

SPRINGER BRIEFS IN COMPUTER SCIENCE

Jenny Benois-Pineau
Frédéric Precioso
Matthieu Cord

Visual Indexing and Retrieval

 Springer

SpringerBriefs in Computer Science

Series Editors

Stan Zdonik

Peng Ning

Shashi Shekhar

Jonathan Katz

Xindong Wu

Lakhmi C. Jain

David Padua

Xuemin Shen

Borko Furht

VS Subrahmanian

For further volumes:

<http://www.springer.com/series/10028>

Jenny Benois-Pineau • Frédéric Precioso
Matthieu Cord

Visual Indexing and Retrieval

 Springer

Jenny Benois-Pineau
Univ. Bordeaux
LABRI UMR 5800
Talence, France
jenny.benois@labri.fr

Frédéric Precioso
UNS
I3S UMR 7271/UNS_CNRS
Sophia Antipolis Cedex, France
precioso@polytech.unice.fr

Matthieu Cord
UPMC
LIP6 UMR 7606
matthieu.cord@lip6.fr

ISSN 2191-5768
ISBN 978-1-4614-3587-7
DOI 10.1007/978-1-4614-3588-4
Springer New York Heidelberg Dordrecht London

ISSN 2191-5776 (electronic)
ISBN 978-1-4614-3588-4 (eBook)

Library of Congress Control Number: 2012935105

© The Author(s) 2012

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed. Exempted from this legal reservation are brief excerpts in connection with reviews or scholarly analysis or material supplied specifically for the purpose of being entered and executed on a computer system, for exclusive use by the purchaser of the work. Duplication of this publication or parts thereof is permitted only under the provisions of the Copyright Law of the Publisher's location, in its current version, and permission for use must always be obtained from Springer. Permissions for use may be obtained through RightsLink at the Copyright Clearance Center. Violations are liable to prosecution under the respective Copyright Law.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

While the advice and information in this book are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

Preface

Visual Indexing and Retrieval is a wide-scope research domain which unites researchers from image analysis and computer vision, to information management. For nearly twenty years, this research has been extremely active world-wide. With the advent of social networks conveying huge amount of visual information to the user, the ever increasing capacities of broadcast transmission of content via heterogeneous networks, popularization of hand-carried devices, the realm of possibilities has got wide open both for industry leaders and researchers. Indeed, the always increasing size of visual content databases has brought to the fore the need for innovative content understanding, retrieval and classification methods. Such methods are henceforth of paramount importance to let people the ability to exploit such huge databases.

The book is the result of joint efforts of the French research community joined through the GDR CNRS ISIS, the national academic network in the field. Thanks to this network, the research results gathered and explained in this book have received global recognition, and are gaining more and more success in technology transfer. The authors hope that the most recent results and fruitful trends in Visual Indexing and Retrieval presented in this book will be helpful to young and experienced researchers willing to put the ideas forward and needing a solid understanding of the state-of-the art to do so, as well as to industry people willing to find an adequate algorithmic solution for their application area.

Acknowledgments: the authors acknowledge the governing board of GDR CNRS ISIS for supporting the project of this book.

Paris, France
January 2012

*Jenny Benois-Pineau
Frédéric Precioso
Matthieu Cord*

Contents

1	Introduction	1
1.1	Context and motivations	1
1.2	Outline of the book	2
2	Visual feature extraction and description	5
2.1	Introduction	5
2.2	Visual primitive detection	5
2.2.1	Point-based detectors	6
2.2.2	Region-based detectors	9
2.2.3	Spatio-temporal extension	12
2.3	Descriptors	13
2.3.1	Feature spaces used by descriptors	13
2.3.2	Scale Invariant Feature Transform	16
2.3.3	Speeded Up Robust Features	18
2.3.4	Global GIST descriptor	19
2.4	Evaluation of feature detectors and descriptors	19
2.5	Conclusion	20
3	Machine learning approaches for visual information retrieval	21
3.1	Bag-of-Feature representations and similarities	21
3.1.1	Bag-of-Visual-Words approaches (BoVW)	22
3.1.2	Vector distances and kernel functions	24
3.1.3	Bag-of-Features (BoF) similarity and retrieval	27
3.2	Learning algorithms	31
3.2.1	Support Vector Machines	31
3.2.2	Multiple Kernel Learning for feature combination	33
3.2.3	Boosting	35
3.2.4	Linear Programming Boosting (LPBoost) for Multiple Kernel Learning (MKL)	36
3.2.5	Interactive learning	38
3.3	Conclusion	40

- 4 Spatial and multi-resolution context in visual indexing 41**
 - 4.1 Introduction 41
 - 4.2 Incorporating spatial context 43
 - 4.2.1 Local histograms of visual words 43
 - 4.2.2 Context-matching kernels 44
 - 4.2.3 Graph-matching 46
 - 4.2.4 Graph Words 46
 - 4.3 Multi-resolution in visual indexing 55
 - 4.3.1 Low resolution and Rough Indexing Paradigm 55
 - 4.3.2 Multi-resolution and multiscale in image indexing 56
 - 4.3.3 Multi-resolution features and visual dictionaries 57
 - 4.4 Conclusion 62

- 5 Scalability issues in visual information retrieval 65**
 - 5.1 Introduction 65
 - 5.2 Scalable Retrieval and Mining: A Typology of Problems 66
 - 5.3 Principled Approaches to Scalability 67
 - 5.3.1 Scalable Retrieval 68
 - 5.3.2 Scalable Mining 71
 - 5.3.3 How to Evaluate Scalability 74
 - 5.4 Trends in Scalable Visual Information Retrieval 75
 - 5.4.1 Complex Data 75
 - 5.4.2 Optimized Representations and Scalability 77
 - 5.4.3 Distributed Data and Resources 79
 - 5.5 Conclusion 80

- 6 Evaluation of visual information indexing and retrieval 83**
 - 6.1 Introduction 83
 - 6.2 Organizing an evaluation campaign 84
 - 6.2.1 Organization 85
 - 6.2.2 Terminology 85
 - 6.2.3 Agenda 86
 - 6.2.4 Copyrights and legal matters on the content set distribution
and exploitation 87
 - 6.2.5 Ground truth and annotation process 88
 - 6.2.6 Evaluation metrics 88
 - 6.3 Main evaluation campaigns overview 90
 - 6.3.1 TRECVID 90
 - 6.3.2 PASCAL VOC and ILSVRC 93
 - 6.3.3 Other evaluation campaigns 94
 - 6.4 Conclusion 95

- References 97**

Chapter 1

Introduction

Jenny Benois-Pineau, Frédéric Precioso, Matthieu Cord

1.1 Context and motivations

The research in visual information indexing and retrieval has become one of the most popular directions in the broad area of information technologies. The reasons for that are the technological maturity of capture, storage and network infrastructures, that allow for common daily-life capturing of images and recording of video with professional equipment and personal mobile devices. According to Internet sources the British Broadcasting Corporation set up team dedicated to process user-generated content, as an experimental group in April 2005 with 3 employees in the staff. The team was then made durable and got expanded, unveiling the integration of “citizen journalist” in the (broadcast) news mainstream. The same concept has been put in place by CNN that launched CNN iReport in 2006. This project was meant to allow CNN to collect user-generated news. So did the american Fox News with their ‘uReport’ project and the french BFM-TV broadcast channels. YouTube, FaceBook, FileMobile, DailyMotion, host and supply facilities for accessing a tremendous amount of professional and user-generated content for educational and entertainment purposes. The areas of societal activity such as video surveillance and security also generate thousands of tera-bytes of video content with specific issues to be tackled. Finally, the digitization and storage of cultural heritage be it Byzantine frescoes, Medieval miniatures, old manuscripts or feature films, documentaries and broadcasting programs or web-sites, lead to the production of a mass of visual data which has to be accessed and searched both by professionals for re-mastering and production of new visual content and by common users for various humanities research.

Thus, visual information indexing and retrieval has attracted a lot of research efforts since the early nineties [24]. However nowadays, owing to the size of large scale databases, complexity of visual interpretation tasks, the famous “semantic gap” between the (semantic) concept(s) the user is looking for in images and the digital representation of the (semantic) visual content, the research directions are widely open. This research field has been dramatically enriched by the last achievements

in machine learning methods, the standardization of content descriptors and large-scale evaluation campaigns providing researchers with access to amounts of visual data associated to benchmarks for testing and evaluation. In this book, we propose a detailed review of the cutting-edge approaches in Visual Information Indexing, along with the original contributions of the authors of chapters.

1.2 Outline of the book

The book structure has naturally emerged from the key components of the visual indexing and retrieval problems. The next chapters, – Spatial and multi-resolution descriptions in visual indexing and – Machine learning approaches for visual information retrieval, are presenting the ingredients of the whole process. The next ones are dedicated to advanced considerations for the image representations, the scalability issues, and the evaluation.

Visual indexing and retrieval tasks always start with feature extraction from the raw data. Many approaches in several contexts have been developed since more than 30 years for now. Many primitives as Points of interest, regions, lines and so on, have been studied. Additionally, efficient descriptors are required. The Chapter 2 provides a deep overview of the basic feature extraction and description methods in the literature.

Representing these features at a higher semantic level description is the second stage of the process. The main steps for deriving image representation from visual local descriptors are described in Chapter 3. Well-known Bag-of-Visual-Words including recent extensions as sparse coding and spatial pooling methods are detailed. Higher semantic level data representation and similarity design are not independent but strongly related processes. Similarity measures between histograms, but also more complex functions such as kernels are presented. In order to address classification, retrieval or detection tasks, these similarities must be integrated into machine learning frameworks. We chose to focus on two major contributions from the Machine Learning community, namely Support Vector Machines and Boosting, very successful in multimedia retrieval applications.

Spatial structure is a key point in the building of the image representation. This information is usually ignored in basic representations. Advanced approaches try to overcome this drawback by adding spatial arrangements in the process. Furthermore, the natural willing to use information according to its visual importance and compressed-stream analysis yields multi-resolution approaches. In Chapter 4, integration of spatial context into visual description is deeply studied. Two trends on how to account for such information are considered. The first is the design of structural descriptors and integration of the spatial context in the signature matching. The second relates to the multi-resolution visual indexing.

Since image and video databases are becoming ever larger, scalability is definitely a critical requirement for visual information retrieval. The chapter 5 is dedicated to these scalability issues. The nature of the problems both for content-based

retrieval and for mining are described. Main ideas and recent advances are presented, like the use of approximation or of shared-neighbour similarity. We also highlight some prospective directions, like embeddings, filtering based on simplified descriptions, optimization of content representations and distributed processing.

Finally, it is crucial to evaluate the visual indexing and retrieval methods and systems against common benchmark corpora. Having a deep understanding of the evaluation process is the best way to identify the strengths and weaknesses of the different approaches, and to trace promising ways to investigate. The chapter 6 gives an overview of the major evaluation campaigns and benchmarks for visual indexing and retrieval tasks. Data collection, relevance judgements, performance measures and experimentation protocols are discussed. The State-of-the-Art performance in recent campaigns and the lessons learnt from these campaigns are also presented.