

图书基本信息

书名：<<大规模多媒体信息管理与检索基础>>

13位ISBN编号：9787302249764

10位ISBN编号：7302249768

出版时间：2011-5

出版时间：清华大学出版社

作者：张智威

页数：300

版权说明：本站所提供下载的PDF图书仅提供预览和简介，请支持正版图书。

更多资源请访问：<http://www.tushu007.com>

<<大规模多媒体信息管理与检索基础>>

内容概要

大规模多媒体信息管理与检索面临着两大类艰巨的技术挑战。

首先，这一工程问题的研究在本质上是多领域、跨学科的，涉及信号处理、计算机视觉、数据库、机器学习、神经科学和认知心理学；其次，一个有效的解决方案必须能解决高维数据和网络规模数据的可扩展性问题。

《大规模多媒体信息管理与检索基础（英）：模拟人类感知数学方法》第一部分（第1~8章）着重介绍如何采用多领域、跨学科算法来解决特征提取及选择、知识表示、语义分析、距离函数的制定等问题；第二部分（第9~12章）对解决高维数据和网络规模数据的扩展性问题提出了有效的处理方法。此外，《大规模多媒体信息管理与检索基础（英）：模拟人类感知数学方法》的附录还给出了作者开发的开源软件的下载地址。

《大规模多媒体信息管理与检索基础（英）：模拟人类感知数学方法》是作者在美国加州大学从事多年的教学科研及在google公司工作多年的基础上编写的。

《大规模多媒体信息管理与检索基础（英）：模拟人类感知数学方法》适合多媒体、计算机视觉、机器学习、大规模数据处理等领域的研发人员阅读，也可作为高等院校计算机专业本科生及研究生的教材或教学参考书。

作者简介

张智威, Dr. Edward Y. Chang was a professor at the Department of Electrical & Computer Engineering, University of California at Santa Barbara, before he joined Google as a research director in 2006. Dr. Chang received his M.S. degree in Computer Science and Ph.D degree in Electrical Engineering, both from Stanford University.

书籍目录

- 1 introduction - key subroutines of multimedia data management
 - 1.1 overview
 - 1.2 feature extraction
 - 1.3 similarity
 - 1.4 learning
 - 1.5 multimodal fusion
 - 1.6 indexing
 - 1.7 scalability
 - 1.8 concluding remarks
- references
- 2 perceptual feature extraction
 - 2.1 introduction
 - 2.2 dmd algorithm
 - 2.2.1 model-based pipeline
 - 2.2.2 data-driven pipeline
 - 2.3 experiments
 - 2.3.1 dataset and setup
 - 2.3.2 model-based vs. data-driven
 - 2.3.3 dmd vs. individual models
 - 2.3.4 regularization tuning
 - 2.3.5 tough categories
 - 2.4 related reading
 - 2.5 concluding remarks
- references
- 3 query concept learning
 - 3.1 introduction
 - 3.2 support vector machines and version space
 - 3.3 active learning and batch sampling strategies
 - 3.3.1 theoretical foundation
 - 3.3.2 sampling strategies
 - 3.4 concept-dependent learning
 - 3.4.1 concept complexity
 - 3.4.2 limitations of active learning
 - 3.4.3 concept-dependent active learning algorithms
 - 3.5 experiments and discussion
 - 3.5.1 testbed and setup
 - 3.5.2 active vs. passive learning
 - 3.5.3 against traditional relevance feedback schemes
 - 3.5.4 sampling method evaluation
 - 3.5.5 concept-dependent learning
 - 3.5.6 concept diversity evaluation
 - 3.5.7 evaluation summary
 - 3.6 related reading
 - 3.6.1 machine learning

- 3.6.2 relevance feedback
- 3.7 relation to other chapters
- 3.8 concluding remarks
- references
- 4 similarity
- 4.1 introduction
- 4.2 mining image feature set
 - 4.2.1 image testbed setup
 - 4.2.2 feature extraction
 - 4.2.3 feature selection
- 4.3 discovering the dynamic partial distance function
 - 4.3.1 minkowski metric and its limitations
 - 4.3.2 dynamic partial distance function
 - 4.3.3 psychological interpretation of dynamic partial distance function
- 4.4 empirical study
 - 4.4.1 image retrieval
 - 4.4.2 video shot-transition detection
 - 4.4.3 near duplicated articles
 - 4.4.4 weighted dpf vs. weighted euclidean
 - 4.4.5 observations
- 4.5 related reading
- 4.6 concluding remarks
- references
- 5 formulating distance functions
- 5.1 introduction
- 5.2 dfa algorithm
 - 5.2.1 transformation model
 - 5.2.2 distance metric learning
- 5.3 experimental evaluation
 - 5.3.1 evaluation on contextual information
 - 5.3.2 evaluation on effectiveness
 - 5.3.3 observations
- 5.4 related reading
 - 5.4.1 metric learning
 - 5.4.2 kernel learning
- 5.5 concluding remarks
- references
- 6 multimodal fusion
- 6.1 introduction
- 6.2 related reading
 - 6.2.1 modality identification
 - 6.2.2 modality fusion
- 6.3 independent modality analysis
 - 6.3.1 pca
 - 6.3.2 ica
 - 6.3.3 img

- 6.4 super-kernel fusion
- 6.5 experiments
 - 6.5.1 evaluation of modality analysis
 - 6.5.2 evaluation of multimodal kernel fusion
 - 6.5.3 observations
- 6.6 concluding remarks
- references
- 7 fusing content and context with causality
 - 7.1 introduction
 - 7.2 related reading
 - 7.2.1 photo annotation
 - 7.2.2 probabilistic graphical models
 - 7.3 multimodal metadata
 - 7.3.1 contextual information
 - 7.3.2 perceptual content
 - 7.3.3 semantic ontology
 - 7.4 influence diagrams
 - 7.4.1 structure learning
 - 7.4.2 causal strength
 - 7.4.3 case study
 - 7.4.4 dealing with missing attributes
 - 7.5 experiments
 - 7.5.1 experiment on learning structure
 - 7.5.2 experiment on causal strength inference
 - 7.5.3 experiment on semantic fusion
 - 7.5.4 experiment on missing features
 - 7.6 concluding remarks
 - references
- 8 combinational collaborative filtering , considering personalization
 - 8.1 introduction
 - 8.2 related reading
 - 8.3 ccf : combinational collaborative filtering
 - 8.3.1 c-u and c-d baseline models
 - 8.3.2 ccf model
 - 8.3.3 gibbs & em hybrid training
 - 8.3.4 parallelization
 - 8.3.5 inference
 - 8.4 experiments
 - 8.4.1 gibbs + em vs. em
 - 8.4.2 the orkut dataset
 - 8.4.3 runtime speedup
 - 8.5 concluding remarks
 - references
- 9 imbalanced data learning
 - 9.1 introduction
 - 9.2 related reading

- 9.3 kernel boundary alignment
 - 9.3.1 conformally transforming kernel k
 - 9.3.2 modifying kernel matrix k
- 9.4 experimental results
 - 9.4.1 vector-space evaluation
 - 9.4.2 non-vector-space evaluation
- 9.5 concluding remarks
- references
- 10 psvm : parallelizing support vector machines on distributed computers
 - 10.1 introduction
 - 10.2 interior point method with incomplete cholesky factorization
 - 10.3 psvm algorithm
 - 10.3.1 parallel icf
 - 10.3.2 parallel ipm
 - 10.3.3 computing parameter b and writing back
 - 10.4 experiments
 - 10.4.1 class-prediction accuracy
 - 10.4.2 scalability
 - 10.4.3 overheads
 - 10.5 concluding remarks
- references
- 11 approximate high-dimensional indexing with kernel
 - 11.1 introduction
 - 11.2 related reading
 - 11.3 algorithm spheredex
 - 11.3.1 create - building the index
 - 11.3.2 search - querying the index
 - 11.3.3 update - insertion and deletion
 - 11.4 experiments
 - 11.4.1 setup
 - 11.4.2 performance with disk ios
 - 11.4.3 choice of parameter g
 - 11.4.4 impact of insertions
 - 11.4.5 sequential vs. random
 - 11.4.6 percentage of data processed
 - 11.4.7 summary
 - 11.5 concluding remarks
 - 11.5.1 range queries
 - 11.5.2 farthest neighbor queries
- references
- 12 speeding up latent dirichlet allocation with parallelization and pipeline strategies
 - 12.1 introduction
 - 12.2 related reading
 - 12.3 ad-lda : approximate distributed lda

- 12.3.1 parallel gibbs sampling and allreduce
- 12.3.2 mpi implementation of ad-lda
- 12.4 plda+
- 12.4.1 reduce bottleneck of ad-lda
- 12.4.2 framework of plda+
- 12.4.3 algorithm for pw processors
- 12.4.4 algorithm for pd processors
- 12.4.5 straggler handling
- 12.4.6 parameters and complexity
- 12.5 experimental results
- 12.5.1 datasets and experiment environment
- 12.5.2 perplexity
- 12.5.3 speedups and scalability
- 12.6 large-scale applications
- 12.6.1 mining social-network user latent behavior
- 12.6.2 question labeling (ql)
- 12.7 concluding remarks
- references

章节摘录

版权页：插图：Feature extraction is fundamental to all multimedia computing tasks. Features can be classified into two categories, content and context. Content refers directly to raw imagery, video, and music data such as pixels, motions, and tones, respectively, and their representations. Context refers to metadata collected or associated with content when a piece of data is acquired or published. For instance, EXIF camera parameters and GPS location are contextual information that some digital cameras can collect. Other widely used contextual information includes surrounding texts of an image/photo on a Web page, and social interactions on a piece of multimedia data instance. Context and content ought to be fused synergistically when analyzing multimedia data.

编辑推荐

《大规模多媒体信息管理与检索基础:模拟人类感知数学方法》: Foundations of Large-Scale Multimedia Information Management and Retrieval Mathematics of Perception covers knowledge representation and semantic analysis of multimedia data and scalability in signal extraction, data mining, and indexing. The book is divided into two parts: Part I - Knowledge Representation and Semantic Analysis focuses on the key components of mathematics of perception as it applies to data management and retrieval. These include feature selection/reduction, knowledge representation, semantic analysis, distance function formulation for measuring similarity, and multimodal fusion. Part II - Scalability Issues presents indexing and distributed methods for scaling up these components for high-dimensional data and Web-scale datasets. The book presents some real-world applications and remarks on future research and development directions. The book is designed for researchers, graduate students, and practitioners in the fields of Computer Vision, Machine Learning, Large-scale Data Mining, Database, and Multimedia Information Retrieval.

版权说明

本站所提供下载的PDF图书仅提供预览和简介，请支持正版图书。

更多资源请访问:<http://www.tushu007.com>