Digital Image Processing and Analysis: Human and Computer Vision Applications with CVIPtools, Second Edition

Scott E. Umbaugh, 977 pp., ISBN 978-1-4398-0205-2, CRC Press (2011), \$119.95 hardcover.

Reviewed by Jeffrey Snyder and Elena Fedorovskaya, Kodak Research Laboratories, Rochester, New York

Digital image processing and analysis is an immense field with applications in robotics, space exploration, medicine, entertainment, photography, defense, and other areas, and is growing larger every year. This leads to an increased demand for publications presenting the foundations of digital image processing to a broad audience in a systematic and practical way. Meeting such a demand in the second edition of his easy-to-read book, Scott Umbaugh provides an introduction to this field for students, researchers, and engi-

neers. The book covers a wide breadth of topics, starting with a high-level description of the author's view of the field, and ending with a detailed description of programming tools and a practical guide for the development of various applications. Although the book has nearly 1000 pages, it would be impossible to cover each of these topics in depth. Instead, the author focuses on providing the reader with a basic understanding of each topic area, illustrations of what can be done with the various algorithms presented, and a means to experiment with their new-found knowledge.

A unique feature of the book is the inclusion of the software CVIPtools, provided on an attached CD-ROM. CVIPtools has two major pieces. The first is the CVIPtools application, which has a graphical user interface and allows the user to read an image, select algorithms from a menu, change the parameters of algorithms, and apply the algorithms to the image. The second component is CVIPlab, which consists of the source code for the image-processing algorithms in CVIPtools and a simple console application. With the help of Microsoft Visual Studio or some other development environment, the user can modify the source code of CVIPlab and create new applications and algorithms. The CVIPtools package is used to demonstrate concepts and analysis algorithms introduced throughout the book, and to enable the reader to experiment with existing image-processing algorithms and develop new algorithms, then use their applications to explore practical problems.

The text, which is richly illustrated with both drawings and example images, and the CVIPtools package are highly integrated. This creates a learning environment that is particularly suited to students, computer scientists, and application developers. A significant portion (at least 30%) of the book is devoted to the detailed description of the CVIPtools package. The intent is for the reader to expand their learning experience by installing and

using the package both while studying the material and working the exercises at the end of each chapter.

A typical chapter in the book begins with an introduction and an overview of a selected topic, such as image restoration and reconstruction. To facilitate the overall understanding of a topic, the author then presents a system model, which includes a block diagram and description of the component blocks. Next, algorithms satisfying the requirements for each component are introduced. At this point, images are used to illustrate the behavior and tradeoffs for each algorithm. The author provides usage instructions for CVIPtools, enabling the reader to experiment and enhance their understanding. Following the main exposition, each chapter contains a key points section that summarizes the concepts defined in the chapter. A set of pen and paper problems are followed by programming exercises. A list of references and suggestions for further reading completes the chapter.

The book is organized into five sections. The first section, comprising chapters one and two, is titled "Introduction to Digital Image Processing and Analysis." Chapter one is a concise description of the digital image-processing field that the author divides into computer-vision applications and human-vision applications. Historically, digital image processing grew from electrical engineering as an extension of the signal-processing branch, and the computer science discipline was largely responsible for developments in computer vision. Umbaugh refers to digital image processing, or computer imaging, as the general field, with separate computer- and human-vision application areas. Computer-vision applications process an image for later use by a computer, while human-vision applications are motivated by a desire to improve image quality for a user or to reduce the storage burden on the user by involving a human being in the visual loop. Both types of applications employ image analysis, defined to include image segmentation, feature extraction, pattern classification, and transforms. Human-vision applications include image restoration, enhancement, and compression.

Chapter two contains an overview of imaging systems. It starts by outlining the basic model for visible light imaging, including the geometry of the lens, object plane, and imaging plane, then proceeds with the introduction of CVIPtools, and finishes with the discussion of image representation in the form of digital files of various types and formats. The exercises of the section focus on the basic principles of imaging systems.

Section 2, "Digital Image Analysis and Computer Vision," includes chapters three through six. Chapter three is an introduction to digital image analysis. Image analysis is presented as a sequence of steps that consists of preprocessing, data reduction, and feature analysis. Each of the steps is further broken down in several block diagrams. Convolutions, translation and rotation, arithmetic and logical operations, and spatial filters are introduced as ways to process a region of interest in the image. Quantization and thresholding are discussed as methods

of image data reduction and basic image analysis that can be performed on pixel values or their spatial coordinates. K-means clustering is presented as an algorithm to find a threshold in an image with good object-to-background contrast. Connectivity and labeling are mentioned in the context of more than one object labeling. The chapter ends with an example, using CVIPtools to show how a simple decision tree-based object classification tool can be created.

Chapter four presents segmentation and edge detection as techniques to take vast amounts of low-level pixel values and extract useful data that represent higher level information. The standard gradient edge-detection operators, those based on the first and second derivative of the luminance channel, are defined and demonstrated. The discussion of these operators is very clear and the results are amply illustrated with many high-quality images. The book truly excels in showing the reader the effects of various operations on the image. The author discusses more complex edge-detection algorithms, such as the Canny algorithm. He also includes a nice section on the performance evaluation of edge-detection techniques. Once you have found the edges in your image, what are you going to do with them? You could use the venerable Hough transform to detect lines or the Harris corner detection algorithm to find corners. Each of these is described and evaluated in the text. This leads to a discussion of image-segmentation methods, where the three different categories of segmentation algorithms are explored: region growing and shrinking; clustering; and boundary detection, the latter being an extension of the edge-detection techniques. The chapter ends with a discussion of morphological filtering as filtering of objects in the spatial domain, wherein binary, monochrome, and color images are considered.

Chapter five offers a foray into discrete transforms aiming to provide information about the image spatial, frequency, and sequency domains. The Fourier transform is defined and illustrated. The properties of the Fourier transform, which makes it so useful in linear space-invariant systems, are formulated but not derived. Sampling and aliasing are briefly mentioned. In addition to the Fourier transform, the discrete cosine, Walsh-Hadamard, and wavelet transforms are reviewed. The author emphasizes the multiresolution decomposition aspect of wavelet transforms, such as the Haar transform, that allow them to retain both spatial and frequency information. The principal components transform (PCT) is also described. PCT use in decorrelating the components of a complex signal, such as a multiband image, and its relationship to compression, are discussed. The usefulness of filtering for selectively modifying and analyzing various components of the spectral information is exemplified by considering the effects of lowpass, highpass, bandpass, and bandreject filters. The difference between ideal versus nonideal filters is explained.

Feature analysis and pattern classification, often considered as the final stages in the image-analysis process, are the topics of chapter six. The author discusses the steps involved in feature extraction, analysis, and pattern classification, introducing shape, histogram, color, spectral, and texture-related features and their description and similarity measures. The utility of CVIPtools is once again asserted by presenting examples of the processes involved in feature extraction and pattern classification. Still, the widely popular SIFT, GIST, and SURF features are not mentioned. While some common pattern-classification techniques, such as linear discriminant analysis and neural nets are discussed, descriptions of other common approaches, such as decision trees and support vector machines, are missing.

Chapters seven through ten are included in Section 3, "Digital Image Processing and Human Vision." In this section, the author describes elements of the human visual system (HVS)

and discusses three types of image-processing applications that are used to produce images for human viewing. Chapter seven provides a short overview of the basic functions of human visual perception, understanding of which is required for the development of the relevant image-processing applications. The anatomical and physiological structure and light sensitivity of the human eye are mentioned. The phenomena related to spatial frequency sensitivity, temporal resolution of the HVS, as well as brightness adaptation, simultaneous contrast, and color metamerism are briefly described. A relationship between these phenomena and corresponding digital image properties is pointed out. Additionally, a section on image quality gives examples of subjective as well as objective image quality measures, utilized when the effects of image enhancement, restoration, or compression need to be evaluated.

Image enhancement is the topic of chapter eight. The author groups enhancement techniques into the categories of point, mask (i.e., region), or global operations, depending on which pixels influence the pixel being modified. These same algorithms can also be classified as either spatial or frequency/sequence domain operations. The majority of the chapter is devoted to histogram-based techniques and various types of filtering, including sharpening and smoothing. Both linear and nonlinear approaches, including adaptive techniques, are discussed. CVIPtools is used to generate many illustrations for the various algorithms.

Image-restoration methods are greatly benefitted by having an accurate model of the process used to create the image. In chapter nine, the author presents a standard system model for this process and then explores what can be done to improve the image under various additional assumptions. In the author's model, the image is subject to a point-spread function degradation of some sort (e.g., motion blur), and then degraded by additional noise. The author discusses assumptions and various techniques to mitigate the noise. Restoration techniques involving both linear spatially invariant filters and adaptive filters are presented.

Image compression involves data reduction, mapping, quantization, and coding processes. In chapter 10, the author describes each of these. Lossy and lossless coding techniques are discussed. The reader is exposed to ideas from both the spatial and transform domains. Near the end of the chapter, the author provides a brief introduction to the JPEG and JPEG2000 standards.

Section 4, encompassing chapters 11 through 13, is devoted to programming and application development using CVIPtools. The source code and Visual Studio project for the console-based application CVIPlab are provided on the CD. Chapter 11 discusses how to compile the software and extend it so that an interested party could add their own features. The data structures used in CVIPtools are also defined. Chapter 12 describes algorithm development using CVIP-ATAT, a tool used to analyze the performance of a sequence of algorithms on a population of images, and CVIP-FEPC, a tool used to develop pattern classification algorithms. Several real-world examples are included to illustrate the usage of these tools. Chapter 13 is a programmer's guide to the C functions available in CVIPtools.

The final section of the book, Sec. 5, contains a set of appendices, which provide brief instructions for the installation and usage of the CVIPtools software.

Altogether, the author has produced a book that presents both the practical and theoretical aspects of digital image processing in a way that is efficient for learning, as well as research- and application-oriented development. He does an excellent job of showing how algorithms work and what the results look like. The drawback for the more sophisticated reader lies in the fact that he does not spend a lot of time on theory and some topics are

not covered in depth. For example, during the discussion of binary morphology in chapter four, the author does not present the set notation for binary morphology and does not spend much time talking about gray-level morphology. This is understandable, however, as he is attempting to cover “all areas of digital image processing and analysis, both human and computer vision applications” in one volume. Many of the chapter topics, such as transforms and pattern-classification analysis methods, have been expanded in books by other authors. While the author presents a very good introduction to the main areas of digital image processing, the reader will want to look elsewhere for more detailed theoretical discussions about specific topics.

The CVIPtools application and libraries are very useful. Compared with the first edition, two newly added software tools, the Computer Vision and Image Processing Algorithm Test and Analysis Tool (CVIP-ATAT) and the CVIP Feature Extraction and Pattern Classification Tool (CVIP-FEPC) offer a more complete set of tools for developing and testing algorithms. However, a few minor problems can be mentioned. First, the user interface of the CVIPtools application is somewhat unconventional and different than that in PhotoShop, IrfanView, and GIMP, three other well-known packages that have overlapping functionality. One distinction is that, unlike the other applications, CVIPtools uses a prefix approach to operate on the image. This requires the user to select an operation first and then select the data. A second problem involves a bug in the CVIPtools application. False edges sometimes appear when the user attempts to display an image with a higher resolution than the display. A final minor critique can be attributed to the design of the CVIPtools C function libraries. Although the author provides a very useful set of functions, they all allocate the result image inside the function. Some software developers may prefer to have more control over memory allocations by externally allocating memory to the function. One would not expect an image-processing application developed as a se-

ries of academic projects to be as polished as a commercially developed tool. Even with these minor problems, the CVIPtools package provides an outstanding value to the reader who is motivated to explore.

Overall, this is an excellent text for the student or professional looking for a practical introduction to or an overview of the wide world of digital image processing and analysis.

Jeffrey Snyder is a member of the computational science group at the Kodak Research Labs in Rochester, New York. He received his MS in computer science from the University of Rochester and has been with Kodak since 1979. His work has included algorithm development and software engineering for some of Kodak's successful imaging applications, such as Kodak Perfect Touch. His current interests are in computer vision, image processing, and computer-assisted storytelling. He is a member of the IEEE Computer Society and the ACM.

Elena Fedorovskaya has been a research scientist at Kodak Research Labs in Rochester, New York, since 1997. She received a PhD in psychophysiology and a MSc degree in applied mathematics, both from Lomonosov Moscow State University, Russia. Her work at Kodak focuses on human-centered computing involving modeling perceptual, cognitive, and emotional aspects of human experience in relation to images and imaging systems. She is a member of IS&T and SIGCHI.