In Search of the Right Image: Recognition and Tracking of Images in Image Databases, Collections, and The Internet

Neil F. Johnson
Center for Secure Information Systems
George Mason University
http://www.ise.gmu.edu/~csis
njohnson@gmu.edu

Abstract:
Owners of digital media (images) are cautious about making their work available to the public due to the risk of illicit copying and distribution. Withholding images protects the author's work but also prevents the author from gaining interest and recognition from the public. Making samples of the work available raises public awareness but also makes that work subject to theft. If the theft of an image is suspected, then how can one locate the image over a distributed system such as the Internet? We consider several methods for image recognition: methods applied by image database systems, digital watermarking techniques, and an alternative method of image recognition. In doing so, we provide classifications of image database systems and digital watermarking methods. We then introduce an alternative method and illustrate how it fills in the void of image recognition between image databases and digital watermarks.

1 INTRODUCTION

With the explosion of web page development, the availability of color scanners, printers, and digital media, people now have access to hundreds of thousands of images. This trend is likely to continue which will provide more and more people with access to increasingly large image databases. The Internet provides an inexpensive means for authors of digital media to distribute their works to a growing audience. Many authors are leery of distributing their work in fear that it may be copied illegally or represented as another's work. If this occurs then how can one find his work in a sea of images?

We have several issues to contend with in searching for an image over the Internet. The scope of filtering through images available on the Internet is daunting. Some authors of digital images may wish to track their works or identify if any have been copied illegally. We are interested in finding an image – the right image. This image may be distorted from the original, yet we want to identify it as being a descendant of some original image.

Much research has been conducted in the topics of image databases and digital watermarking. Both have techniques that assist in our search for the right image. Section 2 examines image database techniques and attempts to classify them in their image recognition approaches and assesses them with respect to our specific query. Section 3 considers digital watermarking for image recognition as an alternative to methods proposed for images databases. Section 4 introduces an alternative for image recognition and discuss how this approach compares to image database or digital watermarking techniques. Section 5 summarizes the major points of this paper.

2 IMAGE RECOGNITION AND IMAGE DATABASES

Images, graphics, animations, and videos are being published on the Web at an increasing rate. Many search engines help us filter through many pages of text, but these search engines fall short in filtering through this visual data [1]. Image database research provides methods that sift through myriad images and can reduce our search population for the right image. For an introduction to multimedia databases principles and concepts, see [2] and [3].

2.1 Finding Images in Databases and Collections

As the size of image databases grows, traditional methods of image searching break down. For example, it is relatively easy for a person to look over a few hundred "thumbnail" images to find a specific image, it is much harder to locate that image among several thousand. Exhaustive searching breaks down as an effective strategy when the database becomes sufficiently large. For this reason, various techniques have been developed to assist in querying images.

Image database techniques attempt to provide a simpler representation of image data. The representation contains some information that is unique to each image (*features*). Features capture salient aspects and objects of importance in the data [4]. Examples of features may be color, edges, edge density, textures, gradient magnitude, and intensity [5]. These features represent "interesting areas" of an image and may be identifying points within image objects. In any case, these features can be used to identify an image [6, 12]. Using many features provides better recognition, but a less efficient process. Applying fewer features produces faster results but also increases the number of images retrieved. Similar ideas are explored in pattern recognition problems [7]. Finding a balance between too many or too few measures is necessary to effectively handle large image databases [24].

2.2 Classifying Image Database Techniques

Identifying and finding images can be simplified into two phases as described in [8]. The first is the *image summary* where every image in the database is "summarized" with identifying features computed prior to retrieval. These features are used in the query process (*summary comparison*) when the user presents a query, a comparison measure is used to retrieve some number of the most similar images based on their feature match.

A variety of papers attempt to classify these two phases in image databases. Some of the same concepts are classified under differing terminology. Some authors identify these classifications based on "image identification" properties, and others define them as being "image query" properties. Two examples are from [9] and [10]. In [9], information about images is broken into three categories: content dependent, content-descriptive, and content-independent. Content dependent features are those that depend on the content of the image, such as color. Content-descriptive features are those that may describe the scene, such as mountain, car, or face. Content-independent features are those that do not rely on the image scene but properties of the image, such as scale and image file format.

In [10], queries are classified into five areas: retrieval by browsing, retrieval by objective attributes, retrieval by spatial constraints, retrieval by shape similarity, and retrieval by semantic attributes. Retrieval by browsing (RBR) is an example of a thumbnail search by a user, looking for a match. Retrieval by objective attributes (ROA) attempts to retrieve images based on matching the attribute values. Retrieval by spatial constraints (RSC) considers the spatial relationship of objects within an image, such as overlap, adjacency, multiples, or groups of objects. Retrieval by shape similarity (RSS) matches images based on similar shapes. Retrieval by semantic attributes (RSA) is based on the user's perception and understanding about the image.

Images may be indexed or categorized based on visual features, terms and key-terms, assigned subjects, or image types [11]. A lot of overlap exists in the classification of images and image queries. This section will attempt to classify image database techniques into several broad categories and explain what is entailed with each. These are *annotation, image properties, image contents*, and *semantics*.

2.2.1 Annotation

A common approach in searching for images is to index the image database with keywords [11]. In many cases, this is a manual procedure that relies on some human intervention and interpretation (a time-consuming task). Some visual aspects of the image may be inherently difficult to describe; while other images may be quite easy to describe in different ways (different descriptions for the same image) [13]. In addition, the user may have difficulty in guessing which visual aspects have been indexed.

The annotation process may be automated, somewhat, by evaluating the context of the image. For example, in searching and indexing images over the Internet, associated text from the web page is parsed and used to classify the images into subject areas [14]. The text gathered may be the image file name, captions, web page titles, and other text near the image tags. Annotating images for indexing is quite demanding. An alternative is to use image properties that are less likely to require intervention.

2.2.2 Image Properties

Image properties are those, such as colors, textures, and general size, that can be processed without human intervention or interpretation. These features are also referred to as *salient features* [11] and *surface properties* [12] contained within visual scenes. An example may the color histogram of an image or color regions within an image [13,14]. Some approaches make the assumption that these properties are less likely to change and are sufficient for image recognition. Such representations of images are: compact, capable of supporting search over orientation, scale and changes in lighting, quickly matched, and capable of distinguishing textures from a large database [12, 15]. Color Histograms are popular for image searching because the are relatively invariant to small changes in viewpoint (rotation, skewing, perspective, and scale) [16, 15]. A variety of image database systems use color histograms as features in the query process [15, 17, 18, 19]. Color histograms only capture the color distribution in an image and do not include any spatial correlation information [20]. However, image objects and regions may have unique color properties that can be use for image identification.

2.2.3 Image Contents

Colors can be used to provide spatial segmentation, scene breaks, and color grouping within an image [20]. This approach can be used to identify image objects and specify the spatial arrangement of color regions. The way objects within an image relate to one another in space and color can provide recognition features [26, 27]. This image segmentation or object grouping can be automated [21] and these areas can be used for spatial queries [22]. A variety of techniques and tools apply similar techniques [23, 24, 25, 26, 27, 34, 36]. Techniques for identifying images by content demand more perception and interaction than reliance on image properties.

Other image contents include structure and composition of the image scene. The relationship between these can also be use as image recognition features [28]. However, when one starts to consider the relationship between objects and interpretation of the image scene and composition, the system must adapt or at least accept input that becomes far more subjective and perceptual.

2.2.4 Semantics

Semantics is the understanding of the image and information about the image. Categories under this heading may be perceptual similarity and image or picture metadata. An example may be browsing and navigation by content through an image and video archive [29] (similar to retrieval by browsing in [10]). This type of information is difficult to capture and highly subjective. The authors of [30] and [31] state that no single model of similarity combines perceptual similarity and matching and psychological component matching. Some applications attempt to capture this interactively by logging the user's preferences and search habits with respect to the image objects and spatial relationships [32]. The user can specify the general shape and placement of objects as a query for an image (query by example). These inaccurate measures are used to build an approximate match for the image [33].

2.3 Finding Images

A variety of tools use combinations of these classifications in building queries and searching for images. Techniques such as content-based retrieval, or query by example are typically based on color, image content (objects), spatial relationships, and annotation of image objects [13, 17, 19, 34, 35, 36, 37, 38]. Some so called *content-based queries* still rely on associated text to initiate the query process [11].

As image database systems evolve, the queries must be developed to cope with human perception where the similarity of two items is measured by the end-user [39]. The basic approaches to image querying has been referred to as *query by content*, *query by example*, and *similarity retrieval.* The common end result in with any of these approaches is the retrieval of *similar* not the exact image.

2.4 Short Falls of Image Database Concepts

Digital images are complex and although computers are good at representing and manipulating this type of information, decoding their contents is still a research issue. Methods of image and query classifications provide means to find collections of similar images. As introduced in this section, such classifications may include annotation. However, manual classification is time-consuming and potentially error-prone [40]. Collecting text from web pages and file names may incorrectly identify and index images.

Color histograms are used to compare images. However, color histograms lack spatial information, so images with very different appearances can have similar histograms. Colors may also change without changing the content, scene, or objects in the image (e.g. convert to gray-scale).

Various image database approaches assume that all images are scaled to contain the same number of pixels (are of the same dimensions) [15, 23, 24]. Or only a small variation is present in the size, position, and orientation of the objects in images [41]. Several factors make such restrictions difficult in image databases. The query image is typically very different from the target image, so the retrieval method must allow for some distortions. If the query is scanned, it may suffer artifacts such as color shift, poor resolution, and dithering effects. In order to match such imperfect queries more effectively, the image database system must accommodates these distortions and yet distinguishes the target image from the rest of the database.

Since the input is only approximate, the approach taken, by image database systems, is to present the user with a small set of the most promising target images as output, rather than with a single "correct" match. In our problem, images may go through drastic color shifts and cropping, yet we wish to find a correct match - regardless of the scale, resampling, or cropping. We are not interested in finding similar images, but an exact match.

3 DIGITAL WATERMARKING: IDENTIFYING AND TRACKING THE RIGHT IMAGE

Digital watermarks play a role by placing information within digital media [42]. This information may constitute registration of ownership for copyright or a means to locate an image that has been distributed [43]. Some commercial applications search web sites for images that contain watermarked images. When watermarked images are found, the information is reported back to the registered owners of the images [44, 45]. The focus of watermarking in this section is in identifying and tracking digital images.

3.1 Digital Watermarks vs. Image Databases Methods

Digital watermarks have several desirable advantages over image database techniques of identifying images. Many digital watermarks are invariant to scale, changes in color, and image format. A digital watermark is integrated with the image content so it cannot be removed easily without severely degrading the image. Watermarks provide information embedded within the image content that may relate to the owner, license, or tracking of an image. This embedded information may be a code that can be used to identify an image. Instead of searching for image properties, contents, or similarity measures, one can simply search for the code. The result of finding a matching code is the exact image containing that code. If multiple images contain the same code (author information), then the set of images containing that code is returned. In image database terms, a query for an image containing an embedded watermark, should yield an exact image match as opposed to "similar" images. Using an embedded code frees system resources from storing and processing image metadata (color, scale, content, objects, etc.).

Digital watermarks have been explored and various techniques are presented in [46, 47] and classified in [43, 48, 49, 53].

3.2 Watermark Weakness – Now What?

However, embedded watermarks may fail due to accidental corruption or attack [50, 51]. When a watermark fails, the reading mechanism cannot detect the existence of a watermark and the task of finding the illicit copies becomes daunting, especially so when the owner may have tens of thousands of digital images (this becomes a problem similar to image database queries).

Many techniques for watermarking of digital images have appeared recently. Most of these techniques are sensitive to cropping and/or to affine distortions (e.g., rotations and scaling - see [53] of an explanation of affine transforms and invariants and their relationship to digital watermarks). Some image database techniques are invariant to these types of changes. However, watermarks are invariant to color shifts and other manipulations where image databases fail to recognize the appropriate images.

So far we saw that image database techniques provide means to find similar images or classes of images, but not to find a specific image. Many image database techniques are fairly rigid in the similarity measures used and may match well if an image goes through several transformations, color shifts, warping, cropping, or resizing.

Watermarks provide means to identify images fairly independent of image format, size, and color. However, since identification and recognition relies on the survivability of embedded features, then watermarks become vulnerable to distortions that make the embedded codes unreadable.

Disabling a watermark or embedded message is fairly easy [50, 51] and software is available that automates the image processing techniques require to make enough subtle changes to the image as to disable the watermark [52]. If watermarks are used to recognize and locate images, then other means must be devised to find the images.

4 Finding the Right Image: An Alternative Image Recognition Processes (Image Fingerprinting)

How then, can one track digital images over a broad network such as the Internet when the tracking signal (the watermark) is damaged? An alternative to embedding a watermark is to use salient characteristics of an image for identification as a "fingerprint" [43, 53]. In [53], we describe a method for recognizing images based on the concept of *identification watermark* that requires only a small number of *salient feature points*. We show that, using our method, it is possible to recognize distorted images and recover their "original" appearance. Once the image is recognized we use a second technique based on the normal flow to fine-tune image parameters. The restored image can be used to recover the watermark that had been embedded in the image by its owner.

4.1 Recognition Strategy

In image matching, a common technique for image matching is sum-of-squares (SSD) differences [54, 55, 56]. This technique is useful if only slight changes to the image occur. Even changes to color will reduce the effectiveness of an SSD method for image matching. This approach will not work in our specific problem. We are interested in using a "good image" as the query and finding a target may be distorted to the point that SSD will not find a match.

We have a variation of the term *salient features* than that presented in Section 2 and described in [12]. Like the previous definition, we wish to express these features as objects native to the image, however, these must be invariant to changes in color, file formats, and texture. A *salient feature* or *feature point* should be one that can be use to uniquely identify an image even if colors change and various transformations occur that distort the image. The salient features we use in image recognition include points and lines. These points are so significant to the image, that the only way to disable them is to destroy the image [53].

Our method of feature point selection is based on considering the gradient over an image and select candidate points as those containing high values of gradient magnitude, which typically correspond to edges. We can get similar points by applying an edge detector and using the points along these edges as candidate feature points [57] (see Figure1).

(a) (b)

Figure 1: Illustration of candidate feature points from edges.

Candidate feature points along these lines are compared with other candidates. Those that are relatively unique are saved as feature points and used in recognition (see Figure 2).

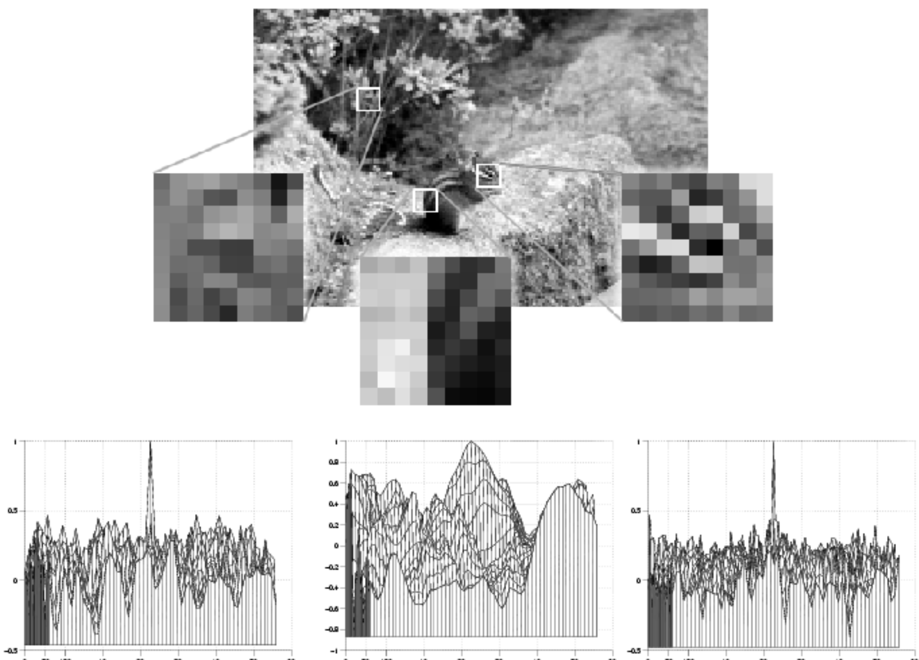

Figure 2: Feature point selection. Candidate points along the selected edges (Figure 1a) are compared with neighboring neighborhoods and compared for uniqueness. The left and right points are good candidates while the center is not.

The graphs in Figure 2 illustrate the correlation of points within the feature point neighborhood. The center point in the graphs corresponds to the center points in the feature point neighborhoods. The center point has the highest value of 1.0 since it is self-similar. Other points are compared with that point and the other peaks in the graphs show their corresponding similarity measure. Note in the center graph, several points are similar to the center point. Determining weather a point should be selected is based on how unique we wish the feature points to be. In all the graphs many points have similarity measures greater than zero. This means that some similarity is detected. However, placing our threshold at 50% is sufficient to select the left and right points as features since no other points within the neighborhood meet or exceed a similarity measure of 50% of the central point.

Feature points are collected for the image at various resolutions. This is to counter scaling and blurring so the image can still be recognized. These sets of feature points provide a summary of the original image and can be use to identify variations of the same image [53]. When we search for matches with other images, we begin with the lowest resolution and try to find matches. If no matches are found then we can say the image does not match our original. If a number of points match, then we have a candidate image that matches our original. (Note: very few points are actually required to match an image. Only three points are required to calculate an inverse affine transform that will recover some of the features lost if any distortion has taken place. Details and examples of this recognition and recovery process are in [53].)

4.2 Advantages of this Recognition Method

The type features we use in image recognition survive transforms and distortions that many image database and watermark techniques cannot. Our recognition method is robust against color, scale, cropping, and image format. Since our method does not embed information into images, it survives manipulation by tools designed to disable watermarking techniques [58, 59]. Image recognition of this type will identify images derived from another image, given a set of feature points. This is not a similarity measure, but an exact match. However, we can manipulate the way features are processed during recognition and identify “similar” images. In this setting, a “similar image” may not have similar content but similar structure. So we get a set of images with few matching points. By using a higher resolution of feature points, the selection process can be refined until we have an exact match.

4.3 Related Research in the Center for Secure Information Systems

This alternative method of image recognition is introduced in [43] and described in [53]. This method relies on a set of unique feature points, within an image, along edges with high gradient values. Even though the points used are unique, an image such as the one shown in Figure 1a can have over one thousand feature points to process at 50% resolution (200x160 pixels). How many feature points are required for positive image identification? We may be able to stop at thirty. Even so, these points are now processed sequentially, so stopping at thirty may only have a small portion of the image recognizable and would thus be vulnerable to cropping attacks.

Currently we are investigating methods to improve the efficiency of this system and look into the refinement of image recognition [60]. Only a few points are required for image recognition. However, these points should be distributed over the image. We can consider point strength as a discriminator and only use the strongest of points. This, too, can result in a large number of feature points.

Methods under investigation include, using corners of edges as feature points. This will greatly reduce the number of points required for recognition and the corners are more resilient to change than lines and edges [57] (see Figure 3).

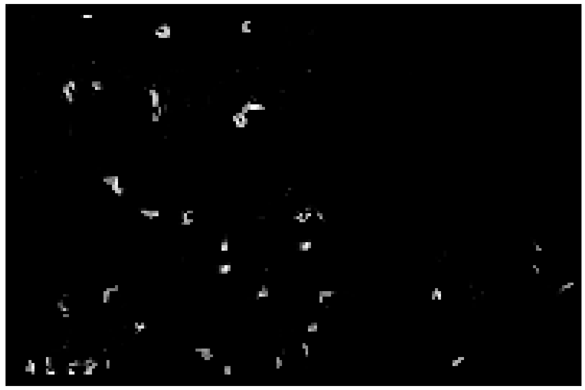

Figure 3: example of feature points of the image in Figure 1a and based on the cornerness measure

Figure 3 is an example of possible feature points based on the cornerness measure. Using corners as feature points, we naturally have far fewer points than with the current method described in [53]. We are not so concerned with the uniqueness of the feature points and symmetry is allowed. This has an advantage over our other method. In the points based on edges and lines, recognition of line drawings such as clip art or cartoons is poor since many lines are similar. However, the corner point approach allows us to recognize these types of images. Corners provide better spread of points and fewer points than lines (see Figure 4).

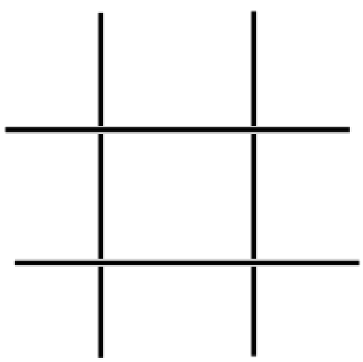

(a)

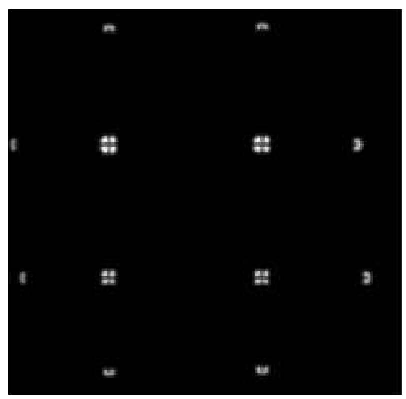

(b)

Figure 4: the image shown in (a) has no feature points if processed with the edge and line method. The image in (b) displays a possible set of feature candidates using the cornerness measure approach to feature point selection.

The corner approach is a variant on the current edge and line approach that yields fewer points (faster pattern matching) and provides a mechanism for recognizing simple images such as line drawings.

5 CONCLUSION

A central task to multimedia information systems is the management of images (storage and retrieval). Research in the area of image databases has focused on retrieval based on objects within images and based in matching algorithms for image similarities or in annotation. Such methods provide a means to reduce the searchable universe in locating the right image

Digital watermarks further reduce the scope and provide a means of tracking for images. Watermarks can be used to locate a specific image; however, watermarks are dependent on survivability of the embedded information and are vulnerable to attacks. Our methods of image recognition (fingerprinting) do not rely on embedded information and can be use to recognize a specific image or images distorted by various transformations. Digital watermarking and image fingerprinting do not prevent illicit copying nor apply any enforcement. Its effectiveness depends on providing evidence of illicit copying and dissemination by extracting watermarks from stolen images. Individual examination of a suspicious image is very costly, although it is unavoidable in some cases where human interference is needed to restore a seriously distorted image in order to detect watermarks from it. Future work is required to develop solutions to these problems. We have started working in the area of watermark recovery from damaged images [43, 53] and continue to investigate alternatives for image recognition and recovery.

Image database, digital watermarking, and image-fingerprinting techniques can be use together in a cooperative means of image querying, recognition, and recovery. Each has strengths and weaknesses that can be leveraged to build a better search for the right image.

6 Bibliography and Referen

1 Marcus, S.: Querying Multimedia Databases in SQL. In [3], pp. 263—277, 1996.

2 Subrahmanian, V.S.: *Principles of Multimedia Database Systems*, Morgan Kaufmann Publishers, San Francisco, CA, 1998.

3 Subrahmanian, V. S., and Jajodia, S. (eds.): *Multimedia Database Systems: Issues and Research Directions*, Springer-Verlag, Berlin, Heidelberg, New York, 1996.

4 Marcus, S. and Subrahmanian, V.S: Towards a Theory of Multimedia Database Systems. In [3] pp. 1—31, 1996.

5 Pass, G. and Zabih, R.: Comparing Images Using Joint Histograms. *Journal of Multimedia Systems*, 1998.

6 Jacobs, C.E., Finkelstein, A., and Salesin, D. H.: Fast Multiresolution Querying. *Proceedings of the ACM SIGGRAPH 95: Computer Graphics Proceedings*, August 1995.

7 Baker, S. and Nayar, S.: Pattern Rejection. *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pp. 544—549, 1996.

8 Pass, G. and Zabih, R.: Comparing Images Using Joint Histograms. *Journal of Multimedia Systems*, 1998.

9 Kashyap, V., Shah, K., and Sheth, A.: Metadata for Building the Multimedia Patch Quilt. In [3], pp. 297—319, 1996.

10 Gudivada, V.N., Raghavan, V.V., and Vanapipat, K.: A Unified Approach to data Modeling and Retrieval for a Class of Image Database Applications. In [3], pp. 36—78, 1996.

11 Smith, J. R. and Chang, S-F.: Searching for Images and Videos on the World-Wide Web. Center for Telecommunication Research Technical Report #459-96-25, Columbia University, 1996.

12 Swain, M.J., Frankel, C.H., and Lu, M.: View-based Techniques for Searching for Objects and Textures. *Proceedings of the Asian Conference on Computer Vision* (ACCV), December 1995.

13 Niblack, W. et al.: The QBIC Project: Querying Images by Content Using Color, Texture, and Shape. *Storage and Retrieval for Image and Video Databases*, SPIE vol. 1908, February 1993.

14 Chang, S-F. et al.: Finding Images/Video in Large Archives, D-Lib Magazine, February, 1997.
URL: http://www.dlib.org/dlib/february97/columbia/02chang.html

15 Swain, M.: Interactive Indexing into Image Databases. *Storage and Retrieval for Images and Video Databases*, SPIE, San Jose, CA, February, pp. 95—103, 1993.

16 Swain, M.J. and Ballard, D. H.: Color Indexing. *International Journal of Computer Vision*, 7:1 1991.

17 Flickner, M. et al.: Query by Image and Video Content: The QBIC System. *IEEE Computer*, 28(9): 23—32, 1995.

18 Ogle, V. and Stonebraker, M.: Chabot - Retrieval from a Relational Database of Images. *IEEE Computer*, 28(9): 40—48, 1995.

19 Pentland, A., Picard, R., and Sclaroff, J.: Photobook - Content-based Manipulation of Image Databases. *International Journal of Computer Vision*, 18(3): 233—254, 1996.

20 Huang, J. et al.: Image Indexing Using Color Correlograms. *Proceedings of 16th IEEE Conference on Computer Vision and Pattern Recognition*, pp. 762—768, 1997.

21 Zhong, D., Zhang, H. J., and Chang, S-F: Clustering Methods for Video Browsing and Annotation. *Symposium on Electronic Imaging: Science and Technology – Storage & Retrieval for Image and Video Databases IV*, vol. 2670, (IS&T/SPIE) San Jose, CA, February 1996.

22 Belussi, A. et al.: A Data Access Structure for Filtering Distance Queries in Image Retrieval. In [3], pp. 185—213, 1996.

23 Pass, G. and Zabih, R.: Histogram Refinement for Content-based Image Retrieval. *IEEE Workshop on Applications of Computer Vision*, pp. 96—102, 1996.

24 Pass, G., Zabih, R., and Miller, J.: Comparing Images Using Color Coherence Vectors. *ACM Conference on Multimedia*, Boston, Massachusetts, November 1996.

25 Rickman, R. and Stonham, J.: Content-based Image Retrieval Using Color Tuple Histograms. *SPIE Proceedings*, 2670: 2—7, 1996.

26 Smith, J. R. and Chang, S-F: Tools and Techniques for Color Image Retrieval. *SPIE Proceedings*, 2670: 1630—1639, 1996.

27 Stricker, M. and Dimai, A.: Color Indexing with Weak Spatial Constraints. *SPIE Proceedings*, 2670: 29—40, 1996.

28 Bach, J. R. et al.: Virage Image Search Engine - An Open Framework for Image Management. *Symposium on Electronic Imaging: Science and Technology – Storage and Retrieval for Image and Video Databases IV*, pp. 76—87, 1996.

29 Smith, J. R. and Chang, S-F: Tools and Techniques for Color Image Retrieval. *Symposium on Electronic Imaging: Science and Technology – Storage & Retrieval for Image and Video Databases IV*, vol. 2670, (IS&T/SPIE) San Jose, CA, February 1996.

30 Santini, S. and Jain, R.: Similarity Queries in Image Databases. Proceedings of *IEEE International Conference on Computer Vision and Pattern Recognition*, (CVPR 96), San Francisco, June 1996.

31 Santini, S.: Similarity Matching. To appear in *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1999.

32 Arya, M. et al.: Design and Implementation of QBISM, a 3D Medical Image Database System. In [3], pp. 79—100, 1996.

33 Sistla, A.P. and Yu, C.: Retrieval of Pictures Using Approximate Matching. In [3], pp. 101—112, 1996.

34 Wu, J.K. et al.: CORE – A Content-based Retrieval Engine for Multimedia Information Systems. *Multimedia Systems*, February vol. 2, pp. 25—41, 1995.

35 Smith, J. R. and Chang, S-F: Querying by Color Regions using the VisualSEEK content-based visual query system. In Maybury, M. T. (ed.), *Intelligent Multimedia Information Retrieval*. IJCAI, 1996.

36 Smith, J. R. and Chang, S-F: VisualSEEK - A Fully Automated Content-based Image Query System. *ACM Multimedia Conference*, Boston, MA, November 20, 1996.

37 Swain, M.J., Frankel, C., and Athitsos, V.: WebSeer – An Image Search Engine for the World Wide Web. IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR), 17-19 June 1997.

38 Frankel, C., Swain, M.J., and Athitsos, V.: WebSeer – An Image Search Engine for the World Wide Web. University of Chicago, Computer Science Department, Technical Report 96-14, 1996.

39 Barros, J. et al.: Using the Triangle Inequality to Reduce the Number of Comparisons Required for Similarity-based Retrieval. *Symposium on Electronic Imaging: Science and Technology – Storage & Retrieval for Image and Video Databases IV*, vol. 2670, (IS&T/SPIE) San Jose, CA, February 1996.

40 Eakins, J. P., Boardman, J. M., and Graham, M. E.: Similarity Retrieval of Trademark Images. *IEEE Multimedia*, 5(2): 53—63, April-June 1998.

41 Swets, D. L. and Weng, J.: Using Discriminant Eigenfeatures for Image Retrieval. *IEEE Transactions on Pattern Analysis and Machine Intelligence,* vol. 18, no. 8, pp. 831—836, 1996.

42 Johnson, N.F. and Jajodia, S.: Exploring Steganography: Seeing the Unseen. *IEEE Computer,* 31(2): 26—34, 1998.

43 Johnson, N.F., Duric, Z., and Jajodia, S.: A Role for Digital Watermarking in Electronic Commerce. To appear in ACM Computer Surveys, 1999.

44 Digimarc Corporation: Picturemarc, MarcSpider, http://www.digimarc.com

45 Signum Technologies: SureSign, http://www.signumtech.com/

46 Anderson, R., (ed.), Information Hiding: First International Workshop, Cambridge, UK. *Lecture Notes in Computer Science*, vol. 1174. Springer-Verlag, 1996.

47 Aucsmith, D. (ed), Second Information Hiding Workshop, Portland, Oregon. *Lecture Notes in Computer Science*, vol. 1525. Springer-Verlag, 1998.

48 Craver, S., Memon, N., Yeo, B., and Yeung, N.M.: Resolving Rightful Ownerships with Invisible Watermarking Techniques. Research Report RC 20755 (91985), Computer Science/Mathematics, IBM Research Division, 1997. An updated version of this paper with the same title is in *IEEE Journal on Selected Areas in Communications*, 16(4): 573—586, 1998.

49 Bender, W. et al.: Techniques for Data Hiding. *IBM Systems Journal*, 35(3-4): 313—336, 1996.

50 Johnson, N.F. and Jajodia, S.: Steganalysis of Images Created using Current Steganography Software. In Aucsmith, D. (ed.) Second Information Hiding Workshop, Portland, Oregon. *Lecture Notes in Computer Science*, vol. 1525. Springer-Verlag, 1998.

51 Petitcolas, F., Anderson, R., and Kuhn, M.: Attacks on Copyright Marking Systems. In Aucsmith, D. (ed) Second Information Hiding Workshop, Portland, Oregon. *Lecture Notes in Computer Science,* vol. 1525. Springer-Verlag, 1998.

52 Kuhn, M. and Petitcolas, F.: StirMark, Tool for evaluating watermarks. http://www.cl.cam.ac.uk/~fapp2/watermarking/image_watermarking/stirmark, 1997.

53 Duric, Z., Johnson, N.F., and Jajodia, S.: Recovering Watermarks from Images. Technical Report: ISE-TR-99-04 in the Center for Secure Information Systems, George Mason University, 1999.

54 Anandan, P.: A Computational Framework and an Algorithm for the Measurement of Visual Motion. *International Journal of Computer Vision*, 2: 283—310, 1989.

55 Turk, M. and Pentland, A.: Eigenfaces for Recognition. *Journal of Cognitive Neuroscience*, 3: 71—86, 1991.

56 Woodfill, J.: *Motion Vision and Tracking for Robots in Dynamic, Unstructured Environments*. PhD Thesis, Stanford University, 1992.

57 Trucco, E. and Verri, A.: *Introductory Techniques for 3-D Computer Vision*. Prentice-Hall, New Jersey, 1998.

58 unZign: last URL: http://altern.org/watermark/, 1997.

59 Kuhn, M. and F. Petitcolas: StirMark, http://www.cl.cam.ac.uk/~fapp2/watermarking/image_watermarking/stirmark, 1997.

60 Johnson, N.: *Image Recognition using Corners and Triangulation of Feature Points*. Term paper in INFS 835: Computational Vision, George Mason University, Spring Semester 1999.