

Advances in Computer Vision and Pattern Recognition

Founding editor

Sameer Singh, Rail Vision, Castle Donington, UK

Series editor

Sing Bing Kang, Microsoft Research, Redmond, WA, USA

Advisory Board

Horst Bischof, Graz University of Technology, Austria

Richard Bowden, University of Surrey, Guildford, UK

Sven Dickinson, University of Toronto, ON, Canada

Jiaya Jia, The Chinese University of Hong Kong, Hong Kong

Kyoung Mu Lee, Seoul National University, South Korea

Yoichi Sato, The University of Tokyo, Japan

Bernt Schiele, Max Planck Institute for Computer Science, Saarbrücken, Germany

Stan Sclaroff, Boston University, MA, USA

More information about this series at <http://www.springer.com/series/4205>

Gabriela Csurka
Editor

Domain Adaptation in Computer Vision Applications

Editor
Gabriela Csurka
Naver Labs Europe
Meylan
France

ISSN 2191-6586 ISSN 2191-6594 (electronic)
Advances in Computer Vision and Pattern Recognition
ISBN 978-3-319-58346-4 ISBN 978-3-319-58347-1 (eBook)
DOI 10.1007/978-3-319-58347-1

Library of Congress Control Number: 2017939548

© Springer International Publishing AG 2017

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.

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.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, express or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Printed on acid-free paper

This Springer imprint is published by Springer Nature
The registered company is Springer International Publishing AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

To Gabriel, Elisabeth and Mikhaël

Preface

While the proliferation of sensors being deployed in cell phones, vehicles, buildings, roadways, and computers allows for larger and more diverse information to be collected, the cost of acquiring labels for all these data remains extremely high. To overcome the burden of annotation, alternative solutions have been proposed in the literature to learn decision making models by exploiting unlabeled data from the same domain (data acquired in similar conditions as the targeted data) or also data from related but different domains (different datasets due to different conditions or provided by different customers). In many real-world machine learning scenarios, using only the data from the same domain might be insufficient and data or models borrowed from similar domains can significantly improve the learning process. Such a process, referred to as *domain adaptation*, aims to leverage labeled data in one or more related domains (sources), in order to build models for a target domain.

Domain adaptation is particularly critical for service companies, where all machine learning components deployed in a given service solution should be customized for a new customer either by annotating new data or, preferably, by calibrating the models in order to achieve a contractual performance in the new environment. While adaptation across domains is a challenging task for many applications, in this book, we focus on solutions for *visual applications*.

The aim of the book is to give a relatively broad view of the field by selecting a diverse set of methods which made different advances in the field. The book begins with a comprehensive survey of domain adaptation and transfer learning, including historical shallow methods, more recent methods using deep architectures, and methods addressing computer vision tasks beyond image categorization, such as detection, segmentation or visual attributes. Then, Chap. 2 gives a deeper look at dataset bias in existing datasets when different representations including features extracted from deep architectures are used. The rest of the book is divided into four main parts, following the same structure as the survey presented in Chap. 1.

Part I is dedicated to shallow domain adaptation methods, beginning with the widely used Geodesic Flow Kernel (Chap. 3) and Subspace Alignment (Chap. 4). Both chapters propose solutions for selecting landmark samples in the source dataset. Chapter 5 presents domain-invariant embedding methods and Chap. 6

describes the Transductive Transfer Machine, a method that combines local feature space transformation with classifier selection and parameter adaptation. The first part ends with Chap. 7 that addresses domain adaptation cases where the access to the source data is constrained.

Part II is dedicated to deep adversarial discriminative domain adaptation methods. The first two methods presented use a confusion loss as an adversarial objective to adapt the source network towards the target data. The deep CORAL (Chap. 8) learns a nonlinear transformation that aligns correlations of activation layers of the deep model. The Deep Domain Confusion network (Chap. 9) uses a Maximum Mean Discrepancy based domain confusion loss to induce domain invariant representations. In contrast, Chap. 10 presents the domain-adversarial neural network that integrates a Gradient Reversal Layer to promote the emergence of features discriminative for the main learning task and non-discriminate with respect to the domain shift.

Part III is a collection of contributions addressing domain adaptation problems different from classical image categorization. As such, Chap. 11 focuses on Fisher vector based patch encoding adaptation in the context of vehicle re-identification. Chapter 12 explores the adaptation of semantic segmentation models trained on synthetic images to correctly operate in real scenarios. Chapter 13 addresses the challenge of pedestrian detection by adapting a Deformable Part-Based Model trained on virtual-world data to real world data using structure-aware adaptive structural SVMs. Finally, Chap. 14 proposes a method to generalize semantic part detectors across domains.

Part IV concludes the book with unifying perspectives. On the one hand, Chap. 15 proposes to use multi-source domain generalization techniques for the purpose of learning cross-category generalizable attribute detectors. On the other hand, Chap. 16 proposes a common framework that unifies multi-domain and multi-task learning which can be flexibly applied also to zero-shot learning and zero-shot domain adaptation.

Overall, this comprehensive volume, designed to form and inform professionals, young researchers, and graduate students, is the first collection dedicated to domain adaptation for visual applications. In this book I wanted not only to address historically shallow and recent deep domain adaptation methods, but also contributions focused on object or object part detection, re-identification, image segmentation, attribute detection as well to present generic frameworks that unify domain adaptation with multi-domain, multi-task and zero-shot learning.

To give such a broad view, I brought together leading experts in the field to showcase their techniques. I would like to thank them specially for accepting my invitation and for their dedicated effort to share in this book their valuable experiences in the various chapters. Finally, I would also like to thank our Springer editors, Wayne Wheeler and Simon Rees, for their advice and their help in guiding me through the book production process.

Contents

1	A Comprehensive Survey on Domain Adaptation for Visual Applications	1
	Gabriela Csurka	
2	A Deeper Look at Dataset Bias	37
	Tatiana Tommasi, Novi Patricia, Barbara Caputo and Tinne Tuytelaars	
Part I Shallow Domain Adaptation Methods		
3	Geodesic Flow Kernel and Landmarks: Kernel Methods for Unsupervised Domain Adaptation	59
	Boqing Gong, Kristen Grauman and Fei Sha	
4	Unsupervised Domain Adaptation Based on Subspace Alignment	81
	Basura Fernando, Rahaf Aljundi, Rémi Emonet, Amaury Habrard, Marc Sebban and Tinne Tuytelaars	
5	Learning Domain Invariant Embeddings by Matching Distributions	95
	Mahsa Baktashmotlagh, Mehrtash Harandi and Mathieu Salzmann	
6	Adaptive Transductive Transfer Machines: A Pipeline for Unsupervised Domain Adaptation	115
	Nazli Farajidavar, Teofilo de Campos and Josef Kittler	
7	What to Do When the Access to the Source Data Is Constrained?	133
	Gabriela Csurka, Boris Chidlovskii and Stéphane Clinchant	

Part II Deep Domain Adaptation Methods

- 8 Correlation Alignment for Unsupervised Domain Adaptation** 153
Baochen Sun, Jiashi Feng and Kate Saenko
- 9 Simultaneous Deep Transfer Across Domains and Tasks** 173
Judy Hoffman, Eric Tzeng, Trevor Darrell and Kate Saenko
- 10 Domain-Adversarial Training of Neural Networks** 189
Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain,
Hugo Larochelle, François Laviolette, Mario Marchand
and Victor Lempitsky

Part III Beyond Image Classification

- 11 Unsupervised Fisher Vector Adaptation for Re-identification** 213
Usman Tariq, Jose A. Rodriguez-Serrano and Florent Perronnin
- 12 Semantic Segmentation of Urban Scenes via Domain
Adaptation of SYNTHIA** 227
German Ros, Laura Sellart, Gabriel Villalonga, Elias Maidanik,
Francisco Molero, Marc Garcia, Adriana Cedeño, Francisco Perez,
Didier Ramirez, Eduardo Escobar, Jose Luis Gomez, David Vazquez
and Antonio M. Lopez
- 13 From Virtual to Real World Visual Perception Using Domain
Adaptation—The DPM as Example** 243
Antonio M. López, Jiaolong Xu, José L. Gómez, David Vázquez
and Germán Ros
- 14 Generalizing Semantic Part Detectors Across Domains** 259
David Novotny, Diane Larlus and Andrea Vedaldi

Part IV Beyond Domain Adaptation: Unifying Perspectives

- 15 A Multisource Domain Generalization Approach to Visual
Attribute Detection** 277
Chuang Gan, Tianbao Yang and Boqing Gong
- 16 Unifying Multi-domain Multitask Learning: Tensor
and Neural Network Perspectives** 291
Yongxin Yang and Timothy M. Hospedales

References 311

Index 341