

3D Shape

Its Unique Place in Visual Perception

Zygmunt Pizlo

The MIT Press
Cambridge, Massachusetts
London, England

© 2008 Massachusetts Institute of Technology

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher.

For information about special quantity discounts, please email special_sales@mitpress.mit.edu

This book was set in Stone Sans and Stone Serif by SNP Best-set Typesetter Ltd., Hong Kong.

Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Pizlo, Zygmunt.

3D shape : its unique place in visual perception / Zygmunt Pizlo.

p. cm.

Includes bibliographical references and index.

ISBN 978-0-262-16251-7 (hardcover : alk. paper)

1. Form perception. 2. Visual perception. I. Title.

BF293.P59 2008

152.14'23—dc22

2007039869

10 9 8 7 6 5 4 3 2 1

Preface

This book is the very first devoted exclusively to the perception of shape by human beings and machines. This claim will surely be surprising to many, perhaps most, readers, but it is true nonetheless. Why is this the first such book? I know of only one good reason. Namely, the fact that shape is a unique perceptual property was not appreciated, and until it was, it was not apparent that shape should be treated separately from all other perceptual properties, such as depth, motion, speed, and color. Shape is special because it is both complex and structured. These two characteristics are responsible for the fact that shapes are perceived veridically, that is, perceived as they really are “out there.” The failure to appreciate the unique status of shape in visual perception led to methodological errors when attempts were made to study shape, arguably the most important perceptual property of many objects. These errors resulted in a large conflicting literature that made it impossible to develop a coherent theoretical treatment of this unique perceptual property. Even a good working definition of shape was wanting. What got me interested in trying to understand this unique, but poorly defined, property of objects?

My interest began when I was working on an engineering application, a doctoral project in electrical engineering that involved formulating statistical methods for pattern recognition. Pattern recognition was known to be an important tool for detecting anomalies in the manufacture of integrated circuits. The task of an engineer on a production line is like the task of a medical doctor; both have to diagnose the presence and the nature of a problem based on the pattern of data provided by “signs.” I realized shortly after beginning to work on this problem that it was very difficult to write a pattern recognition algorithm “smart” enough to accomplish what an engineer did very easily just by looking at histograms and scatter

plots. It became obvious to me that before one could make computers discriminate one pattern from another, one might have to understand how humans manage to do this so well. This epiphany came over me on the night before I defended my first doctoral dissertation. My interest in studying human shape perception started during the early morning hours of that memorable day as I tried to anticipate issues likely to come up at my defense.

Studying pattern and shape perception requires more than a cursory knowledge of geometry, both Euclidean and projective. It also requires the ability to apply this knowledge to a perspective projection from a three-dimensional (3D) space to a two-dimensional (2D) image. I had a reasonable background in electrical engineering, but it did not include projective geometry. I had to learn it from scratch. It took both time and effort, but it paid off. At the time I did not realize that this was unusual. It never occurred to me that anyone would try to study shape, the topic that served for my second doctoral degree, without knowing geometry quite well.

My formal study of human shape perception was done in the Sensori-Neural and Perceptual Processes Program (SNAPP) of the Psychology Department at the University of Maryland at College Park where Robert M. Steinman served as my doctoral advisor. My dissertation also benefited a great deal from interactions with several members of the Center for Automation Research and Computer Science at this institution. My independent study of projective geometry was greatly facilitated by numerous discussions with Isaac Weiss. Realize that I was starting from scratch. I was analyzing known properties of geometrical optics simultaneously with learning about groups, transformations, and invariants. Here, my limited formal background in geometry led me to stumble onto some new aspects of projective geometry that had not been explored before. I was encouraged to pursue this path by Azriel Rosenfeld, my second doctoral mentor, who was affiliated with SNAPP. Azriel Rosenfeld, who was well-known for his many contributions to machine vision, was a mathematician by training. He was always interested in exploring the limits of mathematical knowledge and of mathematical formalisms, and he, Isaac Weiss, and I published some of our insights about a new type of perspective invariants that grew out of my dissertation. After mastering what I needed to understand in projective geometry, and after developing the new geometrical tools needed for a model of the perspective projection in the human eye,

I realized that I should also learn regularization theory with elements of the calculus of variations. Learning this part of mathematics was facilitated by interactions with Yannis Aloimonos, who was among the first to apply this formalism in computer vision. He asked me, now almost 20 years ago, whether regularization theory is the right formalism for understanding human vision. I answered then that I was not sure. My answer now is “Yes” for reasons made abundantly clear in this book. My interactions and learning experiences during my graduate education at the University of Maryland at College Park were not limited to geometry and regularization theory. From Azriel Rosenfeld I learned about pyramid models of figure-ground organization, and I learned about computational applications of Biederman’s and Pentland’s theories of shape from Sven Dickinson. Both figure prominently in my treatment of shape presented in this book. Now that the reader knows the circuitous route that led me to study human shape perception, I will explain why I decided to write this book.

The primary motivation for writing it grew out of my teaching obligations. When I began to teach, I tried to present the topic called “shape perception” as if it were a traditional topic within the specialty called “perception.” As such, shape perception, like other topics such as color perception, should be taught on the basis of the accumulation of specialized knowledge. Clearly, the history of a topic in a scientific specialty, such as shape perception, should be more than a collection of names, theories, and experimental results. The history of the topic should reveal progress in our understanding of the relevant phenomena. I found it impossible to demonstrate the accumulation of knowledge in the area called “shape perception.” The existing literature did not allow a coherent story, and I decided to try to figure out what was going on. Knowing this was important for doing productive research, as well as for teaching. How do you decide to take the *next step* toward understanding shape when where the last step left you was unclear? Recognizing that shape is a special perceptual property did the trick. It made both teaching and productive research possible. This book describes how much we currently understand about shape and how we came to reach the point that we have reached. It is a long story with many twists and turns. I found it an exciting adventure and hope that the reader experiences it this way, too.

By trying to maintain the focus of my presentation, I deliberately left out material that ordinarily would have been included if I were writing a

comprehensive review of visual perception, rather than a book on the specialized topic called “shape perception.” Specifically, I did not include a treatment of the neuroanatomy or neurophysiology of shape perception. Little is known about shape at this level of analysis because we are only now in a position to begin to ask appropriate questions. The emphasis of the book is on understanding perceptual mechanisms, rather than on brain localization. For example, the currently available knowledge of neurophysiology cannot inform us about which “cost function” is being minimized when a 3D shape percept is produced. I also did not include a large body of evidence on the perception of 2D patterns and 3D scenes that is only tangentially relevant to our understanding of the perception of 3D shapes.

The text concentrates on the discussion of the main concepts; technical material has been reduced to a minimum. This made it possible to tell the “story of shape” without interruption. A full understanding of the material contained in this book, however, requires understanding the underlying technical details. The appendices provide the basic mathematical and computational information that should be sufficient for the reader to achieve a technical understanding of the infrastructure that provided the basis for my treatment of shape. The references to sources contained in these appendices can also serve as a starting point for more in-depth readings in geometry and computational vision, readings that I hope will encourage individuals to undertake additional work on this unique perceptual property. Much remains to be done.

I had six goals when I began writing this book, namely, I set out to (i) critically review *all* prior research on shape; (ii) remove apparent contradictions among experimental results; (iii) compare several theories, computational and noncomputational, to each other, as well as to dozens of psychophysical results; (iv) present a new theory of shape; (v) show that this new theory is consistent with all prior and new results on shape perception; and (vi) set the stage for *meaningful* future research on shape. My choice of these particular goals and the degree to which I have been successful in reaching each of them can only be evaluated by reading the book. Obviously, my success with each goal is less important than my success in (i) encouraging the reader to think deeply about the nature and significance of shape perception and (ii) stimulating productive research on this fundamental perceptual problem.

The new theory presented in this book shows how a 3D shape percept is produced from a 2D retinal image, assuming only that the image has been *organized* into 2D shapes. One can argue that this new theory is able to solve the most difficult aspect of 3D shape perception. What remains to be done is to explain how the 2D shapes on the retina are *organized*. The process that accomplishes this, called “figure–ground organization” by the Gestalt psychologists, is not dealt with in great detail in this book, simply because not much is known about it at this writing. It is likely, however, that now that I have called attention to the importance of this critical organizing process in shape perception, it will be easier to (i) expand our understanding of how it works and (ii) formulate plausible computational models of the mechanisms that allow human beings to perceive the shapes of objects veridically.

I will conclude this preface by acknowledging individuals who contributed to this book and to the research that made it possible, beginning with the contributions of my students: Monika Salach-Golyska, Michael Scheesele, Moses Chan, Adam Stevenson, and Kirk Loubier worked with me on shape perception and figure–ground organization; Yunfeng Li designed and conducted recent psychophysical experiments on a number of aspects of shape and helped me formulate and test the current computational model; and he, along with Emil Stefanov and Jack Saalweachter, helped prepare the graphical material used in this book.

I also acknowledge the contributions of the late Julie Epelboim, who was a valuable colleague at the University of Maryland, where she served as a subject in my work on pyramid models and perspective invariants. My son, Filip Pizlo, contributed to a number of aspects of my shape research. He helped write programs for our psychophysical experiments and was instrumental in designing demos illustrating many of the key concepts. Interactions with my colleagues, Charles Bouman, Edward Delp, Sven Dickinson, Gregory Francis, Christoph Hoffmann, Walter Kropatsch, Longin Jan Latecki, Robert Nowack, Voicu Popescu, and Karthik Ramani contributed to my understanding of inverse problems, regularization theory, shape perception, geometrical modeling, and figure–ground organization. I also acknowledge the suggestion and encouragement to write a book like this that I received from George Sperling and Misha Pavel after a talk on the history of shape research that I gave at the 25th Annual Interdisciplinary Conference at Jackson Hole in 2000. None of these indi-

viduals are responsible for any imperfections, errors, or omissions present in this book.

I acknowledge support from the National Science Foundation, National Institutes of Health, the Air Force Office of Scientific Research, and the Department of Energy for my research and for writing this book. I thank Barbara Murphy, Kate Blakinger, Meagan Stacey, and Katherine Almeida at MIT Press for editorial assistance.

Finally, I thank my family for their understanding and support while my mind was bent out of shape by concentrating excessively on this unique perceptual property.