

Practice and Challenges in Trademark Image Retrieval

Jan Schietse

AKTOR Knowledge Technology NV
St-Pietersvliet 7
B-2000 Antwerp, Belgium
+32 3 220 73 67

jan.schietse@thomson.com

John P. Eakins

Department of Computer Science
University of York
Heslington
York YO10 5DD, UK

eakins@cs.york.ac.uk

Remco C. Veltkamp

Department of Computer Science
Utrecht University
3584 CH Utrecht
The Netherlands

Remco.Veltkamp@cs.uu.nl

ABSTRACT

In this paper, we outline some of the main challenges facing trademark searchers today, and discuss the extent to which current automated systems are meeting those challenges.

Categories and Subject Descriptors

H.3.3 Information Search and Retrieval: *Search process, Selection process*; H.3.5 Online Information Services: *Commercial services*; I.4.0 Image processing and computer vision (general): *image processing software*

General Terms

Design, Economics, Human Factors, Legal Aspects, Management, Performance.

Keywords

Trademark similarity, Content-Based Image Retrieval, Pattern Matching.

1. INTRODUCTION

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. Techniques for the retrieval of trademark images, and other artificially-produced images such as icons, logos, coats of arms, and clip-art images, have received less attention, even though there is evidence, that these images require different techniques for effective retrieval.

All these artificially-produced images are 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. 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.
- 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 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.

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.

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. 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.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CIVR'07, July 9-11, 2007, Amsterdam, The Netherlands.

Copyright 2007 ACM 978-1-59593-733-9/07/0007 ...\$5.00.

* All trademark images illustrated in this article are UK crown copyright




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

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. 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 there is a need for new techniques for indexing their relative spatial layout. This is true not only for proprietary databases, but also the collection of trademark images on the web.

2. Trademark Image Retrieval

One of the major issues in the Intellectual Property field is trademark infringement. It is very important for a firm to know if there are other firms which are using “confusingly similar” trademark logos with respect to their newly designed trademark logo since this can lead to a (costly) legal battle; this is a search task. Besides that, firms with “strong” trademarks want to monitor all new registrations of trademarks since they do not want to admit trademarks which are similar to their own trademark on the market; this is a watch task.

Thomson Compu-Mark (TCM) is the market leader of trademark research with offices in Antwerp, Milan, Stockholm, Paris, London, Boston and Tokyo. They are offering both search and watch products for textual and graphic trademarks. We will describe in some more detail the 2 basic trademark research

products, i.e. the search and watch product.

Typically a search is performed when one wants to launch a new product or service in the market. A trademark (candidate) is created by a name creation team, but before registration they want to perform a check if there are registered trademarks which are similar to their trademark candidate because the marketing campaign will fail when an existing trademark successfully opposes the registration of the new trademark. To check potential infringement one wants to compare the trademark candidate to all registered trademarks in a set of registers and classes defining both the geographic region and the goods and the services attached to the new product/service.

Trademarks are registered in individual countries (by the Trademark Offices) and by international organisations like OHIM (Office of Harmonization for the Internal Market) for European trademarks, or WIPO (World Intellectual Property Organization) for international trademarks.

In table 1a we list the sizes of the International register (INTE), the Community Trademark register (CTM), the Benelux register (BENE), the French register (FRAN) and the UK register (GBRI). As can be observed, about 30% of all registered trademarks contain next to the textual information graphical elements.

Trademarks are registered for a certain product/service class. This classification (there are 45 different product categories) defines the goods or services you can use your trademark for.




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

Table 1a. Trademark Database sizes

Trademark register	Number of trademarks	Nr. of trademarks with logo
INTE	667659	209700
CTM	447421	151398
BENE	440481	140686
FRAN	961355	298665
GBRI	816807	223125

Examples are:

- Class 32: Beers; mineral and aerated waters and other non-alcoholic drinks; fruit drinks and fruit juices; syrups and other preparations for making beverages.
- Class 13: Firearms; ammunition and projectiles; explosives; fireworks.
- Class 42: Scientific and technological services and research and design relating thereto; industrial analysis and research services; design and development of computer hardware and software; legal services.

The result of a search done in a set of registers and classes is a report containing a list of registered trademarks similar to the “candidate trademark”. Legal experts will use this report to form an opinion about whether it is safe to register (to avoid claims).

The situation is different when you own a registered trademark. To protect your trademark from infringement it is useful to perform a watch, because if you use as trademark logo a triangle with a hand inside to sell hand cream you will want to oppose another producer which registers a logo also containing triangles with a hand for facial cream since an average consumer might be confused about this.

Table 1b. Number of new Trms in 2006

Trademark Register	# new trademarks in 2006
INTE	44.727
CTM	66.653
BENE	27.327
FRAN	69.706
GBRI	32.383

If a watch is performed, every day one compares the watched trademark with all new trademarks published that day. Similar trademarks are reported on a daily basis to the watch client and again legal experts will evaluate these possible infringements and decide if it is appropriate to start a legal action. This is called opposition.

In table 1b we list the number of new trademarks in the year 2006 in the same registers as in table 1a. As you can see for each register you have several hundreds of new trademarks per day.

3. EXISTING TRADEMARK SEARCH TECHNOLOGY

Until now, the principal means of organizing service- and trademark image collections for retrieval has been to use manually assigned classification codes to reflect image content. The most widely used system is the Vienna classification developed by the World Intellectual Property Organization. In principle, this solves the problem of retrieving all images similar to a given logo by ensuring that similar images will receive identical classification codes.

Table 2. Extract from Vienna Codes

3.5.1	Rabbits, hares
3.5.3	Squirrels
3.5.5	Beavers, marmots, badgers, martens, mink
3.5.7	Rats, mice, moles
3.5.9	Hedgehogs, porcupines
3.5.11	Pangolins, anteaters
3.5.15	Kangaroos, koalas
...	...
26.3	TRIANGLES, LINES FORMING AN ANGLE
26.3.1	One triangle
26.3.2	Two triangles, one inside the other
26.3.3	More than two triangles, inside one another
26.3.4	Several triangles, juxtaposed, joined or intersecting
26.3.10	Triangles containing one or more circles, ellipses or polygons (except 26.3.11)
26.3.11	Triangles containing one or more quadrilaterals
26.3.12	Triangles containing one or more other geometrical figures
26.3.23	Lines or bands forming an angle

An extract of the codes can be found in Table2.

Practically, it goes as follows. Every new registered/published trademark logo will be analysed and will be attributed one or more Vienna codes. These codes will be added as indexes in the database. When a logo search has to be carried out, one determines which Vienna codes could be attached to the order (i.e. query) trademark logo. These codes will then be queried and the human expert will be presented a list of trademark logos which will have to be verified visually one by one, and the human expert has to decide whether or not it will be put in the search report as being similar, or at least relevant for the client.

The watch is organized in a similar way. One compares for all device (i.e. drawing) watch orders the attached Vienna codes to the codes of the newly registered and when there is a match, the resulting query logo is compared with the newly registered logo. It is again a human expert who does the final evaluation.

However, this approach suffers from two major drawbacks, both inherent in any retrieval system based on manual classification

codes. Manual classification of images is time-consuming and potentially error-prone, and classification codes are not always helpful for retrieval, particularly for abstract images. Similarity judgments may be based on a number of criteria, including overall shape, the shapes of image components, the spatial configuration of components and the presence of particular types of object (image semantics). No current classification scheme can reflect the full range of such criteria.

As a result, both in search and watch, the human expert is confronted with large sets of logos to inspect. The ranking of the query result is also quasi random. Only by using some combination matching one can influence the ranking of the retrieved trademarks.

For example, for a search order containing a hand and a triangle, one typically queries first the logos containing a hand AND a triangle, and next one would query all logos with a hand and all logos with a triangle. The first category “should” contain the most similar trademarks. The other queries can also contain similar images because the trademark knowledge logic can lead to the conclusion that trademarks with a very dominant hand in their logo are confusingly similar.

Since TCM is confronted with faster and faster delivery times and higher quality constraints together with the fact that the number of trademark logos grows rapidly, it becomes necessary to investigate the possibilities of a system or decision support tool based on content based image retrieval (CBIR) techniques to streamline their device searching ensuring consistency and an acceptable degree of precision and recall. It is especially crucial that no confusingly similar trademarks are missed while doing a trademark search.

4. POSSIBLE FUTURE TECHNOLOGY

By investigating the field for device mark comparison in detail, it is clear that a high level of sophistication is needed to provide a refined similarity comparison and ranking system. In contrast to spotting identical or near-identical images, the challenge for providing refined similarity measurement and judgement comparable is much bigger. The human decision taking in the current device watch and search product lines of TCM is based on *refined image understanding* (decomposing images, recognizing explicit but also more implicit image components and configurations), *refined comparison* (invariance for rotation, scaling, transformation, occlusion and noise), and last but not least on *trademark knowledge* (judging the strength or weakness of used image elements, judging the relative importance of elements, etc). In these judgements, human experts perform an image interpretation based on recognition of shapes, regions, texture, text and spatial configuration.

On a high level an image retrieval system suited for comparing trademark images should fulfil the following constraints:

- One should take into account every possible interpretation of a trademark image.
- It should be possible to search in big sets of images with an acceptable speed (relatively short delivery times).
- Very similar (to the query image) images in the database can not be missed (zero tolerance).

- Trademark images should be compared in great detail (such as shape, contour, and structure) taking into account all sorts of transformations (such as rotation, scaling, inversion, and blurring).

4.1 Scope

Before going into detailed characteristics of a trademark image retrieval systems it is important to elaborate on the scope of an “ideal” trademark image system.

First of all we have to take care of the semantic gap problem. For a search of a logo with a lion, the client will want to receive in his search result all logos containing a lion, even if the “image characteristics” of both lions are completely different. In that case we are dealing with a semantic search and as a result this kind of orders will not be solved by a “traditional” content based image retrieval system which compares contours, shapes, lines or structures. Fortunately, it is easier to perform a search with natural objects than for more geometric order queries (the number of logos to inspect are smaller and the decision is easier).

The added value for abstract orders containing mainly geometric shapes will be higher, since currently with the text retrieval system, for this kind of orders the human expert is confronted with very large collections of logos which have to be inspected. This is simply due to the fact that a very big part of all registered trademark logos are abstract and the fact that in order to retrieve all potential similar trademarks using the Vienna codes, one should enter general/broad categories.

4.2 System Features

The main characteristics of a possible solution for a trademark logo retrieval system are the following:

Order Query Specification

In a 'Query by example' environment the order image is taken as a starting point, image understanding is performed by the system, and the analyst is able to provide additional information. Codification is no longer needed, except for image components with clear semantic meaning (natural objects like pelicans and known artefacts like the statue of liberty). It is clear that segmentation should be an important module in the image analysis component.

Since the human expert can indicate the relative importance of certain elements/shapes, add tags to natural objects, and correct the segmentation results, we will start from an analyzed and enriched order image.

Analysis of target images

The system should provide a (semi-)automatic analysis of new target images. This ensures the incorporation of new (trade)marks for device watching, and also incorporation of existing trademark databases for device searching. Coding should no longer be needed, except for image components with clear semantic meaning (natural objects and known artefacts). As in the case of the order query, it seems likely that the images are segmented.

Robust to noise

The system should be robust to noise in both order image and trademark logo images. Trademark logos with a noise level too

high to do an automatic segmentation should be (semi) automatically cleaned.

Advanced image interpretation

The system should provide image interpretation, and should be able to detect image components like shapes, regions, texture, colour and text components. Also the spatial configuration of all components should be detected. The system should use human-perception-based segmentation to also identify more implicit shapes in the image or partially occluded shapes.

Advanced image comparison

The system should be able to compare device mark images by comparing all elements resulting from the image understanding step, and taking into account their spatial configuration. A query image with a circle in a triangle is more similar to a logo with both shapes present in a deformed way but in the same configuration than to a logo with a circle and triangle in a completely different configuration.

This comparison of the shapes should remain effective under variations like rotation, scaling, transformation, partial occlusion and noise. It is essential for this application that partial matching is supported and on top of that the matching algorithm should reflect in detail how good the partial matching is and which parts of the images are matching.

The fact that text is present on a trademark logo is important. It is not needed to take into account the individual letters of the text, but the fact that a text field is present in both query and target logo in a comparable spatial configuration contributes to the similarity measure.

Colour is also recorded as image element attributes, and can be used as a feature in the comparison.

(Trade)Mark Knowledge Layer

The image comparison results are combined in (trade)mark knowledge rules to provide the similarity judgement and ranking. This provides the flexibility to tune/refine the system based on human expertise. One of the most important trademark features is dominance. The concept of dominance is influenced by the size or other characteristics of a shape or object but also the frequency of occurrence of a certain object can influence the fact that it is dominant. For example: if there are only a very limited number of trademarks which use a certain shape in their logo, then this shape is very distinctive and therefore dominant. Even if there are stars, triangles or circles added to the distinctive shape (for example a swoosh) it will be important to retrieve all trademarks containing this distinctive shape and to rank them high. Trademarks containing stars, triangles or circles can lead also to a similar trademark but the probability is much lower.

It should be possible to tune the system in order to solve quality issues from clients or internal quality checks. Therefore the system should represent all information in an image content graph. Based on comparison results from both graphs, a *tunable (trade)mark logic layer* decides on the overall similarity between order and trademark logo. This knowledge representation approach will enable quality updates and complaint solving.

Acceptance

The similar device marks are presented in an acceptance environment, that provides ranking, and logical groupings. The analyst acceptance is used to refine the proposal even more by using relevance-feedback. The human expert should also be provided with tools supporting consistent selection.

Indexing

The fast delivery times are implying that a trademark image retrieval solution also includes advanced indexing schemes that provide a fast response despite the complexity of the underlying computations.

4.3 Benefits

The benefits of a system such as the one described above would be quite considerable. Such a system should make it possible to deliver in a consistent way high volume logo searches and watches with quality assurance and controllable cost.

5. TECHNOLOGICAL CHALLENGES

While trademark image retrieval has been the subject of considerable research over the last fifteen years [1], no system described in the literature is yet capable of meeting all the criteria set out above. A brief outline of previous research in the field is given below.

5.1 Previous research

One strand of research has concentrated on extracting and comparing features from trademark images *taken as a whole*: The earliest example of the first approach was Kato's TRADEMARK system [2]. It maps normalized trademark images to an 8×8 pixel grid, and calculated a *GF-vector* for each image from frequency distributions of black and edge pixels appearing in each cell of the grid. Query and stored images could then be matched by comparing GF-vectors. Other work following this approach include the following.

- Jain & Vailaya [3] use a two-stage process comprising rapid screening using edge direction histograms and moment invariants followed by template matching;
- Kim & Kim [4] calculate all Zernike moments up to order 17 for each stored and query image, and then select and use the moment with greatest discriminating power for matching;
- Ravela & Manmatha [5] use multi-resolution matching based on histograms of local curvature ratios and gradient orientations computed from Gaussian derivatives.

The second approach regards trademark images as a set of discrete components which are best matched on an individual basis. Overall image similarity can then be computed in a variety of ways from component similarities. The earliest example of this method was the STAR system developed by Wu et al. [6]. This system is based on the principles that perceived trademark similarity is a function of shape, structural and semantic similarity, and that human intervention is essential to achieve acceptable results. The first stage of processing thus involves human indexers, who segment trademark images into perceptually meaningful components. A mixture of human and automated labelling can then be performed, assigning shape features such as

Fourier descriptors and moment invariants, structural features such as the presence of regular patterns of shapes and semantic features such as the presence of particular types of object. The overall similarity between trademarks can then be computed from component feature similarities.

The ARTISAN¹ system developed by Eakins et al. [7] is based on similar principles, though with the important difference that all segmentation and feature extraction is performed automatically. Gestalt principles are used to derive rules allowing individual image components to be grouped into perceptually significant families. Similarity matching can be performed at three levels: whole images, component families or individual image components. More recent versions of the system [8] have incorporated multiresolution analysis to remove texture and group low-level components into higher-level regions, as well as a wider range of shape and structural features.

ARTISAN's use of Gestalt principles has been taken one step further by Alwis and Austin [9], who aim to identify all significant line segments in an image and then cluster these into perceptually significant units according to Gestalt rules. Rather than using conventional similarity matching, their system uses an evidence counting method based on feature values extracted from closed contours in both raw images and "Gestalt" images.

Another technique based on differing views of an image is that of Leung and Chen [10]. They characterize regions as either solid or line-like, extracting boundary contours for the former and skeletons for the latter. After extracting features from line segments derived from both types of representation, overall image similarity is computed by performing a best match between line segments in query and stored images.

5.2 Limitations of current systems

Despite considerable ingenuity by researchers into trademark matching, it is clear that a significant gap remains between the needs of users and the capabilities of current technology. Indeed, it is not immediately apparent that researchers have always been aware of user needs, suggesting that much research may not even have tried to tackle the most pressing problems. Taking the criteria from Section 3 in turn:

1. **One should take into account every possible interpretation of a trademark image.** Studies of the ways by which humans perceive and interpret images confirm that it is a complex process [11], and one that is not straightforward to model in software [12]. However, most research to date has concentrated on matching trademarks purely on the basis of shape or other low-level features. Some researchers have looked at shape features derived from the image as a whole, while others have compared shape features derived from components of segmented images. Some have used regions as the basis for shape comparison, others have used line segments. But with the exception of Alwis & Austin [9] and to a lesser extent Eakins et al. [7], there have so far been few attempts to base image matching on *multiple* views of an image.

2. **It should be possible to search in big sets of images with an acceptable speed (relatively short delivery times).** Most research to date has been conducted on relatively small sets of trademark images, many researchers using collections of little more than a thousand images. Search efficiency has therefore not been a high priority, though the two-stage approach of Jain & Vailaya [3] demonstrates the feasibility of one approach to the problem. In the future, significant improvements in search efficiency will still be needed before any system becomes usable in a commercial environment.
3. **Very similar (to the query image) images in the database can not be missed (zero tolerance).** This is a fundamental requirement of trademark searching, though not necessarily true of image searching in general. For many applications such as journalism and fashion, it does not matter if some relevant images are missed as long as the ones retrieved are acceptable to users. In this context it is important that the retrieval effectiveness of prototype systems should be exhaustively investigated.
4. **Trademark images should be compared in great detail (such as shape, contour, and structure) taking into account all sorts of transformations (such as rotation, scaling, inversion, and blurring).** This requirement is in fact relatively easy for current image matching technology to fulfil. Most, if not all, current feature matching and shape comparison techniques are either inherently invariant to transformations, or can be made so. Multi-resolution matching can handle images at varying levels of detail and blurring. However, this kind of processing is extremely computationally expensive. Hence the more exhaustively query and stored images have to be analysed, the slower the system. Even with the most powerful modern computers, there still needs to be a tradeoff between speed and effectiveness.

5.3 Challenges and prospects for future progress

Perhaps the most serious limitation of current automated systems lies in the area of initial image analysis. Unless all crucial features of target images have been effectively computed and stored, subsequent matching is unlikely to identify all relevant similarities. As indicated above, an ideal system should be able to recognize similarities of shape, structure, and semantics, and to be able to handle (possibly stylised) text – a challenge well beyond the capability of current technology. Even at the level of retrieval by shape or structure, considerable advances will need to be made in modelling human image perception.

The importance of providing alternative representations based on different views of an image has already been mentioned. One possible way to achieve this is follows:

- **Line-based views** of an image can be generated by taking the output from a suitable edge detector and aggregating it into perceptually significant groupings according to Gestalt principles, following the approach pioneered by Alwis and Austin [9].
- **Region-based views** can be generated by multiresolution analysis using techniques derived from those already developed by Eakins et al [8], augmented by texture

¹ Automatic Retrieval of Trademark Images by Shape Analysis

classification and possibly by splitting and merging regions (following rules similar to those proposed by Hoffmann and Richards [13]) to form more perceptually significant groupings.

- **Concept-based views** can be generated by identifying and characterizing familiar (i.e. named) visual concepts within an image. These could include shapes such as circles, triangles, squares, and hexagons, as well as more abstract concepts such as crossover, linear repetition and symmetry, which appear from previous studies [8] to play a crucial role in similarity determination in some contexts.

Developing a whole series of views in this way runs the risk that many of them will represent a nonsense interpretation of the image. This can be avoided by using AI techniques to select those views of a given image most likely to make perceptual sense, or identify the most effective combination of processing methods. However, it may not be possible to train up a machine to perform this task to the exacting standards required by trademark searchers. Hence a hybrid system may be necessary in which human indexers review and if necessary correct machine interpretations of images added to a trademark database. Such indexers could also assign semantic terms to the images (a task which even the best machine learning systems are still incapable of performing reliably), thus bringing such a system one step closer to commercial acceptability.

Conventional image matching techniques, based on 1-to-1 comparison of pairs of image feature vectors, are for the most part far too slow to be acceptable with databases containing up to a million images. Methods based on searching and matching **groups** of lines, regions or feature vectors may be needed before acceptable performance is achieved. Further improvements in search efficiency may be gained by using multidimensional indexes such as the X-tree [14] to organize feature vectors, and Vantage Objects [15] for indexing object space.

Interfacing, both at the query specification and results display stage, is another area that has been relatively neglected by researchers to date. Better methods of search formulation are needed, allowing users to specify:

- whether the search should be based on a complete image, specified parts of an image or a sketch, and
- the most appropriate search parameters for a given image – for example, giving shape and structural features different weights.

Potentially useful improvements at the display stage include:

- two- or even three-dimensional display of retrieved images, allowing searchers to view similarities between them, and
- relevance feedback [16], allowing users to improve system effectiveness by indicating which retrieved images are genuinely relevant to the query.

Many further approaches remain to be explored, and prospects for long-term progress remain good. But the difficulty of finding solutions to the trademark matching problem which are sufficiently robust for commercial use should not be underestimated.

6. ACKNOWLEDGMENTS

Part of this work has been supported by the European Commission under the Sixth Framework Programme project 511572-2 PROFI (Perceptually-relevant Retrieval Of Figurative Images).

7. REFERENCES

- [1] Eakins, J. P. Trademark image retrieval. Ch 13 in *Principles of Visual Information Retrieval* (Lew, M, ed). Springer-Verlag, Berlin, 2001
- [2] Kato, T. Database architecture for content-based image retrieval. *Image Storage and Retrieval Systems*, Proc SPIE 2185, 112-123, 1992
- [3] Jain, A.K., & Vailaya, A. Shape-based retrieval: a case study with trademark image databases. *Pattern Recognition*, **31**(9) 1369-1390, 1998.
- [4] Kim, Y.S., & Kim, W.Y. Content-based trademark retrieval system using a visually salient feature. *Image and Vision Computing*, **16**, 931-939, 1998.
- [5] Ravela, S., & Manmatha, R. Multi-modal retrieval of trademark images using global similarity. Internal Report, University of Massachusetts at Amherst, 1999.
- [6] Wu, J.K. et al. Content-based retrieval for trademark registration. *Multimedia Tools and Appl.* **3**, 245-267, 1996.
- [7] Eakins, J.P., Boardman, J.M., & Graham, M.E. Similarity Retrieval of Trademark Images. *IEEE Multimedia*, **5**(2), 53-63, 1998.
- [8] Eakins, J.P., Edwards, J.D., Riley, J., & Rosin, P.L. A comparison of the effectiveness of alternative feature sets in shape retrieval of multi-component images. *Storage and Retrieval for Media Databases*, Proc SPIE 4315, 196-207, 2001.
- [9] Alwis, S. and Austin, J. Trademark image retrieval using multiple features. Presented at *CIR-99: The Challenge of Image Retrieval*, Newcastle-upon-Tyne, U.K., Feb. 1999.
- [10] Leung, W.H. & Chen, T. Trademark retrieval using contour-skeleton classification. *Proc. IEEE Intl. Conf. on Multimed. and Expo (ICME 2002)*, Lausanne, Switzerland, Aug. 2002.
- [11] Goldmeier, E. "Similarity in visually perceived forms" *Psychological Issues* **8**(1), 1-135, 1972.
- [12] Ren, M., Eakins, J. P. & Briggs, P. Human perception of trademark images: implications for retrieval system design. *Journal of Electronic Imaging*, **9**(4) 564-575, 2000.
- [13] Hoffmann, D. D. & Richards, W. A. Parts of recognition. *Cognition* **18**, 65-96, 1985.
- [14] Berchtold, S., Keim, D.A. & Kriegel, H. P. The X-tree: an index structure for high-dimensional data. *Proceedings of the 22nd Conference on Very Large Databases*, 1996
- [15] Vleugels, J. & Veltkamp, R. C. Efficient Image Retrieval through Vantage Objects, *Pattern Recognition* **35**(1), 69-80, 2002.
- [16] Zhou, X. S. & Huang, T. S. Relevance feedback in image retrieval: a comprehensive review. *Multimedia Systems* **8**, 536-544, 2003.