

Image Databases

Search and Retrieval of Digital Imagery

Edited by

Vittorio Castelli and Lawrence D. Bergman




IMAGE DATABASES


Search and Retrieval of Digital Imagery

Edited by

Vittorio Castelli

Lawrence D. Bergman

IBM T.J. Watson Research Center



JOHN WILEY & SONS, INC.

This book is printed on acid-free paper. ∞

Copyright © 2002 by John Wiley & Sons, Inc., New York. All rights reserved.

Published simultaneously in Canada.

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 750-4744. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 605 Third Avenue, New York, NY 10158-0012, (212) 850-6011, fax (212) 850-6008, E-mail: PERMREQ@WILEY.COM.

For ordering and customer service, call 1-800-CALL-WILEY.

Library of Congress Cataloging-in-Publication Data Is Available

ISBN 0-471-32116-8

Printed in the United States of America.

10 9 8 7 6 5 4 3 2 1

CONTRIBUTORS

LAWRENCE D. BERGMAN, IBM T.J. Watson Research Center, 30 Saw Mill River Rd., Hawthorne, NY 10532

VITTORIO CASTELLI, IBM T.J. Watson Research Center, P.O. Box 128, Yorktown Heights, NY 10598

KAUSHIK CHAKRABARTI, University of California at Irvine, Information & Computer Science, 444 Computer Science Building, Irvine, CA 92697-3425

SHIH-FU CHANG, Columbia University, Department of Electrical Engineering, S.W. Mudd Building, Room 1312, New York, NY 10027

CARLO COLOMBO, Università Di Firenze, Dipartimento di Sistemi E Informatica, Via Santa Marta 3, I-50139 Firenze, Italy

ALBERTO DEL BIMBO, Università di Firenze, Dipartimento di Sistemi E Informatica, Via Santa Marta 3, I-50139 Firenze, Italy

CHRISTOS FALOUTSOS, Carnegie Mellon University, Department of Computer Science, Wean Hall, Pittsburgh, PA 15213-3891

SHEILA S. HEMAMI, Cornell University, School of Electrical Engineering, 332 Rhodes Hall, Ithaca, NY 14853

ALEJANDRO JAIMES, Columbia University, Department of Electrical Engineering, 500 West 120th Street, S.W. Mudd Bldg., Room 1312, New York, NY 10027

BENJAMIN B. KIMIA, Brown University, Department of Engineering, Barus and Holley 320, Box D, Providence, RI 02912

C.-C. JAY KUO, University of Southern California, Department of Electrical Engineering Systems, MC2564, Los Angeles, CA 90089-2564

YING LI, University of Southern California, Department of Electrical Engineering-Systems, EEB 400, Los Angeles, CA 90089

WEI-YING MA, Microsoft Research China, Beijing, China

B.S. MANJUNATH, University of California at Santa Barbara, Department of Electrical and Computer Engineering, Santa Barbara, CA 93106-9560

SHARAD MEHROTRA, University of California at Irvine, Information & Computer Science, 444 Computer Science Building, Irvine, CA 92697-3425

MICHAEL ORTEGA-BINDERBERGER, University of Illinois, Department of Computer Science, Urbana-Champaign, IL

SETHURAMAN PANCHANATHAN, Arizona State University, Department of Computer Sciences and Engineering, P.O. Box 875406, Tempe, AZ 85287-5406

JEFFREY W. PERCIVAL, University of Wisconsin, Space Astronomy Laboratory, 6295 Chamberlain Hall, 1150 University Avenue, Madison, WI 53706

H.K. RAMAPRIYAN, Earth Science Data and Information System Project, Code 423 NASA, Goddard Space Flight Center, Greenbelt, MD 20771

PRASHANT SHENOY, University of Massachusetts, Department of Computer Science, Amherst, MA 01003-4610

JOHN R. SMITH, IBM T.J. Watson Research Center, 30 Saw Mill River Road, Hawthorne, NY 10532

KENT SOO HOO, JR., University of California at San Francisco, Radiology, Box 0628, 505 Parnassus Avenue, San Francisco, CA 94143-0628

PETER TILKE, Schlumberger-Doll Research Center, Old Quarry Rd, Ridgefield, CT 06877

HARRICK M. VIN, University of Texas at Austin, Department of Computer Sciences, Taylor Hall 2.124, Austin, TX 78750

XIA SHARON WAN, 1704 Automation Parkway, San Jose, CA 95131

STEPHEN WONG, University of California at San Francisco, Radiology, P.O.Box 0628, 505 Parnassus Avenue, San Francisco, CA 94143-0628

PREFACE

Image databases pose new and challenging problems to the research community. Over the past 40 years, database technology has matured with the development of relational databases, object-relational databases, and object-oriented databases. Extensions to these databases have been developed to handle nontraditional data types, such as images, video, audio, and maps. The core functionalities of classical databases, however, are tailored toward simple data types and do not extend gracefully to nonstructured information.

As the amount of digital imagery grows, techniques that are specifically tailored to such data types need to be developed and more widely deployed. Over the last few years, standardization efforts have taken place in fields such as medicine, geographic information systems, and video in parallel with the development of large digital archives. Search and indexing technology has developed over the same time frame, leading to a wide variety of research prototypes and ad-hoc commercial solutions for searching imagery, video, and other multimedia types. There still is a lack, however, of large commercial systems that integrate existing techniques for storage, indexing, and retrieval of image databases.

In this volume, we present the state of the art of a number of disciplines that converge in image database technology. We motivate the volume by presenting selected applications, including photographic, remotely sensed, petroleum, and medical imagery.

The technology section is divided into two main areas: a portion on storage and one on retrieval. We start with a chapter on system architecture, detailing hierarchical storage management schemes, the role of caching, storage layout issues, and architectural trade-offs. This is followed by a chapter on applications of traditional databases to indexing image repositories, with emphasis on data modeling and organization. Image compression is an integral part of image storage management, and we devote a chapter to the topic, describing the fundamental concepts, reviewing and comparing the main existing standards, outlining nonstandard compression techniques, and discussing evaluation and trade-offs in selecting appropriate methods. Since delivery of imagery over limited bandwidth channels is a universal problem, we include a chapter describing the transmission of imagery in digital format, including techniques such as progressive transmission.

The second half of the technology section describes search and retrieval techniques. We begin with an overview of the area. Low-level image features, such as color, texture, and shape, are typically used to index image archives, and we

devote a chapter to each of these. This is followed by a chapter on indexing techniques for spatial queries, range queries, similarity, and nearest-neighbor queries. The next chapter discusses the use of multiple abstraction levels and compressed or transformed images for improving search efficiency. The last chapter addresses the automatic extraction of semantic information from imagery.

This book is intended for several audiences. It can be used as a textbook for a graduate-level course on image databases, as it provides a wide range of introductory material and extensive bibliographies that are appropriate for directing further reading. It is also a valuable reference for developers and researchers in the field, as well as an introduction for IT professionals who need to further their understanding of the discipline.

CONTENTS

Contributors	xiii
Preface	xv

INTRODUCTION

1 Digital Imagery: Fundamentals	1
<i>Vittorio Castelli and Lawrence D. Bergman</i>	1
1.1 Digital Imagery	1
1.2 Applications of Digital Images	1
1.3 Technological Factors	6
1.4 Indexing Large Collection of Digital Images	8
1.5 Overview of the Book	9
References	10

SELECTED APPLICATION

2 Visible Image Retrieval	11
<i>Carlo Colombo and Alberto Del Bimbo</i>	11
2.1 Introduction	11
2.2 Image Retrieval and Its Applications	12
2.3 Advanced Design Issues	17
2.4 Visible Image Retrieval Examples	22
2.5 Conclusion	31
Acknowledgments	31
References	31
3 Satellite Imagery in Earth Science Applications	35
<i>H.K. Ramapriyan</i>	35
3.1 Introduction	35
3.2 Historical Background and Remote Sensing Missions	36
3.3 Applications of Remote Sensing	38
3.4 Data Collection Systems	41
	vii

3.5	Errors, Artifacts, and Required Corrections	45
3.6	Processing	49
3.7	Implications for Storage and Access	61
3.8	Examples of Systems that Store or Provide Access to Remotely Sensed Data	66
3.9	Summary and Conclusion	77
3.10	Acknowledgments	78
	References	78
	Additional Reading	81
4	Medical Imagery	83
	<i>Stephen Wong and Kent Soo Hoo, Jr.</i>	83
4.1	Introduction	83
4.2	Applications	85
4.3	Challenges	89
4.4	Enabling Technologies	95
4.5	Standards	98
4.6	Systems Integration	101
4.7	Conclusion	102
	Appendix	103
	References	103
5	Images in the Exploration for Oil and Gas	107
	<i>Peter Tilke</i>	107
5.1	Introduction	107
5.2	Data Characteristics	108
5.3	Selected Application Scenarios	117
5.4	Enabling Technologies	125
5.5	Challenges and Open Problems	132
	Appendix	132
	References	137

STORAGE AND SYSTEM ARCHITECTURE

6	Storage Architectures for Digital Imagery	139
	<i>Harrick M. Vin</i>	139
6.1	Storage Management	140
6.2	Fault Tolerance	145
6.3	Retrieval Techniques	149
6.4	Caching and Batching Issues	152
6.5	Architectural Issues	153
6.6	Conclusion	156
	References	156

7 Database Support for Multimedia Applications	161
<i>Michael Ortega-Binderberger and Kaushik Chakrabarti</i>	161
7.1 Introduction	161
7.2 A Model for Content-Based Retrieval	162
7.3 Overview of Current Database Technology	166
7.4 Image Retrieval Extensions to Commercial DBMSs	172
7.5 Current Research	186
7.6 Conclusion	203
Acknowledgments	204
References	204
8 Image Compression — A Review	211
<i>Sheila S. Hemami</i>	211
8.1 Introduction	211
8.2 Entropy — A Bound on Lossless Compression	213
8.3 Rate-Distortion Theory — Performance Bounds for Lossy Compression	215
8.4 Human Visual System Characteristics	217
8.5 Pixel-Based Redundancy Reduction	220
8.6 Quantization	227
8.7 Lossless Image-Compression Systems	230
8.8 Lossy Image-Compression Systems	231
8.9 Some Comments on JPEG-2000	236
8.10 Conclusion	237
References	237
9 Transmission of Digital Imagery	241
<i>Jeffrey W. Percival</i>	241
9.1 Introduction	241
9.2 Bulk Transmission of Raw Data	242
9.3 Progressive Transmission	243
9.4 Some Examples	251
9.5 Summary	258
References	258

INDEXING AND RETRIEVAL

10 Introduction to Content-Based Image Retrieval — Overview of Key Techniques	261
<i>Ying Li and C.-C. Jay Kuo</i>	261
10.1 Introduction	261
10.2 Feature Extraction and Integration	263
10.3 Similarity Functions	264

10.4	Feature Indexing	271
10.5	Interactive Content-Based Image Retrieval	272
10.6	The MPEG-7 Standard	276
10.7	Conclusion	278
	References	279
11	Color for Image Retrieval	285
	<i>John R. Smith</i>	285
11.1	Introduction	285
11.2	Color Descriptor Extraction	286
11.3	Color Descriptor Metrics	294
11.4	Retrieval Evaluation	300
11.5	Summary	308
	References	309
12	Texture Features for Image Retrieval	313
	<i>B.S. Manjunath and Wei-Ying Ma</i>	313
12.1	Introduction	313
12.2	Texture Features	316
12.3	Texture Features Based on Spatial-Domain Analysis	318
12.4	Autoregressive and Random Field Texture Models	321
12.5	Spatial Frequency and Transform Domain Features	323
12.6	Comparison of Different Texture Features for Image Retrieval	326
12.7	Applications and Discussions	329
12.8	Conclusion	339
	Acknowledgments	339
	References	339
13	Shape Representation for Image Retrieval	345
	<i>Benjamin B. Kimia</i>	345
13.1	Introduction	345
13.2	Where Is Indexing by Shape Relevant?	346
13.3	Image Preparation and Query Formulation	348
13.4	Representation of Shape	349
13.5	Matching, Shape Similarity, and Validation	357
	References	358
14	Multidimensional Indexing Structures for Content-Based Retrieval	373
	<i>Vittorio Castelli</i>	373
14.1	Introduction	373
14.2	Feature-Level Image Representation	374

14.3	Taxonomies of Indexing Structures	390
14.4	The Main Classes of Multidimensional Indexing Structures	391
14.5	Choosing an Appropriate Indexing Structure	396
14.6	Future Directions	398
	Appendix	399
	References	424
15	Multimedia Indexing	435
	<i>Christos Faloutsos</i>	435
15.1	Introduction	435
15.2	Gemini: Fundamentals	437
15.3	1D Time Series	441
15.4	2D Color Images	447
15.5	Extension: Subpattern Matching	452
15.6	Conclusion	458
	Appendix: Survey	459
	References	460
16	Compressed or Progressive Image Search	465
	<i>Sethuraman Panchanathan</i>	465
16.1	Introduction	465
16.2	Image Indexing in the Compressed Domain	466
16.3	Conclusion	492
16.4	Summary	492
	Acknowledgment	493
	References	493
17	Concepts and Techniques for Indexing Visual Semantics	497
	<i>Alejandro Jaimes and Shih-Fu Chang</i>	497
17.1	Introduction	497
17.2	Indexing Visual Information at Multiple Levels	500
17.3	Content-Based Techniques and Object Recognition	510
17.4	Learning Visual Object Detectors: The Visual Apprentice	536
17.5	Conclusion and Future Directions	554
	Appendix	555
	References	556
	Index	567

1 Digital Imagery: Fundamentals

VITTORIO CASTELLI

IBM T.J. Watson Research Center, Yorktown Heights, New York

LAWRENCE D. BERGMAN

IBM T.J. Watson Research Center, Hawthorne, New York

1.1 DIGITAL IMAGERY

Digital images have a predominant position among multimedia data types. Unlike video and audio, which are mostly used by the entertainment and news industry, images are central to a wide variety of fields ranging from art history to medicine, including astronomy, oil exploration, and weather forecasting. Digital imagery plays a valuable role in numerous human activities, such as law enforcement, agriculture and forestry management, earth science, urban planning, as well as sports, newscasting, and entertainment.

This chapter provides an overview of the topics covered in this book. We first describe several applications of digital imagery, some of which are covered in Chapters 2 to 5. The main technological factors that support the management and exchange of digital imagery, namely, acquisition, storage (Chapter 6), database management (Chapter 7), compression (Chapter 8), and transmission (Chapter 9) are then discussed.

Finally, a section has been devoted to content-based retrieval, a large class of techniques specifically designed for retrieving images and video. Chapters 10 to 17 cover these topics in detail.

1.2 APPLICATIONS OF DIGITAL IMAGES

Applications of digital imagery are continually developing. In this section, some of the major ones have been reviewed and the enabling economical and technological factors have been discussed briefly.

1.2.1 Visible Imagery

Photographic images are increasingly being acquired, stored, and transmitted in digital format. Their applications range from personal use to media and advertising, education, art, and even research in the humanities.

Image Databases: Search and Retrieval of Digital Imagery, Edited by Vittorio Castelli and Lawrence D. Bergman.

ISBN 0-471-32116-8 © 2002 John Wiley & Sons, Inc.

In the consumer market, digital cameras are slowly replacing traditional film-based cameras. The characteristics of devices for acquiring, displaying, and printing digital images are improving while their prices are decreasing. The resolution and color fidelity of digital cameras and desktop scanners are improving. Advancements in storage technology make it possible to store large number of pictures in digital cameras before uploading them to a personal computer. Inexpensive color printers can produce good quality reproductions of digital photographs. Digital images are also easy to share and disseminate: they can be posted on personal web sites or sent via e-mail to distant friends and relatives at no cost.

The advertisement and the media industries maintain large collection of images and need systems to store, archive, browse, and search them by content.

Museums and art galleries are increasingly relying on digital libraries to organize and promote their collections [1], and advertise special exhibits. These digital libraries, accessible via the Internet, provide an excellent source of material for the education sector and, in particular, for K-12.

A novel and extremely important application of digital libraries is to organize collections of rare and fragile documents. These documents are usually kept in highly controlled environments, characterized by low light and precise temperature and humidity levels. Only scholars can gain access to such documents, usually for a very limited time to minimize the risk of damage. Technology is changing this scenario: complex professional scanners have been developed that have very good color fidelity and depth (42-bit color or more) as well as high resolution. These scanners can capture the finest details, even those invisible without the aid of a magnifying lens, without risk to the original documents. The resulting digital images can be safely distributed to a wide audience across the Internet, allowing scholars to study otherwise inaccessible documents.

1.2.2 Remotely Sensed Images

One of the earliest application areas of digital imagery was remote sensing. Numerous satellites continuously monitor the surface of the earth. The majority of them measure the reflectance of the surface of the earth or atmospheric layers. Others measure thermal emission in the far-infrared and near-microwave portion of the spectrum, while yet others use synthetic-aperture radars and measure both reflectance and travel time (hence elevation). Some instruments acquire measurements in a single portion of the spectrum; others simultaneously acquire images in several spectral bands; finally, some radiometers acquire measurements in tens or hundreds of narrow spectral bands. Geostationary satellites on a high equatorial orbit are well suited to acquire low-resolution images of large portions of the earth's surface (where each pixel corresponds to tens of square miles), and are typically used for weather prediction. Nongeostationary satellites are usually on a polar orbit—their position relative to the ground depends both on their orbital motion and on the rotation of the earth. Lower-orbiting satellites typically acquire higher-resolution images but require more revolutions to cover the

entire surface of the earth. Satellites used for environmental monitoring usually produce low-resolution images, where each pixel corresponds to surface areas on the order of square kilometers. Other commercial satellites have higher resolution. The Landsat TM instrument has a resolution of about 30 m on the ground, and the latest generation of commercial instruments have resolutions of 1 to 3 m. Satellites for military applications have even higher resolution.

The sheer volume of satellite-produced imagery, on the order of hundreds of gigabytes a day, makes acquisition, preparation, storage, indexing, retrieval, and distribution of the data very difficult.

The diverse community of users often combine remotely sensed images with different types of data, including geographic or demographic information, ground-station observations, and photographs taken from planes. The resulting need for data fusion and interoperability poses further challenges to database and application developers.

1.2.3 Medical Images

Images are used in medicine for both diagnostic and educational purposes. Radiology departments produce the vast majority of medical images, while anatomic photographs and histological microphotographs account for a small fraction of the overall data volume.

Radiological images capture physical properties of the body, such as opacity to X rays (radiographies and CT scans), reflectance to ultrasounds, concentration of water or other chemicals (MRI), and distribution of elements within organs (PET). Medical images can be used to investigate anatomic features (e.g., broken bones or tumors) and physiological functions (e.g., imbalances in the activity of specific organs).

The availability of high-quality, high-resolution displays of sensors with better characteristics (sensitivity, quantum efficiency, etc.) than photographic film, of fast interconnection networks, and of inexpensive secondary and tertiary storage are driving radiology departments toward entirely digital, filmless environments, wherein image databases play a central role.

The main challenges faced by medical image databases are integration with the hospital work flow and interoperability.

1.2.4 Geologic Images

Oil companies are among the main producers and consumers of digital imagery. Oil exploration often starts with seismic surveys, in which large geologic formations are imaged by generating sound waves and measuring how they are reflected at the interface between different strata. Seismic surveys produce data that is processed into two- or three-dimensional imagery.

Data are also routinely acquired during drilling operation. Measurements of physical properties of the rock surrounding the well bore are measured with special-purpose imaging tools, either during drilling or afterwards. Some instruments measure aggregate properties of the surrounding rock and produce a single

measurement every sampling interval; others have arrays of sensors that take localized measurements along the circumference of the well bore. The former measures are usually displayed as one-dimensional signals and the latter measures are displayed as (long and thin) images.

Sections of rock (core) are also selectively removed from the bottom of the well, prepared, and photographed for further analysis. Visible-light or infrared-light microphotographs of core sections are often used to assess structural properties of the rock, and a scanning electron microscope is occasionally used to produce images at even higher magnification.

Image databases designed for the oil industry face the challenges of large data volumes, a wide diversity of data formats, and the need to combine data from multiple sources (data fusion) for the purpose of analysis.

1.2.5 Biometric Identification

Images are widely used for personal-identification purposes. In particular, fingerprints have long been used in law enforcement and are becoming increasingly popular for access control to secure information and identity. checks during firearm sales. Some technologies, such as face recognition, are still in the research domain, while others, such as retinal scan matching, have very specialized applications and are not widespread.

Fingerprinting has traditionally been a labor-intensive manual task performed by highly skilled workers. However, the same technological factors that have enabled the development of digital libraries, and the availability of inkless fingerprint scanners, have made it possible to create digital fingerprint archives[2,3].

Fingerprint verification (to determine if two fingerprints are from the same finger), identification (retrieving archived fingerprints that match the one given), and classification (assigning a fingerprint to a predefined class) all rely on the positions of distinctive characteristics of the fingerprint called *minutiae*. Typical minutiae are the points where a ridge bifurcates or where a ridge terminates. Fingerprint collections are searched by matching the presence and relative positions of minutiae. Facial matching procedures operate similarly—extracting essential features and then matching them against pre-extracted features from images in the database.

The main challenges in constructing biometric databases are the reliable extraction of minutiae and the matching strategy. Matching, in particular, is difficult: it must rely on rotation- and translation-invariant algorithms, it must be robust to missing and spurious data (the extraction algorithm might fail to identify relevant features or might extract nonexistent features, especially if the image quality is poor), and it must account for distortions due to lighting and positioning. The development of efficient indexing methods that satisfy these requirements is still an open research problem.

1.2.6 Astronomical Imagery

Astronomers acquire data in all regions of the electromagnetic spectrum. Although the atmosphere blocks most high-energy waves (UV, X- rays and γ -

rays), as well as large portions of the infrared and lower-frequency waves, it has large transparency windows that allowed the development of visible-light and radio wave astronomy. High-energy astronomy is possible using instruments mounted on high-altitude planes or orbiting satellites. Radio telescopes acquire signals in the long-wave, short-wave and microwave ranges, and are used to produce two-dimensional maps, often displayed as images.

Traditionally, astronomy has heavily relied on plate-based photography for infrared, visual, and ultraviolet studies (in astronomy, glass plates are used instead of photographic film). Long exposures (of up to tens of hours) make it possible to capture objects that are one to two orders of magnitude too dim to be detected by the human eye through the same instrument.

Photographic plates are not the ideal detectors—they are expensive and fragile, often have small defects that can hide important information, are not very sensitive to light, lose sensitivity during exposure, and their reproduction for distribution is labor-intensive. Their main benefits are large size and high resolution.

Starting from the mid-1980s, sensors that acquire images in digital format have become more and more widely used. In particular, charge-coupled devices (CCD) have found widespread application in astronomy. High-resolution sensor arrays, with responses that go beyond the visible spectrum are now commonly available. These instruments are extremely sensitive—when coupled with photomultipliers, they can almost detect the arrival of individual photons. Images that used to require hours of exposure can now be produced in minutes or less. Additionally, techniques exist to digitally reduce the inherent electrical noise of the sensor, further enhancing the quality of the image. Since the images are produced directly in digital format, a photograph is often acquired by collecting a sequence of short-exposure snapshots and combining them digitally. Image-processing techniques exist to compensate for atmospheric turbulence and for inaccuracies in telescope movement. Solid-state devices are also the detectors of choice for orbiting telescopes.

Digital libraries that organize the wealth of astronomical information are growing continuously and are increasingly providing support for communities beyond professional astronomers and astrophysicists, including school systems and amateur astronomers.

1.2.7 Document Management

Digital imagery plays an increasingly important role in traditional office management. Although we are far from the “paperless office” that many have envisioned, more and more information is being stored digitally, much of it in the form of imagery.

Perhaps the best case in point is archiving of cancelled checks. This information in the past was stored on microfilm—a medium that was difficult to manage. Moving this information to digital storage has resulted in enhanced ease of access and reduced storage volume. The savings are even more dramatic when digital imagery is used to replace paper records.