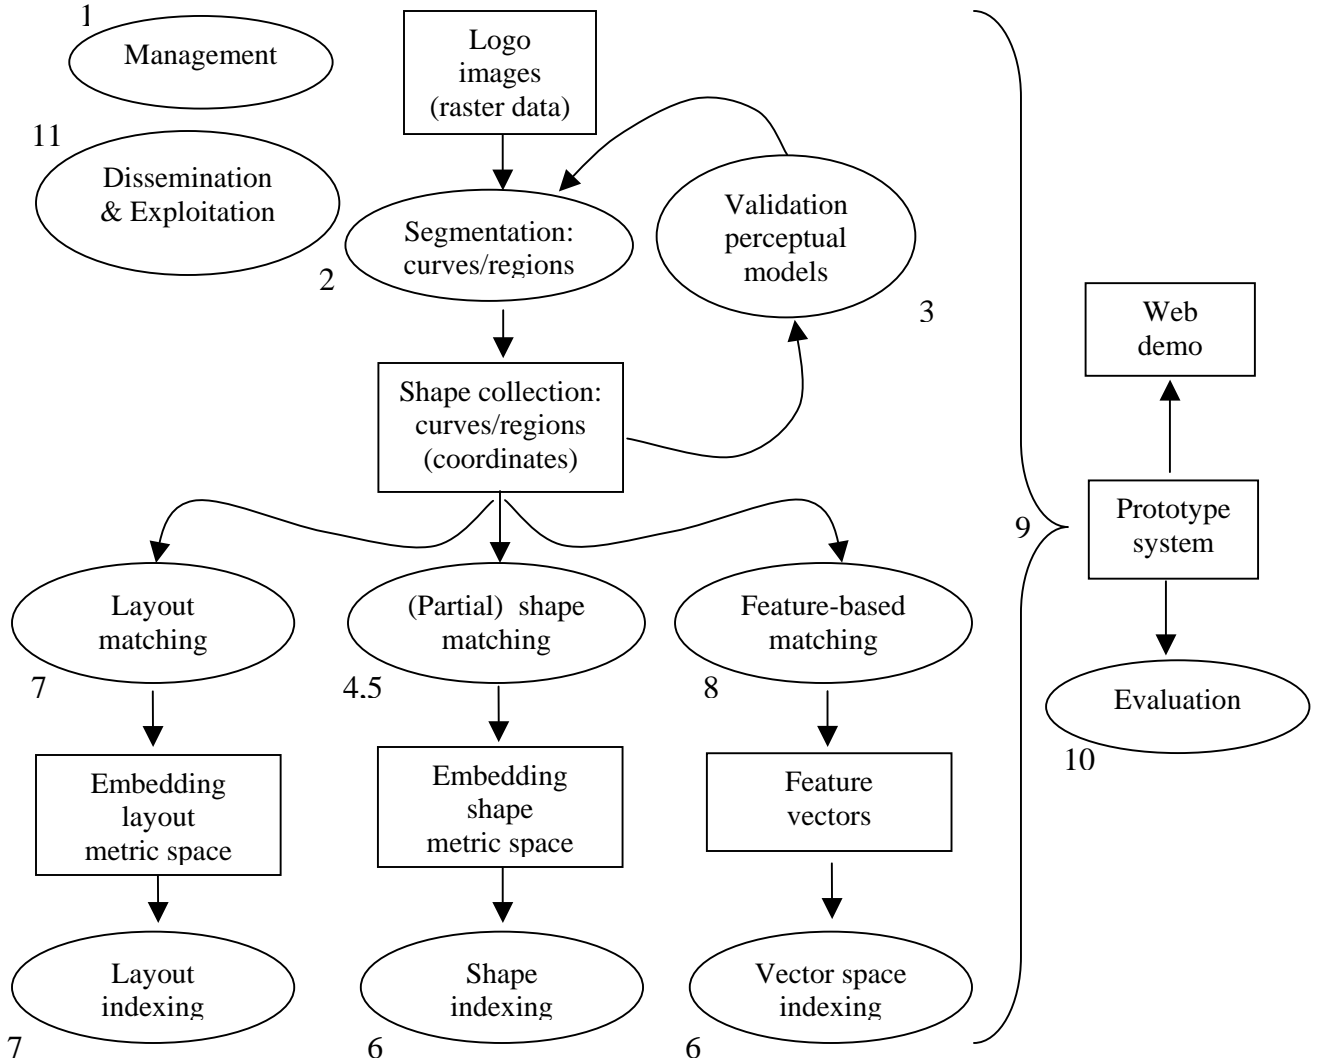
The contribution of each of the partners to each of these workpackages is shown in table 1, below.

Partner	WP2 Developing perceptual segmentation techniques	WP3 Combination of perceptual image models	WP4 Shape Matching	WP5 Partial Matching	WP6 Shape Indexing
UU	+	+	+	+	*
UoY	*	*			
FUB			*	*	+
Aktor					

Partner	WP7 Layout Indexing	WP8 Statistical pattern recognition	WP9 Integra- tion	WP10 Evalua- tion	WP 11 Dissemination & exploitation
UU	*	+	*	+	
UoY	+	*	+	+	+
FUB	+	+	+	+	
Aktor			+	*	*

Table 1. Expertise and contribution of partners to the workpackages. An asterisk (*) indicates the lead contractor, a plus (+) indicates other contributions.

5.2. Graphical presentation of work packages, Pert diagram



Appendix A – Consortium description

A.1 Participants and consortium

A wide range of specialist expertise needs to be brought together for this project, including algorithm design and implementation in the areas of image processing and database indexing, systems integration, evaluation and exploitation. As indicated above, we feel that no single EU country at present contains all the necessary expertise for this project. Hence the consortium consists of three academic partners and one commercial partner, from four different countries, who between them possess all the necessary skills and experience:

1. Partner 1 is the Centre for Geometry, Imaging, and Virtual Environments, Utrecht University (the Netherlands). Utrecht University is the coordinator of this project.
2. Partner 2 is the Advanced Computer Architecture Group, University of York (UK).
3. Partner 3 is the Theoretical Computer Science group, Free University Berlin (Germany).
4. Partner 4 is Aktor-Knowledge Technology (Antwerp, Belgium)

Partner 1 will supply expertise in the area of design and experimental evaluation of indexing schemes for shape features in content-based image retrieval. Partner 2 will contribute its extensive experience in research in figure retrieval, focussing on aspects such as required functionality, perceptual segmentation, and feature extraction. Partner 3 will provide expertise in algorithmic aspects of shape matching. Partner 4 will contribute its ample experience in applying leading-edge technology in business environments.

All partners have been working on shape matching in the past, from different perspectives: trademark image analysis, algorithmic design, shape indexing for content-based image retrieval, and business trademark applications. They contribute complementary knowledge and experience, from different backgrounds, but with similar applications. We have confidence that all necessary skills have been assembled for this project. If a need for additional expertise should become apparent during the project (as is always possible with a highly innovative project such as this), the participating universities all have sufficient depth of expertise at their disposal to rise to additional challenges.

Utrecht University

The Centre for Geometry, Imaging, and Virtual Environments (<http://www.cs.uu.nl/centers/give/>) is headed by Prof. Mark Overmars. The focal topics of research are geometric algorithms (the theory of computational geometry, the implementational aspects, and algorithms for Geographic Information Systems), imaging and multimedia (shape recognition and indexing, and the application in shape-based image retrieval), and virtual environments (motion planning and collision detection, virtual surgery).

In the shape-based retrieval area we have collaborated with Philips Research, and are currently developing a Shape Matching Environment, built on top of the CGAL library (result of the Esprit project CGAL). The developed vantage indexing method resulted from a large national project, Advanced Multimedia Indexing and Searching (AMIS), and is used in our Web demo of the Similarity-based Multimedia Retrieval Framework (SMURF). Shape matching algorithms are applied in a project on multimedia retrieval: Matching and Indexing through Shape Decomposition (MINDSHADE).

Remco Veltkamp obtained his Ph.D. in 1992 from Erasmus University Rotterdam. He is now associate professor at UU, and is principal investigator of the Imaging and Multimedia Group. He has written over 50 refereed papers on indexing, shape matching, shape reconstruction, geometric constraint management, and variational curve and surface design. He is editor of the International

Journal on Shape Modelling, the international journal Pattern Recognition, guest editor of the special issue on shape algorithmics of *Algorithmica* and on multimedia algorithmics of *Multimedia Tools & Applications*, was editor of the Eurographics'95 State-of-the-Art proceedings, programme committee member of various Eurographics Workshops on Programming Paradigms in Graphics, VISUAL 99, 2000, Shape Modeling International 2002, 2003, 2004, and organised the Dagstuhl Seminars on Content-Based Image and Video Retrieval of 1999, 2002, and 2004. He participated in the ESPRIT IV LTR projects CGAL (as the technical manager) GALIA, and SPIRIT, participates in the 6th FW Network of Excellence Aim@Shape.

Mark Overmars received his Ph.D. degree in Computer Science in 1983 from Utrecht University in the Netherlands. Currently he is a full professor at the Department of Computer Science at the same university. His main research interests include computational geometry and application in virtual reality, robotics, and shape matching. Over the past years he published about 175 papers in refereed journals and conferences, and he is author of one of the prime textbooks on computational geometry. He heads a research group of about 25 people that combine theoretical algorithms research with practical implementations and experiments. He headed the ESPRIT IV LTR project CGAL and ESPRIT V project MOVIE, and was involved in 5 other ESPRIT projects (ALCOM, ALCOM2, PROMOTION, MOLOG and GALIA).

Jan Koenderink received his PhD in 1972 (cum laude) from Utrecht University. He is full professor at the Helmholtz school for Autonomous Systems Research of Utrecht University. He received the degree Doctor Honoris Causa in Medicine, awarded by the University of Louvain (Belgium), for his work in Neuroinformatics, was the first “Douglas Holder Fellow” at University College, Oxford (UK), and was appointed a Fellow of the Royal Netherlands Academy of Arts and Sciences (KNAW). His research interests include visual and haptic perception of shape. He has published about 500 refereed articles. Esprit projects he participated in are Insight, REALISE, ECVision and nEurone. Past and present editorial boards memberships include perception journals such as *Vision Research*, *Spatial Vision*, *Perception*, and computer vision journals such as *Computer Vision and Image Understanding*, and *Int. Journal of Computer Vision*.

University of York

The University of York is rated as one of the top ten universities in the United Kingdom. York is rated 6th out of 172 UK Higher Education institutions for research, with 18 of its 23 departments rated 5 or 5* in the 2001 research rating exercise. The Higher Education Funding Council for England (HEFCE) most recently assessed research in the Department of Computer Science in 2003 at grade 6*. This is the highest grade possible and one of only three Computer Science Departments to obtain this grade. The Advanced Computer Architecture Group is one of the largest groups in Computer Science and has many collaborations with industry including Rolls-Royce, EDS, Altos Origin and DS&S. Professor Jim Austin heads the groups 40 researchers, and within the group he leads a team of researchers working in neural networks. Most of group's core research effort has focussed on the development of scalable neural network methods known collectively as AURA. These operate on large datasets, in real time, while maintaining a powerful pattern matching ability, and are supported both in software and in specialist hardware.

Jim Austin is Professor of Neural Computation, and directs the Advanced Computer Architecture Group. He undertakes research in high-performance computer systems based on neural networks for problems in aircraft systems, computer vision and associated areas. He has over 150 publications and is most well known for his work in binary neural networks.

John P Eakins, Professor of Computing, has over ten years' experience of research into storage and retrieval techniques for drawings and images. He is the author of over 20 papers on the subject, including commissioned reviews of the state of the art in CBIR and trademark retrieval.

He has given invited presentations on the topic of image retrieval in Britain, Holland, France, Germany and Italy. He has for four years been co-chair of the *Challenge of Image and Video Retrieval* conferences set up to promote exchange between researchers and practitioners in this field. He led the ARTISAN, SPIRIT and SHREW projects.

Simon O'Keefe has degrees in Natural Science/Management Studies, Operational Research, and Computer Science. Before moving to Computer Science in 1989, he was a management consultant for the Civil Service. Dr O'Keefe has been a lecturer in the Department of Computer Science at the University of York since 1999. His areas of interest cover binary neural networks and image analysis and the practical applications of neural networks to image databases.

Freie Universitaet Berlin

The Theoretical Computer Science research group (<http://www.inf.fu-berlin.de/inst/ag-ti>) at the Institute of Computer Science has done extensive theoretical work in the analysis and matching of geometric shapes and patterns in the last ten years. Several distance measures have been considered for comparing shapes. In particular, it was shown that the Hausdorff-distance between point patterns or shapes consisting of line segments in two dimensions can be computed very efficiently. Computing the Fréchet distance, which is more suitable in certain applications, can still be done efficiently for curves in two and higher dimensions, but turns out to be NP-hard for surfaces in three dimensions. In addition to just computing the distance of two static shapes, the problem of optimal matching of shapes under transformations (translations, rotations, scaling) was considered. The group participated in developing the CGAL software library.

The group has worked on application projects in astronomy and medicine, for which they have also developed software packages. The group has been a partner in several other projects funded by the European Union within the last fifteen years, in particular in the projects ALCOM, ALCOM II, CGAL, GALIA, and now ECG. All these projects were concerned with algorithms and, in particular, computational geometry.

Helmut Alt is professor and the head of the research group. He obtained in 1973 a diploma in mathematics and in 1976 a doctorate in computer science from the University of the Saarland and in 1984 his habilitation from the University of Kaiserslautern. He was a faculty member of the Pennsylvania State University and the Hochschule Hildesheim. Since 1986 he has been a professor of computer science at the Free University of Berlin. His research work is within the fields of algorithms and complexity with an emphasis on computational geometry in the last years, in particular shape matching.

Christian Knauer obtained his diploma in computer science from the University of Erlangen/Nuernberg in 1997 and been a research and teaching associate at the Free University since. His Ph.D. dissertation (2002) is concerned with matching problems of shapes and point patterns.

Lutz Meissner obtained a Masters degree in mathematics from the Free University of Berlin and has been working as a teaching and research associate at the institute of computer science since 1997. His PhD thesis addresses the problem of finding efficient data structures for shapes in order to retrieve for a given shape the most similar one. Similarity again is measured by the distance measures for shapes described previously.

Aktor Knowledge Technology

AKTOR-KT (Antwerp, Belgium) is a Thomson company, operating in the Thomson Legal & Regulatory branch of the Thomson Corporation. AKTOR works in close partnership with Compu-Mark offering integrated knowledge management and powering their trademark research systems with advanced knowledge technology. Compu-Mark delivers trademark research products and services for the European market. Together with their affiliates in UK, Italy,

Sweden, The Netherlands and France, and with their US sister-company Thomson&Thomson, they are world leader in global trademark research.

Guy Coene has degrees in computer science and Master of Science in Artificial Intelligence, and has a consulting background in the field of knowledge technology. He developed knowledge solutions for numerous clients in finance, insurance and industry. Since 1993 he became involved in the knowledge management and technology solutions for Trademark Comparison Systems of Compu-Mark NV. He is founder and managing director of Aktor Knowledge Technology NV. AKTOR was founded at 1998.

Appendix B – References

- Alt H and Guibas L (1999) Discrete Geometric Shapes: Matching, Interpolation, and Approximation. In J.-R. Sack, J. Urrutia, editors, *Handbook of Computational Geometry*, pages 121 - 153. Elsevier, Amsterdam.
- Alt H, Behrends B and Bloemer J, Approximate matching of polygonal shapes, In: Proc. 7th Annu. ACM Sympos. Comput. Geom., pages 186-193, 1991.
- Alwis, S and Austin, J (1998) A novel architecture for trademark image retrieval systems. In: CIR-98: the Challenge of Image Retrieval research workshop, Newcastle upon Tyne, Feb 1998
- Ballard, D. H. (1981). Generalized Hough Transform to Detect Arbitrary Patterns, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(2), 111-122.
- Berchtold, S. and Keim, D. A. and Kriegel, H.-P. (1996). The X-tree: An index structure for higher dimensional data. In: Proc. 22th VLDB Conference, 28—39.
- Berg, M de, Giannopoulos, P, Knauer, C, Oostrum, R van, Veltkamp, R C (2003) The Area of Overlap of two Unions of Convex Objects under Translation. Technical Report UU-CS-2003-25.
- Biederman, I (1987) Recognition-by-components: A theory of human image understanding. *Psychological Review* **94**: 115-147
- Chang, S. K. and Yan, C. W. and Dimitroff, D. and Arudt, T. (1988). Iconic indexing by 2D string. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 9(3), 413-428.
- Chang, S. K. and Jungert, E. (1991). Pictorial data management based upon the theory of symbolic projections. *Journal of Visual Language and Computing* 2,3, 195-215.
- Dimacs (2003) DIMACS Workshop on Discrete Metric Spaces and their Algorithmic Applications, Princeton, August 2003.
- Dohnal, V, Gennaro, C, Saino, P, and Zezalu, P (2003) D-Index: Distance Searchig Index for Metric Data Sets. *Multimedia Tool and Applications* 21, 9-23.
- Eakins J P et al (1998) Similarity retrieval of trademark images. *IEEE Multimedia* **5**(2) 53-63
- Eakins J P (2001) Trademark Image Retrieval. In: *Principles of Visual Information Retrieval*, M. Lew (ed.), Springer, Berlin. ISBN 1-85233-381-2, 319-350.
- Eakins J P (2002) Towards intelligent image retrieval. *Pattern Recognition*, **35**, 3-14
- Eakins, J P et al (2003) Shape feature matching for trademark image retrieval. In: *Proceedings of CIVR2003*, Champaign-Urbana, Illinois, July 2003. *Lecture Notes in Computer Science* **2728**, 28-38 (2003)
- Eakins, J P & Graham, M E (1999) Content-based image retrieval. *JISC Technology Applications Programme Report 39* (<http://www.unn.ac.uk/iidr/CBIR/report.html>).
- El-Kwae, E. A. and Kabuka, M. R. (1999). A robust framework for content-based retrieval by spatial similarity in image databases. *ACM Transactions on Information Systems* 17(2), 174-198.
- Fisher, R. (1936) The use of multiple measures in taxonomic problems. *Annals of Eugenics*, 7: 179-188.
- Folkers, A & Samet, H (2002) Content-Based Retrieval Using Fourier Descriptors on a Logo Database. *Proceedings ICPR '02*.
- Forsyth D A et al (1997) Finding pictures of objects in large collections of images. In: *Digital Image Access and Retrieval: 1996 Clinic on Library Applications of Data Processing*, University of Illinois, 118-139
- Gaede, V and Günther, V (1998) Multidimensional Access Methods. *ACM Computing Surveys*, **30**(2), 170-231.
- Gold, S. (1995). *Matching and Learning Structural and Spatial Representations with Neural Networks*, PhD dissertation Yale University.
- Goldmeier, E (1972) Similarity in visually perceived forms. *Psychological Issues* **8**(1), 1-135
- Gudivada, V. N. (1995). θ R-String: a geometry-based representation for efficient and effective retrieval of images by spatial similarity. TR CS95-02, School of Electrical and Computer Science, Ohio University.

- Gudivada, V. N. and Raghavan, V. V. (1995). Design and Evaluation of algorithms for image retrieval by spatial similarity. *ACM Transactions on Information Systems* 13(2), 115-144.
- Hoffmann D D and Richards W A (1985) Parts of recognition. *Cognition* **18**, 65-96
- Hou , T. -Y. and Lui, P. and Chui, M. Y. (1992). A content-based indexing technique using relative geometry features. In: *Image Storage and Retrieval Systems*. Proceedings of the SPIE, vol. 1662.
- Hu M K (1962) Visual pattern recognition by moment invariants. *IRE Transactions on Information Theory* **IT-8** 179-187
- Huang, P. W. and Jean, Y. R. (1994). Using 2D C⁺ strings as spatial knowledge representation for image database systems. *Pattern Recognition* 27(9), 1249-1257.
- Huttenlocher, D., and S. Ullman, S. Object Recognition using Alignment, *Proceedings of the International Conference on Computer Vision*, London, 102-111.
- Isbell, C L, and Viola, P. (1998) Restructuring Sparse High Dimensional Data for Effective Retrieval. *Advances in Neural Information Processing Systems*.
- Jacobs, C. and Finkelstein, A. and Salesin, D. (1995). Fast Multiresolution Image Querying, *Computer Graphics (Proceedings SIGGRAPH)*, 277-286.
- Jain A K and Vailaya A (1998) Shape-based retrieval: a case study with trademark image databases. *Pattern Recognition* **31**(9), 1369-1390
- Jain A K et al (1996) Object matching using deformable templates. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **18**(3) 267-277
- Katayama, N. and Satoh, S. (1997) The SR-tree: An index structure for high-dimensional nearest neighbor queries. In: *SIGMOD '97*, 369—380.
- Kato, T (1992) Database architecture for content-based image retrieval. In: *Image Storage and Retrieval Systems*, Proc SPIE 2185, 112-123
- Kim, Y S and Kim, W Y (1998) Content-based trademark retrieval system using a visually salient feature. *Image and Vision Computing* **16**, 931-939
- Kimia B B et al (1995) Shapes, shocks and deformations. I: the components of two-dimensional shape and the reaction-diffusion space. *International Journal of Computer Vision* **15**, 189-224
- Kruizinga P. and Petkov N. (1999). Non-linear operator for oriented texture. *IEEE Transactions on Image Processing* **8**(10), 1395-1407.
- Lee C S et al (1999) Information embedding based on users' relevance feedback for image retrieval. In: *Storage and Retrieval for Image and Video Databases VII*, Proc SPIE 3656, 294-304
- Levine M D (1985) *Vision in man and machine*, ch 10. McGraw-Hill, N Y
- Li , J. Z. and Ozsu, T. and Szafron, D. (1996). Spatial reasoning rules in multimedia management systems. TR96-05, University of Alberta, Edmonton, Canada.
- Lin, K. I. And Jagdish, H. V. and Faloutsos, C. (1994). The TV-tree: An index structure for higher dimensional data. *VLDB Journal*, 4:517—542.
- Loncaric, S. (1998). A Survey of Shape Analysis Techniques, *Pattern Recognition*, 31(8), 983-1001.
- Ma W Y and Manjunath B S (1997) Edge flow: a framework of boundary detection and image segmentation. ECE Technical Report #97-02, University of California, Santa Barbara
- Mahmud S et al (2003) Segmentation of multiple salient closed contours from real images. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, in press.
- Manjunath B S et al (2002) *Introduction to MPEG-7*. Wiley, NY
- Michalewicz, Z: (1996) *Genetic algorithms + data structures = evolution programs*. Springer, Berlin
- Mokhtarian, F. and Abbasi, S. and Kittler, J. (1996). Efficient and Robust Retrieval by Shape Content through Curvature Scale Space, In: *Image Databases and Multi-Media Search*, proceedings of the First International Workshop IDB-MMS'96, Amsterdam, The Netherlands, 35-42.

- Oostrum R van and Veltkamp Remco C, Parametric Search Made Practical. In: Proceedings ACM Symposium on Computational Geometry (SoCG 2002), 1-9. Accepted for publication in Computational Geometry, Theory and Applications.
- Osada, R. and Funkhouser, T. and Chazelle, B. and Dobkin, D. (2001) Matching 3D Models with Shape Distributions. International Conference on Shape Modeling and Applications (SMI 2001), 154-166.
- Pham, T D (2003) Unconstrained Logo Detection in Document Images. Pattern Recognition 36, 3023-3025.
- Quinlan, J R (1993) Programs for Machine Learning. Morgan Kaufmann
- Peng H L and Chen S Y (1997) Trademark shape recognition using closed contours. Pattern Recognition Letters **18**, 791-803
- Ranade, S. and Rosenfeld, A. (1980). Point Pattern Matching by Relaxation, Pattern Recognition, 12, 269-275.
- Ravela, S and Manmatha, R (1999) Multi-modal retrieval of trademark images using global similarity. Internal Report, University of Massachusetts at Amherst
- Ren, M et al (2000) Human perception of trademark images: implications for retrieval system design. Journal of Electronic Imaging, **9**(4), 564-575
- Rosin P L (2000) Measuring shape: ellipticity, rectangularity and triangularity. In: Proceedings of 15th International Conference on Pattern Recognition, Barcelona, **1** 952-955 (2000)
- Scassellati, B et al (1994) Retrieving images by 2-D shape: a comparison of computation methods with human perceptual judgements. In: Storage and Retrieval for Image and Video Databases II, Proc SPIE 2185, 2-14
- Scaroff, S. and Pentland, A. (1995). Modal Matching for Correspondence and Recognition, IEEE Transactions on Pattern Analysis and Machine Intelligence, 17(6).
- Siddiqui K et al (1996) Parts of visual form: psychophysical aspects. Perception **25**, 399-424
- Smeulders A W M et al (2000) Content-based image retrieval at the end of the early years. IEEE Transactions on Pattern Analysis and Machine Intelligence, **22**(12), 1349-1380
- Stockman, G. (1987). Object Recognition and Localization via Pose Clustering. Computer Vision, Graphics, and Image Processing, 40(3), 361-387.
- Teh C H and Chin R T (1988) Image analysis by methods of moments. IEEE Transactions on Pattern Analysis and Machine Intelligence **10**(4) 496-513
- Turner, M. and Austin, J. (1999) Graph matching by neural relaxation. *Neural Computing and Application*, 7(3) 238-248
- Uhlmann, J. K. (1991). Satisfying general proximity/similarity queries with metric trees. Inform. Process. Lett., 40:175—179.
- Umeyama, S. (1993). Parameterized Point Pattern Matching and its Application to Recognition of Object Families, IEEE Transactions on Pattern Analysis and Machine Intelligence, 15 (1), 136-144.
- Vailaya A et al (2001) Image classification for content-based indexing. IEEE Transactions on Image Processing **10**(1), 117-129
- Veltkamp, R C and Hagedoorn, M (2001) State-of-the-art in shape matching. In: Principles of Visual Information Retrieval, M. Lew (ed.), Springer, Berlin. ISBN 1-85233-381-2, 87-119
- Veltkamp, R C and Tanase, M (2001) Features in Content-Based Image Retrieval Systems: A Survey. In: State-of-the-Art in Content-Based Image and Video Retrieval. Kluwer, NY
- Vleugels, J and Veltkamp, R (2002) Efficient image retrieval through vantage objects. Pattern Recognition, **35**(1), 69-80
- Webb, A. (2002) *Statistical Pattern Recognition* (2nd Ed.). John Wiley & Sons Ltd, West Sussex, England
- Weber, M et al (2000) Towards automatic discovery of object categories. In: IEEE Conference on Computer Vision and Pattern Recognition, 101-108
- Wertheimer, M (1923) Untersuchungen zur Lehre von der Gestalt. Psychologische Forschung **4**, 301-350

White, D. A. and Jain, R. (1996). Similarity indexing with the SS-tree, In: Proc. 12th IEEE Internat. Conf. Data Eng., 516—523.

Wolfson, H., and Rigoutsos, I. (1997). Geometric Hashing: an Overview, IEEE Computational Science & Engineering, October-December, 10-21.

Wu J K et al (1996) Content-based retrieval for trademark registration. Multimedia Tools and Applications **3**, 245-267

Yianilos, P. N. (1993) Data structures and algorithms for nearest neighbor search in general metric spaces. In: Proc. 4th ACM-SIAM Sympos. Discrete Algorithms, 311—321.

Yi, C. (2002) High-Dimensional Indexing – Transformational approaches to high-dimensional range and similarity searches. LNCS 2341. Springer.

Zahn C T and Roskies C Z (1972) Fourier descriptor for plane closed curves. IEEE Transactions on Computers **C-21** 269-281