

IS THE HUMAN EYE LIKE A CAMERA?
WHAT MAKES YOUR EARS 'POP' ON A PLANE?
WHY DID WOMEN IN THE MIDDLE AGES
PUT BELLADONNA INTO THEIR EYES?

This fully updated 2nd edition of *Sensation and Perception* is an accessible introduction to the field of perception. Covering in detail the perceptual processes related to vision and hearing, taste and smell, touch and pain, as well as the vestibular and proprioceptive systems, this textbook is essential reading for any student of perception.

New material includes:

- 'Applications' feature that connects key content to real-life contexts
- 'Thinking Critically' feature that pushes students beyond the basics
- End-of-chapter essay questions
- An entirely new chapter on Perception and Action.

JOHN HARRIS is Emeritus Professor of Psychology at the University of Reading.
JARED G. SMITH is Senior Research Fellow at the Population Health Research Institute of St George's, University of London.

 **SAGE** www.sagepublishing.com
Los Angeles | London | New Delhi | Singapore | Washington DC | Melbourne

COVER DESIGN BY WENDY SCOTT | COVER IMAGE © GETTY IMAGES

ISBN 978-1-5264-6771-3

9 781526 467713



2E

SENSATION & PERCEPTION HARRIS & SMITH

2ND EDITION

SENSATION & PERCEPTION

JOHN HARRIS & JARED G. SMITH



SENSATION & PERCEPTION

Sara Miller McCune founded SAGE Publishing in 1965 to support the dissemination of usable knowledge and educate a global community. SAGE publishes more than 1000 journals and over 800 new books each year, spanning a wide range of subject areas. Our growing selection of library products includes archives, data, case studies and video. SAGE remains majority owned by our founder and after her lifetime will become owned by a charitable trust that secures the company's continued independence.

Los Angeles | London | New Delhi | Singapore | Washington DC | Melbourne

2ND EDITION

SENSATION & PERCEPTION

JOHN HARRIS & JARED G. SMITH



Los Angeles | London | New Delhi
Singapore | Washington DC | Melbourne



Los Angeles | London | New Delhi
Singapore | Washington DC | Melbourne

SAGE Publications Ltd
1 Oliver's Yard
55 City Road
London EC1Y 1SP

SAGE Publications Inc.
2455 Teller Road
Thousand Oaks, California 91320

SAGE Publications India Pvt Ltd
B 1/I 1 Mohan Cooperative Industrial Area
Mathura Road
New Delhi 110 044

SAGE Publications Asia-Pacific Pte Ltd
3 Church Street
#10-04 Samsung Hub
Singapore 049483

Editor: Donna Goddard
Editorial assistant: Esmé Carter
Production editor: Martin Fox
Copyeditor: Jane Fricker
Proofreader: Neil Dowden
Indexer: Silvia Benvenuto
Marketing manager: Fauzia Eastwood
Cover design: Wendy Scott
Typeset by: C&M Digitals (P) Ltd, Chennai, India
Printed in the UK

© John Harris and Jared G. Smith 2022

Apart from any fair dealing for the purposes of research, private study, or criticism or review, as permitted under the Copyright, Designs and Patents Act, 1988, this publication may not be reproduced, stored or transmitted in any form, or by any means, without the prior permission in writing of the publisher, or in the case of reprographic reproduction, in accordance with the terms of licences issued by the Copyright Licensing Agency. Enquiries concerning reproduction outside those terms should be sent to the publisher.

Library of Congress Control Number: 2021946476


British Library Cataloguing in Publication data

A catalogue record for this book is available from the British Library

ISBN 978-1-5264-6770-6
ISBN 978-1-5264-6771-3 (pbk)

At SAGE we take sustainability seriously. Most of our products are printed in the UK using responsibly sourced papers and boards. When we print overseas we ensure sustainable papers are used as measured by the PREPS grading system. We undertake an annual audit to monitor our sustainability.

For Jake, Noah, Finn and Leo, who might read it one day



CONTENTS

| | |
|---------------------|------|
| Online resources | xiii |
| About the authors | xv |
| Preface | xvii |
| How to use the book | xix |

PART 1 FOUNDATIONS AND INVESTIGATIVE TECHNIQUES 1

1 The Nature of Perception, and Some Ways of Investigating It 3

| | |
|--|----|
| Introduction | 4 |
| Approaches to the study of perception | 4 |
| Studying human perception | 11 |
| Role of past experience in perception | 18 |
| Synaesthesia | 20 |
| Bottom-up and top-down processes in perception | 22 |
| Stimuli and their presentation | 23 |
| Overview of the nature of perception and some methods of investigation | 28 |
| Summary | 28 |

2 Research Methods in Perception 31

| | |
|--|----|
| Introduction | 32 |
| Introspection and verbal report | 33 |
| Psychophysical methods | 34 |
| Brain imaging | 40 |
| Studies of the effects of brain lesions | 46 |
| Psychopharmacology | 47 |
| Single-unit physiology | 47 |
| Transcranial magnetic stimulation (TMS) | 50 |
| Modelling | 50 |
| Combination of techniques | 55 |
| Overview of research methods in perception | 56 |
| Summary | 56 |

| | |
|---|------------|
| PART 2 VISUAL AND AUDITORY PERCEPTION | 59 |
| 3 Mechanisms of Early and Middle Visual Processing | 61 |
| Introduction | 62 |
| The eye | 62 |
| Receptive fields | 70 |
| Retino-geniculo-striate pathway | 73 |
| V1 and V2 | 78 |
| Extra-striate visual streams | 79 |
| Functions of V4 | 81 |
| Grouping and segmentation | 83 |
| Functions of the dorsal stream | 87 |
| Two streams or a network? | 88 |
| Overview of studies of early and middle vision | 89 |
| Summary | 90 |
| 4 Seeing in Colour | 93 |
| Introduction | 94 |
| Colour and wavelength | 94 |
| Colour mixing and trichromacy | 97 |
| Photoreceptors | 100 |
| Opponent processes | 106 |
| Physiology of colour vision | 109 |
| Colour contrast, assimilation and spreading | 112 |
| Colour classification | 114 |
| Colour constancy | 115 |
| What is colour vision for? | 116 |
| Lightness perception | 118 |
| Overview of studies of colour vision | 121 |
| Summary | 121 |
| 5 Seeing Pattern and Motion | 123 |
| Introduction | 124 |
| Perceiving form | 124 |
| Perceiving motion | 137 |
| Overview of the perception of pattern and of motion | 149 |
| Summary | 150 |
| 6 Hearing | 153 |
| Introduction | 154 |
| The ear | 154 |
| Neural processes in hearing | 160 |
| Encoding frequency | 163 |
| Perceiving pitch | 165 |

| | |
|----------------------------|-----|
| Encoding loudness | 169 |
| Perceiving loudness | 170 |
| Auditory streaming | 171 |
| Perception of speech | 173 |
| Hearing and spatial vision | 180 |
| Overview of hearing | 180 |
| Summary | 181 |

PART 3 THE CHEMICAL SENSES AND SOMATOSENSORY PERCEPTION 183

7 Taste and Smell 185

| | |
|---|-----|
| Introduction | 186 |
| Taste | 186 |
| Smell | 197 |
| Interactions between taste and smell (and other modalities) | 212 |
| Overview of taste and smell | 213 |
| Summary | 213 |

8 Touch and Pain 215

| | |
|--|-----|
| Introduction | 216 |
| Stimuli and methods in tactile research | 216 |
| Tactile receptors in the skin | 217 |
| Spatial resolution of the skin | 219 |
| Temporal resolution of the skin | 220 |
| Tactile perception of textures | 221 |
| Neural basis of tactile perception | 221 |
| Tactile recognition by passive touch | 225 |
| Tactile recognition by active touch | 229 |
| Pain | 230 |
| Stimuli in experiments on pain | 230 |
| Pain receptors in the skin | 231 |
| Neural pathways for pain and temperature | 231 |
| Sex differences in pain perception | 233 |
| Pain processing in the brain | 234 |
| Central control of pain | 236 |
| Central sensitisation and chronic pain | 239 |
| Overview of research on touch and pain | 242 |
| Summary | 243 |

9 Vestibular and Proprioceptive Systems 247

| | |
|--------------------------------------|-----|
| Introduction | 248 |
| Vestibular system | 248 |
| Organisation of cilia and hair cells | 250 |

| | |
|--|-----|
| How to stimulate the vestibular system | 252 |
| Neural vestibular pathways | 253 |
| The vestibulo-ocular reflex (VOR) | 255 |
| Calibration of the vestibular system by vision | 256 |
| Vestibular acuity | 256 |
| Vestibular illusions | 258 |
| Motion sickness | 261 |
| Vestibular contributions to cognition and affect | 263 |
| Proprioception | 264 |
| Neural pathways for proprioception | 265 |
| Proprioceptive acuity | 266 |
| Control of limb movement | 268 |
| Proprioceptive contributions to body-ownership | 269 |
| Interactions between proprioception and vision | 269 |
| Interactions between the vestibular and other systems | 271 |
| Effects of space flight on the vestibular and proprioceptive systems | 273 |
| Overview of vestibular and proprioceptive functions | 274 |
| Summary | 274 |

PART 4 PERCEIVING THE WORLD AROUND US 277

10 Visual and Auditory Localisation 279

| | |
|--|-----|
| Introduction | 280 |
| How is space represented in the brain? | 280 |
| Information about depth and distance | 282 |
| Task-dependence of localisation | 293 |
| Auditory localisation | 297 |
| Combination of cues | 304 |
| Overview of visual and auditory localisation | 305 |
| Summary | 305 |

11 Perception and Action 309

| | |
|---|-----|
| Introduction | 310 |
| Reaching and grasping | 310 |
| Maintaining posture | 321 |
| Perceiving the direction of self-motion | 322 |
| Avoiding collisions | 324 |
| Interceptive timing | 325 |
| The relationship between perceptual and motor coding in the brain | 327 |
| Overview of perception and action | 330 |
| Summary | 330 |

| | |
|--|------------|
| 12 Eye Movements and Perception of Natural Scenes | 333 |
| Introduction | 334 |
| Categorising natural scenes | 334 |
| Statistical properties of natural scenes | 335 |
| Perceiving the gist of scenes | 341 |
| Scene analysis | 343 |
| Eye movements and vision | 345 |
| Trans-saccadic perception | 349 |
| Gaze control in natural scenes | 356 |
| Attention in natural scenes | 357 |
| Perceiving natural sounds | 358 |
| Statistics of natural sounds | 359 |
| Auditory salience | 360 |
| Overview of eye movements and the perception of natural scenes | 360 |
| Summary | 361 |
| | |
| 13 Recognising Faces | 363 |
| Introduction | 364 |
| Structural descriptions | 364 |
| Some theories of visual object recognition | 365 |
| The nature of face processing | 368 |
| Expertise and face recognition | 369 |
| Representation of faces in the brain: norm-based coding | 374 |
| Average faces, caricatures and anti-faces | 376 |
| Face aftereffects and face-specific mechanisms | 381 |
| Variability and ensemble coding | 385 |
| The neural basis of face identification | 387 |
| Controversies in studies of face recognition | 393 |
| Overview of studies of face recognition | 394 |
| Summary | 394 |
| | |
| 14 Attention and Awareness | 397 |
| Introduction | 398 |
| What is attention? | 398 |
| Perception and awareness | 410 |
| Perception of ambiguous figures | 410 |
| Blindsight | 414 |
| Binocular rivalry | 419 |
| What is awareness, or consciousness, for? | 423 |
| Overview of studies of attention and awareness | 423 |
| Summary | 424 |

| | | |
|---------------|--|------------|
| PART 5 | CHANGES IN PERCEPTION | 427 |
| 15 | Changes in Perception Through the Life-span | 429 |
| | Introduction | 430 |
| | Development of vision | 430 |
| | Development of hearing | 440 |
| | Changes in perception with ageing | 445 |
| | Overview of studies of changes of perception through the life-span | 455 |
| | Summary | 456 |
| 16 | Pathologies of Perception | 459 |
| | Introduction | 460 |
| | Visual perception in Autistic Spectrum Disorder (ASD) | 460 |
| | Agnosia | 467 |
| | Impairments involving the control of actions | 472 |
| | Hemispatial neglect | 474 |
| | Overview of studies of perceptual abnormality | 486 |
| | Summary | 487 |
| | References | 491 |
| | Index | 547 |

ONLINE RESOURCES



Visit <https://study.sagepub.com/SensationPerception2e> to find a range of additional resources for both students and lecturers, to aid study and support teaching.

FOR LECTURERS

A **Teaching Guide** outlining the key points covered in each chapter and providing you with further reading and suggested activities to use in class or for assignments. This guide includes links to the Key Notes relating to each chapter.

A **Testbank** containing questions related to the key concepts in each chapter which can be downloaded and used in class, as homework or exams.

PowerPoint Decks featuring figures and tables from the book, which can be downloaded and customised for use in your own presentations.

FOR STUDENTS

Key Notes relating to each chapter which will expand on material covered in the book.

Demonstrations show important perceptual phenomena, or provide links to websites where these may be found.



ABOUT THE AUTHORS

John Harris is Emeritus Professor of Psychology at the University of Reading. He has published more than 70 papers on perception and its abnormalities in psychiatric and neurological illness. For many years, he was associated with the international journal *Perception*, initially as an associate editor, and later as a member of the editorial board.

Jared G. Smith is a Senior Research Fellow in the Population Health Research Institution at St George's, University of London. He is an experienced methodologist with an extensive publication record of research evaluating psychological-based health and healthcare interventions in clinical populations, including perceptual and cognitive function in people with psychiatric and neurological illness.



PREFACE

As with the first edition of this book, authored by JH, this second edition is aimed at students who are studying perception beyond introductory classes. It has been extensively revised and updated. The chapter on the perception of emotion has been dropped and replaced by a new chapter on the control of action, an already vibrant area of research, whose importance is likely to grow in the future. Although the structure of the chapters in the first edition has been largely maintained, much research published over the past few years has been added.

We have adopted a similar approach to the earlier book in producing this second edition, with the following ideas in mind. First, as a student, you may sometimes have difficulty in relating the method of an experimental paper to its conclusions, and may tend to jump from the introductory question to the discussion with little regard to how the data were obtained. We have tried to address this by giving more detail of many studies than is found in most textbooks, since understanding the stimuli will help you to evaluate the claims made by authors towards the end of journal articles. Second, in British universities, most students in psychology are required to undertake a research project. However, the précis of experiments presented in lectures does not always seem to transfer well to the design of one's own undergraduate experiments. The experimental detail in the book is also intended to help with this. Third, the general approach in the book is that perception is an information processing problem – what information in the stimulus does the perceiver need to do a particular task, and how is this used? This also seems to be the (usually) implicit approach in most of the articles referred to in the text.

Experimental methods are emphasised throughout the book. Chapter 2 is devoted to ways of making behavioural measures of perception, together with a description of various physiological techniques. A recent development in perceptual research has been the increasing use of a variety of brain imaging techniques, and most chapters include descriptions of research employing functional imaging. Understanding the results of these studies, and of the effects of brain lesions, requires some knowledge of neuroanatomy, and often of neurophysiological work on animals, which is provided where relevant. Studies which include genetic analysis are also likely to become increasingly important, and some examples are given in the chapters on 'Seeing in Colour' (Chapter 4) and on 'Taste and Smell' (Chapter 7). The specialised techniques which have been developed for investigating the range of sensory modalities are also described.

Each chapter begins with a list of questions which will be discussed (and to which summary answers are highlighted in the text). Each section concludes with a list of the key points which have been made. At the end of each chapter are an overview and a summary, together with suggestions for further reading, either to expand on the material which has been covered, or to point you to related topics for which there was no space. A list of key terms is given, which

you should be able to define in your own words after reading the chapter, together with some possible essay questions, which you can answer in essay form. Various inserts will be found throughout the chapters. 'Thinking about Research' will help you to design your own experiments. 'Applications' show how laboratory studies of perception can help to alleviate problems in everyday life. 'Thinking Critically' aims to get you thinking a little more deeply about issues raised in the text.

Perception is a very active area of research. Inevitably, we have had to be selective in choosing topics and articles to describe, and in how to organise the material, and this means that many important and interesting studies have been omitted. We apologise to those who feel that their work should have been included. As in the first edition, the book is dominated by studies of vision, and to a lesser extent hearing, because those are the research areas in which most work is being done.

The list of people who have helped us is long. Michael Carmichael at Sage suggested the first edition, and so is the grandfather of the second. Dianne Berry, Andrew Glennerster, Eugene McSorley and Johannes Zanker, and some anonymous reviewers arranged by the publisher, gave advice on the first edition, and we have continued to make use of it here. Other anonymous reviewers commented on several revised chapters of this edition, and provided many helpful comments, much of which we have acted on. Two chapters were quite extensively changed from their first drafts, as a result of detailed reviews (which, given the increasing pressures on academics, were beyond the call of duty). Joel Winter and Roger Edelman informed us about the trigeminal nerve, and Kate Barnes about drugs which affect the pupil. Robert O'Shea and Olivia Diane Vatmanides have pointed out errors in the first edition (and we would like to know of any others in this one). Many people at Sage have advised and nudged us during writing: initially, Amy Jarrold, Robert Patterson, Marc Barnard, Katie Rabot, and more recently Donna Goddard and Esmé Carter, who have tolerated our delays and provided useful advice and support. We thank them all, though none is responsible for any misuse that we have made of their advice, or for any remaining shortcomings. Finally, but not least, JH thanks Dianne Berry for her continued support and encouragement (and her suggestion that the time before the morning coffee break is often the most productive), and JGS thanks Rebecca Abayakoon for her enduring support (and eternal patience!).

HOW TO USE THE BOOK

You will probably be using the book in connection with a course of lectures or teaching seminars. Sometimes a lecture will map on to one of the chapters, so that Chapter 13 would be relevant to a lecture or essay on face recognition. However, this might not always be the case. For example, if your essay assignment required an assessment of various stages in visual perception, you would find it worthwhile to read Chapters 3, 5 as well as 13. Chapter 16 would be a good start, if you were asked to assess how studies of perceptual abnormality throw light on the processes of normal perception, but other relevant material will be found in most chapters of the book. The chapter summaries (see below) will help you to locate material within a chapter.

Introductory questions at the start of each chapter signpost many of the topics that will be covered. At the end of each topic you will find a short passage in bold text (usually, one or two sentences) which gives a brief summary answer to one of the questions.

Key points at the end of sections give a brief summary of the information in that section in the form of bullet points. You may find it useful to read these before looking at each section, so that you read with the ‘take-home message’ in mind.

Key studies and **Case studies** are important studies which highlight ideas discussed in the text.

Thinking about research in each chapter suggests a possible study to address a question raised by some of the material, and asks you to design it (and sometimes to think about possible ethical issues). Of course, each of these is only one of many possible examples: when you have finished reading a section, ask yourself what questions are left unanswered in the current state of research.

Applications explain how some of the research mentioned in the chapter is (sometimes potentially) useful in solving real-world problems.

Thinking critically looks a little deeper into some of the issues raised by the material in the chapter.

Overviews at the end of each chapter sketch out the current state of knowledge about a topic or a perceptual system. They may suggest where there are gaps in our knowledge.

Chapter summaries present the main points made in the chapter in the form of a series of numbered statements. They usually indicate which aspects of a question are undecided, though you will need to read the text to understand the nature of the uncertainty.

Key terms: usually, of one or two words, which have been used in the chapter, and which you should be able to define in your own words.

Essay questions: short essay questions, to which you should be able to sketch an outline answer as a series of short notes.

Suggestions for further reading refer you to articles (or occasionally a book) which cover some of the topics in that chapter in greater depth. Sometimes there is a reference which addresses an aspect of the topic which is mentioned only briefly in the chapter. Many of the references are to review articles that summarise and discuss relevant literature.

PART I

FOUNDATIONS AND INVESTIGATIVE TECHNIQUES

Part Contents

| | | |
|---|---|----|
| 1 | The Nature of Perception, and Some Ways of Investigating It | 3 |
| 2 | Research Methods in Perception | 31 |



THE NATURE OF PERCEPTION, AND SOME WAYS OF INVESTIGATING IT

Chapter Summary

| | | | |
|--|----|--------------------------------------|----|
| Introduction | 4 | Stimuli and their presentation | 23 |
| Approaches to the study of perception | 4 | Experiments in perception | 23 |
| What is there to explain? | 4 | The nature of sound | 23 |
| Theoretical approaches | 6 | Presenting sounds in experiments | 25 |
| Studying human perception | 11 | Measuring luminance | 25 |
| Measuring thresholds | 11 | Presenting visual stimuli | 26 |
| Selective adaptation | 13 | Overview of the nature of perception | |
| Visual search | 16 | and some methods of | |
| Role of past experience in perception | 18 | investigation | 28 |
| Perceptual learning in early vision | 18 | Summary | 28 |
| Effects of vision on auditory learning | 19 | Key terms | 29 |
| Synaesthesia | 20 | Essay questions | 29 |
| Bottom-up and top-down processes in perception | 22 | Suggestions for further reading | 30 |

INTRODUCTION

Research into perception studies how information about the world is acquired through the senses. Since the only other way in which information about the world can be acquired is through our genes, perception is an important aspect of cognition. This chapter aims to set the scene for the remainder of the book in three main ways. First, it asks what an explanation for a perceptual process might look like, and sketches the ways in which some influential theorists have viewed the nature of perception, and a few of the phenomena which led them to their views. Second, it looks at some examples of the kinds of measurements that have been made, and touches on the conclusions that can be drawn from them. These examples serve to introduce some key ideas about studying perception:

- 1 The information available from measurements of thresholds.
- 2 The importance of selective adaptation as a technique for understanding perception.
- 3 The distinction between pre- and post-attentive processes, as revealed by visual search tasks.
- 4 The effects of experience on perception, and the calibration of one sense by another.
- 5 The importance of brain imaging in current studies of perception, using synaesthesia as an example.
- 6 The distinction between ‘top-down’ and ‘bottom-up’ processes.

Third, it emphasises the importance of understanding the stimuli and the apparatus to present them, using hearing and vision as examples. The chapter includes discussion of the following questions:

- Are percepts direct or constructed?
- What is a threshold?
- Why does staring at something cause distortions in vision?
- When do we need attention to see something?
- Can one learn to see better?
- When can vision improve our hearing of where things are?
- Are headphones better than loudspeakers?
- When is it useful to move farther away from a visual display?

APPROACHES TO THE STUDY OF PERCEPTION

What is there to explain?

The motivation to study mental processes is often triggered by their shortcomings. We have all forgotten names, or the planned visit to a shop on the way home. It is easy to demonstrate the limitations of memory in the laboratory, and to begin to suggest how these can be accounted for. New parents (and developmental psychologists, who often are, or have been, new parents) wonder what a baby’s mental life is like, and how it changes over the years.

For most of us, thinking is hard work. In contrast, adult perception almost always works well. We cross busy roads safely, walk down crowded streets without collisions, recognise our friends, read their writing, and understand their speech, efficiently and without apparent effort. The very efficiency of perception makes it less obvious that there is something to explain, and also makes it difficult to investigate.

Why should one bother? Firstly, as the perceptual psychologist Richard Gregory (1997) pointed out, a leap of imagination is needed to see that understanding vision presents us with problems. The input is the two tiny retinal images, subject to several distortions, and the output is the perceived world of objects. Gregory described this ‘as nothing sort of a miracle’. Studying this ‘miracle’ is a difficult but interesting and attractive intellectual exercise. Secondly, it can have important practical implications. For example, some individuals are born with a pronounced strabismus, or squint (of which one type is shown in Figure 1.1), which means that their eyes are not correctly aligned and they would not develop normal binocular vision. Ophthalmic surgeons can correct such a squint by tightening or repositioning the muscles that control the position of the eyes, but at what age should this best be done? The older the child, the easier it is to explain the procedure, and reassure them that their vision will be improved once healing is complete. On the other hand, if the operation is delayed too long, it may be too late for normal binocular vision to develop. A study is described in Chapter 15, which suggests that the optimal time may be before 4 months of age but that some benefit may occur up to the age of 5 years (and later work has suggested benefits from correction in adulthood). Thirdly, understanding how biological perception works may help in designing artificial systems to carry out the same job, for example, in guiding a vehicle, or picking out a suspected criminal in a crowd, activities which humans do rather well.

Unfortunately, understanding how we perceive has proved less easy than perceiving itself, as the rest of this book will show, though there has been much progress, especially in understanding the initial stages of the different sensory systems. One way to measure progress comes from computer simulations of perceptual processes, since one criterion for measuring how well we understand a biological system is our ability to mimic its activity with a machine. We can be confident that the heart is a pump because it can be replaced temporarily in an operating theatre with a mechanical pump, and the patient continues to live. Making machines that perceive as we do has proved less easy. Current machines can certainly differentiate between carefully designed patterns, such as the high contrast bar codes on goods in a shop, or the characters on a cheque. Humans are not very good at reading bar codes, but they can read hand-written scrawl which has been partially obliterated, and use the context to hear otherwise ambiguous sounds correctly. Abilities such as these, which involve the application of stored knowledge (often unconsciously), have proved difficult to mechanise.



Figure 1.1 The actor and comedian Marty Feldman had a pronounced uncorrected squint, in which one eye pointed outwards (so called ‘exotropic strabismus’). © David Levy, Wikimedia Commons.

Key Points

- Studying perception is not easy, but it can have important practical implications.
- The speed and efficiency of perception might suggest that there is nothing to explain. The failure of attempts to make machines perceive as adult humans do shows that this is not so.

What might an explanation for a perceptual process look like? A computer program is a list of operations to be carried out on data, written in some language that the computer can understand, such as BASIC or C++. However, the description of the program, which the systems analyst gives to the programmer, could be much more general and applicable to any computer language or machine. For example, the specification of a program to find the square root of a number could start as follows:

- 1 Input the number (x), whose square root is to be found.
- 2 Generate a number (y) which is less than x.
- 3 If y^2 is less than x, increase y.
- 4 If y^2 is greater than x, decrease y.

and so on....

Researchers into perception aim to give a description like this. Often, their description is informed by knowledge of the physiology and anatomy of perceptual systems; for example, the properties of retinal cones or auditory nerve fibres.

Theoretical approaches




Figure 1.2 Grouping by similarity. The dots appear to be grouped into horizontal rows, even though vertical rows are equally possible. © Wikimedia Commons.

The kind of explanation one might give, and even the kinds of experiment one might do, reflect a view of the nature of perception. Here, we consider briefly some influential accounts. The Gestalt psychologists (chiefly Wertheimer, Koffka and Koehler) emphasised the organising power of perception, the combination of isolated elements into new forms ('Gestalt' is a German word meaning 'form' or 'figure'). An example is that of apparent motion (see Chapter 5) in which two spots, spatially separated, and turned on and off in alternation, may appear to be a single spot in smooth oscillatory motion. To use the Gestalt description, 'the whole is other than the sum of the parts': one sees the whole (a spot in motion), not the parts (two flashing spots). Perceptual phenomena like this were explained by a series of laws. For example, in Figure 1.2, the grouping of the dots illustrates the Law of Similarity. Unfortunately, the Gestalt laws have little explanatory power: the Law of Similarity states that similar elements are grouped

perceptually, rather than explaining why, or how, and so is a redescription of the phenomenon it seeks to explain. Nevertheless, the perceptual effects which the Gestalt school (and others working in that tradition) discovered remain influential challenges to perceptual theories.

A second influential approach is the constructivist, of which the most influential recent exponent is perhaps Richard Gregory (1980). **In his view, percepts are constructed in the brain, as implied by the illusory triangle in Figure 1.3.** Here, perhaps the simplest explanation for the gaps in the brown elements is that these elements are in fact discs, masked by a white triangle, which is then created by the visual system. Gregory suggested that percepts are like hypotheses in science, based on, but going beyond, sensory data. The idea that percepts are hypotheses (or opinions) about the world fits well with the changes of perception of images like that in Figure 1.4. When two opinions are similarly supported by sensory data, the prevailing opinion alternates.

Note that, though Gregory may be correct about the nature of percepts, his account is incomplete, and does not explain when and why the perception of ambiguous figures should change, questions discussed in Chapter 14. The idea of perceptual construction goes back to von Helmholtz (2005), who suggested that percepts, which result from what he called ‘unconscious inference’, are symbols which can be manipulated mentally.




Figure 1.3 An illusory triangle, an effect devised by Kanisza (1976). The ‘whiter-than-white’ figure has the same luminance as the surrounding area.

Go Online

Visit <https://study.sagepub.com/SensationPerception2e> to find the Key Notes for this chapter:

- 1A Shows you more Gestalt demonstrations, which pit various ‘rules’ behind perceptual grouping against each other.
- 1B Describes Richard Gregory’s Hollow Face illusion, and provides a link to a website showing a video of the effect, along with some other demonstrations.

Gregory often contrasted his view with that of James Gibson (1950), the so-called ecological approach, from its emphasis on natural scenes, rather than impoverished laboratory stimuli. Gibson was especially interested in a different kind of task from that of perceiving objects, namely the perceptual guidance of movements.

Figure 1.5 shows a static view of a runway as seen by a pilot landing an aircraft. As the aircraft moves forward, the image of the world on the retina expands. As Gibson




Figure 1.4 The duck/rabbit figure. Is this a picture of a duck, facing to the left, or of a rabbit, facing to the right? Typically, one sees these possibilities in alternation during prolonged viewing of the figure.



Figure 1.5 Pilot's view of an airport runway. As the aircraft moves forward, to take off or land, an expanding pattern of optic flow, like that illustrated in Figure 11.14, is generated on the pilot's retina. Photo by Brent Cox on Unsplash.

suggested, and as we shall see in Chapter 11, this optic flow provides a powerful cue to the direction in which the pilot is heading. Where Gregory differed from Gibson was not in the nature of the visual information, but in how it is perceived. **In the constructivist approach, a representation of the world is built by following rules which incorporate the probability of occurrence of objects and situations, as in the perception of the illusory triangle in Figure 1.3. Instead, Gibson proposed that the scene layout is perceived 'directly', without the need for any intervening representation or computation.** Pointing out that most visual illusions occur with impoverished displays, he argued that the perceiver 'picks up' the detailed information available in natural scenes. This is done through some kind of resonance, in which the observer's nervous system resonates with aspects of the scene, much as the components of a radio resonate with aspects of the electromagnetic signal from the transmitter (without the radio building a representation of the information in the signal). An apparent difficulty for the notion of direct perception is that the optical information about surfaces does not specify them unambiguously (see Figure 1.6). Any retinal image could be produced by a large number of scenes. As we shall see in Chapter 2, in connection with Bayesian modelling, the perception of surfaces appears to depend to some extent on the application of prior knowledge. Critics of Gibson would say that the idea that their

perception is direct is to underestimate the complexity of the problem solved by the visual system. Supporters of Gibson would argue that, as an observer moves around the world, there is in fact a wealth of information about surfaces in their retinal images, which is why we rarely make errors in perceiving them. For example, if the trapezoids in Figure 1.6 were covered in a uniform texture, the upper texture elements on Trapezoid A would be larger in the retinal image than the lower elements (and vice versa for Trapezoid B), which would suggest that the surfaces were not vertical. Of course, the change of size of the elements might actually occur on a vertical object, but this seems less likely than that the surface is tilted.

Though the notion of direct perception has not been universally accepted, Gibson's emphasis on the information available in natural scenes and its relationship to actions has been influential. One of his ideas was that of 'affordance', the suggestion implicit in the structure of an object about how our motor systems should respond to it (see Figure 1.7). Thus the handle of a cup suggests 'grasp me', a steering wheel suggests 'turn me', and so on. Norman (2002) has used the idea in his discussion of the principles of good design: it should be obvious to the potential user from the appearance (or affordance) of say a door handle what one needs to do to open the door.




Figure 1.6 Upper: Side view of two flat trapezoids (A, B) of different sizes and orientations, and a rectangle (C), which could produce the same 2D retinal image. Lower: Front view of the trapezoids and rectangle.



A



B

Figure 1.7 Examples of Gibsonian affordances. The door handle (A) tells the perceiver to 'grasp and twist', the steering wheel (B) to 'rotate'. A © J.Dncsn CC-BY-SA-3.0, B © Øyvind Holmstad CC-BY-SA-3.0.

Earlier, we asked what an explanation for a perceptual process might look like, and suggested a series of operations, like the specification for a computer program. This suggestion was informed by the computational approach, of which perhaps the best-known exponent is David Marr (1982). He thought of the visual system as carrying out information processing: performing operations on an input (the retinal image) to produce an output (perceptual experience), by analogy with a computer program which performs a series of operations on an input to give an output. One of his enduring ideas is that one can describe a computer program, or by analogy the visual system, at three levels:

- 1 Computational level: what does the system do and why? (e.g. produce an internal representation of the world in a sufficiently detailed form to guide actions and enable object recognition).
- 2 Algorithmic level: how does the system operate on the input? (e.g. what aspects of the retinal image does it use, and what calculations does it perform on these?).
- 3 Implementation level: how do nerve cells carry out the algorithms? (e.g. how are edges in the retinal image identified?).

Our description of the program to find a square root would be at the algorithmic level. At the computational level, one would refer to a calculation in statistics, performed by a method using successive approximations to find a square root. The finished program, written in say C++ and taking into account the hardware of a particular computer set-up, would be the implementation. Note that the implementation in a digital computer, in which the elements can be in one of two states, must be very different from that in the brain, in which the firing rates of neurons can vary over a much larger range (though brains and computers would share the computational and algorithmic levels of description).

Computational theorists test their theories with computer programs that operate on images. A strength of the approach is the need to incorporate all aspects of the model in the program, or to produce a program which can learn to carry out a perceptual task, as we shall see in Chapter 2. To a human audience, one can skate over details in a verbal description (perhaps you can identify where this is done in the other approaches above), but this is not possible in a working computer program. The computational approach is very influential in modern research on perception.

Key Points

- An explanation for a process in perception would be a set of rules which transforms a sensory input into a percept.
- Perception involves the organisation of sensory input.
- Percepts are constructed by bringing past experience to bear on sensory data.
- Theories of direct perception overlook the ambiguity of retinal images.
- The computational approach distinguishes three levels of explanation: computational, algorithmic, implementation.
- A computational model needs to specify all aspects of the perceptual process which it seeks to explain.

THINKING CRITICALLY

In describing various approaches to perception, we pointed out apparent flaws in each. The constructivist approach, as exemplified by Gregory, lacks detailed accounts of how perceptual hypotheses are arrived at and tested during perception. Nevertheless, phenomena like the illusory triangle and the perceptual alternation of ambiguous figures suggest that the broad approach may be right, at least for one job carried out by perceptual systems, namely object recognition. Similarly, the Gestalt laws, though not themselves explanatory, point to a range of phenomena which need explanation, again largely relevant to object recognition. In contrast, Gibson's suggestion

of direct perception is concerned with the perceptual control of action, a job which need not involve object recognition at all. To avoid a lamp-post while walking down the street, one does not need to recognise that it is not a tree. One can criticise aspects of the theory, pointing out the ambiguity of retinal images, but it tries to address a different set of problems to the constructivist approach. As we shall see later in the book, object recognition and the control of actions seem to be mediated by separate anatomical pathways within the visual system. It seems that the ideas of Gibson and Gregory may be complementary, rather than antagonistic.

STUDYING HUMAN PERCEPTION

Perception, luckily for us as perceivers, is usually fast and accurate. It should not come as a surprise that much of the brain is devoted to it, so that, for example, perhaps half the cerebral cortex is devoted to aspects of vision. To understand what all this tissue does, a range of behavioural techniques have been developed to challenge the speed and accuracy of perception, and force the perceptual systems into error. As with memory, it is by exploring the limitations of perception that we can begin to understand how it works. Some common types of behavioural task, and how results from them have led to views of perception, are described below. How to make measurements in these tasks, and other techniques for investigating perception, are described in Chapter 2.

Measuring thresholds

Thresholds are of two types: absolute and difference (or discrimination) thresholds. One measures an absolute threshold by finding the weakest stimulus (the dimmest light, the faintest sound) which can just be detected. One example of a set of absolute threshold measurements is the dark adaptation curve (see Figure 3.6), repeated measurements of the faintest light which can be detected at various times after the observer has been immersed in total darkness. Another is the contrast sensitivity function (Figure 5.4), measurements of the lowest contrast needed to detect stripes of differing bar widths.

Snowden and Kavanagh (2006) were interested in whether the perception of moving objects changes as one gets older, and they addressed the question in several ways. In one experiment, they measured the lower absolute threshold for motion: how slow movement




Figure 1.8 Motion discrimination in two age groups, older (filled symbols) and younger (open symbols). Discrimination improves as speed increases, but is consistently worse in the older group.

Source: Modified from Snowden RJ, Kavanagh E (2006) Motion perception in the ageing visual system: minimum motion, motion coherence, and speed discrimination thresholds. *Perception* 35: 9–24. Pion Ltd, London. www.pion.co.uk.

could be made and yet still be detected. Their participants had to judge the direction (up or down) of slowly moving patterns of stripes. The older group (mean age 61.5 years) required higher velocities than the younger (23.2 years) for four different stripe widths (coarse to fine) to make the discrimination, and were also worse when the stimulus was a pattern of randomly positioned dots. The authors also measured motion coherence thresholds for random dot patterns drifting at one of four speeds. Motion coherence is defined as follows: the motion in a pattern of randomly positioned dots in which all dots move at the same velocity in the same direction is 100% coherent; a pattern in which all dots move in different directions at different velocities (a ‘snowstorm’) is 0% coherent. If the number of the dots in the snowstorm moving in the same direction at the same velocity is gradually increased, at some point an observer will correctly report the direction of that motion. This is the coherence threshold. Snowden and Kavanagh found that coherence thresholds

were higher (motion perception was worse) in the older group at the two lower speeds they tested, but not at two higher speeds.

Discrimination thresholds are measures of the smallest difference between two stimuli needed to reliably detect that they are different. As well as absolute thresholds, Snowden and Kavanagh also measured speed discrimination thresholds. In these experiments, two drifting gratings (patterns of stripes) were presented in quick succession, and the task was to report whether the first or second was moving faster. One of the gratings moved at a fixed speed, and the other was different from it by a variable amount. The minimum difference between two stimuli which allows discrimination is sometimes called the just noticeable difference (or JND). The older group needed a greater difference between the velocities, which was similar at all three fixed speeds, to make the discrimination (see Figure 1.8). These results have some practical importance, since they suggest that older people are worse at judging, say, vehicle speed than younger people, a difference which could matter when driving, or crossing a busy road. Other visual changes in ageing are discussed in Chapter 15.

Key Points

- Thresholds are of two types: absolute (or detection) and discrimination (or difference).
- Motion thresholds are generally higher, and so motion perception is worse, in older people.

Selective adaptation

Thresholds tell us about the sensitivity of a perceptual system (a less sensitive mechanism requires a more intense stimulus to achieve detection, or a larger difference between two stimuli to achieve discrimination). All perceptual systems change their sensitivity if continuously stimulated, and this change of sensitivity is often selective. In vision, for example, staring at a high contrast pattern of stripes for a while raises the absolute threshold (contrast has to be higher) for detecting faint versions of that pattern, but not that for much coarser or finer patterns. Similarly, listening to a pure tone for a time increases the loudness needed to detect a faint version of the tone, but not that needed to detect tones which are very different in pitch. The extended periods of gazing and listening are known as adaptation, and the increases in threshold as after-effects. As well as increasing thresholds, adaptation can change the appearance of high contrast patterns which are similar but not identical to the adapting pattern. For example, after one has adapted to a pattern of stripes tilted a few degrees away from vertical, stripes that are physically vertical appear to be tilted away from vertical in the opposite direction from those to which one has adapted. This effect is known as the tilt aftereffect (see Figure 1.9).

How is the tilt aftereffect to be explained? It is generally supposed that the orientation of a line or edge is encoded by activity in several visual mechanisms, each responding maximally to a different orientation, but also in a graded fashion to other nearby orientations (see Figure 1.10). Perceived orientation is given by the mechanism which gives the strongest response. **Before adaptation, perception is veridical: perceived and displayed orientation correspond. But, after adaptation, the most active mechanism is that signalling is not vertical, but 5 degrees anticlockwise from vertical, which is now the perceived orientation.**




Figure 1.9 Demonstration of the tilt aftereffect. First, look at the black square in the middle of the right-hand panel (B), noting that the patterns of stripes above and below the squares are vertical and parallel. Now, for about 30 seconds, let your eyes run from side to side along the horizontal black line in the centre of the left hand panel (A), then glance back at the centre of Panel B. The stripes above and below should appear tilted away from vertical in the opposite directions to those in Panel A.

Source: Shepherd AJ (2001) Increased visual after-effects following pattern adaptation in migraine: a lack of intracortical excitation? *Brain* 124: 2310–2318.

How might this explanation for the tilt aftereffect be incorporated into the computational approach to perception? At the computational level, one would point to the importance of edges and their orientation in the retinal image for identifying and recognising objects. Several mechanisms could be involved in the perception of an edge because individual mechanisms can be unreliable and subject to random variations in activity. At the algorithmic level, one would specify a comparison of several mechanisms, each responding to the edge, because this is likely to give a more accurate answer. In addition, the nature of the comparison would need to be given (perhaps, finding the more active of one pair of mechanisms and comparing that to a further mechanism, and so on). At the implementation level, one could postulate interactions between orientation-sensitive neurons of the type identified in the visual cortex by physiologists, and described in Chapter 2. At least part of the loss of sensitivity caused by adaptation may come from inhibition, that is, the action of the most active mechanism on less active mechanisms, which makes them even less active. This would make the difference in relative activity, suggested in Panel A, Figure 1.10, even greater than that caused directly by the stimulus. Such inhibition can explain illusions such as the simultaneous tilt illusion shown in Figure 5.7.

Other dimensions or aspects of vision which give aftereffects include colour (see Figure 4.7), size, movement and distance (from binocular disparity – see Chapter 10). In hearing, adapting to a tone of a particular frequency causes higher tones to appear even higher, and lower tones even lower. Such aftereffects are usually explained in terms of mechanisms responding to the physical characteristics of the stimuli, and occurring early in the chain of processes which leads to conscious perception. Other aftereffects have been discovered which appear to involve adaptation to more complex stimulus properties, such as the configuration which forms an individual's face (see Chapter 13), and probably occur at a later stage of processing.

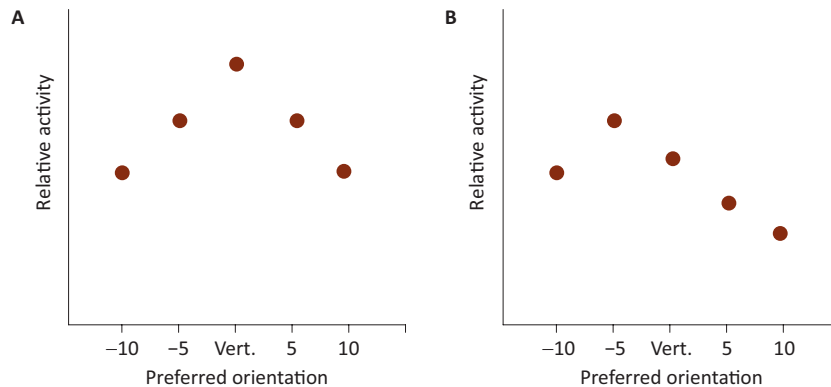


Figure 1.10 Explanation for the tilt aftereffect. Panel A (before adaptation): Vertical stripes excite several orientation-tuned mechanisms, most strongly that responding maximally to vertical, but also, less strongly, those responding maximally to orientations 5 degrees clockwise or anticlockwise from vertical, and even less those responding maximally to orientations 10 degrees from vertical. Panel B (after adaptation): Adaptation has reduced the activity in mechanisms which would normally respond maximally to vertical and orientations clockwise from vertical. Now the greatest response to vertical stripes comes from a mechanism whose preferred orientation is anticlockwise from vertical.

Figures 1.9 and 5.7 can demonstrate the tilt aftereffect and tilt illusion, respectively, but by themselves they are of little use if one wants to investigate the properties of the underlying mechanisms. For example, to determine whether the effects are specific to the width of the stripes, one would measure their strength in several conditions in which the adapting/inducing stripes were identical to, or differed by varied amounts from, the test stripes. One way to do this would be to vary the orientation of the test stripes, and measure how different from vertical they had to be in order to appear vertical. This method is called nulling. Another method is to have a set of stripes or a line of adjustable orientation near to the test stripes, and to measure the orientation at which this appears parallel to the test stripes, a method known as matching. Matching may be superior to nulling, since it does not involve direct interference with the effect that one is trying to measure. Ways in which to make the measurements are discussed in Chapter 2.

THINKING ABOUT RESEARCH

Here are two aftereffects of touch and temperature to try out yourself, using apparatus you probably have at home. Rub the finger-tips of the left hand over fine sandpaper, and those of the right hand over coarse sandpaper for a minute or so. Then rub medium sandpaper with both sets of finger-tips simultaneously. The medium sandpaper should feel coarser to the left hand than to the right. For the second experiment, fill three mugs with water: one with cold, one with warm, and one with lukewarm water. Put the index fingers of the left hand into the cold water, and of the right hand into the warm water. Wait a while, then put both index fingers into the lukewarm water. The water should feel warm to the left index finger and cold to the right (but lukewarm to both hands if the index fingers are removed and other fingers inserted).

Aftereffects decay over time. Consider how you might measure the decay of the temperature aftereffect. Presumably, at

some time after you have inserted both adapted fingers into the lukewarm water, its temperature will feel the same to both fingers. You could investigate whether this time increases with increases in adapting time (insertion into the warm and cold water mugs). Another question is whether decay of the aftereffect is hastened by insertion into the lukewarm water, or does it decay at the same rate regardless of what happens to the fingers. How could you answer this question? (Hint: you don't have to put your fingers into the lukewarm water immediately after adaptation.) Is the decay faster if you insert an unadapted finger as well as an adapted finger into the lukewarm water? If so, why might this be?

For an extended experiment, rather than a simple demonstration, you need to ensure that the water temperature in the mugs is roughly constant. It could be checked every five minutes or so with a thermometer and adjusted as necessary from the hot or cold tap.

Key Points

- Adaptation causes increases in detection thresholds for identical (and to a lesser extent for similar) patterns.
- Adaptation causes distortions in the perception of similar high contrast patterns (e.g. tilt aftereffect).
- Aftereffects suggest that aspects of stimuli are encoded by a range of mechanisms which vary in their sensitivity to the adapting stimulus.

- Although some aftereffects occur in mechanisms which detect low-level physical characteristics of stimuli (e.g. in vision, orientation, colour, direction of motion and size), others appear to involve the processing of more complex configurations, such as faces.

Visual search




Figure 1.11 A: Searching for a horizontal red bar in an array of vertical red bars is easy. B: Finding the vertical red bar in an array of red horizontal and green vertical and horizontal bars is harder.




Figure 1.12 Typical patterns of reaction times for serial and parallel visual search. Visual search tasks have been used to study aspects of the perception of natural scenes (Chapter 12). As we shall see in Chapter 12, it is likely that higher-level cognitive processes, as well as features in the image, guide visual search.

A common thread in accounts of perception is that it takes place in stages. For example, all sense organs convert electromagnetic, mechanical or chemical stimulation into electrical activity that can be processed by the nervous system. This conversion is known as transduction, and it is found in many artificial devices, such as microphones and light meters. Sometimes, identification of later stages can be based on anatomy and physiology. For example, it is likely that the neural basis of the tilt aftereffect lies in the visual cortex, because neurons in earlier stages of the main visual pathway are not orientation-selective. Even when anatomy and physiology are of little help, operations can be localised with respect to some perceptual process, although its physiological locus may be uncertain. An example of this comes from studies of visual search, in which an observer is required to scan arrays like those in Figure 1.11 to find a target. For the left-hand panel A, no effort is needed – the target ‘pops out’. If the task is to press a switch as soon as the target is located, reaction times are short, and do not increase as the number of other elements (distractors) in the array is increased.

It is as though the observer can process all regions of the display simultaneously (or ‘in parallel’). In contrast, finding the target in the right-hand panel B is much harder. Each item has to be attended to and scrutinised in turn (‘serially’). Thus reaction times to find the target rise with the number of distractors (see Figure 1.12). **Treisman and Gelade (1980) suggested that these differences arose because, early in visual processing, at a stage before attention was deployed, individual features, such as orientation, can be identified in parallel. However, when search for a conjunction of two such features is needed, attention is required to bind the features together.** Later work (Wolfe, 1994) has suggested that information from the pre-attentive

stage can be used to guide attention-driven serial search. Note that the distinction between pre- and post-attentive stages of perception is independent of where in the brain attention operates. In terms of the computational approach, this is an algorithmic, not an implementation, distinction.

CASE STUDY

Camouflage: hindering detection and recognition during visual search

Evolution generates a sort of arms race between predators and prey. Predators such as hawks have visual acuity superior to that of humans, and so can detect tiny creatures at long distances. Big cats such as cheetahs run at high speeds to overtake antelopes. One way in which prey can fight back is to develop camouflage, which makes them more difficult to pick out from their surroundings. If the predator cannot see them, its sharp teeth and strong limbs are useless. Thus the stripes on the zebra make its detection in the long grass harder for the lioness, a type of camouflage known as background matching, of which an example is shown in Figure 1.13.

A commonly held view is that camouflage is of no use once the potential prey has moved. It has been suggested that the system which detects motion in random-dot kinematograms (see Figure 5.19 and associated text) evolved to detect a moving object against an otherwise identical background. Hall et al. (2013) investigated how well observers could detect and recognise artificial patterns camouflaged in different ways from their background.



Figure 1.13 A mantis amongst leaves. This is an example of background matching. © Wikimedia Commons

When object and background were stationary, camouflage was very effective in preventing detection, until the object moved, when it was quickly detected. However, camouflage hindered the identification of moving objects when the background contained similar objects (distractors). Thus camouflage may hinder the predator in identifying say the smallest or weakest member of the herd.

Key Points

- Visual search reveals both pre-attentive and attentive processes.
- Search for simple features can be in parallel (pre-attentive), whereas search for conjunctions of features requires attentive scrutiny and is serial.
- Parallel search times do not vary with number of distractors, whereas serial search times do.
- Camouflage impairs visual search, unless the target moves, and continues to impair identification, when distractors also move.

ROLE OF PAST EXPERIENCE IN PERCEPTION

Perceptual learning in early vision

Figure 1.14 Stimulus to measure Vernier acuity. The task is to decide whether the lower line is offset to the left or to the right of the upper line.

It used to be thought that, after an initial period early in life during which learning could take place, sensory systems in the human brain became set, and could no longer be affected by experience. An example of early learning is the oblique effect, discussed in Chapter 12, which demonstrates that humans are more sensitive to vertical and horizontal than to oblique contours, probably because the former are more frequent in natural environments. Learning could still take place, of course, but it was thought to occur in higher, more cognitive, processes. Recently, this view has begun to change, as a result of studies suggesting that training can affect early stages of vision. One set of studies has investigated learning in Vernier acuity, the ability to discriminate the direction of the offset of one line from another (see Figure 1.14), showing that performance continues to improve (thresholds fall) with practice. It might be that discrimination of Vernier stimuli is based on activity in orientation-specific neurons like those thought to

underlie the tilt aftereffect. Thus the stimulus in Figure 1.14 would maximally excite a neuron whose preferred orientation was a degree or so anticlockwise from vertical. If the lower line was on the left of the upper line, the stimulus would excite a neuron whose preferred orientation was clockwise from vertical. The improvement in discrimination with training might come from an increase in the inhibitory effects of such neurons on other neurons sensitive to similar orientations. Particularly interesting studies in this area are those which have looked for transfer of learning: how different the stimulus has to be from that on which the observer was trained before its discrimination is no longer affected by the learning.

If one eye is trained, thresholds when the other eye makes the discrimination are the same as if there has been no training (Fahle, 2009), suggesting that the effects occur early in vision, perhaps in the primary visual cortex (Area V1 – see Chapter 3). If the observer practises on a vertical stimulus, there is no transfer of training to a horizontal stimulus. Tilting the stimulus by only 3 degrees from vertical reduces the effects of training by about half (Fahle, 2009). Performance rapidly improves during the early stages of training, then slows so that additional improvements take longer. Ahissar and Hochstein (1997) trained observers to detect the odd target in arrays comparable to Figure 1.11A (with the orientation of the target made progressively more similar to that of the distractors). **They found that during the early stages of learning, when the task was easy, learning to detect the target transferred to other regions of the array. However, in the later stages of learning, when the task was more difficult, learning did not transfer to other regions of the display. The authors suggested that the early learning took place in more central cognitive processes, whereas the later learning took place in specific early visual mechanisms.**

APPLICATION

Treating amblyopia

The suggestion from the studies described in the text that the adult brain may be plastic has potentially important practical implications. Some individuals have a condition called amblyopia (or 'lazy eye'), one type of which is caused by a defocus of the optics so that the retinal image in that eye is extremely blurred. This is often treated by covering the good eye with a patch, which can lead to improvements in vision of the amblyopic eye. However, such improvements do not always occur and, when they do, are not always complete. The laboratory studies of perceptual learning suggest that it may be possible to improve amblyopic vision by suitable training even during adulthood. Such improvements were demonstrated in a study by Huttunen et al. (2018) on two groups of participants in their

mid to late thirties. One group was given a drug fluoxetine, which is thought to increase neural plasticity, the other a placebo. There were no significant differences in improvement between the two groups, but both groups showed significant improvements in visual acuity and contrast sensitivity as a result of training. Participants were asked to spend 30 minutes in training in each of a 70-day period (or 35 hours if the instructions were faithfully carried out). The training consisted of seven computer games, in which the participant tracked one or more moving objects, and had to report when the target changed in some way, using foveal or peripheral vision. This is a nice example of how laboratory studies of visual perception can lead to improvements in the quality of life of the wider community.

Key Points

- Performance on a Vernier acuity task improves with practice, initially rapidly, then more slowly.
- At least some of the learning is specific to the eye which is trained and the orientation of the stimulus.
- The initial fast phase probably reflects changes in more cognitive processes; the slower later phase changes in mechanism, an early stage of perceptual processing.
- The suggestion from laboratory studies of plasticity in adult sensory systems has led to improved treatments for amblyopia in adulthood.

Effects of vision on auditory learning

Although the different perceptual systems are often described and investigated separately, in many situations they provide correlated information about the world. For example, we both see and hear the location of someone walking towards us on a hard surface. Sometimes the locations given by sound and vision are not perfectly correlated, as when we watch a movie or TV: the sound appears to come from the actors, not from its physical origin (loudspeakers in other positions). This is the basis of ventriloquism: we hear the sound produced by the ventriloquist as coming from the dummy's

moving lips. Alais and Burr (2004) showed that the effect depends on the quality of the visual information, since location can normally be determined precisely from the position of images on the retina. Their participants were required to locate a combination of clicks and briefly flashed blurred blobs, which they were asked to treat as a single event, such as a ball hitting the screen. However, in some conditions the locations suggested by the visual and auditory information were different. When the blobs were small and so easy to localise, it was the visual information which governed perceived location (the ventriloquism effect). However, when the blobs were large and so harder to localise, perceived location was biased more towards that suggested by the auditory information. As we shall see in Chapter 10, locating objects by hearing alone can be very accurate. However, as implied by the ventriloquism effect, it turns out that this accuracy is influenced during development and maintained during adulthood by visual information. **For example, putting moulds into the ears of adult humans impairs their ability to localise sounds. Over the next few weeks, however, the ability to localise gradually improves, as the auditory system is recalibrated by the visual system (King, 2009).**

Studies of the interactions between the senses are likely to become of increasing importance in the future. They are touched on throughout the book; for example, in the effects of vision on the perception as well as the location of sounds (Chapter 6), and the role of the vestibular system in vision (Chapter 9).

Key Points

- Visual and auditory information about location are usually correlated.
- When they are not correlated, visual information usually dominates, unless it is degraded.
- When auditory localisation is impaired, by putting moulds in the ears, the ability is gradually recovered, via the use of visual information.

SYNAESTHESIA

Synaesthesia is a set of intriguing perceptual effects in which one stimulus attribute (the inducer) always leads to the conscious perception of another attribute (the concurrent). The inducer and concurrent can be in the same sensory modality, as when a printed number or letter (a grapheme) always evokes the colour red, or in different modalities, as when the sound of a name can evoke a particular taste. People without synaesthesia have questioned its reality, which has led to studies in which the consistency of the inducer–concurrent association has been tested. When synaesthetes are given a list of inducers and asked to describe the concurrents, their consistency on retest is 80% or more, whereas that of controls who were given a list of possible inducers and asked to free associate their concurrents is 50% or less (Mattingley et al., 2001). The incidence of different types of synaesthesia, measured by screening large groups, ranges from 2 to 200 per thousand. Different types of synaesthesia often occur in the same individual, suggesting a common developmental origin (Ward, 2013; Speed and Majid, 2018).

Although the reality of synaesthesia is widely accepted, the reasons for it are uncertain. There is evidence that it requires attention to the inducer. For example, as a group, grapheme–colour

synaesthetes are better than controls at discriminating shapes formed from arrays of letters which induce colours from a background of other letters (Ward et al., 2010). However, some participants did not report colour effects, and those who did experienced colour effects on only about one-third of the letters. The authors concluded that the results were more consistent with the allocation of attention to some letters than with the kind of pre-attentive pop-out described above.

Synaesthetes perform better on laboratory memory tests than non-synaesthetes (Rothen et al., 2012), and also show enhanced autobiographical memory (Chin and Ward, 2018). In childhood, synaesthetes perform better on a letter cancellation task, and also in recalling lists of letters spoken aloud by the experimenter (Simner and Bain, 2018). Synaesthesia can occur between sensory modalities, as well as between different dimensions within a modality, and again may be associated with superior cognitive performance. For example, Speed and Majid (2018) tested colour–odour synaesthetes (individuals in whom odours evoked colours). Their participants were superior to controls in discriminating both colours and odours.

It seems likely that synaesthesia is associated with brain structures that are different from those of non-synaesthetes, for example, connections between different areas which are not present or not active in other brains. Modern methods of brain imaging, described in Chapter 2, have supported this view, implicating two regions in the parietal lobe, that around the intra-parietal sulcus (IPS) and the pre-cuneus (see Figure 1.15). There is increased grey matter (more cell bodies) in the IPS in grapheme–colour synaesthetes, and more activity is found there when graphemes are presented. The region is thought to be involved in binding (i.e. linking grapheme and colour into a perceptual whole), and may also be involved in the attentional aspects of synaesthesia (see Chapter 14 for further discussion of attention). The pre-cuneus is involved in imagery in several sensory modalities, and also shows increased activity during synaesthesia (Ward, 2013). Studies of synaesthesia provide an example of the importance of brain imaging in understanding perception. Other examples will be found throughout the book.




Figure 1.15 A: Lateral view of the right hemisphere, showing the intra-parietal sulcus. B: Medial view, showing the pre-cuneus in red.

Key Points

- Synaesthesia is the evocation by one stimulus attribute of another attribute, sometimes in a different modality.
- It appears to require attention to the inducing stimulus, and to be associated with increased activity in regions of the parietal lobe.
- Synaesthetes are often superior to non-synaesthetes in discriminations in the sensory modalities which are linked in their brains.




Figure 1.16 Contextual effects in perception. The identity (and perhaps the appearance) of the central character(s) depends on whether one is reading the horizontal or vertical string.

BOTTOM-UP AND TOP-DOWN PROCESSES IN PERCEPTION

Throughout this chapter, we have contrasted two kinds of process, one driven by information from the sense organs, the other influenced by past experience. The former are often referred to as ‘bottom-up’ and the latter as ‘top-down’ processes. In terms of broad theoretical approaches, Gibson’s is bottom-up, in that the information in the retinal image is directly perceived. On the other hand, the Gestalt psychologists, and certainly Gregory, emphasise the importance of top-down processes. Most contemporary researchers accept that both types of process are involved in perception. It would be inefficient not to allow past experience to have some influence on seeing or hearing. At the same time, it would be dangerous to allow past experience to dominate current sensory input. Presumably, the hallucinations which occur in some psychiatric and neurological illnesses, and can produce maladaptive behaviour, are examples of the latter.

Features of the image in Figure 1.4 take on a different meaning, depending on how the figure is being perceived. For example, one region may appear as a beak or as a pair of ears, depending on whether a bird or animal is being perceived, a clear example of top-down processing. When one gazes at Figure 1.4, perception changes although the stimulus does not. Figure 1.16 illustrates an effect of a simple change of context. How a character is read depends whether it is embedded in a list of letters or numbers. However, this would only work for a literate and numerate participant familiar with the Roman alphabet and Arabic numerals. Someone familiar only with Arabic script and Roman numerals would find the demonstration underwhelming, depending as it does on top-down processes programmed during a particular set of past experiences.

Although, as we noted earlier, perception proceeds in stages, the information flow between the stages is not one-way. Contextual or top-down effects in perception are probably mediated at least to some extent by feedback from later to earlier stages. As mentioned in Chapter 3, the interconnections between different cortical areas devoted to vision often run in both directions. The lateral geniculate nucleus in the thalamus receives input from the retina, and sends output to the visual cortex. But more than half the inputs to this nucleus come not from the retina but from the cortex, suggesting that top-down control occurs even at early stages of vision. As we shall see in Chapter 6, top-down effects also occur in hearing.

Key Points

- Both bottom-up and top-down processes play roles in perception.
- Top-down processing is mediated by feedback connections from later to earlier stages of perception.

STIMULI AND THEIR PRESENTATION

Experiments in perception

Experiments in perception depend on presenting stimuli in an appropriate manner, which requires some knowledge of the sensory system under investigation, the range of stimuli to which it responds, and of the proposed apparatus. To see why this matters, consider an experiment in which we wish to compare absolute thresholds for a high- and low-pitched tone in hearing. We have an oscillator to generate the tones, an amplifier, and a loudspeaker. Our listener adjusts the amplitude of the two tones until she can just hear them. We find that the amplitude of the high-pitched tone is greater and conclude that our listener is less sensitive to it. The problem with this experiment is that we cannot be sure why the difference occurred. It might indeed be due to the listener's auditory system, but it could also occur because the amplifier/loudspeaker combination was less efficient at producing the higher-pitched tone, and so attenuated its amplitude. We would need to measure the amplitude of the sounds leaving the loudspeaker with some suitable instrument (in other words, to calibrate it) to be sure that the effect originated in the listener, not in the apparatus. Here we consider some techniques for stimulating the auditory and visual systems. Stimulation of the other sensory modalities is considered in later chapters.

Key Point

- It is important to ensure that the apparatus is appropriate for the experimental question.

The nature of sound

Sound waves are fluctuations in air pressure caused by small movements or vibrations of objects, such as a guitar string or someone's vocal cords. As the object moves to and fro, so the pressure in the wave increases and decreases. A microphone is a device which contains a diaphragm, a flexible sheet which vibrates in time with the sound wave, and whose vibrations are converted into an electrical signal by some form of transducer (and this is also not a bad initial description of the ear). The time-varying electrical signal can be displayed, for example, on an oscilloscope, to show the characteristics of the wave form.

Some sound sources, such as a tuning fork or many whistles, produce a sinusoidally varying sound wave (see Figure 1.17A). Changes in the amplitude of the wave are perceived as changes of loudness, and in its frequency as changes of pitch, though loudness and pitch, as we shall see in Chapter 6, are not perfectly predicted by amplitude and frequency. Low frequencies correspond to low-pitched sounds, such as that produced by a trombone, whereas high frequencies correspond to high-pitched sounds, as produced by a piccolo. Humans can detect the range of frequencies from about 20 Hz to about 20 kHz. As examples of the frequencies of familiar sounds, the lowest note on a piano is about 27 Hz, and the highest note over 4 kHz.

The intensity of sounds is usually expressed in decibels (dB), a scale which expresses by how many times a sound is louder than some reference sound. In hearing research, intensity is usually expressed as a sound pressure level (SPL – the extent of the change in air pressure produced by the sound). In the decibel sound pressure scale, the reference is 0.0002 dynes/cm², the smallest pressure which can be detected with normal hearing for a tone of 1 kHz. Because there are other scales, measurements of this type are designated as being in dB (SPL). The formula for calculating the loudness of a sound in dB (SPL) is:

$$20 \times \log_{10}(P_s / P_r)$$

where P_s is the pressure of the sound to be expressed in dB and P_r is that of the reference level. A dB scale is used because of the enormous range of sound pressures (1:10 million or so) to which the ear is sensitive. A logarithmic scale compresses the range and makes the larger numbers more manageable. By definition, a barely audible sound (that is, one at threshold) is 0 dB (SPL). Other typical values in dB (SPL) are: rustling leaves, 20; conversation, 40; jet aircraft at take-off, 140. Sounds of above about 120 dB (SPL) cause pain and impair hearing.

Most sounds are not simple sinusoids. Figure 1.17B depicts white noise, a sound in which all audible frequencies occur with random amplitudes, which can be thought of as the complementary sound to that of a pure tone (single frequency, constant amplitude). Between these two extremes are sounds composed of several frequencies, often of a fundamental frequency and its harmonics (frequencies which are multiples of the fundamental frequency). Figure 5.3, which represents a distribution of luminance across space, could equally well represent the fundamental frequency of a sound together with its third and fifth harmonics.




Figure 1.17 Graphical representations of two sound waves, depicted as changes of amplitude against time. A: Simple sinusoid or pure tone. B: White noise (all audible frequencies, each of a randomly chosen amplitude). A © Pluke/Wikimedia Commons, B © Cralize/Wikimedia Commons.

Presenting sounds in experiments

Sounds are presented via loudspeakers or headphones. The ideal output device for research purposes should have a flat frequency response (i.e. for a fixed input voltage, the output should have the same level for all frequencies), and should not introduce distortions (components of the output that were not in the input), or phase shifts between different frequencies (say, delay the output of high frequencies with respect to low frequencies). Moore (2012) discusses how well various types of headphone and loudspeaker meet this ideal specification (and most do not). **Headphones have the advantage that they are cheaper, and the sounds they produce are not influenced by room acoustics.** Moore mentions the Eymotic Research ER2 earphone, which fits into the ear and has a flat frequency response (varying by less than ± 2 dB) over the range 100–14,000 Hz. Even when thresholds for sounds are being measured, it may not matter that a headphone does not have a flat frequency response, provided that this is known, so that losses in the equipment, as well as in the listener, can be calculated.

Increasingly, auditory stimuli are presented by digital computer. Depending on their nature, they may be generated by the computer as the experiment proceeds, or recorded offline and stored on the computer's hard disc in digital form. To produce the highest audible frequency (say 20 kHz), each cycle of the wave needs to be sampled at least twice (to capture the peak and trough), so that the rate at which the sound is sampled needs to be greater than 40 kHz. Digital recorders and computer sound cards often have 16 bit resolution, so that 65,536 voltage levels for each sample can be stored. Processing of the signal (e.g. filtering to remove or reduce the amplitude of some frequencies) can also be done by computer.

Key Points

- Although some sounds, such as that produced by a tuning fork, are pure tones, most sounds are a mixture of frequencies.
- The range of audible frequencies runs from about 20 to about 20,000 Hz.
- Sound intensity is usually measured on a scale of sound pressure level (SPL) in decibels.
- Headphones have advantages over loudspeakers as devices to present sounds in experiments, such as a flatter frequency response.
- Sound may be filtered by computer and stored on a hard disc for presentation via a sound card.

Measuring luminance

Light is an example of electromagnetic radiation, and can be thought of as a wave, or as a stream of particles, or 'photons'. As with sound, it varies in intensity and wavelength. The specification and measurement of wavelength are important in studies of colour vision, as we shall see in Chapter 4. A measure of the intensity of light reflected from or emitted by a surface is its luminance. The SI unit of luminance is the candela per square metre (cd m^{-2}), though older publications express luminance in other ways, such as the foot lambert (one of which is equal to

3.43 cd m^{-2}). Luminance is measured with a photometer, such as the Minolta Chromameter, whose output is given in candelas per square metre.

From measurements on a set of images of natural scenes, the average luminance of the sky is around 4000 cd m^{-2} , and that of the ground about 300 cd m^{-2} (Frazor and Geisler, 2006). Depending on adjustments, the luminance of the brighter parts of the display on a computer monitor could be around 150 cd m^{-2} . The sun at noon on a clear summer's day has a luminance of more than 1,000,000 cd m^{-2} , an intensity which can cause retinal damage when looked at directly.

Presenting visual stimuli

As with hearing, most experiments on vision are now run by a computer. Whether a standard computer monitor is used to present the stimuli depends on the nature of the experiment. Two things govern performance: the 'graphics adaptor', the piece of hardware in the computer which generates the electronic signals that produce the image on the monitor; and the monitor

itself. Two important measures of the performance of the graphics adaptor are its spatial and temporal resolution. Spatial resolution is given in picture elements, or 'pixels' (e.g. 1024×768). A pixel is one of the tiny rectangles into which the display screen is divided, and whose colour and brightness can be changed individually, rather like the squares of a crossword puzzle. Conventionally, the horizontal number is given first (1024), then the vertical number (768). For a given display size, the larger these numbers, the smaller is each pixel. One type of experiment in which this is important involves the measurement of Vernier acuity (see Figure 1.14), and requires tiny spatial offsets of one line from another. The smallest offset possible is governed by the pixel size. Size in experiments in vision is usually expressed in terms of the angle subtended by the object at the eye (see Figure 1.18 for the associated geometry and how to calculate visual angle).

If a monitor is 34 cm wide, and this width is occupied by 1024 pixels, each centimetre of screen width is occupied by about 30 pixels, and so each pixel measures about 0.33 mm. At a viewing distance of 1 metre, each pixel subtends an angle of inverse tan ($0.33/1000$) or 0.01891 degrees, or 0.01891×3600 , or about 68 seconds of arc. Would such a set-up be appropriate? The answer is probably not: the best observers have a Vernier acuity of about 10 seconds of arc. Figure 1.18B suggests one answer.




Figure 1.18 Illustration of angle subtended at the eye by two stimuli, identical in physical size(s). In B, the viewing distance ($2d$) is twice that in A. The tangent of the visual angle is the object size divided by the viewing distance. Thus the visual angle subtended by the stimulus in A (α) is twice that in B (β).

Each time the viewing distance is doubled, the visual angle is halved. Increasing viewing distance to 8 m would decrease pixel size to less than 10 seconds, which would probably be adequate.

A second important property of a graphics adaptor is the frequency at which it draws a new image on the screen. This is known as the frame rate, or refresh rate. It varies between computers, but is often 60, 72 or 75 Hz (or 'refreshes per second'). If the frame rate is 60 Hz, the duration of each frame is 16.67 (1000/60) milliseconds. This means that the duration of a stimulus must be a multiple of 16.67 ms. It is irrelevant whether, for example, the computer program running the experiment asks for a stimulus duration of 10 or 15 ms – in both cases, stimulus duration will be the same, namely, a single frame. When desired exposure durations are relatively long (say 100 ms or more), one can ignore the imprecision imposed by the frame duration. Another limitation set by the frame rate is the maximum frequency at which the display can be made to flicker. If luminance is set to be light and dark on alternate frames, the maximum flicker rate is 30 Hz when the frame rate is 60 Hz. This is not adequate for measuring human sensitivity to flicker, which can exceed 60 Hz in ideal circumstances. One answer to this difficulty is to arrange for the computer to program a signal generator which can produce the required flicker frequencies, and drive, say, an array of light-emitting diodes, whose luminance can be varied at very high rates.

A third limitation of a graphics adaptor is the number of levels of luminance at which each pixel can be set. This is often 256 (8-bit resolution), which may not be adequate for measuring contrast thresholds, when very small differences between neighbouring regions of the display may be required. Specialist graphics adaptors, intended for vision research, have better resolution, so that 12-bit resolution would give 4,096 luminance levels.

Computer monitors are of two types, namely CRTs (cathode ray tubes) and flat screen LCDs (liquid crystal displays). CRTs have existed for many years, and much is known about their properties. An important aspect of their performance is the relationship between the voltage supplied by the graphics adaptor and the luminance on the screen. For some applications, such as measuring thresholds, one would like the relationship to be linear, so that doubling the voltage also doubles the luminance. For most CRT monitors, this is not the case, and screen luminance rises as a power function of input voltage. The exponent in the function is referred to by the Greek letter γ (gamma). This varies between monitors, but is around 2.2, so that luminance typically rises with rather more than the square of the voltage. For any monitor, it can be measured with a photometer, and corrected during an experiment from sets of values stored in the computer (a look-up table). Of course, this is often not necessary. If an experiment requires high contrast stimuli, such as black letters on a white background, and contrast is never varied, it would not be necessary to gamma-correct the monitor.

Most CRT monitors come with in-built electronics which mean that their frame rate is fixed, matched to the computer's graphics adaptor, and cannot be altered by the user. One type of CRT is the display oscilloscope, whose refresh rate can be selected by the user, who needs to provide appropriate electronics. Such devices can be driven at high frame rates (e.g. 200 Hz) and often do not need gamma correction over most of their luminance range.

LCD displays are now common in domestic TV sets and computers. They are adequate for some types of vision research, in which fast changes of luminance are not critical (such as studies of Vernier acuity). But when they are asked to switch repeatedly between white and black, there is a short delay before the new luminance is reached, during which an intermediate

luminance is displayed. This makes them unsuitable for some applications in which it is important that average luminance remains constant; for example, studying visual evoked potentials (VEPs) (changes in voltage produced by brain activity, in response to sensory stimulation, and recorded on the surface of the scalp – see Chapter 2). When comparing various LCD displays, Ghodrati et al. (2015) reported that, because the luminance of LCDs varies with viewing angle, the brightness of peripherally viewed stimuli was reduced by up to 80% compared with centrally viewed stimuli of the same nominal luminance. This was the case even for two displays designed for vision research (VPixx and Display++), though these displays were less affected by problems when images were rapidly switched, provided that image duration was at least two frames.

Key Points

- The intensity of light as emitted by or reflected from a surface (luminance) is expressed in candelas per square metre.
- The performance of a computer-controlled monitor depends on the spatial resolution (number of pixels) of the graphics adaptor, its refresh rate, and the number of different luminance levels which it can produce.
- The size of the images of pixels at the eye can be halved by doubling the viewing distance.
- It may be necessary to gamma-correct a monitor (e.g. when measuring contrast thresholds).
- LCD displays may not reproduce fast changes of luminance accurately, and perceived luminance differs for different screen regions.

OVERVIEW OF THE NATURE OF PERCEPTION AND SOME METHODS OF INVESTIGATION

Despite earlier claims that perception is direct and depends solely on sensory information, there is now general agreement that it also involves stored knowledge. The computational approach to explaining perception, which includes stored knowledge in the form of constraints, has become very influential. Because perception is usually very fast and accurate, a range of special behavioural techniques are needed to investigate it. These include measurement of thresholds, selective adaptation and search. Perceptual systems can change their state through experience at various levels, and different systems interact with each other. Studies of search, adaptation, and learning have strengthened the view that perception proceeds in stages, often with influences of later stages on earlier ones. Presenting auditory and visual stimuli appropriately requires an understanding of apparatus.

SUMMARY

- 1 The difficulties of making a machine which perceives suggest that perception is a complex process, though this may be obscured by its speed and accuracy.
- 2 The Gestalt psychologists emphasised the organising power of perceptual systems, and the constructivist theorist Gregory, the creative aspect of perception.

- 3 The notion of direct perception is not consistent with the ambiguity of retinal images.
- 4 Marr's computational approach postulated three levels of explanation: computational, algorithmic and implementation.
- 5 Measurement of absolute and discrimination thresholds is an important technique, which has shown, for example, that the visual perception of motion is impaired in older people.
- 6 Selective adaptation produces aftereffects in all perceptual systems, raising detection thresholds and producing perceptual distortions.
- 7 Visual search has revealed pre-attentive processes in addition to processes which require attention.
- 8 Perceptual learning can take place in early perceptual mechanisms in adults as well as later more cognitive processes.
- 9 Visual and auditory information interacts in localising objects, with visual information dominant, unless it is degraded.
- 10 Synaesthesia appears to require attention and to be associated with abnormal activity in parts of the parietal lobe.
- 11 Bottom-up and top-down processes interact in perception.
- 12 Intensity of sound is usually measured in db (SPL) and of luminance in cd m^{-2} .
- 13 It is important to select appropriate apparatus, and to understand its limitations, in experiments in perception.

Key Terms

You should now be able to say in your own words what each of the following terms means:

| | | |
|----------------------|------------------|------------------|
| Law of Similarity | Serial search | Decibel |
| Affordance | Pop-out | Refresh rate |
| Absolute threshold | Vernier acuity | Gamma correction |
| Selective adaptation | Top-down process | |

Essay Questions

- 1 Describe three accounts of the nature of perception. Does each have weaknesses as well as strengths?
- 2 Can studies of perceptual learning influence rehabilitation programmes?
- 3 What problems may be encountered in using a computer display in vision research?

Go Online

Visit <https://study.sagepub.com/SensationPerception2e> to find the Key Notes for this chapter that expand on material covered in the book.

Suggestions for Further Reading

- Gregory RL (1997) *Eye and brain: the psychology of seeing* (5th edition). Oxford: Oxford University Press. An introduction to visual perception, from a constructivist perspective.
- Kanisza G (1976) Subjective contours. *Scientific American* 234(4): 48–52. A popular account of interesting phenomena in vision.
- Spillmann L (1997) Colour in a larger perspective: the rebirth of Gestalt psychology. *Perception* 26: 1341–1352. Considers Gestalt ideas in a more modern context.

RESEARCH METHODS IN PERCEPTION

Chapter Summary

| | | | |
|---|----|-----------------------------------|----|
| Introduction | 32 | Psychopharmacology | 47 |
| Introspection and verbal report | 33 | Single-unit physiology | 47 |
| Psychophysical methods | 34 | Transcranial magnetic stimulation | |
| Method of Adjustment | 34 | (TMS) | 50 |
| Signal detection | 34 | Modelling | 50 |
| Forced choice methods | 36 | Bayesian interpretation | 50 |
| Adaptive methods | 38 | Computational modelling | 53 |
| Magnitude estimation | 39 | Deep learning | 53 |
| Brain imaging | 40 | Combination of techniques | 55 |
| Electroencephalography (EEG) | 40 | Overview of research methods in | |
| Magnetoencephalography (MEG) | 41 | perception | 56 |
| Positron emission tomography (PET) | 42 | Summary | 56 |
| Functional magnetic resonance | | Key terms | 57 |
| imaging (fMRI) | 42 | Essay questions | 57 |
| Studies of the effects of brain lesions | 46 | Suggestions for further reading | 57 |

INTRODUCTION

A variety of techniques have been developed for investigating perception. Many of them involve the presentation of simple stimuli, such as pure tones or patches of colour. More complex images are less often used: they may be more interesting for the participants, but do not usually lead to firm conclusions about the detail of the processes underlying their perception. This has led some researchers to question how far research with simple stimuli bears on the perception of natural scenes, an issue touched on in Chapter 12, though complex scenes may be composed of simple features. The availability of cheap powerful computers has changed the nature of research into visual and auditory perception especially, since they provide new ways of generating and presenting stimuli, running experiments, collecting responses, and analysing data (especially the large amounts of data collected in brain imaging). The chapter starts by considering psychophysical methods (in effect, ways of asking people what they see, hear, etc.), then moves on to brain imaging. We briefly consider neurophysiology and the effects of brain lesions (both permanent and temporary), before touching on two approaches to modelling perception.

All methods have advantages and disadvantages. Different psychophysical methods have different trade-offs between speed and accuracy, some being quick but imprecise, others giving greater precision but also taking more time. This trade-off, and reasons for choosing a particular method, are briefly discussed. In principle, the results of a psychophysical experiment could be discussed without reference to the brain, mentioning only the computational and algorithmic aspects of the perceptual task. In practice, this is rare, and knowledge of the anatomy and physiology of perceptual systems informs the design of studies and interpretation of psychophysical results. Applying such knowledge involves a hypothesis which links perceptual performance to brain function. For example, it is known that parts of the brain which process visual information contain neurons which are orientation-selective: they respond most strongly to bars of a particular orientation (see Figure 2.8A). Such neurons adapt – their response reduces with continual stimulation. Perhaps the tilt aftereffect described in Chapter 1 results from a similar effect in similar neurons in the human visual system. If so, then one can ask, for example, whether the aftereffect increases with prolonged adaptation in a similar way to that in which the response of single neurons declines.

Increasingly, information about brain processes comes from functional imaging – records of brain activity during perception. Different methods of brain imaging have different strengths and weaknesses, which are touched on after a description of each method. These methods measure activity in populations of neurons, and their interpretation, in humans, is often informed by knowledge of the activity of single cells in corresponding brain areas in animals. Generally, single-cell recordings are made in humans for clinical reasons, in connection with proposed neurosurgery, so that any research value is incidental (and rare).

Brain damage can produce bizarre changes to perception, such as the inability to recognise familiar faces, discussed in Chapter 13, or to attend to part of external space, discussed in Chapter 16. Identification of the exact nature of such impairments is often difficult because, for example, the damage may affect several areas which subserve different functions. Again, knowledge of the physiology of the affected areas in animals can be useful.

Finally, we describe two approaches to explaining perception. One, the Bayesian approach, attempts to formalise and quantify the application of stored knowledge in perception. The other, deep learning in a neural network, is one type of computational modelling of perception, which is becoming increasingly influential.

The chapter includes discussion of the following questions:

- To what extent can we investigate perception by asking people to reflect on their perceptual experiences?
- Some people are more cautious than others. How can we measure the effects of this in perceptual tasks?
- Can we minimise the effort needed to make perceptual measurements?
- How can we measure the perception of something which is hard to specify, such as beauty?
- How can we best measure brain activity during perception?
- Can we interfere with brain activity during perception?
- How does prior knowledge affect perception?
- Why are computational models worthwhile?

INTROSPECTION AND VERBAL REPORT

A possible way to investigate perception is to reflect on one's perceptual experiences. Perhaps, like an art critic examining the brush strokes in a painting, one could work out how one's conscious experience of a scene was arrived at. This approach has not turned out to work well. It might help in identifying which aspects of a glorious sunset give the scene its aesthetic appeal, but not in understanding how the brain creates its internal representation of that sunset. One reason for this is that some perceptual processes are not available to consciousness. **However, in some circumstances, introspection and verbal report can give useful information. As we shall see in Chapter 14, some patients with damage to V1 (the first cortical visual area) have no or only limited awareness of visual stimuli, yet can make some visual discriminations at better than chance level.** Our only information about the detail of this impairment of awareness comes from their introspections and verbal reports. In larger populations, these can be elicited by questionnaires. For example, Kline et al. (1992) found that older drivers often report problems in detecting other vehicles coming in from the side when two streams of traffic are merging (as might be expected from the laboratory finding that older people have difficulty in detecting targets in peripheral vision in visual search tasks; Sekuler and Ball, 1986). The value of such verbal reports is limited. They can help to identify some differences between groups or individuals, but understanding why such differences arise requires different techniques. One of these is the presentation of well-specified stimuli in conditions which require a simple response from the participant, and to which we now turn. The development of methods in psychophysics has generated a large literature which cannot all be covered in detail in this chapter, in which some representative methods and their costs and benefits will be described.

Key Point

- Questionnaires can yield valuable information about differences in perception between groups or individuals, but introspection has revealed little of value about the underlying processes.

PSYCHOPHYSICAL METHODS

Method of Adjustment

Suppose that you wish to measure how high a contrast an observer needs to reliably detect a pattern (their contrast threshold). You have a device for displaying a grating (a repetitive pattern of light and dark stripes) and some electronics to generate the required signals. The electronics are controlled by a computer. One way to measure a threshold is to give the participant a means (say, two switches) with which they can instruct the computer to increase or decrease the contrast (the difference in luminance between the light and dark stripes) until it is just visible, and a way of signalling that they are satisfied with their setting (perhaps, a third switch). Such a method is fast and meaningful to the participant, but unfortunately the method gives ambiguous results, since a very cautious participant might want to be absolutely sure that the grating was visible, whereas a less cautious participant might be satisfied with just a hint. These two participants would produce different thresholds, even if their visual systems were identical, since the more cautious would set their threshold contrast to be higher than that of the less cautious.

Although the Method of Adjustment can be criticised, there are situations in which it can give acceptable results. Many studies have been published in which experienced observers make repeated measures on each of a range of stimuli. Presumably, their level of caution would be similar for all the stimuli. We noted in Chapter 1 that some individuals have an amblyopic or lazy eye, with which vision is much worse than with the other eye. A clinical researcher might wish to compare contrast thresholds from each eye during a visit by a patient to a hospital, when only a short time is available for testing. Since the same individual is setting the level of caution for each eye, the relationship between the thresholds is likely to be similar to that obtained by more rigorous methods, outlined below.

Signal detection

Signal detection theory (SDT) makes explicit and separable two components in the measurement of threshold: the participant's sensitivity to the stimulus, and their criterion, or level of caution, which are implicit or hidden in the Method of Adjustment. In a typical experiment, the participant would gaze at the display, and attend to a series of presentations. Each presentation would be signalled, perhaps by an audible tone at its start and end. On half the presentations, chosen at random, a grating would appear at a contrast drawn from a range which spans the expected threshold, and, on the other half, the screen would remain blank. The task would be to signal whether or not a grating had been presented. The four possible outcomes, and their names in SDT, are shown in Table 2.1. Note that if the proportions of 'hits' and 'misses' and of 'false alarms' and 'correct rejections' were added, both sums would be 1. The ideal observer, someone who always made use of all available sensory information and never made an error in responding, would have a hit rate of 1.0 and a false alarm rate of 0.0. Real observers do not achieve this, but the better their performance, the larger is the difference between scores for hits and false alarms. SDT provides a principled way of treating the difference, whose assumptions are illustrated in Figure 2.1. The theory assumes that an output is produced in a sensory system,