

Signal Processing for Computer Vision

Gösta H. Granlund
Hans Knutsson

Published by Kluwer Academic Publisher
P.O. Box 17, 3300 AA Dordrecht, The Netherlands

1995

ISBN: 0-7923-9530-1
ISBN: 978-0-7923-9530-0

CONTENTS

| | |
|--|----|
| PREFACE | v |
| 1 INTRODUCTION AND OVERVIEW | 1 |
| 1.1 Hierarchical Computing Structures | 3 |
| 1.2 Low level representation and operations | 12 |
| 1.3 Description in terms of symmetry | 20 |
| 1.4 Cascading of operations | 24 |
| 1.5 Compatibility representation | 25 |
| 1.6 Description of size and scale | 27 |
| 1.7 Model based processing | 33 |
| 1.8 Representation of volumes and time sequences | 38 |
| 1.9 Classification and response generation | 39 |
| 2 BIOLOGICAL VISION | 41 |
| 2.1 Motivation | 41 |
| 2.2 Overview of the visual system | 42 |
| 2.3 Properties of neurons | 43 |
| 2.4 The retina | 49 |
| 2.5 Color vision | 59 |
| 2.6 The visual pathways and the lateral geniculate nucleus | 67 |
| 2.7 The primary visual cortex | 69 |
| 2.8 Columnar organization and the layers of the cortex | 75 |
| 2.9 Possible implementations of filters | 80 |
| 2.10 Channel organization of features | 83 |
| 2.11 The visual pathways beyond the primary visual cortex | 89 |

| | | |
|----------|---|-----|
| 3 | LOW LEVEL OPERATIONS | 97 |
| 3.1 | Information representation | 97 |
| 3.2 | Complex valued convolution functions | 102 |
| 3.3 | Compact representation | 105 |
| 3.4 | Examples of processing | 109 |
| 3.5 | More on compatibility representation | 112 |
| 3.6 | General properties of a useful information representation | 114 |
| 4 | FOURIER TRANSFORMS | 117 |
| 4.1 | Introduction | 119 |
| 4.2 | Basics | 122 |
| 4.3 | Three aspects of the Fourier transform | 143 |
| 4.4 | Separability | 153 |
| 4.5 | Analytic signals | 164 |
| 4.6 | Examples of statistical calculations | 174 |
| 4.7 | Transform pairs | 177 |
| 4.8 | The Fourier transform and discrete functions | 180 |
| 4.9 | The discrete Fourier transform | 192 |
| 5 | KERNEL OPTIMIZATION | 199 |
| 5.1 | Spatial restrictions | 199 |
| 5.2 | Distance measure | 200 |
| 5.3 | Minimizing the distance | 201 |
| 5.4 | The weighting function | 201 |
| 5.5 | Optimization results | 203 |
| 5.6 | Kernel evaluation | 209 |
| 6 | ORIENTATION AND VELOCITY | 219 |
| 6.1 | Tensors - A short introduction | 222 |
| 6.2 | Representing orientation | 224 |
| 6.3 | Orientation estimation | 230 |
| 6.4 | Tensor Construction | 239 |
| 6.5 | Interpretation of the orientation tensor | 249 |
| 6.6 | Time sequences - velocity | 253 |
| 6.7 | Performance measures | 257 |

| | | |
|-----------|---|-----|
| 7 | LOCAL PHASE ESTIMATION | 259 |
| 7.1 | What is local phase? | 259 |
| 7.2 | Local Phase in Scale-space | 266 |
| 7.3 | Phase in higher dimensions | 271 |
| 7.4 | Applications using Local Phase Estimates | 275 |
| 8 | LOCAL FREQUENCY | 279 |
| 8.1 | Frequency estimation | 281 |
| 8.2 | Wide range frequency estimation | 288 |
| 8.3 | Experimental results | 290 |
| 9 | REPRESENTATION AND AVERAGING | 297 |
| 9.1 | Background | 297 |
| 9.2 | Invariance and equivariance | 298 |
| 9.3 | Signals and certainties | 299 |
| 9.4 | Vector and tensor representation of local orientation | 301 |
| 9.5 | Averaging | 303 |
| 9.6 | Normalized averaging | 304 |
| 10 | ADAPTIVE FILTERING | 309 |
| 10.1 | Adaptive filtering | 311 |
| 10.2 | Tensor controlled filter | 313 |
| 10.3 | The control tensor | 313 |
| 10.4 | Adaptive filter synthesis | 320 |
| 10.5 | Examples of filtering | 333 |
| 11 | VECTOR AND TENSOR FIELD FILTERING | 343 |
| 11.1 | Introduction | 343 |
| 11.2 | Orientation in vector and tensor fields | 343 |
| 11.3 | Rotational symmetries | 350 |
| 11.4 | Curvature estimation | 361 |
| 12 | CLASSIFICATION AND RESPONSE GENERATION | 367 |
| 12.1 | Information needed for classification | 368 |
| 12.2 | Linear discriminant functions | 371 |

| | | |
|-----------|---|------------|
| 12.3 | Minimum distance classification | 378 |
| 12.4 | Convolution operations | 380 |
| 12.5 | Generalized discriminant functions | 385 |
| 12.6 | Training of linear discriminant classifiers | 389 |
| 12.7 | Perceptrons and feedforward neural networks | 390 |
| 12.8 | Clustering and unsupervised classification | 396 |
| 13 | TEXTURE ANALYSIS | 399 |
| 13.1 | Introduction | 399 |
| 13.2 | Issues in texture analysis | 401 |
| 13.3 | Human aspects of texture perception | 402 |
| 13.4 | Feature-based texture analysis | 405 |
| 13.5 | Structural approach using certainty gating | 412 |
| | REFERENCES | 419 |
| | INDEX | 433 |

PREFACE

Interest in computer vision has increased considerably as other disciplines of computation and control have developed to require better flexibility and precision in their controlling information. Despite considerable development over the past few decades, the vision field is still in its infancy. Many fundamental questions remain of how to handle and represent spatial information in an effective way.

The intent of writing this book has been to give an introduction to robust low-level methods for computer vision. The field has seen a variety of methods originating from different researchers over a few decades. This book is not a review of published material in the field, but mainly an account of methods developed and tested at the Computer Vision Laboratory (CVL) of Linköping University in Sweden. The book includes as well certain methods developed at other sites; methods that conform to the requirements of robustness and modularity that are deemed essential for vision systems.

Although this book is attributed to two principal authors, both the work described and its documentation is the collective effort of a number of people who have contributed in various ways over the years.

Chapter 1, written by Gösta Granlund, establishes the motivation and use of hierarchical operation structures to provide a systematic organization for the implementation of complicated models. The chapter gives an intuitive treatment of most aspects that are considered in the later chapters.

Chapter 2, written by Gösta Granlund and Jörgen Karlholm, with contributions by Carl-Johan Westelius and Carl-Fredrik Westin, gives an overview of important biological vision mechanisms. Although a great deal is known about neural processing of visual information, most essential questions about biological vision remain as yet unanswered. Nonetheless, the knowledge available has already provided useful guidance to the organization of effective machine vision systems.

Chapter 3, written by Gösta Granlund with contributions by Johan Wiklund, gives an introductory treatment of operations and representations for low-level features in multi-dimensional spaces. An important issue is how to combine contributions from several filters to provide robust statements in accordance with certain low-level models. This chapter gives an introduction to the problems of unambiguous mappings in multi-dimensional spaces.

These introductory chapters establish the framework and motivate a more detailed treatment, which begins in Chapter 4. This chapter, written by Klas Nordberg, is on Fourier methods, with a particular emphasis on definitions and theorems essential to the understanding of filtering procedures in multi-dimensional spaces. This is a central issue in computer vision.

Chapter 5, written by Hans Knutsson, presents a method for obtaining an optimal n -dimensional set of filter coefficients for any given frequency response. An optimality criterion is defined that enables different frequencies to be given individual weights. Appropriate forms of frequency weight functions are discussed and a number of optimization examples are given.

Chapter 6, written by Hans Knutsson, introduces the use of tensors in estimation of local structure and orientation. The tensor representation is shown to be crucial to unambiguous and continuous representation of local orientation in multiple dimensions. In addition to orientation the tensor representation also conveys the degree and type of local anisotropy. The orientation estimation approach is developed in detail for two, three and four dimensions and is shown to be extendable to higher dimensions. Examples and performance measures are given for processing of images, volumes and time sequences. The processing and tests in three dimensions are the work of Mats Andersson, Leif Haglund and Johan Wiklund.

Chapter 7, written by Carl-Johan Westelius deals with the concept of phase and phase representation in multiple dimensions. Phase is an important concept, which emerges in several contexts in vision. The chapter provides a detailed treatment of phase properties in various situations and deals with how to maintain continuity in phase representation. An example is given of how local phase in difference scales can be used for disparity estimation. The chapter contains original material on the representation of phase for signals in three dimensions and higher dimensions.

Chapter 8, written by Hans Knutsson with contributions by Carl-Fredrik Westin, deals with the estimation of local frequency and bandwidth. Local frequency is an important concept which provides an indication of the appro-

appropriate range of scales for subsequent analysis. A number of one-dimensional and two-dimensional examples of local frequency and bandwidth estimation are given.

Chapter 9, written by Carl-Fredrik Westin with contributions by Hans Knutsson, considers what the important properties are for an information representation to behave well in various transformations. There is an extended discussion on the necessity to separate between class membership and certainty of a signal.

Chapter 10, written by Hans Knutsson, presents a computationally efficient technique for adaptive filtering of n -dimensional signals. The approach is based on the local signal description given by the orientation tensor discussed in Chapter 6. The adaptive filter output is synthesized as a tensor-controlled weighted summation of shift-invariant filter outputs. Several examples of adaptive filtering in two and three dimensions are given. Most of the results in the example section were produced by Leif Haglund. The chapter contains original results on the extension of the techniques to n dimensions.

Chapter 11, written by Carl-Fredrik Westin discusses techniques for processing of higher order data such as vector and tensor fields. As abstraction implies a more complex descriptor, developing methods for processing of higher order data is an essential part of any hierarchical or layered approach to vision. The chapter focuses on models for extracting local symmetries and discontinuities in higher order fields.

Chapter 12, written by Gösta Granlund with contributions by Jörgen Karlholm, is not original, but presents methods for linear classification in the tradition of N. J. Nilsson as well as R. O. Duda and P. E. Hart. Part of the motivation for including this well-known material is to allow the vision structure to be brought to a logical conclusion in which feature properties are combined to form responses or class statements. Another motivation developed here is to display the similarity in structure between convolution operations and linear discriminant functions. This brings all operations for feature extraction and classification to the use of a common component, linear discriminants. This is also illustrated in the form of perceptrons, which allows a transition to the modern theory of neural networks.

Chapter 13, written by Morgan Ulvklo, deals with texture analysis, an important application of the methods described in earlier chapters. It introduces ideas from preattentive vision, which gives clues for the extraction of texture primitives. There is also a discussion on how to handle features whose significance varies with spatial position.

In addition to writing Chapter 4, Klas Nordberg has orchestrated the production of the book with its detail work on organization, layout and figures and he is responsible for making the L^AT_EX compilations run. Catharina Holmgren has contributed with entry of text and pictorial material, as well as proof-reading of the manuscript. Finally, the authors want to acknowledge contributions made by all other members, past and present, of the CVL over the years since the start of its activities in the late seventies.

Matt Fichtenbaum of GenRad, Inc., Concord Massachusetts has reviewed the manuscript and suggested numerous improvements to the language and the description of technical topics. Dr. Elizabeth Liebson of Tufts - New England Medical Center in Boston, Massachusetts has reviewed the chapter on biological vision.

There are a number of persons who have visited CVL for shorter or longer periods, persons with whom we have discussed various issues of common interest, persons who have in one way or other been influential in the course of our research. Among those we would like to mention: Prof. Edward Adelson, MIT, Cambridge, USA; Dr. Andrew Calway, Univ. of Wales, Cardiff, Great Britain; Prof. C. Forbes Dewey, MIT, Cambridge, USA; Prof Ramesh Jain, Univ. of Michigan, Ann Arbor, USA; Prof. Bernd Jähne, Univ. of California, San Diego, California; Prof. Murat Kunt, EPFL, Lausanne, Switzerland; Prof. Steven Tanimoto, Univ. of Washington, Seattle, USA; Prof. Torsten N. Wiesel, Rockefeller University, New York; Dr. Roland Wilson, Univ. of Warwick, Coventry, Great Britain. The authors would also like to acknowledge the beneficial influence from partners in the ESPRIT Basic Research Action projects 3038/7108: Vision as Process.

We are grateful for the permission to reproduce copyright figures as referenced in the text. In a few cases it has proved impossible to reach copyright holders. We regret this and hope that no offence has been given.

The major part of the research work described in this book has been supported by research grants from The Swedish National Board for Industrial and Technical Development (NUTEK).

Gösta H. Granlund
Hans Knutsson