

1

Introduction

Look in, let not either the proper quality, or the true worth of anything pass thee, before thou hast fully apprehended it.

—MARCUS AURELIUS *Meditations*, 170–180 AD
(Translated by Meric Casaubon, 1634)

This book presents selected object detection and recognition methods in computer vision, joining theory, implementation as well as applications. The majority of the selected methods were used in real automotive vision systems. However, two groups of methods were distinguished. The first group contains methods which are based on tensors, which in the last decade have opened new frontiers in image processing and pattern analysis. The second group of methods builds on mathematical statistics. In many cases, object detection and recognition methods draw from these two groups. As indicated in the title, equally important is the explanation of the main concepts of the methods and presentation of their mathematical derivations, as their implementations and usage in real applications. Although object detection and recognition are strictly connected, to some extent both domains can be seen as pattern classification and frequently detection precedes recognition, we make a distinction between the two. Object detection in our definition mostly concerns answering a question about whether a given type of object is present in images. Sometimes, their current appearance and position are also important. On the other hand, the goal of object recognition is to tell its particular type. For instance, we can detect a face, or after that identify a concrete person. Similarly, in the road sign recognition system for some signs, their detection unanimously reveals their category, such as “Yield.” However, for the majority of them, we first detect their characteristic shapes, then we identify their particular type, such as “40km/h speed limit,” and so forth.

Detection and recognition of objects in the observed scenes is a natural biological ability. People and animals perform this effortlessly in daily life to move without collisions, to find food, avoid threats, and so on. However, similar computer methods and algorithms for scene analysis are not so straightforward, despite their unprecedented development. Nevertheless, biological systems after close observations and analysis provide some hints for their machine realizations. A good example here are artificial neural networks which in their diversity resemble biological systems of neurons and which – in their software realization – are frequently used by computers to recognize objects. This is how the branch of computer science, called

computer vision (CV), developed. Its main objective is to make computers see as humans, or even better. Sometimes it becomes possible.

Due to technological breakthroughs, domains of object detection and recognition have changed so dynamically that preparation of even a multivolume publication on the majority of important subjects in this area seems impossible. Each month hundreds of new papers are published with new ideas, theorems, algorithms, etc. On the other hand, the fastest and most ample source of information is Internet. One can easily look up almost all subjects on a myriad of webpages, such as Wikipedia. So, nowadays the purpose of writing a book on computer vision has to be stated somewhat differently than even few years ago. The difference between an ample set of information versus knowledge and experience starts to become especially important when we face a new technological problem and our task is to solve it or design a system which will do this for us. In this case we need a way of thinking, which helps us to understand the state of nature, as well as a methodology which takes us closer to a potential solution. This book grew up in just this way, alongside my work on different projects related to object recognition in images. To be able to apply a given method we need first to understand it. At this stage not just a final formula summarizing a method, but also its detailed mathematical background, are of great use. On the other hand, bare formulas don't yet solve the problem. We need their implementations. This is the second stage, sometimes requiring more time and work than the former. One of the main goals of this book is to join the two domains on a selected set of useful methods of object detection and recognition. In this respect I hope this book will be of practical use, both for self study and also as a reference when working on a concrete problem. Nevertheless, we are not able to go through all stages of all the methods, but I hope the book will provide at least a solid start for further study and development in this fascinating and dynamically changing area.

As indicated in the title, one of my goals was to join theory and practice. My experience is that such composition leads to an in-depth understanding of the subject. This is further underpinned by case studies of mostly automotive applications of object detection and recognition. Thus, sections of this book can be grouped as follows:

- Presentations of methods, their main concepts, and mathematical background.
- Method implementations which contain C++ code listings (sections of this type are indicated with word IMPLEMENTATION).
- Analysis of special applications (their names start with CASE STUDY).

Apart from this we have some special entries which contain brief explanations of some mathematical concepts with examples which aim is to help in understanding the mathematical derivation in the surrounding sections.

A comment on code examples. I have always been convinced that in a book like this we should not spoil pages with an introduction to C, C++ or other basic principles of computer science, as sometimes is the case. The reasons are at least twofold: the first is that for computer science there are a lot of good books available, for which I provide the references. The second reason, is so to not divert a Reader from the main purpose of this book, which is an in-depth presentation of the modern computer vision methods and algorithms. On the other hand, Readers who are not familiar with C++ can skip detailed code explanations and focus on implementation in other platforms. However, there is no better way of learning the method than through practical testing and usage in applications.

This book is based on my experience gathered while working on many scientific projects. Results of these were published in a number of conference and journal articles. In this respect, two previous books are special. The first, *An Introduction to 3D Computer Vision Techniques and Applications*, written together with J. Paul Siebert, was published by Wiley in 2009 [1]. The second is my habilitation thesis [2], also issued in 2009 by the AGH University of Science and Technology Press in Kraków, Poland. Extended parts of the latter are contained in different sections of this book, permission for which was granted by the AGH University Press.

Most of all, I have always found being involved in scientific and industry projects real fun and an adventure leading to self-development. I wish the same to you.

1.1 A Sample of Computer Vision

In this section let us briefly take a look at some applications of computer vision in the systems of driver monitoring, as well as scene analysis. Both belong to the on-car Driver Assisting System aimed at facilitating driving, for example by notifying drivers of incoming road signs, and most of all by preventing car accidents, for example due to the driver falling asleep.

Figure 1.1 depicts a system of cameras mounted in a test car. The cameras can observe the driver and allow the system to monitor his or her state. Cameras can also observe the front of the car for pedestrian detection or road sign recognition, in which case they can send an image like the one presented in Figure 1.2.

What type of information can we draw from such an image? This depends on our goal, certainly. In the real traffic situation depicted we are mainly interested in driving the car safely, avoiding pedestrians and other vehicles in motion or parked, as well as spotting and reacting to



Figure 1.1 System of cameras mounted in a car. The cameras can observe a driver to monitor his/her state. Cameras can also observe the front of a car for pedestrian detection or road sign recognition. Such vision modules will probably soon become standard equipment, being a part of the on-board Driver Assisting System of a car.




Figure 1.2 A traffic scene. A car-mounted computer with cameras can provide information on the road scene to help safe driving. But computer vision can also help you identify where the picture was taken.

traffic signals and signs. However, in a situation where someone sent us this image we might be interested in finding out the name of that street, for instance. What can computer vision do for us? To some extent all of the above, and soon driving a car, at least in special conditions. Let us look at some stages of processing by computer vision methods, details of which are discussed in the next chapters.

Let us first observe that even a single color image has three dimensions, as shown in Figure 1.3(a). In the case of multiple images or a video stream, dimensions grow. Thus, we need tools to analyze such structures. As we will see, tensors offer new possibilities in



(a)



(b)

Figure 1.3 A color image can be seen as a 3D structure (a). Internal properties of such multidimensional signals can be analyzed with tensors. Local structures can be detected with the structural tensor (b). Here different colors encode orientations of areas with strong signal variations, such as edges. Areas with weak texture are in black. These features can be used to detect pedestrians, cars, road signs and other objects. (For a color version of this figure, please see the color plate section.)

this respect. Also, their recently developed decompositions allow insight into information contained in such multidimensional structures, as well as their compression or extraction of features for further classification. Much research into computer vision and pattern recognition is on feature detection and their properties. In this respect such transformations are investigated which change the original intensity or color pixels into some new representation which provides some knowledge about image contents or is more appropriate for finding specific objects. An example of an application of the structural tensor to image in Figure 1.2 for detection of areas with strong local structures is shown in Figure 1.3(b). Found structures are encoded with color – their orientation is represented by different colors, whereas strength is by color saturation. Let us observe that areas with no prominent structures show no response of this filter – in Figure 1.3(b) they are simply black. As will be shown, such representation proves very useful in finding specific figures in images, such as pedestrians, cars, or road signs, and so forth.

Let us now briefly show the possible steps that lead to detection of road signs in the image in Figure 1.2. In this method signs are first detected with fast segmentation by specific colors characteristic to different groups of expected signs. For instance, red color segmentation is used to spot all-red objects, among which could also be the red rims of the prohibitive signs, and so on for all colors of interest.

Figure 1.4 shows binary maps obtained of the image in Figure 1.2 after red and blue segmentations, respectively. There are many segmentation methods which are discussed in this book. In this case we used manually gathered color samples which were used to train the support vector classifiers.

From the maps in Figure 1.4 we need to find a way of selecting objects whose shape and size potentially correspond to the road signs we are looking for. This is done by specific methods which rely on detection of salient points, as well as on fuzzy logic rules which define the potential shape and size of the candidate objects.

Figure 1.5 shows the detected areas of the signs. These now need to be fed to the next classifier which will provide a final response, first if we are really observing a sign and not for instance a traffic light, and then what the type of particular sign it is. However, observed signs can be of any size and can also be rotated. Classifiers which can cope with such patterns are for instance the cooperating groups of neural networks or the decomposition of tensors of deformed prototypes. Both of the aforementioned classifiers respond with the correct type of signs visible in Figure 1.2. These, as well as many other methods of object detection and recognition, are discussed in this book.



Figure 1.4 Segmentation of image in Figure 1.2. Red (a), blue color segmentation (b).




Figure 1.5 Circular signs are found by outlining all red objects detected in the scene. Then only those which fulfill the definition and relative size expected for a sign are left (a). Triangular and rectangular shapes are found based on their corner points. The points are checked for all possible rectangles and again only those which comply with fuzzy rules defining sought figures are left (b). (For a color version of this figure, please see the color plate section.)

1.2 Overview of Book Contents

Organizing a book is not straightforward due to many the interrelations between the topics discussed. Such relations are not linear, and in this respect electronic texts with inner links show many benefits. The printed version has its own features. On the one hand, the book can be read linearly, from the beginning to its end. On the other, selected topics can be read independently, especially when looking for a specific method or its implementation. The book is organized into six chapters, starting with the Introduction.

Chapter 2 is entirely devoted to different aspects of tensor methods applied to numerous tasks of computer vision and pattern recognition. We start with basic explanations of what tensors are, as well as their different definitions. Then basic properties of tensors, and especially their distances, are discussed. The next section provides some information on filtering of tensor data. Then structural tensor is discussed, which proves very useful in many different tasks and different types of images. A further important topic is tensor of inertia, as well as statistical moments, which can be used at different stages of object detection and recognition. Eigendecomposition of tensors, as well as their invariants, are discussed next. A separate topic are multi-focal tensors which are used to represent relations among corresponding points in multiple views of the same scene.

The second part of Chapter 2 is devoted to multilinear methods. First the most important concepts are discussed, such as k -mode product, tensor flattening, as well as different ranks




Figure 1.6 Organization of the book.

of tensors. These are followed by the three main important tensor decompositions, namely Higher Order Singular Value Decomposition, best rank-1, as well as best rank- $(R_1, \dots, R_P)$ where R_1 to R_P represent desired ranks of each of the P dimension of the tensor. The chapter ends with a discussion of subspace data representation, as well as nonnegative decompositions of tensors.

Chapter 3 presents an overview of classification methods. We start with a presentation of subspace methods with one of the most important data representation methods – Principal Component Analysis. The majority of the methods have their roots in mathematical statistics, so the next chapters present a concise introduction to the statistical framework of object recognition. Not surprisingly the key concept here is the Bayes theorem. Then we discuss the parametric methods as well as the Kalman filter, frequently used in tracking systems but whose applications reach far beyond this. A discussion on the nonparametric follows, starting with simple, but surprisingly useful, histogram methods. Then the Parzen approach is discussed with its connections to nearest-neighbor methods. Mean shift methods are discussed in the consecutive parts of Chapter 3. Then the probabilistic, Hamming, as well as morphological neural networks are presented.

A separate topic within Chapter 3 concerns kernel processing. These are important novel classification methods which rely on smart data transformation into a higher dimensional space in which linear classification is possible. From this group come Support Vector Machines, one of the most important types of data classifier.

The last part of Chapter 3 is devoted to the family of k -means data clustering methods which find broad application in many areas of data processing. They are used in many of the discussed applications, for which special attention to ensembles of classifiers is deserved, such as the one discussed at the end of Chapter 3.

Chapter 4 deals with object detection and tracking. It starts with a discussion on the various methods of direct pixel classification, used mostly for fast image segmentation, as shown with the help of two applications. Methods of detection of basic shapes and figures follow. These are discussed mostly in the context of automotive applications. Chapter 4 ends with a brief overview of the recent methods of pedestrian detection.

Object recognition is discussed in Chapter 5. We start with recognition methods that are based on analysis of phase histograms of objects which come from the structural tensor. Discussion on scale-space template matching in the log-polar domain follows. This technique has found many applications in CV. From these, two are discussed. Two very important topics are discussed next. The first is the idea of object recognition in the domain of deformable prototypes. The second concerns ensembles of classifiers. As was shown, these show superior results even compared to very sophisticated but single classifiers.

Chapter 5 concludes with a presentation of the road sign classification systems based on ensembles of classifiers and deformable patterns, but realized in two different ways. The first employs Hamming neural networks. The second is based on decomposition of a tensor of deformable prototype patterns. The latter is also shown in the context of handwritten digit recognition.

A very specific topic discussed at the end of Chapter 5 is eye recognition, used for monitoring the driver's state to prevent dangerous situations arising from the driver falling asleep. Chapter 5 concludes with a discussion on the recent methods of object category recognition.

Appendix A discusses a number of auxiliary topics. It starts with a presentation of the morphological scale-space. Then a domain of morphological tensors operators is briefly discussed. Next, the geometry of quadratic forms is provided. Then the problem of testing classifiers is discussed. This section gathers different approaches to classifier testing, as well as containing a list of frequent parameters and measures used to assess classifiers. The rest of Appendix A briefly presents the OpenMP library used to convert serial codes into functionally corresponding but concurrent versions. In the last section some useful MATLAB[®] functions for matrix and tensor processing are presented.

As already mentioned, the majority of the presented topics are accompanied by their full C++ implementations. Their main parts are also discussed in the book. The full implementation in the form of a software library can be downloaded from the book webpage [3]. This webpage also contains some additional materials, such as the manual to the software platform, color images, and other useful links.

Last but not least, I will be very grateful to hear your opinion of the book.

References

- [1] Cyganek B., Siebert J.P.: An Introduction to 3D Computer Vision Techniques and Algorithms, Wiley, 2009.
- [2] Cyganek B.: Methods and Algorithms of Object Recognition in Digital Images. Habilitation Thesis. AGH University of Science and Technology Press, 2009.
- [3] <http://www.wiley.com/go/cyganekobject>