The Maximum Consensus Problem

Recent Algorithmic Advances

Synthesis Lectures on Computer Vision

Synthesis Lectures on Computer Vision is edited by Gérard Medioni of the University of Southern California and Sven Dickinson of the University of Toronto. The series publishes 50- to 150 page publications on topics pertaining to computer vision and pattern recognition. The scope will largely follow the purview of premier computer science conferences, such as ICCV, CVPR, and ECCV. Potential topics include, but not are limited to:

- Applications and Case Studies for Computer Vision
- Color, Illumination, and Texture
- Computational Photography and Video
- Early and Biologically-inspired Vision
- Face and Gesture Analysis
- Illumination and Reflectance Modeling
- Image-Based Modeling
- Image and Video Retrieval
- Medical Image Analysis
- Motion and Tracking
- Object Detection, Recognition, and Categorization
- Segmentation and Grouping
- Sensors
- Shape-from-X
- Stereo and Structure from Motion
- Shape Representation and Matching

- Statistical Methods and Learning
- Performance Evaluation
- Video Analysis and Event Recognition

The Maximum Consensus Problem: Recent Algorithmic Advances
Tat-Jun Chin and David Suter
2017

Extreme Value Theory-Based Methods for Visual Recognition
Walter J. Scheirer
2017

Data Association for Multi-Object Visual Tracking
Margrit Betke and Zheng Wu
2016

Ellipse Fitting for Computer Vision: Implementation and Applications
Kenichi Kanatani, Yasuyuki Sugaya, and Yasushi Kanazawa
2016

Computational Methods for Integrating Vision and Language
Kobus Barnard
2016

Background Subtraction: Theory and Practice
Ahmed Elgammal
2014

Vision-Based Interaction
Matthew Turk and Gang Hua
2013

Camera Networks: The Acquisition and Analysis of Videos over Wide Areas
Amit K. Roy-Chowdhury and Bi Song
2012

Deformable Surface 3D Reconstruction from Monocular Images
Mathieu Salzmann and Pascal Fua
2010

Boosting-Based Face Detection and Adaptation
Cha Zhang and Zhengyou Zhang
2010

Image-Based Modeling of Plants and Trees
Sing Bing Kang and Long Quan
2009

The Maximum Consensus Problem: Recent Algorithmic Advances

Tat-Jun Chin and David Suter

ISBN: 978-3-031-00690-6 paperback
ISBN: 978-3-031-01818-3 ebook

DOI 10.1007/978-3-031-01734-6

A Publication in the Springer series
SYNTHESIS LECTURES ON COMPUTER VISION

Lecture #11
Series Editors: Gérard Medioni, *University of Southern California*
Sven Dickinson, *University of Toronto*
Series ISSN
Print 2153-1056 Electronic 2153-1064

The Maximum Consensus Problem

Recent Algorithmic Advances

Tat-Jun Chin and David Suter
The University of Adelaide

SYNTHESIS LECTURES ON COMPUTER VISION #11

ABSTRACT

Outlier-contaminated data is a fact of life in computer vision. For computer vision applications to perform reliably and accurately in practical settings, the processing of the input data must be conducted in a robust manner. In this context, the maximum consensus robust criterion plays a critical role by allowing the quantity of interest to be estimated from noisy and outlier-prone visual measurements. The *maximum consensus problem* refers to the problem of optimizing the quantity of interest according to the maximum consensus criterion. This book provides an overview of the algorithms for performing this optimization. The emphasis is on the basic operation or "inner workings" of the algorithms, and on their mathematical characteristics in terms of optimality and efficiency. The applicability of the techniques to common computer vision tasks is also highlighted. By collecting existing techniques in a single article, this book aims to trigger further developments in this theoretically interesting and practically important area.

For updates, errata, demo programs, and other information, please visit:

http://cs.adelaide.edu.au/~tjchin/maxcon/

We welcome contributions to the errata list.

KEYWORDS

robust fitting, maximum consensus, algorithms, optimization

Always try the problem that matters most to you.

Andrew Wiles

Contents

Preface

The book is organized into four chapters. The first chapter is primarily concerned with defining the maximum consensus problem as an optimization problem, and what is meant by the term "solution." An outline of the computational hardness of the problem is also given. This sets up the kinds of algorithms and efficiencies that one can realistically expect, with respect to theoretical complexity limitations.

The second chapter describes approximate algorithms for maximum consensus. These include the popular randomized sample-and-test heuristics, as well as more recent algorithms that employ some form of convex optimization as a subroutine. A large part of Chapter 2 is devoted to Chebyshev approximation (ℓ_∞ minimization) and linear program (LP)-type problems. Though they seem tangential to robust fitting or maximum consensus at the onset, these topics provide a consistent mathematical framework to talk about maximum consensus estimation for a useful class of nonlinear models.

The third chapter describes exact algorithms for maximum consensus. The fundamental intractability of maximum consensus means that all exact algorithms conduct some form of search. Thus, in a nutshell, Chapter 3 is about how to conduct the search efficiently (as least more efficiently than brute-force search). To this end, the underlying geometric structure of the model plays a crucial role, and it is thus given attention in the chapter.

The fourth chapter is about a relatively new idea in maximum consensus optimization—preprocessing, or data reduction. Different from a simple culling of input data, the preprocessing must preserve the optimal solution in the original input data. This implies retaining the maximum consensus set and removing only true outliers. How to efficiently identify true outliers and remove them is the main topic in this chapter.

The intention of this book is *not* to propose the "best-performing method"—such an endeavor should be carried out via comprehensive benchmarking, taking into account the requirements of the specific application (available time budget, desired repeatability of estimates, minimum solution quality, etc.) and the expected variability in the data. The main aim of this book is to explore the underlying concepts of the various algorithms and to establish properties related to efficiency and optimality. In some sections, however, sample results and brief comparisons are provided to illustrate the operation of the method.

Tat-Jun Chin and David Suter
February 2017

Acknowledgments

Throughout the authors' careers, they have been fortunate enough to work with many talented and dedicated individuals, including former and current students, postdocs under (and sometimes not under) our supervision, fellow academic staff in our immediate vicinity and more broadly in the community, as well as the researchers and R&D engineers from industry and nonacademic research organizations. These collaborations have had direct and indirect impact on making this book a reality.

We are indebted specifically to Álvaro Parra Bustos, Jin Yu, Huu Minh Le, Pulak Purkait, Trung Thanh Pham, Quoc Huy Tran, Hoi Sim Wong, and Yang Heng Kee, who co-contributed to the development of several of the algorithms and/or their predecessors, to some of the sample results, and to a number of the illustrative diagrams. We would also like to thank Alireza Khosravian, who proofread parts of the book. Special mention is given to Anders Eriksson, Wojciech Chojnacki, and Frank Neumann, who have been careful listeners and sanity testers of some of our initial ideas—it is increasingly rare to find brilliant researchers who are also generous enough with their time to think about "other people's problems."

Certainly the ideas in this book do not exist in a vacuum, as we have leveraged the hard work of a large number of researchers in computer vision and beyond. We can only express our appreciation by citing their excellent work (as far as the local context permits) in our book. We apologize in advance if we left out important works, and we would love to hear about them via email or the next time we cross paths at the conferences.

We thank the editors and editorial staff for their encouragement, kind understanding, and timely reminders to complete the book. We also record our appreciation of the reviewers who provided useful comments to improve the book.

Lastly, we thank the Australian Research Council (ARC), whose funding has been instrumental in enabling the necessary basic research.

Tat-Jun Chin and David Suter
February 2017