

Preface

Aims and Scope

This book is both an introductory textbook and a research monograph on modeling the statistical structure of natural images. In very simple terms, “natural images” are photographs of the typical environment where we live. In this book, their statistical structure is described using a number of statistical models whose parameters are estimated from image samples.

Our main motivation for exploring natural image statistics is computational modeling of biological visual systems. A theoretical framework which is gaining more and more support considers the properties of the visual system to be reflections of the statistical structure of natural images because of evolutionary adaptation processes. Another motivation for natural image statistics research is in computer science and engineering, where it helps in development of better image processing and computer vision methods.

While research on natural image statistics has been growing rapidly since the mid-1990s, no attempt has been made to cover the field in a single book, providing a unified view of the different models and approaches. This book attempts to do just that. Furthermore, our aim is to provide an accessible introduction to the field for students in related disciplines.

However, not all aspects of such a large field of study can be completely covered in a single book, so we have had to make some choices. Basically, we concentrate on the neural modeling approaches at the expense of engineering applications. Furthermore, those topics on which the authors themselves have been doing research are inevitably given more emphasis.

Targeted Audience and Prerequisites

The book is targeted for advanced undergraduate students, graduate students and researchers in vision science, computational neuroscience, computer vision, and image processing. It can also be read as an introduction to the area by people with a background in mathematical disciplines (mathematics, statistics, theoretical physics).

Due to the multi-disciplinary nature of the subject, the book has been written so as to be accessible to an audience coming from very different backgrounds such as psychology, computer science, electrical engineering, neurobiology, mathematics, statistics, and physics. Therefore, we have attempted to reduce the prerequisites to a minimum. The main thing needed are basic mathematical skills as taught in introductory university-level mathematics courses. In particular, the reader is assumed to

know the basics of

- univariate calculus (e.g. one-dimensional derivatives and integrals)
- linear algebra (e.g. inverse matrix, orthogonality)
- probability and statistics (e.g. expectation, probability density function, variance, covariance)

To help readers with a modest mathematical background, a crash course on linear algebra is offered at Chap. 19, and Chap. 4 reviews probability theory and statistics on a rather elementary level.

No previous knowledge of neuroscience or vision science is necessary for reading this book. All the necessary background on the visual system is given in Chap. 3, and an introduction to some basic image processing methods is given in Chap. 2.

Structure of the Book and Its Use as a Textbook

This book is a hybrid of a monograph and an advanced graduate textbook. It starts with background material which is rather classic, whereas the latter parts of the book consider very recent work with many open problems. The material in the middle is quite recent but relatively established.

The book is divided into the following parts:

Introduction which explains the basic setting and motivation.

Part I which consists of background chapters. This is mainly classic material found in many textbooks in statistics, neuroscience, and signal processing. However, here it has been carefully selected to ensure that the reader has the right background for the main part of the book.

Part II starts the main topic, considering the most basic models for natural image statistics. These models are based on the statistics of linear features, i.e. linear combinations of image pixel values.

Part III considers more sophisticated models of natural image statistics, in which dependencies (interactions) of linear features are considered, which is related to computing non-linear features.

Part IV applies the models already introduced to new kinds of data: color images, stereo images, and image sequences (video). Some new models on the temporal structure of sequences are also introduced.

Part V consists of a concluding chapter. It provides a short overview of the book and discusses open questions as well as alternative approaches to image modeling.

Part VI consists of mathematical chapters which are provided as a kind of an appendix. Chapter 18 is a rather independent chapter on optimization theory. Chapter 19 is background material which the reader is actually supposed to know; it is provided here as a reminder. Chapters 20 and 21 provide sophisticated supplementary mathematical material for readers with such interests.

Dependencies of the parts are rather simple. When the book is used as a textbook, **all readers should start by reading the first seven chapters** in the order they are

given (i.e. Introduction, Part I, and Part II except for the last chapter), unless the reader is already familiar with some of the material. After that, it is possible to jump to later chapters in almost any order, except for the following:

- Chapter 10 requires Chap. 9, and Chap. 11 requires Chaps. 9 and 10.
- Chapter 14 requires Sect. 13.1.

Some of the sections are marked with an asterisk *, which means that they are more sophisticated material which can be skipped without interrupting the flow of ideas.

An introductory course on natural image statistics can be simply constructed by going through the first n chapters of the book, where n would typically be between 7 and 17, depending on the amount of time available.

Referencing and Exercises

To keep the text readable and suitable for a textbook, the first 11 chapters do not include references in the main text. References are given in a separate section at the end of the chapter. In the latter chapters, the nature of the material requires that references are given in the text, so the style changes to a more scholarly one. Likewise, mathematical exercises and computer assignments are given for the first 10 chapters.

Code for Reproducing Experiments

For pedagogical purposes as well as to ensure the reproducibility of the experiments, the Matlab™ code for producing most of the experiments in the first 11 chapters, and some in Chap. 13, is distributed on the Internet at

www.naturalimagestatistics.net

This web site will also include other related material.

Acknowledgements

We would like to thank Michael Gutmann, Asun Vicente, and Jussi Lindgren for detailed comments on the manuscript. We have also greatly benefited from discussions with Bruno Olshausen, Eero Simoncelli, Geoffrey Hinton, David Field, Peter Dayan, David Donoho, Pentti Laurinen, Jussi Saarinen, Simo Vanni, and many others. We are also very grateful to Dario Ringach for providing the reverse correlation results in Fig. 3.7. During the writing process, the authors were funded by the University of Helsinki (Department of Computer Science and Department of Mathematics and Statistics), the Helsinki Institute for Information Technology, and the Academy of Finland.

Helsinki

Aapo Hyvärinen
Jarmo Hurri
Patrik Hoyer

Natural Image Statistics

A Probabilistic Approach to Early Computational Vision.

Hyvärinen, A.; Hurri, J.; Hoyer, P.O.

2009, XIX, 448 p., Hardcover

ISBN: 978-1-84882-490-4