

Reading in the Brain

The New Science
of How We Read

STANISLAS DEHAENE



PENGUIN BOOKS

PENGUIN BOOKS

Published by the Penguin Group

Penguin Group (USA) Inc., 375 Hudson Street, New York, New York 10014,
U.S.A.

Penguin Group (Canada), 90 Eglinton Avenue East, Suite 700, Toronto, Ontario,
Canada M4P 2Y3 (a division of Pearson Penguin Canada Inc.)

Penguin Books Ltd, 80 Strand, London WC2R 0RL, England

Penguin Ireland, 25 St Stephen's Green, Dublin 2, Ireland (a division of Penguin
Books Ltd)

Penguin Group (Australia), 250 Camberwell Road, Camberwell, Victoria 3124,
Australia (a division of Pearson Australia Group Pty Ltd)

Penguin Books India Pvt Ltd, 11 Community Centre, Panchsheel Park, New Delhi
– 110 017, India

Penguin Group (NZ), 67 Apollo Drive, Rosedale, North Shore 0632, New Zealand
(a division of Pearson New Zealand Ltd)

Penguin Books (South Africa) (Pty) Ltd, 24 Sturdee Avenue, Rosebank,
Johannesburg 2196, South Africa

Penguin Books Ltd, Registered Offices: 80 Strand, London WC2R 0RL, England

First published in the United States of America by Viking Penguin, a member of
Penguin Group (USA) Inc. 2009

Published in Penguin Books 2010

Copyright © Stanislas Dehaene, 2009

All rights reserved

Illustration credits appear [here](#).

THE LIBRARY OF CONGRESS HAS CATALOGED THE HARDCOVER EDITION AS FOLLOWS:

Dehaene, Stanislas.

Reading in the brain : the science and evolution of a human invention / Stanislas
Dehaene.

p. cm.

Includes bibliographical references and index.

ISBN: 978-1-101-15240-9

1. Reading, Psychology of. 2. Reading—Physiological aspects. I. Title.

BF456.R2D36 2009

418'4019—dc22 2009009389

Except in the United States of America, this book is sold subject to the condition
that it shall not, by way of trade or otherwise, be lent, resold, hired out, or
otherwise circulated without the publisher's prior consent in any form of binding or
cover other than that in which it is published and without a similar condition
including this condition being imposed on the subsequent purchaser.

CONTENTS

[Cover](#)

[Praise for *Reading in the Brain*](#)

[Title Page](#)

[Copyright](#)

[Dedication](#)

[INTRODUCTION](#)

[**The New Science of Reading**](#)

[From Neurons to Education](#)

[Putting Neurons into Culture](#)

[The Mystery of the Reading Ape](#)

[Biological Unity and Cultural Diversity](#)

[A Reader's Guide](#)

[CHAPTER 1](#)

[**How Do We Read?**](#)

[The Eye: A Poor Scanner](#)

[The Search for Invariants](#)

[Amplifying Differences](#)

[Every Word Is a Tree](#)

[The Silent Voice](#)

[The Limits of Sound](#)

[The Hidden Logic of Our Spelling System](#)

[The Impossible Dream of Transparent Spelling](#)

[Two Routes for Reading](#)

[Mental Dictionaries](#)

[An Assembly of Daemons](#)

[Parallel Reading](#)

[Active Letter Decoding](#)

[Conspiracy and Competition in Reading](#)

[From Behavior to Brain Mechanisms](#)

[CHAPTER 2](#)

[The Brain's Letterbox](#)

[Joseph-Jules Déjerine's Discovery](#)

[Pure Alexia](#)

[A Lesion Revealed](#)

[Modern Lesion Analysis](#)

[Decoding the Reading Brain](#)

[Reading Is Universal](#)

[A Patchwork of Visual Preferences](#)

[How Fast Do We Read?](#)

[Electrodes in the Brain](#)

[Position Invariance](#)

[Subliminal Reading](#)

[How Culture Fashions the Brain](#)

[The Brains of Chinese Readers](#)

[Japanese and Its Two Scripts](#)

[Beyond the Letterbox](#)

[Sound and Meaning](#)

[From Spelling to Sound](#)

[Avenues to Meaning](#)

[A Cerebral Tidal Bore](#)

[Brain Limits on Cultural Diversity](#)

[Reading and Evolution](#)

[CHAPTER 3](#)

[The Reading Ape](#)

[Of Monkeys and Men](#)

[Neurons for Objects](#)
[Grandmother Cells](#)
[An Alphabet in the Monkey Brain](#)
[Proto-Letters](#)
[The Acquisition of Shape](#)
[The Learning Instinct](#)
[Neuronal Recycling](#)
[Birth of a Culture](#)
[Neurons for Reading](#)
[Bigram Neurons](#)
[A Neuronal Word Tree](#)
[How Many Neurons for Reading?](#)
[Simulating the Reader's Cortex](#)
[Cortical Biases That Shape Reading](#)

CHAPTER 4

Inventing Reading

[The Universal Features of Writing Systems](#)
[A Golden Section for Writing Systems](#)
[Artificial Signs and Natural Shapes](#)
[Prehistoric Precursors of Writing](#)
[From Counting to Writing](#)
[The Limits of Pictography](#)
[The Alphabet: A Great Leap Forward](#)
[Vowels: The Mothers of Reading](#)

CHAPTER 5

Learning to Read

[The Birth of a Future Reader](#)
[Three Steps for Reading](#)
[Becoming Aware of Phonemes](#)
[Graphemes and Phonemes: A Chicken and Egg Problem](#)
[The Orthographic Stage](#)

[The Brain of a Young Reader](#)

[The Illiterate Brain](#)

[What Does Reading Make Us Lose?](#)

[When Letters Have Colors](#)

[From Neuroscience to Education](#)

[Reading Wars](#)

[The Myth of Whole-Word Reading](#)

[The Inefficiency of the Whole-Language Approach](#)

[A Few Suggestions for Educators](#)

CHAPTER 6

The Dyslexic Brain

[What Is Dyslexia?](#)

[Phonological Trouble](#)

[The Biological Unity of Dyslexia](#)

[A Prime Suspect: The Left Temporal Lobe](#)

[Neuronal Migrations](#)

[The Dyslexic Mouse](#)

[The Genetics of Dyslexia](#)

[Overcoming Dyslexia](#)

CHAPTER 7

Reading and Symmetry

[When Animals Mix Left and Right](#)

[Evolution and Symmetry](#)

[Symmetry Perception and Brain Symmetry](#)

[Dr. Orton's Modern Followers](#)

[The Pros and Cons of a Symmetrical Brain](#)

[Single-Neuron Symmetry](#)

[Symmetrical Connections](#)

[Dormant Symmetry](#)

[Breaking the Mirror](#)

[Broken Symmetry . . . or Hidden Symmetry?](#)

Symmetry, Reading, and Neuronal Recycling
A Surprising Case of Mirror Dyslexia

CHAPTER 8

Toward a Culture of Neurons

Resolving the Reading Paradox

The Universality of Cultural Forms

Neuronal Recycling and Cerebral Modules

Toward a List of Cultural Invariants

Why Are We the Only Cultural Species?

Uniquely Human Plasticity?

Reading Other Minds

A Global Neuronal Workspace

CONCLUSION

The Future of Reading

ACKNOWLEDGMENTS

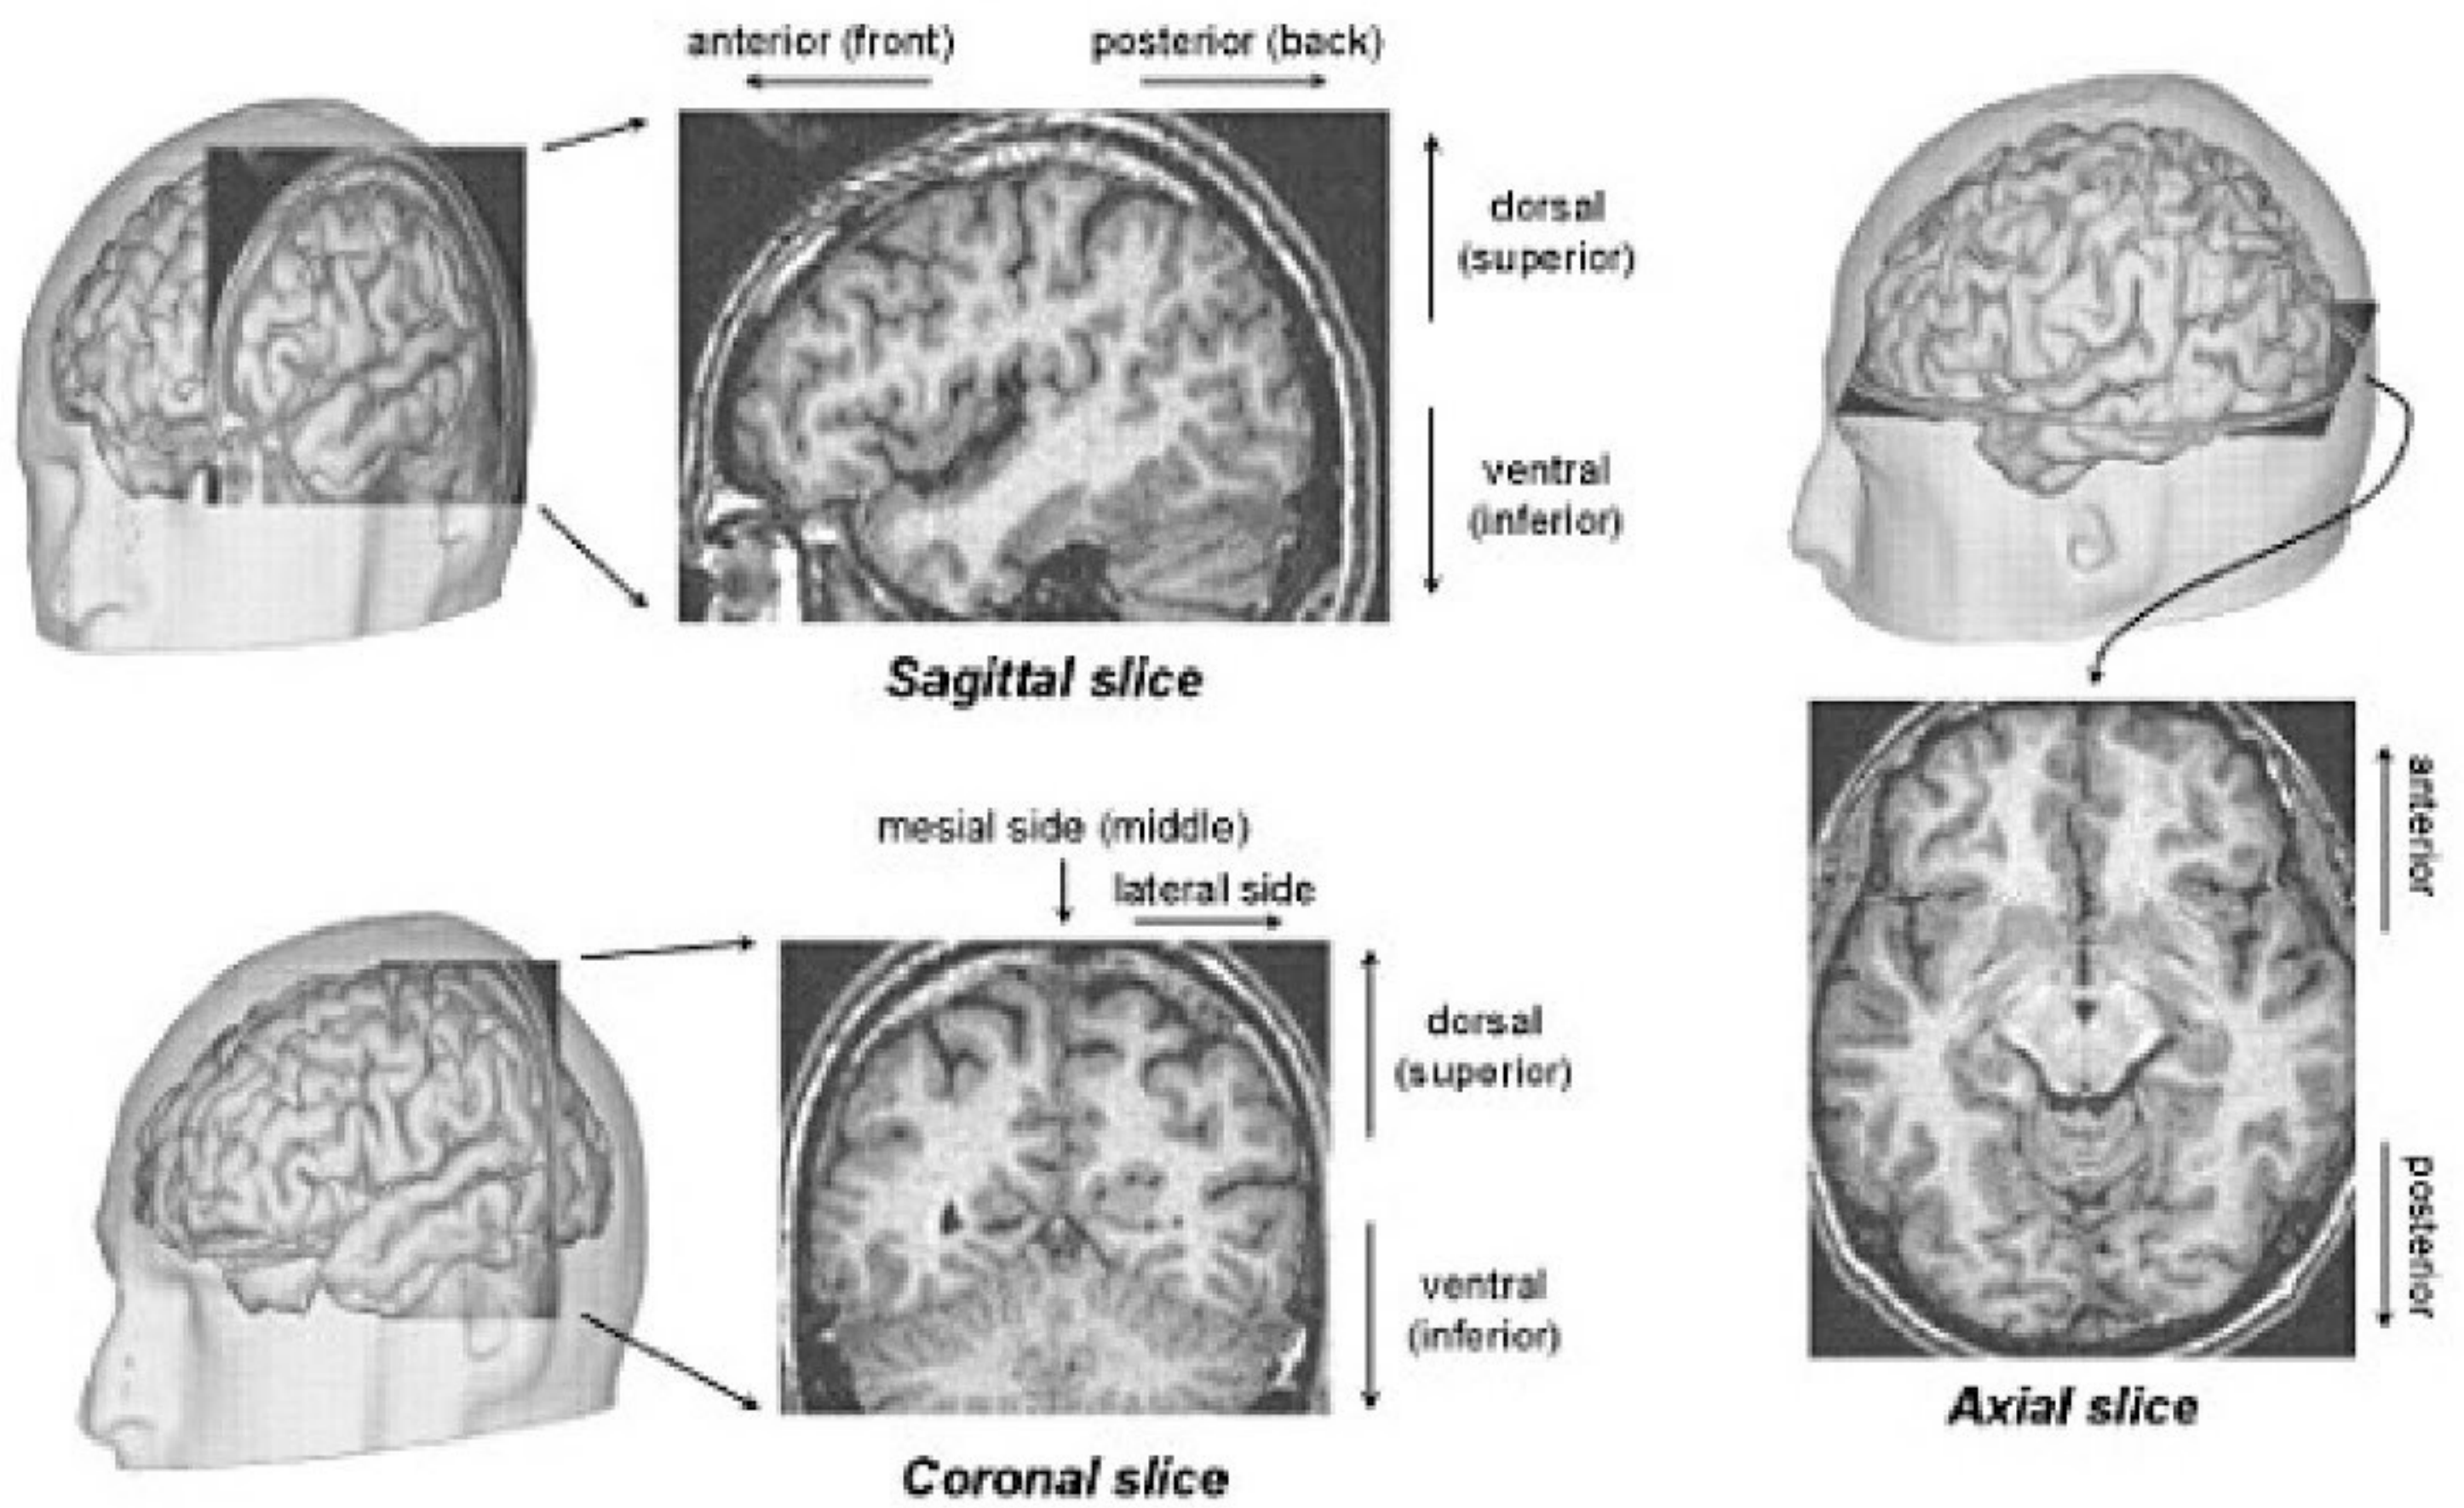
NOTES

BIBLIOGRAPHY

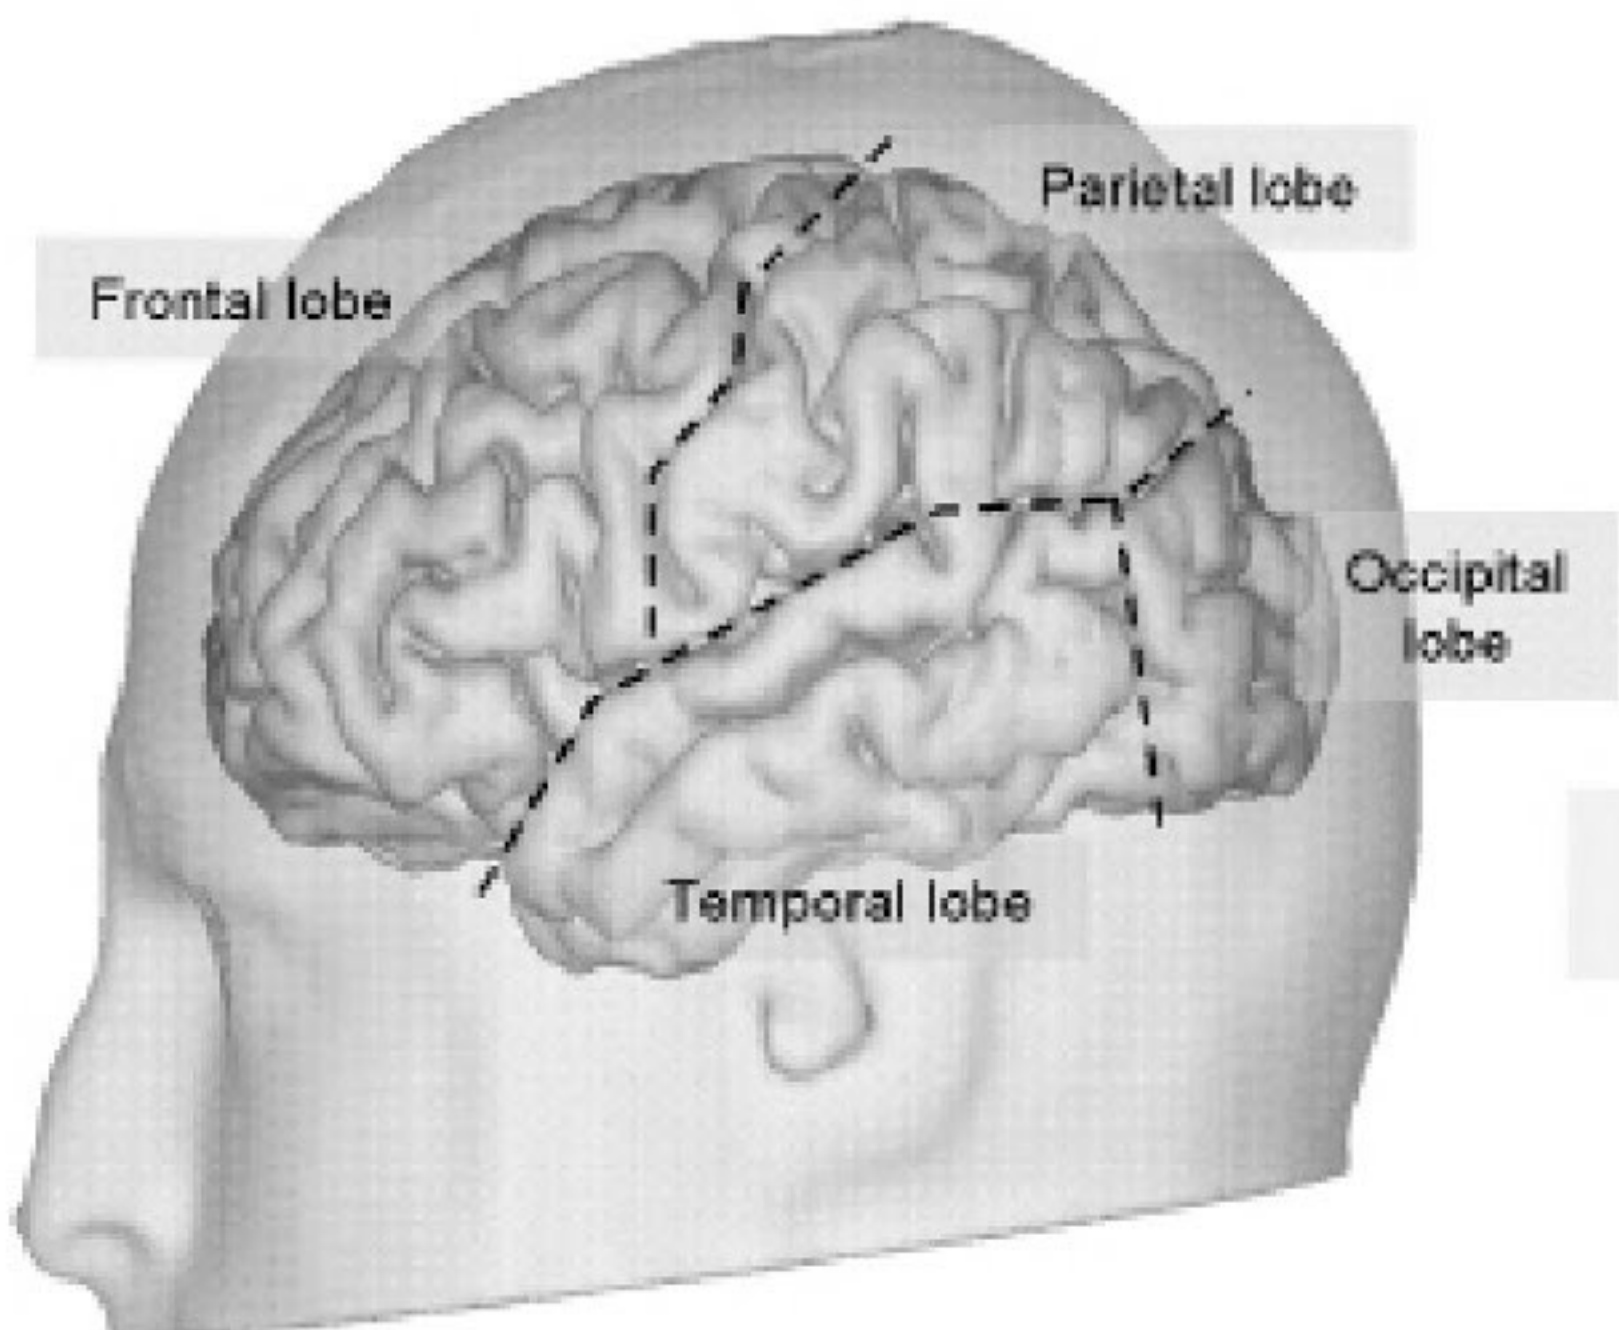
INDEX

FIGURE CREDITS

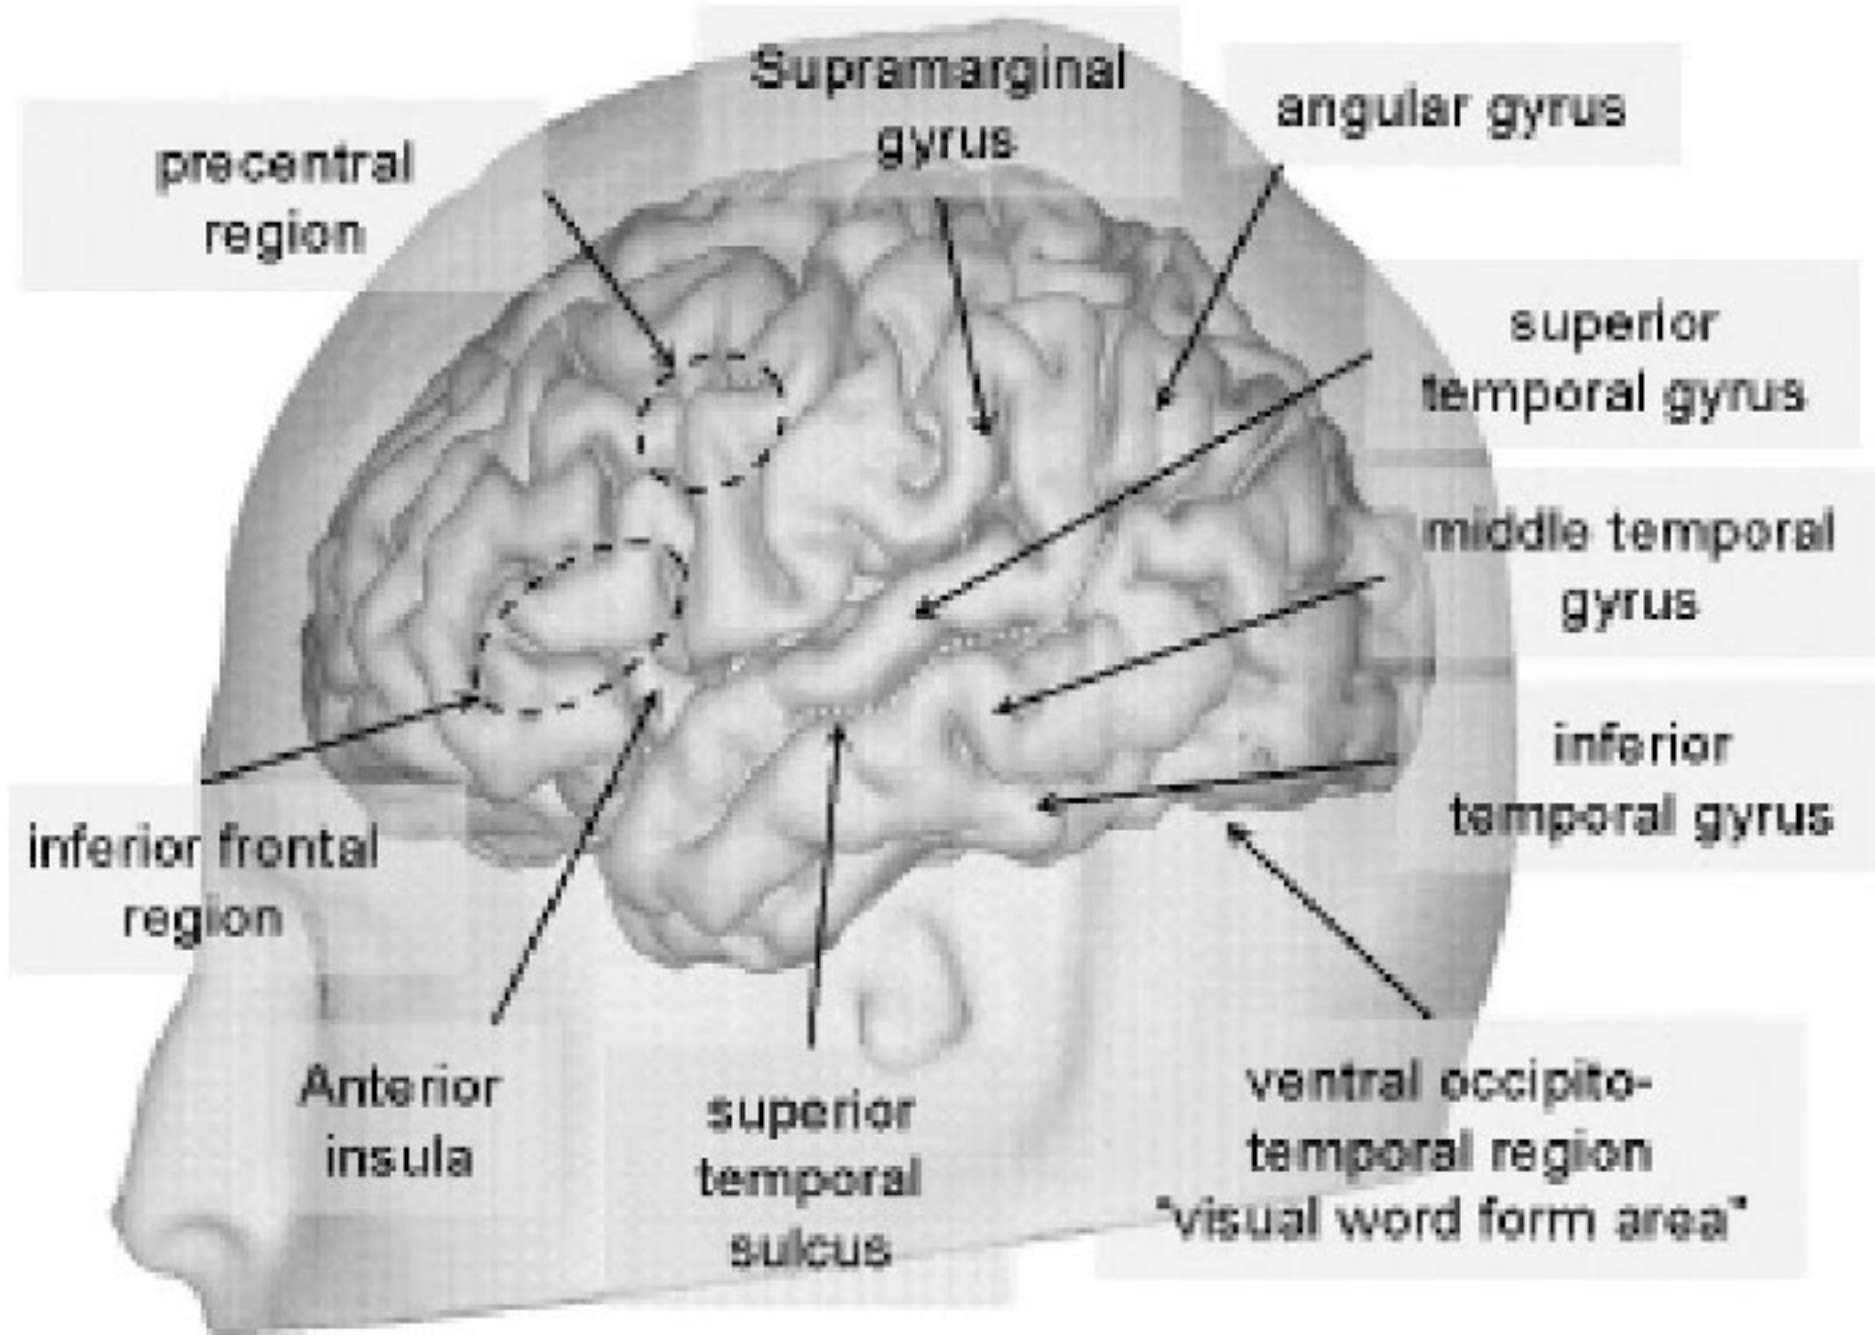
Getting oriented in the brain




The major lobes



Some reading areas



The brain seen from underneath



The New Science of Reading

Withdrawn into the peace of this desert, along with some books, few but wise, I live in conversation with the deceased, and listen to the dead with my eyes

—FRANCISCO DE QUEVEDO

At this very moment, your brain is accomplishing an amazing feat—reading. Your eyes scan the page in short spasmodic movements. Four or five times per second, your gaze stops just long enough to recognize one or two words. You are, of course, unaware of this jerky intake of information. Only the sounds and meanings of the words reach your conscious mind. But how can a few black marks on white paper projected onto your retina evoke an entire universe, as Vladimir Nabokov does in the opening lines of *Lolita*:

Lolita, light of my life, fire of my loins. My sin, my soul.
Lo-lee-ta: the tip of the tongue taking a trip of three
steps down the palate to tap, at three, on the teeth. Lo.
Lee. Ta.

The reader's brain contains a complicated set of mechanisms admirably attuned to reading. For a great many centuries, this talent remained a mystery. Today, the brain's black box is cracked open and a true science of reading is coming into being. Advances in psychology and neuroscience over the last twenty years have begun to unravel the principles underlying the brain's reading circuits. Modern brain imaging methods now reveal, in just a matter of minutes, the brain areas that activate when we

decipher written words. Scientists can track a printed word as it progresses from the retina through a chain of processing stages, each of which is marked by an elementary question: Are these letters? What do they look like? Are they a word? What does it sound like? How is it pronounced? What does it mean?

On this empirical ground, a theory of reading is materializing. It postulates that the brain circuitry inherited from our primate evolution can be co-opted to the task of recognizing printed words. According to this approach, our neuronal networks are literally “recycled” for reading. The insight into how literacy changes the brain is profoundly transforming our vision of education and learning disabilities. New remediation programs are being conceived that should, in time, cope with the debilitating incapacity to decipher words known as dyslexia.

My purpose in this book is to share my knowledge of recent and little-known advances in the science of reading. In the twenty-first century, the average person still has a better idea of how a car works than of the inner functioning of his own brain—a curious and shocking state of affairs. Decision makers in our education systems swing back and forth with the changing winds of pedagogical reform, often blatantly ignoring how the brain actually learns to read. Parents, educators, and politicians often recognize that there is a gap between educational programs and the most up-to-date findings in neuroscience. But too frequently their idea of how this field can contribute to advances in education is only grounded in a few color pictures of the brain at work. Unfortunately, the imaging techniques that allow us to visualize brain activity are subtle and occasionally misleading. The new science of reading is so young and fast-moving that it is still relatively unknown outside the scientific community. My goal is to provide a simple introduction to this exciting field, and to increase awareness of the amazing capacities of our reading brains.

From Neurons to Education

Reading acquisition is a major step in child development. Many children initially struggle with reading, and surveys indicate that

about one adult in ten fails to master even the rudiments of text comprehension. Years of hard work are needed before the clockwork-like brain machinery that supports reading runs so smoothly that we forget it exists.

Why is reading so difficult to master? What profound alterations in brain circuitry accompany the acquisition of reading? Are some teaching strategies better adapted to the child's brain than others? What scientific reasons, if any, explain why phonics—the systematic teaching of letter-to-sound correspondences—seems to work better than whole-word teaching? Although much still remains to be discovered, the new science of reading is now providing increasingly precise answers to all these questions. In particular, it underlines why early research on reading erroneously supported the whole-word approach—and how recent research on the brain's reading networks proves it was wrong.

Understanding what goes into reading also sheds light on its pathologies. In our explorations of the reader's mind and brain, you will be introduced to patients who suddenly lost the ability to read following a stroke. I will also analyze the causes of dyslexia, whose cerebral underpinnings are gradually coming to light. It is now clear that the dyslexic brain is subtly different from the brain of a normal reader. Several dyslexia susceptibility genes have been identified. But this is by no means a reason for discouragement or resignation. New intervention therapies are now being defined. Intensive retraining of language and reading circuits has brought about major improvements in children's brains that can readily be tracked with brain imaging.

Putting Neurons into Culture

Our ability to read brings us face-to-face with the singularity of the human brain. Why is *Homo sapiens* the only species that actively teaches itself? Why is he unique in his ability to transmit a sophisticated culture? How does the biological world of synapses and neurons relate to the universe of human cultural inventions? Reading, but also writing, mathematics, art, religion, agriculture, and city life have dramatically increased the native capacities of

our primate brains. Our species alone rises above its biological condition, creates an artificial cultural environment for itself, and teaches itself new skills like reading. This uniquely human competence is puzzling and calls for a theoretical explanation.

One of the basic techniques in the neurobiologist's toolkit consists of "putting neurons in culture"—letting neurons grow in a petri dish. In this book, I call for a different "culture of neurons"—a new way of looking at human cultural activities, based on our understanding of how they map onto the brain networks that support them. Neuroscience's avowed goal is to describe how the elementary components of the nervous system lead to the behavioral regularities that can be observed in children and adults (including advanced cognitive skills). Reading provides one of the most appropriate test beds for this "neurocultural" approach. We are increasingly aware of how writing systems as different as Chinese, Hebrew, or English get inscribed in our brain circuits. In the case of reading, we can clearly draw direct links between our native neuronal architecture and our acquired cultural abilities—but the hope is that this neuroscience approach will extend to other major domains of human cultural expression.

The Mystery of the Reading Ape

If we are to reconsider the relation between brain and culture, we must address an enigma, which I call the *reading paradox*: Why does our primate brain read? Why does it have an inclination for reading although this cultural activity was invented only a few thousand years ago?

There are good reasons why this deceptively simple question deserves to be called a paradox. We have discovered that the literate brain contains specialized cortical mechanisms that are exquisitely attuned to the recognition of written words. Even more surprisingly, the same mechanisms, in all humans, are systematically housed in identical brain regions, as though there were a cerebral organ for reading.

But writing was born only fifty-four hundred years ago in the Fertile Crescent, and the alphabet itself is only thirty-eight hundred

years old. These time spans are a mere trifle in evolutionary terms. Evolution thus did not have the time to develop specialized reading circuits in *Homo sapiens*. Our brain is built on the genetic blueprint that allowed our hunter-gatherer ancestors to survive. We take delight in reading Nabokov and Shakespeare using a primate brain originally designed for life in the African savanna. Nothing in our evolution could have prepared us to absorb language through vision. Yet brain imaging demonstrates that the adult brain contains fixed circuitry exquisitely attuned to reading.

The reading paradox is reminiscent of the Reverend William Paley's parable aimed at proving the existence of God. In his *Natural Theology* (1802), he imagined that in a deserted heath, a watch was found on the ground, complete with its intricate inner workings clearly designed to measure time. Wouldn't it provide, he argued, clear proof that there is an intelligent clockmaker, a designer who purposely created the watch? Similarly, Paley maintained that the intricate devices that we find in living organisms, such as the astonishing mechanisms of the eye, prove that nature is the work of a divine watchmaker.

Charles Darwin famously refuted Paley by showing how blind natural selection can produce highly organized structures. Even if biological organisms at first glance seem designed for a specific purpose, closer examination reveals that their organization falls short of the perfection that one would expect from an omnipotent architect. All sorts of imperfections attest that evolution is not guided by an intelligent creator, but follows random paths in the struggle for survival. In the retina, for example, blood vessels and nerve cables are situated *in front of* the photoreceptors, thus partially blocking incoming light and creating a blind spot—very poor design indeed.

Following in Darwin's footsteps, Stephen Jay Gould provided many examples of the imperfect outcome of natural selection, including the panda's thumb.¹ The British evolutionist Richard Dawkins also explained how the delicate mechanisms of the eye or of the wing could only have emerged through natural selection or are the work of a "blind watchmaker."² Darwin's evolutionism seems to be the only source of apparent "design" in nature.

When it comes to explaining reading, however, Paley's parable

is problematic in a subtly different way. The clockwork-like brain mechanisms that support reading are certainly comparable in complexity and sheer design to those of the watch abandoned on the heath. Their entire organization leans toward the single apparent goal of decoding written words as quickly and accurately as possible. Yet neither the hypothesis of an intelligent creator nor that of slow emergence through natural selection seems to provide a plausible explanation for the origins of reading. Time was simply too short for evolution to design specialized reading circuits. How, then, did our primate brain learn to read? Our cortex is the outcome of millions of years of evolution in a world without writing—why can it adapt to the specific challenges posed by written word recognition?

Biological Unity and Cultural Diversity

In the social sciences, the acquisition of cultural skills such as reading, mathematics, or the fine arts is rarely, if ever, posed in biological terms. Until recently, very few social scientists considered that brain biology and evolutionary theory were even relevant to their fields. Even today, most implicitly subscribe to a naïve model of the brain, tacitly viewing it as an infinitely plastic organ whose learning capacity is so broad that it places no constraints on the scope of human activity. This is not a new idea. It can be traced back to the theories of the British empiricists John Locke, David Hume, and George Berkeley, who claimed that the human brain should be compared to a blank slate that progressively absorbs the imprint of man's natural and cultural environment through the five senses.

This view of mankind, which denies the very existence of a human nature, has often been adopted without question. It belongs to the default “standard social science model”³ shared by many anthropologists, sociologists, some psychologists, and even a few neuroscientists who view the cortical surface as “largely equipotent and free of domain-specific structure.”⁴ It holds that human nature is constructed, gradually and flexibly, through cultural impregnation. As a result, children born to the Inuit, to the

hunter-gatherers of the Amazon, or to an Upper East Side New York family, according to this view, have little in common. Even color perception, musical appreciation, or the notion of right and wrong should vary from one culture to the next, simply because the human brain has few stable structures other than the capacity to learn.

Empiricists further maintain that the human brain, unhindered by the limitations of biology and unlike that of any other animal species, can absorb any form of culture. From this theoretical perspective, to talk about the cerebral bases of cultural inventions such as reading is thus downright irrelevant—much like analyzing the atomic composition of a Shakespeare play.

In this book, I refute this simplistic view of an infinite adaptability of the brain to culture. New evidence on the cerebral circuits of reading demonstrates that the hypothesis of an equipotent brain is wrong. To be sure, if the brain were not capable of learning, it could not adapt to the specific rules for writing English, Japanese, or Arabic. This learning, however, is tightly constrained, and its mechanisms themselves rigidly specified by our genes. The brain's architecture is similar in all members of the *Homo sapiens* family, and differs only slightly from that of other primates. All over the world, the same brain regions activate to decode a written word. Whether in French or in Chinese, learning to read necessarily goes through a genetically constrained circuit.

On the basis of these data, I propose a novel theory of neurocultural interactions, radically opposed to cultural relativism, and capable of resolving the reading paradox. I call it the “neuronal recycling” hypothesis. According to this view, human brain architecture obeys strong genetic constraints, but some circuits have evolved to tolerate a fringe of variability. Part of our visual system, for instance, is not hardwired, but remains open to changes in the environment. Within an otherwise well-structured brain, visual plasticity gave the ancient scribes the opportunity to invent reading.

In general, a range of brain circuits, defined by our genes, provides “pre-representations”⁵ or hypotheses that our brain can entertain about future developments in its environment. During

brain development, learning mechanisms select which pre-representations are best adapted to a given situation. Cultural acquisition rides on this fringe of brain plasticity. Far from being a blank slate that absorbs everything in its surroundings, our brain adapts to a given culture by minimally turning its predispositions to a different use. It is not a *tabula rasa* within which cultural constructions are amassed, but a very carefully structured device that manages to convert some of its parts to a new use. When we learn a new skill, we recycle some of our old primate brain circuits—insofar, of course, as those circuits can tolerate the change.

A Reader's Guide

In forthcoming chapters, I will explain how neuronal recycling can account for literacy, its mechanisms in the brain, and even its history. In the first three chapters, I analyze the mechanisms of reading in expert adults. [Chapter 1](#) sets the stage by looking at reading from a psychological angle: how fast do we read, and what are the main determinants of reading behavior? In [chapter 2](#), I move to the brain areas at work when we read, and how they can be visualized using modern brain imaging techniques. Finally, in [chapter 3](#), I come down to the level of single neurons and their organization into the circuits that recognize letters and words.

I tackle my analysis in a resolutely mechanical way. I propose to expose the cogwheels of the reader's brain in much the same way as the Reverend Paley suggested we dismantle the watch abandoned on the heath. The reader's brain will not, however, reveal any perfect clockwork mechanics designed by a divine watchmaker. Our reading circuits contain more than a few imperfections that betray our brain's compromise between what is needed for reading and the available biological mechanisms. The peculiar characteristics of the primate visual system explain why reading does not operate like a fast and efficient scanner. As we move our eyes across the page, each word is slowly brought into the central region of our retina, only to be exploded into a myriad of fragments that our brain later pieces back together. It is only because these processes have become automatic and

unconscious, thanks to years of practice, that we are under the illusion that reading is simple and effortless.

The reading paradox expresses the indisputable fact that our genes have not evolved in order to enable us to read. My reasoning in the face of this enigma is quite simple. If the brain did not evolve for reading, the opposite must be true: writing systems must have evolved within our brain's constraints. [Chapter 4](#) revisits the history of writing in this light, starting with the first prehistoric symbols and ending with the invention of the alphabet. At each step, there is evidence of constant cultural tinkering. Over many millennia, the scribes struggled to design words, signs, and alphabets that could fit the limits of our primate brain. To this day, the world's writing systems still share a number of design features that can ultimately be traced back to the restrictions imposed by our brain circuits.

Continuing on the idea that our brain was not designed for reading, but recycles some of its circuits for this novel cultural activity, [chapter 5](#) examines how children learn to read. Psychological research concludes that there are not many ways to convert a primate brain into that of an expert reader. This chapter explores in some detail the only developmental trajectory that appears to exist. Schools might be well advised to exploit this knowledge to optimize the teaching of reading and mitigate the dramatic effects of illiteracy and dyslexia.

I will also go on to show how a neuroscientific approach can shed light on the more mysterious features of reading acquisition. For instance, why do so many children often write their first words from right to left? Contrary to the accepted idea, these mirror inversion errors are not the first signs of dyslexia, but a natural consequence of the organization of our visual brain. In a majority of children, dyslexia relates to another, quite distinct anomaly in processing speech sounds. The description of the symptoms of dyslexia, their cerebral bases, and the most recent discoveries concerning its genetic foundations are covered in [chapter 6](#), while [chapter 7](#) provides an insight into what mirror errors can tell us about normal visual recognition.

Finally, in [chapter 8](#), I will return to the astonishing fact that only our species is capable of cultural inventions as sophisticated

as reading—a unique feat, unmatched by any other primate. In total opposition to the standard social science model, where culture gets a free ride on a blank-slate brain, reading demonstrates how culture and brain organization are inextricably linked. Throughout their long cultural history, human beings progressively discovered that they could reuse their visual systems as surrogate language inputs, thus arriving at reading and writing. I will also briefly discuss how other major human cultural traits could be submitted to a similar analysis.

Mathematics, art, music, and religion might also be looked on as evolved devices, shaped by centuries of cultural evolution, that have encroached on our primate brains.

One last enigma remains: if learning exists in all primates, why is *Homo sapiens* the only species with a sophisticated culture? Although the term is sometimes applied to chimpanzees, their “culture” barely goes beyond a few good tricks for splitting nuts, washing potatoes or fishing ants with a stick—nothing comparable to the seemingly endless human production of interlocking conventions and symbols systems, including languages, religions, art forms, sports, mathematics or medicine. Nonhuman primates can slowly learn to recognize novel symbols such as letters and digits—but they never think of inventing them. In my conclusion, I propose some tentative ideas on the singularity of the human brain. The uniqueness of our species may arise from a combination of two factors: a theory of mind (the ability to imagine the mind of others) and a conscious global workspace (an internal buffer where an infinite variety of ideas can be recombined). Both mechanisms, inscribed in our genes, conspire to make us the only cultural species. The seemingly infinite variety of human cultures is only an illusion, caused by the fact that we are locked in a cognitive vicious circle: how could we possibly imagine forms other than those our brains can conceive? Reading, although a recent invention, lay dormant for millennia within the envelope of potentialities inscribed in our brains. Behind the apparent diversity of human writing systems lies a core set of universal neuronal mechanisms that, like a watermark, reveal the constraints of human nature.

How Do We Read?

Written word processing starts in our eyes. Only the center of the retina, called the fovea, has a fine enough resolution to allow for the recognition of small print. Our gaze must therefore move around the page constantly. Whenever our eyes stop, we only recognize one or two words. Each of them is then split up into myriad fragments by retinal neurons and must be put back together before it can be recognized. Our visual system progressively extracts graphemes, syllables, prefixes, suffixes, and word roots. Two major parallel processing routes eventually come into play: the phonological route, which converts letters into speech sounds, and the lexical route, which gives access to a mental dictionary of word meanings.

The existence of the text is a silent existence, silent until the moment in which a reader reads it. Only when the able eye makes contact with the markings on the tablet does the text come to active life. All writing depends on the generosity of the reader.

—ALBERTO MANGUEL, *THE HISTORY OF READING*

At first sight, reading seems close to magical: our gaze lands on a word, and our brain effortlessly gives us access to its meaning and pronunciation. But in spite of appearances, the process is far from simple. Upon entering the retina, a word is split up into a myriad of fragments, as each part of the visual image is recognized by a distinct photoreceptor. Starting from this input, the real challenge consists in putting the pieces back together in order to decode what letters are present, to figure out the order in which they appear, and finally to identify the word.

Over the past thirty years, cognitive psychology has worked on

analyzing the mechanics of reading. Its goal is to crack the “algorithm” of visual word recognition—the series of processing steps that a proficient reader applies to the problem of identifying written words. Psychologists treat reading like a computer science problem. Every reader resembles a robot with two cameras—the two eyes and their retinas. The words we read are painted onto them. They first appear only as splotches of light and dark that are not directly interpretable as linguistic signs. Visual information must be recoded in an understandable format before we can access the appropriate sounds, words, and meanings. Thus we must have a deciphering algorithm, or a processing recipe akin to automatic character recognition software, which takes the pixels on a page as input and produces the identity of the words as output. To accomplish this feat, unbeknownst to us, our brain hosts a sophisticated set of decoding operations whose principles are only beginning to be understood.

The Eye: A Poor Scanner

The tale of reading begins when the retina receives photons reflected off the written page. But the retina is not a homogeneous sensor. Only its central part, called the fovea, is dense in high-resolution cells sensitive to incoming light, while the rest of the retina has a coarser resolution. The fovea, which occupies about 15 degrees of the visual field, is the only part of the retina that is genuinely useful for reading. When foveal information is lacking, whether due to a retinal lesion, to a stroke having destroyed the central part of the visual cortex, or to an experimental trick that selectively blocks visual inputs to the fovea, reading becomes impossible.⁶

The need to bring words into the fovea explains why our eyes are in constant motion when we read. By orienting our gaze, we “scan” text with the most sensitive part of our vision, the only one that has the resolution needed to determine letters. However, our eyes do not travel continuously across the page.⁷ Quite the opposite: they move in small steps called saccades. At this very moment, you are making four or five of these jerky movements

every second, in order to bring new information to your fovea.

Even within the fovea, visual information is not represented with the same precision at all points. In the retina as well as in the subsequent visual relays of the thalamus and of the cortex, the number of cells allocated to a given portion of the visual scene decreases progressively as one moves away from the center of gaze. This causes a gradual loss of visual precision. Visual accuracy is optimal at the center and smoothly decreases toward the periphery. We have the illusion of seeing the whole scene in front of us with the same fixed accuracy, as if it were filmed by a digital camera with a homogeneous array of pixels. However, unlike the camera, our eye sensor accurately perceives only the precise point where our gaze happens to land. The surroundings are lost in an increasingly hazy blurriness ([figure 1.1](#)).⁸


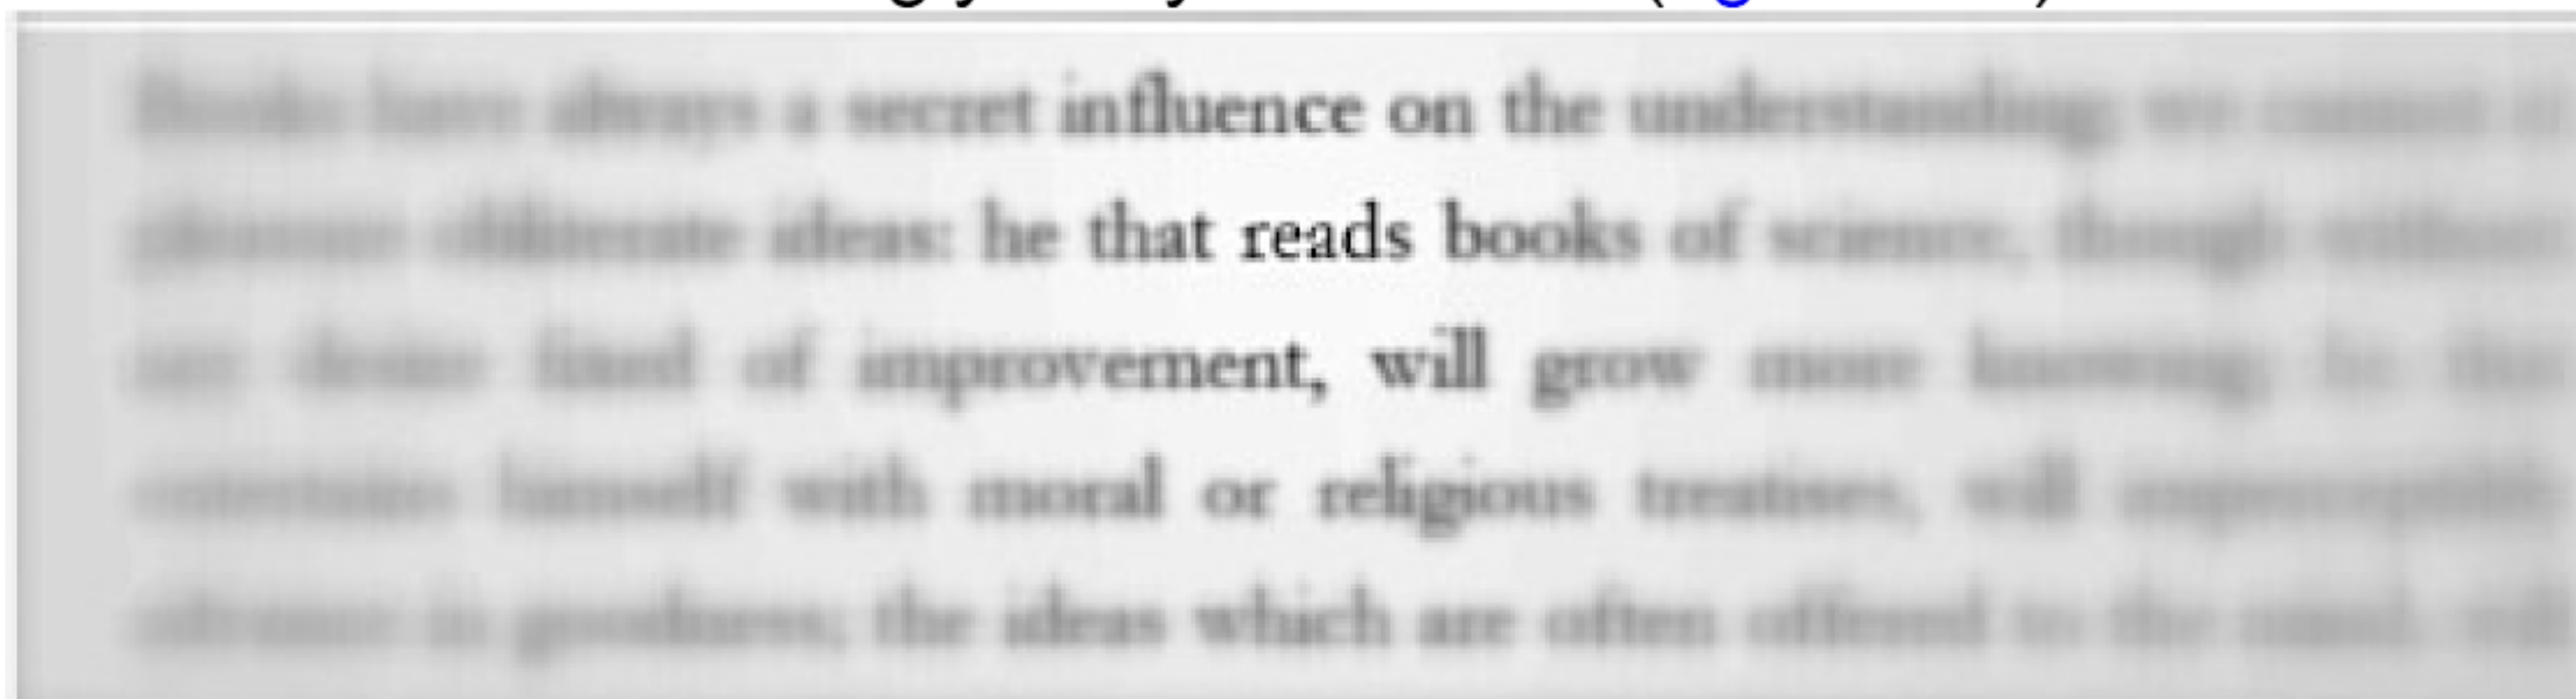


Figure 1.1 The retina stringently filters what we read. In this simulation, a page from Samuel Johnson's *The Adventurer* (1754) was filtered using an algorithm that copies the decreasing acuity of human vision away from the center of the retina. Regardless of size, only letters close to fixation can be identified. This is why we constantly explore pages with jerky eye movements when we read. When our gaze stops, we can only identify one or two words.

One might think that, under these conditions, it is the absolute size of printed characters that determines the ease with which we can read: small letters should be harder to read than larger ones.

Oddly enough, however, this is not the case. The reason is that the larger the characters, the more room they use on the retina. When a whole word is printed in larger letters, it moves into the periphery of the retina, where even large letters are hard to discern. The two factors compensate for each other almost exactly, so that an **enormous** word and a minuscule one are essentially equivalent from the point of view of retinal precision. Of course, this is only true provided that the size of the characters remains larger than an absolute minimum, which corresponds to the maximal precision attained at the center of our fovea. When visual acuity is diminished, for instance in aging patients, it is quite logical to recommend books in large print.

Because our eyes are organized in this way, our perceptual abilities depend exclusively on the number of letters in words, not on the space these words occupy on our retina.⁹ Indeed, our saccades when we read vary in absolute size, but are constant when measured in numbers of letters. When the brain prepares to move our eyes, it adapts the distance to be covered to the size of the characters, in order to ensure that our gaze always advances by about seven to nine letters. This value, which is amazingly small, thus corresponds approximately to the information that we can process in the course of a single eye fixation.

To prove that we see only a very small part of each page at a time, George W. McConkie and Keith Rayner developed an experimental method that I like to call the “Cartesian devil.” In his *Metaphysical Meditations*, René Descartes imagined that an evil genius was playing with our senses:

I shall then suppose, not that God who is supremely good and the fountain of truth, but some evil genius not less powerful than deceitful, has employed all his energy to deceive me; I shall consider that the heavens, the earth, colors, figures, sound, and all other external things are naught but the illusions and dreams of which this genius has availed himself in order to lay traps for my credulity. I shall consider myself as having no hands, no eyes, no flesh, no blood, nor any senses, yet falsely believing myself to possess all these things.

Much like the supercomputer in the *Matrix* movies, Descartes' evil genius produces a pseudo-reality by bombarding our senses with signals carefully crafted to create an illusion of real life, a virtual scene whose true side remains forever hidden. More modestly, McConkie and Rayner designed a "moving window" that creates an illusion of text on a computer screen.¹⁰ The method consists in equipping a human volunteer with a special device that tracks eye movements and can change the visual display in real time. The device can be programmed to display only a few characters left and right of the center of gaze, while all of the remaining letters on the page are replaced with strings of x's:

We the pxxxxx xx xxx xxxxxxxx xxxxxxxx, xx xxxxxxx xx
↑

As soon as the eyes move, the computer discreetly refreshes the display. Its goal is to show the appropriate letters at the place where the person is looking, and strings of x's everywhere else:

Xx xxx people of txx xxxxxxx xxxxxxx, xx xxxxxxx xx
↑

Xx xxx xxxxxxx xx xhe United xxxxxxx, xx xxxxxxx xx
↑

Xx xxx xxxxxxx xx xxx Xxxxxed States, ix xxxxxxx xx
↑

Xx xxx xxxxxxx xx xxx Xxxxxx Xxxxxx, in order to
↑

Using this device, McConkie and Rayner made a remarkable and paradoxical discovery. They found that the participants did not notice the manipulation. As long as enough letters are presented left and right of fixation, a reader fails to detect the trick and believes that he is looking at a perfectly normal page of text.

This surprising blindness occurs because the eye attains its

maximum speed at the point when the letter change occurs. This trick makes the letter changes hard to detect, because at this very moment the whole retinal image is blurred by motion. Once gaze lands, everything looks normal: within the fovea, the expected letters are in place, and the rest of the visual field, on the periphery, cannot be read anyway. McConkie and Rayner's experiment thus proves that we consciously process only a very small subset of our visual inputs. If the computer leaves four letters on the left of fixation, and fifteen letters on the right, reading speed remains normal.¹¹ In brief, we extract very little information at a time from the written page. Descartes' evil genius would only have to display twenty letters per fixation to make us believe that we were reading the Bible or the U.S. Constitution!

Twenty letters is, in fact, an overestimate. We identify only ten or twelve letters per saccade: three or four to the left of fixation, and seven or eight to the right. Beyond this point, we are largely insensitive to letter identity and merely encode the presence of the spaces between words. By providing cues about word length, the spaces allow us to prepare our eye movements and ensure that our gaze lands close to the center of the next word. Experts continue to debate about the extent to which we extract information from an upcoming word—perhaps only the first few letters. Everyone agrees, however, that the direction of reading imposes asymmetry on our span of vision. In the West, visual span is much greater toward the right side, but in readers of Arabic or Hebrew, where gaze scans the page from right to left, this asymmetry is reversed.¹² In other writing systems such as Chinese, where character density is greater, saccades are shorter and visual span is reduced accordingly. Each reader thus adapts his visual exploration strategy to his language and script.

Using the same method, we can also estimate how much time is needed to encode the identity of words. A computer can be programmed so that, after a given duration, all of the letters are replaced by a string of x's, even in the fovea. This experiment reveals that fifty milliseconds of presentation are enough for reading to proceed at an essentially normal pace. This does not mean that all of the mental operations involved in reading are completed in one-twentieth of a second. As we shall see, a whole

pipeline of mental processes continues to operate for at least one-half second after the word has been presented. However, the initial intake of visual information can be very brief.

In summary, our eyes impose a lot of constraints on the act of reading. The structure of our visual sensors forces us to scan the page by jerking our eyes around every two or three tenths of a second. Reading is nothing but the word-by-word mental restitution of a text through a series of snapshots. While some small grammatical words like “the,” “it,” or “is” can sometimes be skipped, almost all content words such as nouns and verbs have to be fixated at least once.

These constraints are an integral part of our visual apparatus and cannot be lifted by training. One can certainly teach people to optimize their eye movement patterns, but most good readers, who read from four hundred to five hundred words per minute, are already close to optimal. Given the retinal sensor at our disposal, it is probably not possible to do much better. A simple demonstration proves that eye movements are the rate-limiting step in reading.¹³ If a full sentence is presented, word by word, at the precise point where gaze is focalized, thus avoiding the need for eye movements, a good reader can read at staggering speed—a mean of eleven hundred words per minute, and up to sixteen hundred words per minute for the best readers, which is about one word every forty milliseconds and three to four times faster than normal reading! With this method, called rapid sequential visual presentation, or RSVP, identification and comprehension remain satisfactory, thus suggesting that the duration of those central steps does not impose a strong constraint on normal reading. Perhaps this computerized presentation mode represents the future of reading in a world where screens progressively replace paper.

At any rate, as long as text is presented in pages and lines, acquisition through gaze will slow reading and impose an unavoidable limitation. Thus, fast reading methods that advertise gains in reading speed of up to one thousand words per minute or more must be viewed with skepticism.¹⁴ One can no doubt broaden one’s visual span somewhat, in order to reduce the number of saccades per line, and it is also possible to learn to

avoid moments of regression, where gaze backtracks to the words it has just read. However, the physical limits of the eyes cannot be overcome, unless one is willing to skip words and thus run the risk of a misunderstanding. Woody Allen described this situation perfectly: “I took a speed-reading course and was able to read *War and Peace* in twenty minutes. It involves Russia.”

The Search for Invariants

Can you read, Lubin?

Yes, I can read printed letters, but I was never able to read handwriting.

—MOLIÈRE, *GEORGES DANDIN*

Reading poses a difficult perceptual problem. We must identify words regardless of how they appear, whether in print or handwritten, in upper- or lowercase, and regardless of their size. This is what psychologists call the *invariance problem*: we need to recognize which aspect of a word does not vary—the sequence of letters—in spite of the thousand and one possible shapes that the actual characters can take on.

If perceptual invariance is a problem, it is because words are not always in the same location, in the same font, or in the same size. If they were, just listing which of the cells on the retina are active and which are not would suffice to decode a word, much like a black-and-white computer image is defined by the list of its pixels. In fact, however, hundreds of different retinal images can stand for the same word, depending on the form in which it is written ([figure 1.2](#)). Thus one of the first steps in reading must be to correct for the immense variety of those surface forms.

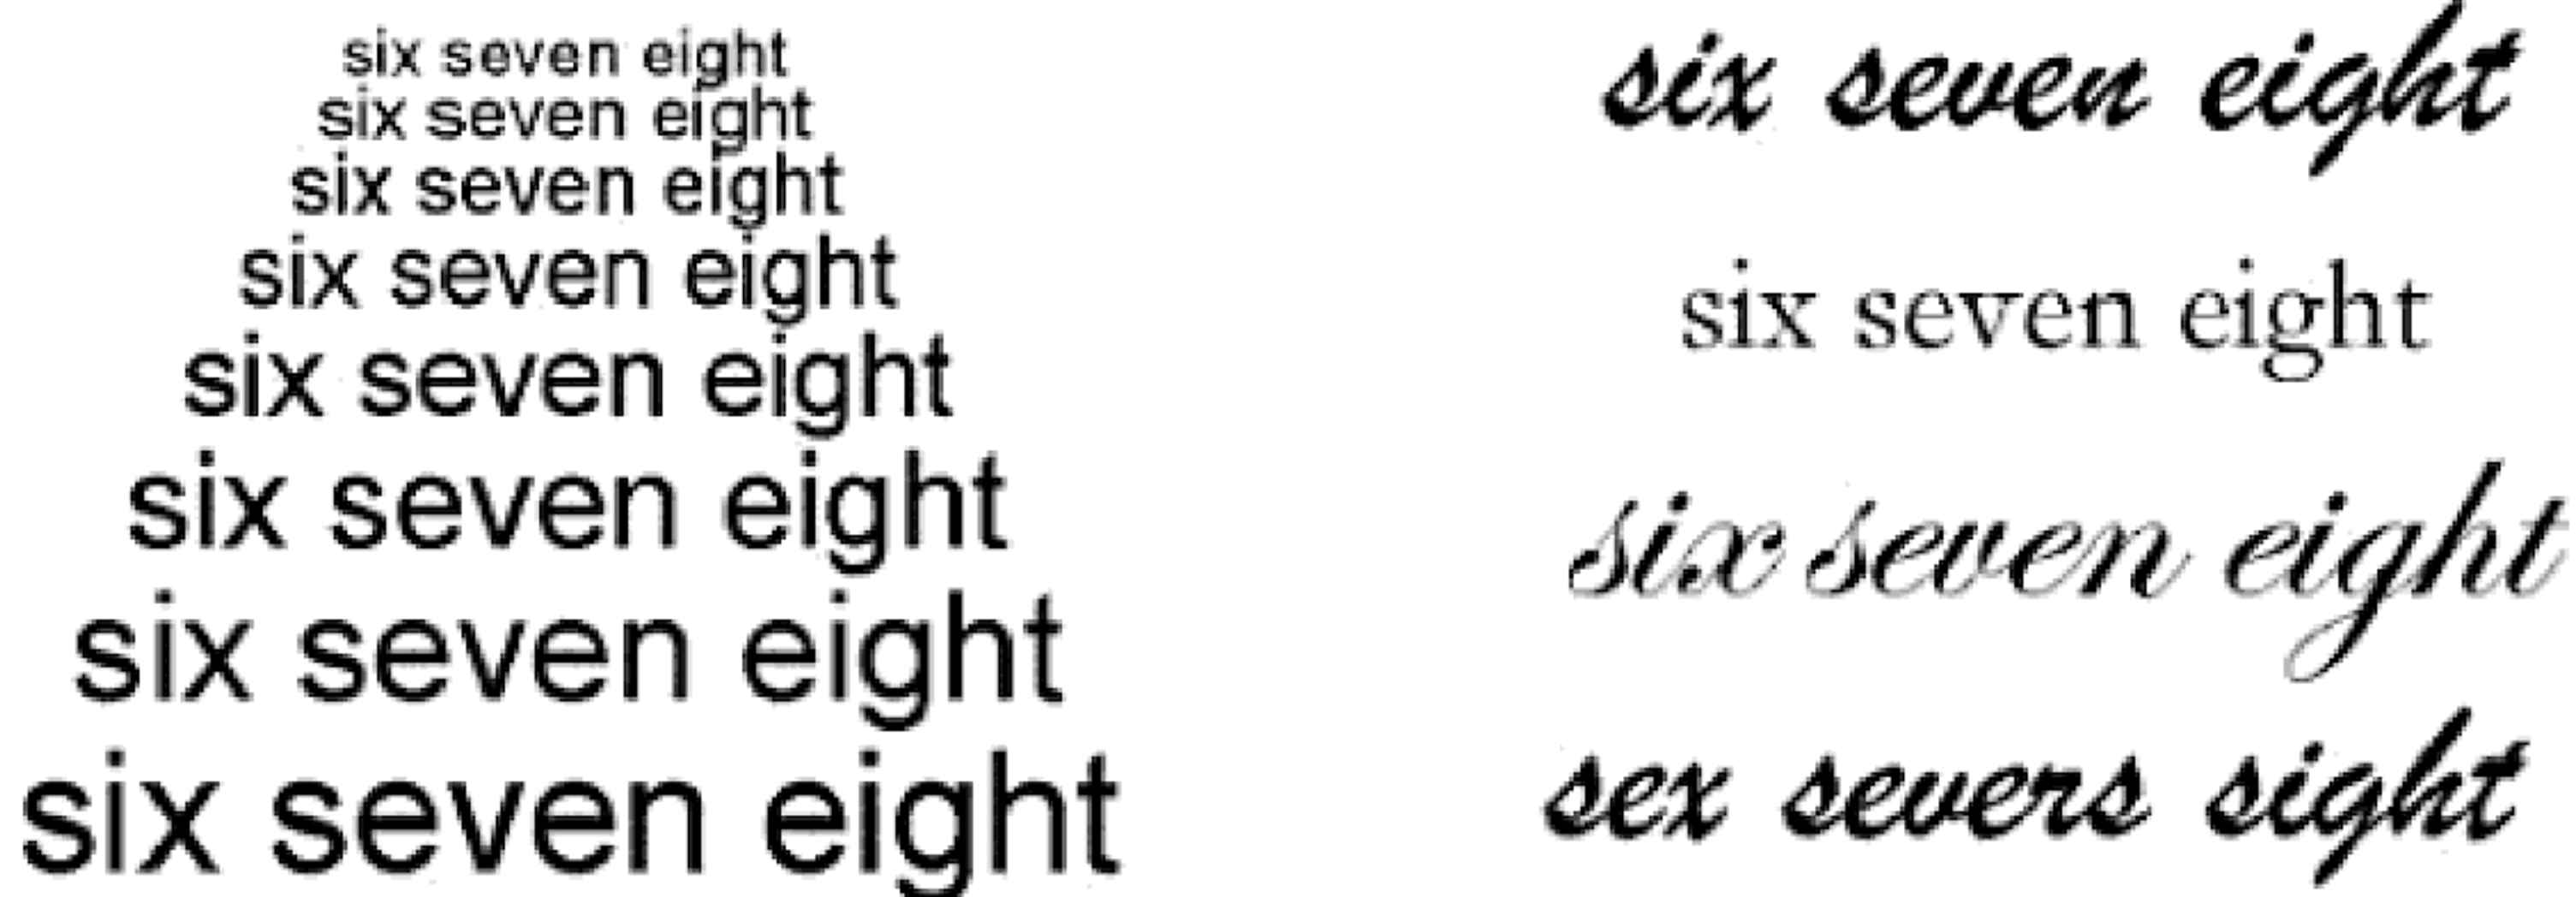


Figure 1.2 Visual invariance is one of the prime features of the human reading system. Our word recognition device meets two seemingly contradictory requirements: it neglects irrelevant variations in character shape, even if they are huge, but amplifies relevant differences, even if they are tiny.



Unbeknownst to us, our visual system automatically compensates for enormous variations in size or font. Yet it also attends to minuscule changes in shape. By turning an “s” into an “e,” and therefore “sight” into “eight,” a single mark drastically reorients the processing chain toward entirely distinct pronunciations and meanings.

Several cues suggest that our brain applies an efficient solution to this perceptual invariance problem. When we hold a newspaper at a reasonable distance, we can read both the headlines and the classified ads. Word size can vary fiftyfold without having much impact on our reading speed. This task is not very different from that of recognizing the same face or object from a distance of two feet or thirty yards—our visual system tolerates vast changes in scale.

A second form of invariance lets us disregard the location of words on the page. As our gaze scans a page, the center of our retina usually lands slightly left of the center of words. However, our targeting is far from perfect, and our eyes sometimes reach the first or last letter without this preventing us from recognizing the word. We can even read words presented on the periphery of our visual field, provided that letter size is increased to compensate for the loss of retinal resolution. Thus size constancy goes hand in hand with normalization for spatial location.

Finally, word recognition is also largely invariant for character shape. Now that word processing software is omnipresent, technology that was formerly reserved to a small elite of

typographers has become broadly available. Everyone knows that there are many sets of characters called “fonts” (a term left over from the time when each character had to be cast in lead at a type foundry before going to the press). Each font also has two kinds of characters called “cases,” the UPPERCASE and the lowercase (originally the case was a flat box divided into many compartments where lead characters were sorted; the “upper case” was reserved for capital letters, and the “lower case” for the rest). Finally, one can choose the “weight” of a font (normal or **bold** characters), its inclination (*italics*, originally invented in Italy), whether it is underlined or not, and any combination of *these options*. These well-calibrated variations in fonts, however, are nothing compared to the enormous variety of writing styles. Manuscript handwriting obviously takes us to another level of variability and ambiguity.

In the face of all these variations, exactly how our visual system learns to categorize letter shapes remains somewhat mysterious. Part of this invariance problem can be solved using relatively simple means. The vowel “o,” for instance, can easily be recognized, regardless of size, case, or font, thanks to its unique closed shape. Thus, building a visual o detector isn’t particularly difficult. Other letters, however, pose specific problems. Consider the letter “r,” for instance. Although it seems obvious that the shapes r, R,  and  all represent the same letter, careful examination shows that this association is entirely arbitrary—the shape e, for instance, might serve as well as the lowercase version of the letter “R.” Only the accidents of history have left us this cultural oddity. As a result, when we learn to read, we must not only learn that letters map onto the sounds of language, but also that each letter can take on many unrelated shapes. As we shall see, our capacity to do this probably comes from the existence of abstract letter detectors, neurons that can recognize the identity of a letter in its various guises. Experiments show that very little training suffices to DeCoDe, At An EsSeNtIaLly NoRmAl SpEeD, EnTiRe SeNtEnCes WhOsE LeTtErS HaVe BeEn PrInTeD AlTeRnAtEIY iN uPpErCaSe aNd In LoWeRcAsE.¹⁵ In the McConkie and Rayner “evil genius” computer, this letter-case alternation can be changed in between every eye saccade, totally

unbeknownst to the reader!¹⁶ In our daily reading experience, we never see words presented in alternating case, but our letter normalization processes are so efficient that they easily resist such transformation.

In passing, these experiments demonstrate that global word shape does not play any role in reading. If we can immediately recognize the identity of “words,” “WORDS,” and “WoRdS,” it is because our visual system pays no attention to the contours of words or to the pattern of ascending and descending letters: it is only interested in the letters they contain. Obviously, our capacity to recognize words does not depend on an analysis of their overall shape.

Amplifying Differences

Although our visual system efficiently filters out visual differences that are irrelevant to reading, such as the distinction between “R” and “r,” it would be a mistake to think that it always discards information and simplifies shapes. On the contrary, in many cases it must preserve, and even amplify, the minuscule details that distinguish two very similar words from each other. Consider the words “eight” and “sight.” We immediately access their very distinct meanings and pronunciations, but it is only when we look more closely at them that we realize that the difference is only a few pixels. Our visual system is exquisitely sensitive to the minuscule difference between “eight” and “sight,” and it amplifies it in order to send the input to completely different regions of semantic space. At the same time, it pays very little attention to other much greater differences, such as the distinction between “eight” and “EIGHT.”

As with invariance for case, this capacity to attend to relevant details results from years of training. The same reader who immediately spots the difference between letters “e” and “o,” and the lack of difference between “a” and “a,” may not notice that the Hebrew letters א and א differ sharply, a fact that seems obvious to any Hebrew reader.

Every Word Is a Tree

Our visual system deals with the problem of invariant word recognition using a well-organized system. As we shall see in detail in [chapter 2](#), the flow of neuronal activity that enters into the visual brain gets progressively sorted into meaningful categories. Shapes that appear very similar, such as “eight” and “sight,” are sifted through a series of increasingly refined filters that progressively separate them and attach them to distinct entries in a mental lexicon, a virtual dictionary of all the words we have ever encountered. Conversely, shapes like “eight” and “EIGHT,” which are composed of distinct visual features, are initially encoded by different neurons in the primary visual area, but are progressively recoded until they become virtually indistinguishable. Feature detectors recognize the similarity of the letters “i” and “l.” Other, slightly more abstract letter detectors classify “e” and “E” as two forms of the same letter. In spite of the initial differences, the reader’s visual system eventually encodes the very essence of the letter strings “eight” and “EIGHT,” regardless of their exact shape. It gives these two strings the same single mental address, an abstract code capable of orienting the rest of the brain toward the pronunciation and meaning of the word.

What does this address look like? According to some models, the brain uses a sort of unstructured list that merely provides the sequence of letters E-I-G-H-T. In others, it relies on a very abstract and conventional code, akin to a random cipher by which, say, [1296] would be the word “eight” and [3452] the word “sight.” Contemporary research, however, supports another hypothesis. Every written word is probably encoded by a hierarchical tree in which letters are grouped into larger-sized units, which are themselves grouped into syllables and words—much as a human body can be represented as an arrangement of legs, arms, torso, and head, each of which can be further broken down into simple parts.

A good example of the mental decomposition of words into relevant units can be found if we dissect the word “unbuttoning.” We must first strip off the prefix “un” and the familiar suffix or grammatical ending “ing.” Both frame the central element, the

word inside the word: the root “button.” All three of these components are called “morphemes”—the smallest units that carry some meaning. Each word is characterized, at this level, by how its morphemes are put together. Breaking down a word into its morphemes even allows us to understand words that we have never seen before, such as “reunbutton” or “deglochingization” (we understand that this is the undoing of the action of “gloching,” whatever that may be). In some languages, such as Turkish or Finnish, morphemes can be assembled into very large words that convey as much information as a full English sentence. In those languages, but also in ours, the decomposition of a word into its morphemes is an essential step on the path that leads from vision to meaning.

A lot of experimental data show that, very quickly and even downright unconsciously, our visual system snips out the morphemes of words. For instance, if I were to flash the word “departure” on a computer screen, you would later say the word “depart” slightly faster when confronted with it. The presentation of “departure” seems to preactivate the morpheme [depart], thus facilitating its access. Psychologists speak of a “priming” effect—the reading of a word primes the recognition of related words, much as one primes a pump. Importantly, this priming effect does not depend solely on visual similarity: words that look quite different but share a morpheme, such as “can” and “could,” can prime each other, whereas words that look alike but bear no intimate morphological relation, such as “aspire” and “aspirin,” do not. Priming also does not require any resemblance at the level of meaning; words such as “hard” and “hardly,” or “depart” and “department,” can prime each other, even though their meanings are essentially unrelated.¹⁷ Getting down to the morpheme level seems to be of such importance for our reading system that it is willing to make guesses about the decomposition of words. Our reading apparatus dissects the word “department” into [depart] + [ment] in the hope that this will be useful to the next operators computing its meaning.¹⁸ Never mind that this does not work all the time—a “listless” person is not one who is waiting for a grocery list, nor does sharing an “apartment” imply that you and your partner will soon live apart. Such parsing errors will have to be

caught at other stages in the word dissection process.

If we continue to undress the word “unbuttoning,” the morpheme [button] itself is not an indivisible whole. It is made up of two syllables, [bʌ] and [ton], each of which can be broken down into individual consonants and vowels: [b] [ʌ] [t] [o] [n]. Here lies another essential unit in our reading system: the grapheme, a letter or series of letters that maps onto a phoneme in the target language. Note that in our example, the two letters “tt” map onto a single sound *t*.¹⁹ Indeed, the mapping of graphemes onto phonemes isn’t always a direct process. In many languages, graphemes can be constructed out of a group of letters. English has a particularly extensive collection of complex graphemes such as “ough,” “oi,” and “au.”

Our visual system has learned to treat these groups of letters as bona fide units, to the point where we no longer pay attention to their actual letter content. Let us do a simple experiment to prove this point. Examine the following list of words and mark those that contain the letter “a”:

garage
metal
people
coat
please
meat

Did you feel that you had to slow down, ever so slightly, for the last three words, “coat,” “please,” and “meat”? They all contain the letter “a,” but it is embedded in a complex grapheme that is not pronounced like an “a.” If we were to rely only on letter detectors in order to detect the letter “a,” the parsing of the word into its graphemes would not matter. However, actual measurement of response times clearly shows that our brain does not stop at the single-letter level. Our visual system automatically regroups letters into higher-level graphemes, thus making it harder for us to see that groups of letters such as “ea” actually contain the letter “a.”²⁰

In turn, graphemes are automatically grouped into syllables. Here is another simple demonstration of this fact. You are going to

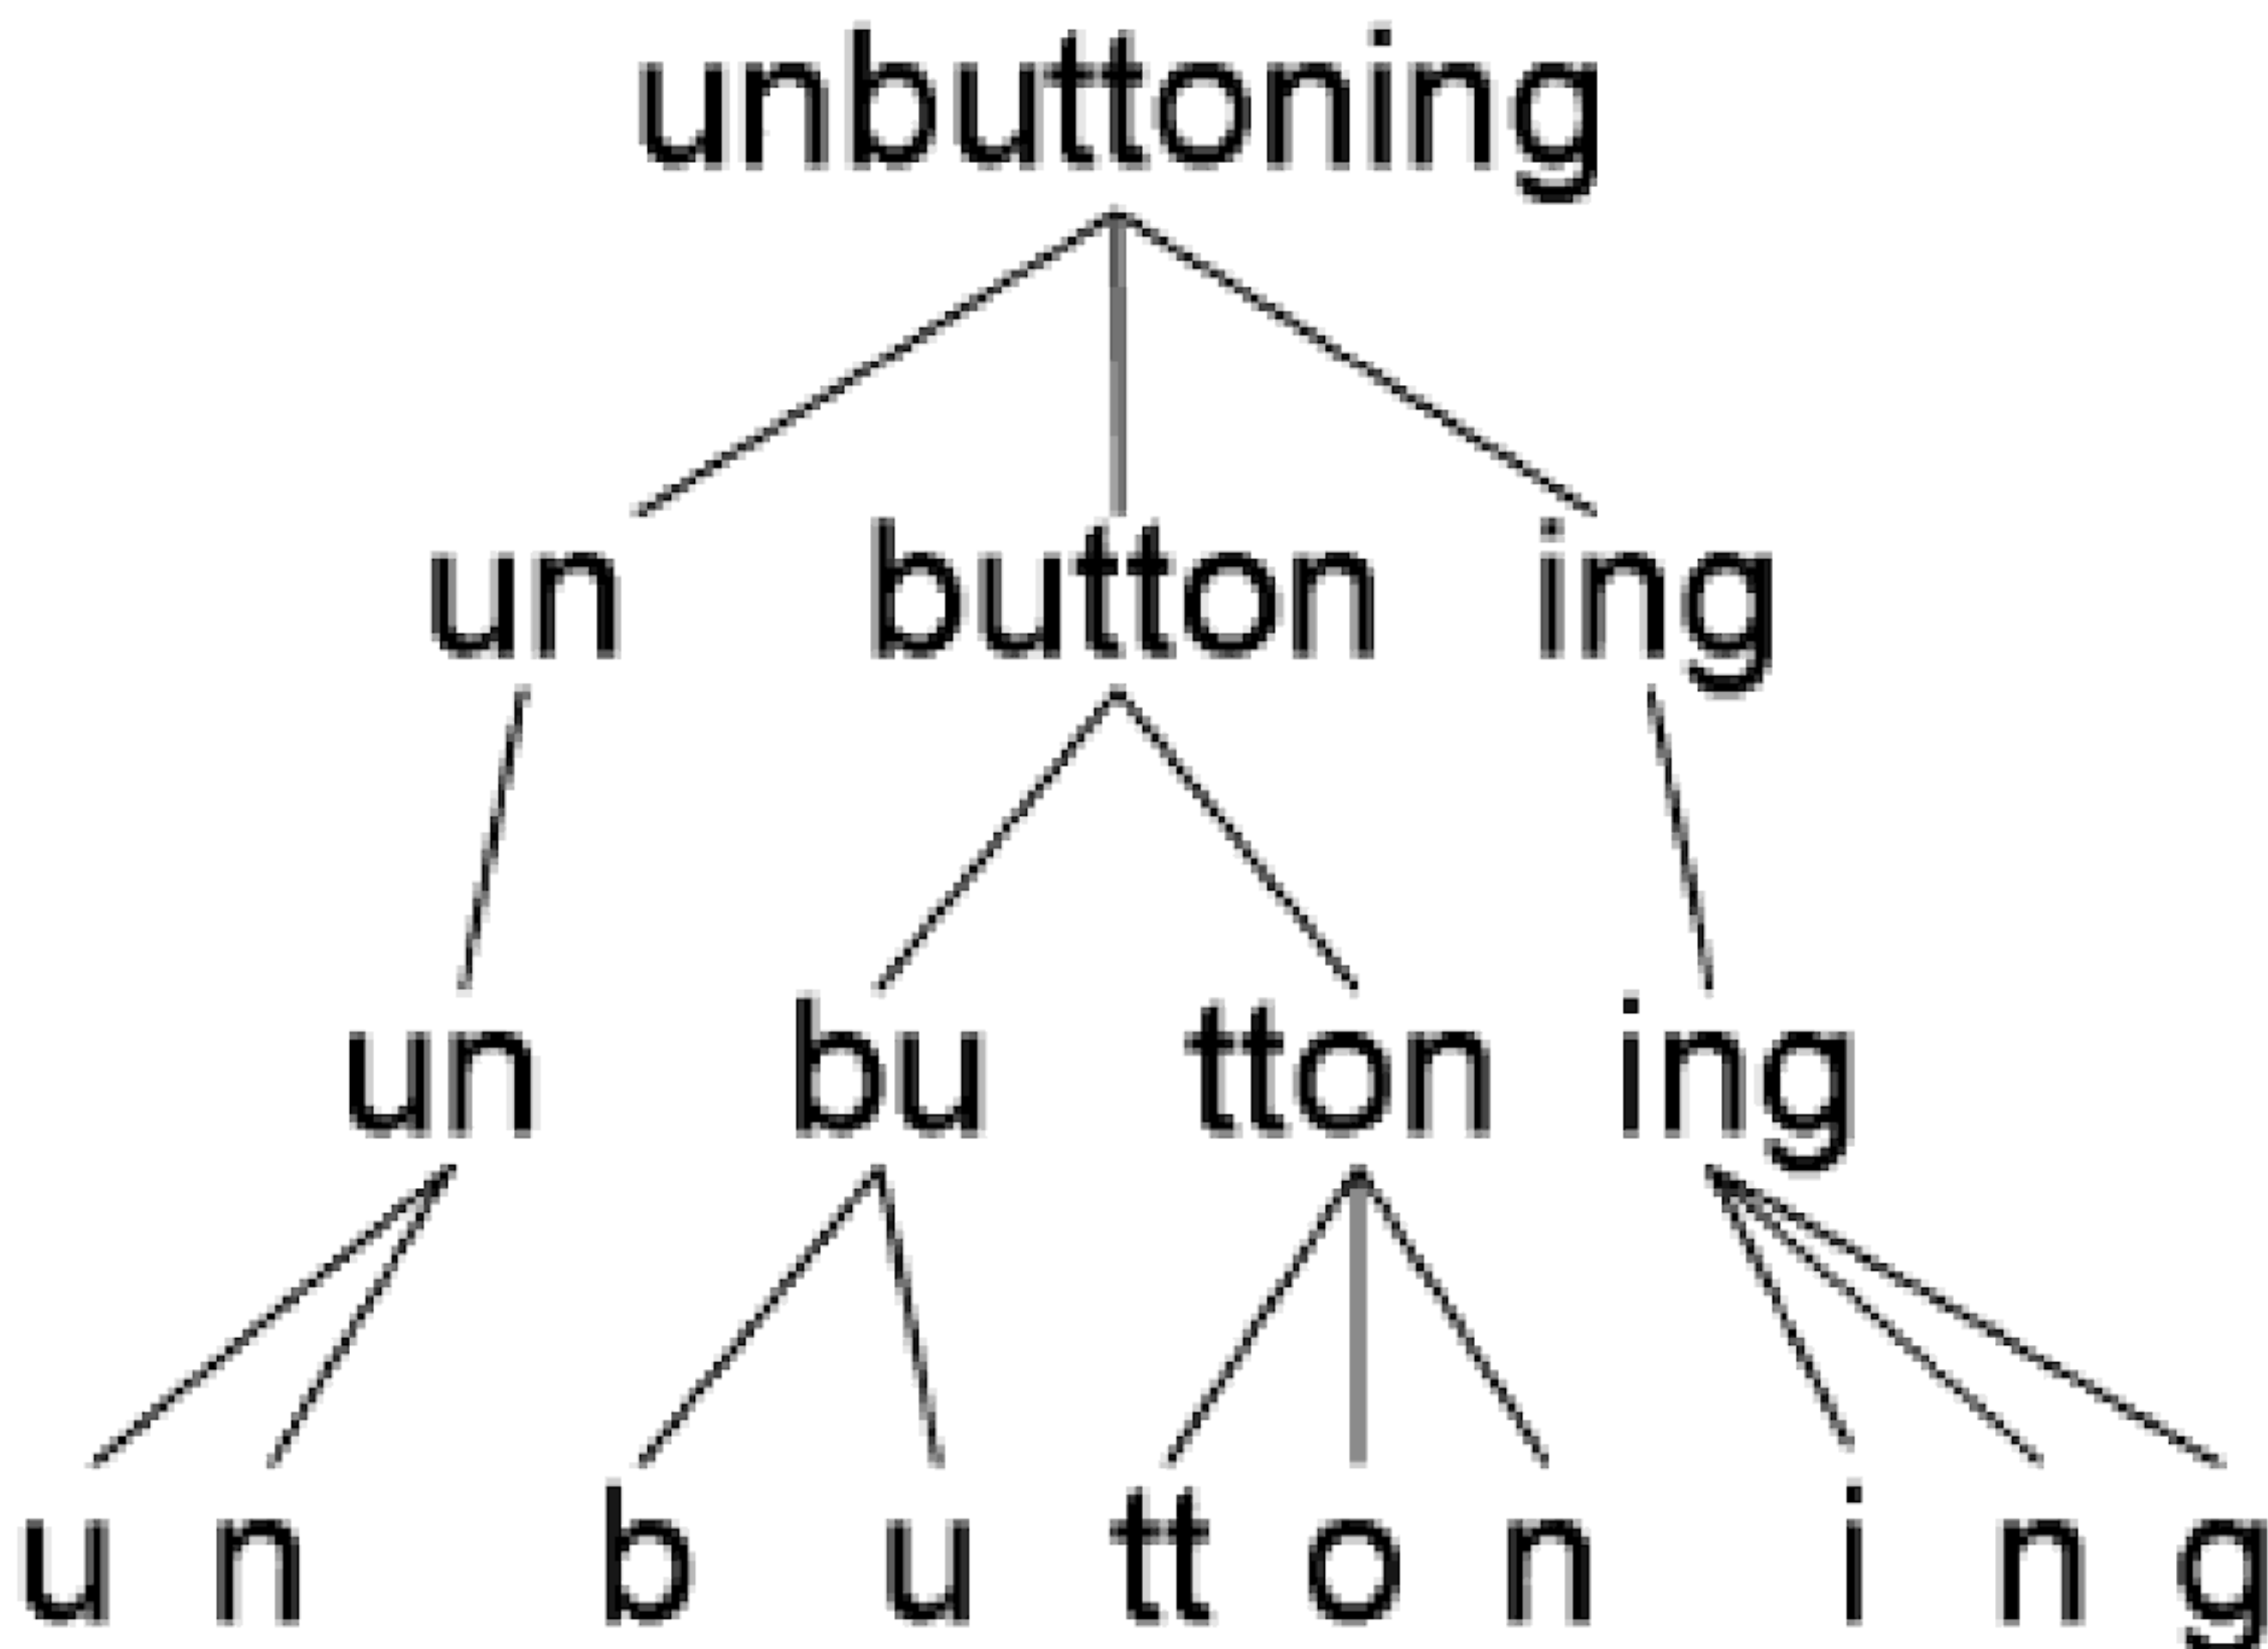
see five-letter words. Some letters are in **bold** type, others in a normal font. Concentrate solely on the middle letter, and try to decide if it is printed in normal or in bold type:

List 1: HORNY RIDER GRAVY FILET

List 2: VODKA METRO HANDY SUPER

Did you feel that the first list was slightly more difficult than the second? In the first list, the bold characters do not respect syllable boundaries—in “**RIDER**,” for instance, the “D” is printed in bold type while the rest of the syllable is in normal type. Our mind tends to group together the letters that make up a syllable, thus creating a conflict with the bold type that leads to a measurable slowing down of responses.²¹ This effect shows that our visual system cannot avoid automatically carving up words into their elementary constituents even when it would be better not to do so.

The nature of these constituents remains a hot topic in research. It would appear that multiple levels of analysis can coexist: a single letter at the lowest level, then a pair of letters (or “bigram,” an important unit to which we will return later), the grapheme, the syllable, the morpheme, and finally the whole word. The final point in visual processing leaves the word parsed out into a hierarchical structure, a tree made up of branches of increasing sizes whose leaves are the letters.



Reduced to a skeleton, stripped of all its irrelevant features like font, case, and size, the letter string is thus broken down into the elementary components that will be used by the rest of the brain to compute sound and meaning.

The Silent Voice

Writing—this ingenious art to paint words and speech for the eyes.

—GEORGES DE BRÉBEUF (FRENCH POET, 1617–1661)

When he paid a visit to Ambrose, then bishop of Milan, Augustine observed a phenomenon that he judged strange enough to be worth noting in his memoirs:

When [Ambrose] read, his eyes scanned the page and his heart sought out the meaning, but his voice was silent and his tongue was still. Anyone could approach him freely and guests were not commonly announced,

so that often, when we came to visit him, we found him reading like this in silence, for he never read aloud.²²

In the middle of the seventh century, the theologian Isidore of Seville similarly marveled that “letters have the power to convey to us silently the sayings of those who are absent.” At the time, it was customary to read Latin aloud. To articulate sounds was a social convention, but also a true necessity: confronted with pages in which the words were glued together, without spaces, in a language that they did not know well, most readers had to mumble through the texts like young children. This is why Ambrose’s silent reading was so surprising, even if for us it has become a familiar experience: we can read without articulating sounds.

Whether our mind ever goes straight from the written word to its meaning without accessing pronunciation or whether it unconsciously transforms letters into sound and then sound into meaning has been the topic of considerable discussion. The organization of the mental pathways for reading fueled a debate that divided the psychological community for over thirty years. Some thought that the transformation from print to sound was essential—written language, they argued, is just a by-product of spoken language, and we therefore have to sound the words out, through a *phonological route*, before we have any hope of recovering their meaning. For others, however, phonological recoding was just a beginner’s trait characteristic of young readers. In more expert readers, reading efficiency was based on a direct *lexical route* straight from the letter string to its meaning.

Nowadays, a consensus has emerged: in adults, both reading routes exist, and both are simultaneously active. We all enjoy direct access to word meaning, which spares us from pronouncing words mentally before we can understand them. Nevertheless, even proficient readers continue to use the sounds of words, even if they are unaware of it. Not that we articulate words covertly—we do not have to move our lips, or even prepare an intention to do so. At a deeper level, however, information about the pronunciation of words is automatically retrieved. Both the lexical and phonological pathways operate in parallel and reinforce each other.

There is abundant proof that we automatically access speech sounds while we read. Imagine, for instance, that you are presented with a list of strings and have to decide whether each one is a real English word or not. Mind you, you only have to decide if the letters spell out an English word. Here you go:

rabbit
bountery
culdolt
money
dimon
karpit
nee

You perhaps hesitated when the letters sounded like a real word—as in “demon,” “carpet,” or “knee.” This interference effect can easily be measured in terms of response times. It implies that each string is converted into a sequence of sounds that is evaluated like a real word, even though the process goes against the requested task.²³

Mental conversion into sound plays an essential role when we read a word for the first time—say, the string “Kalashnikov.” Initially, we cannot possibly access its meaning directly, since we have never seen the word spelled out. All we can do is convert it into sound, find that the sound pattern is intelligible, and, through this indirect route, come to understand the new word. Thus, sounding is often the only solution when we encounter a new word. It is also indispensable when we read misspelled words. Consider the little-known story by Edgar Allan Poe called “The Angel of the Odd.” In it, a strange character mysteriously intrudes into the narrator’s apartment, “a personage nondescript, although not altogether indescribable,” and with a German accent as thick as British fog:

“Who are you, pray?” said I, with much dignity, although somewhat puzzled; “how did you get here? And what is it you are talking about?”

“Az vor ow I com’d ere,” replied the figure, “dat iz

none of your pizzness; and as vor vat I be talking apout, I be talk apout vat I tink proper; and as vor who I be, vy dat is de very ting I com'd here for to let you zee for yourzef. . . . Look at me! Zee! I am te Angel ov te Odd.”

“And odd enough, too,” I ventured to reply; “but I was always under the impression that an angel had wings.”

“Te wing!” he cried, highly incensed, “vat I pe do mit te wing? Mein Gott! Do you take me vor a shicken?”

In reading this passage we return to a style we had long forgotten, one that dates back to our childhood: the phonological route, or the slow transformation of totally novel strings of letters into sounds that miraculously become intelligible, as though someone were whispering them at us.

What about everyday words, however, that we have already met a thousand times? We do not get the impression that we slowly decode through mental enunciation. However, clever psychological tests show that we still activate their pronunciation at a nonconscious level. For instance, suppose that you are asked to indicate which of the following words refer to parts of the human body. These are all very familiar words, so you should be able to focus on their meaning and neglect their pronunciation. Try it:

knee
leg
table
head
plane
bucket
hare

Perhaps you felt the urge to respond to the word “hare,” which sounds like a body part. Experiments show that we slow down and make mistakes on words that sound like an item in the target category.²⁴ It is not clear how we could recognize this homophony if we did not first mentally retrieve the word’s pronunciation. Only an internal conversion into speech sounds can explain this type of

error. Our brain cannot help but transform the letters “h-a-r-e” into internal speech and then associate it with a meaning—a process that can go wrong in rare cases where the string sounds like another well-known word.

Of course, this imperfect design is also what grants us one of the great pleasures of life: puns, or the “joy of text,” as the humorist Richard Lederer puts it. Without the gift of letter-to-sound conversion, we would not be able to enjoy Mae West’s wisecrack (“She’s the kind of girl who climbed the ladder of success wrong by wrong”) or Conan Doyle’s brother-in-law’s quip (“there’s no police like Holmes”). Without Augustine’s “silent voice,” the pleasure of risqué double entendres would be denied us:

An admirer says to President Lincoln, “Permit me to introduce my family. My wife, Mrs. Bates. My daughter, Miss Bates. My son, Master Bates.”

“Oh dear!” replied the president.²⁵

Further proof that our brain automatically accesses a word’s sound patterns derives from subliminal priming. Suppose that I flash the word “LATE” at you, immediately followed by the word “mate,” and ask you to read the second word as fast as you can. The words are shown in a different case in order to avoid any low-level visual resemblance. Nevertheless, when the first word sounds and spells like the second, as in this example, we would observe a massive acceleration of reading time, in comparison with a situation where two words are not particularly related to one another (“BOWL” followed by “mate”). Part of this facilitation clearly arises from similarity at the level of spelling alone. To flash “MATH” eases recognition of “mate,” even though the two strings sound quite different. However, crucially, even greater facilitation can be found when two words share the same pronunciation (“LATE” followed by “mate”), and this sound-based priming works even when spelling is completely different (“EIGHT” followed by “mate”). Thus, pronunciation seems to be automatically extracted. As one might expect, however, spelling and sound are not encoded at the same time. It takes only twenty or thirty milliseconds of word viewing for our brain to automatically activate

“choir” and “chair,” “friend” and “fiend.” For most irregular words, the recovery of pronunciation, far from being the source of word comprehension, seems to depend on its outcome: it is only after we have recognized the identity of the word “dough” that we can recover its sound pattern.

The Hidden Logic of Our Spelling System

One may wonder why English sticks to such a complicated spelling system. Indeed, Italians do not meet with the same problems. Their spelling is transparent: every letter maps onto a single phoneme, with virtually no exceptions. As a result, it only takes a few months to learn to read. This gives Italians an enormous advantage: their children’s reading skills surpass ours by several years, and they do not need to spend hours of schooling a week on dictation and spelling out loud. Furthermore, as we shall discuss later, dyslexia is a much less serious problem for them. Perhaps we should follow Italy’s lead, burn all our dictionaries and design a new spelling system that even a three-year-old child could easily read.

There is no doubt that English spelling could be simplified. The weight of history explains a lot of its peculiarities—today’s pupils should lament the loss of the battle of Hastings, because the mixture of French and English that ensued is responsible for many of our spelling headaches, such as the use of the letter “c” for the sound *s* (as in “cinder”). Centuries of academic conservatism, sometimes bordering on pedantry, have frozen our dictionary. Well-meaning academics even introduced spelling absurdities such as the “s” in the word “island,” a misguided Renaissance attempt to restore the etymology of the Latin word *insula*. Worst of all, English spelling failed to evolve in spite of the natural drift of oral language. The introduction of foreign words and spontaneous shifts in English articulation have created an immense gap between the way we write and the way we speak, which causes years of unnecessary suffering for our children. In brief, reason

calls for a radical simplification of English spelling.

Nevertheless, before any revisions can be made, it is essential to fully understand the hidden logic of our spelling system.

Spelling irregularities are not just a matter of convention. They also originate in the very structure of our language and of our brains. The two reading routes, either from spelling to sound or from spelling to meaning, place complex and often irreconcilable constraints on any writing system. The linguistic differences between English, Italian, French, and Chinese are such that no single spelling solution could ever suit them all. Thus the abominable irregularity of English spelling appears inevitable. Although spelling reform is badly needed, it will have to struggle with a great many restrictions.

First of all, it is not clear that English spelling, like Italian, could attribute a single letter to each sound, and a fixed sound to each letter. It would not be a simple thing to do because the English language contains many more speech sounds than Italian. The number of English phonemes ranges from forty to forty-five, depending on speakers and counting methods, while Italian has only thirty. Vowels and diphthongs are particularly abundant in English: there are six simple vowels (as in *bat*, *bet*, *bit*, *but*, *good*, and *pot*), but also five long vowels (as in *beef*, *boot*, *bird*, *bard*, and *boat*) and at least seven diphthongs (as in *bay*, *boy*, *toe*, *buy*, *cow*, *beer*, *bear*). If each of those sounds were granted its own written symbol, we would have to invent new letters, placing an additional burden on our children. We could consider the addition of accents to existing letters, such as ã, õ, or ü. However, it is entirely utopian to imagine a universal alphabet that could transcribe all of the world's languages. Such a spelling system does exist: it is called the International Phonetic Alphabet and it plays an important role in technical publications by phonologists and linguists. However, this writing system is so complex that it would be ineffective in everyday life. The International Phonetic Alphabet has 170 signs, some of which are particularly complex (e e, b, ɔ, ʒ, or ɲ). Even specialists find it very hard to read fluently without the help of a dictionary.

In order to avoid learning an excessive number of symbol shapes, languages with a great many phonemes, such as English

and French, all resort to a compromise. They indicate certain vowels or consonants using either special characters such as ü, or groups of letters like “oo” or “oy.” Such peculiarities, which are unique to any given language, are far from being gratuitous embellishments: they play an essential role in the “mental economy” of reading, and have to find their place in any kind of spelling reform.

Although we cannot easily assign a single letter shape to each speech sound, we could perhaps try the opposite. Many spelling errors could be avoided if we systematically transcribed each sound with a fixed letter. For instance, if we were to avoid writing the sound *f* with both the letter “f” and with “ph,” life would be much simpler. There is little doubt that we could easily get rid of this and many other useless redundancies whose acquisition eats up many years of childhood. In fact, this is the timid direction that American spelling reform took when it simplified the irregular British spellings of “behaviour” or “analyse” into “behavior” and “analyze.” Many more steps could have been taken along the same lines. As expert readers, we cease to be aware of the absurdity of our spelling. Even a letter as simple as “x” is unnecessary, as it stands for two phonemes *ks* that already have their own spelling. In Turkey, one takes a “taksi.” That country, which in the space of one year (1928–29) adopted the Roman alphabet, drastically simplified its spelling, and taught three million people how to read, sets a beautiful example of the feasibility of spelling reform.

Yet here again, great caution is needed. I suspect that any radical reform whose aim would be to ensure a clear, one-to-one transcription of English speech would be bound to fail, because the role of spelling is *not* just to provide a faithful transcription of speech sounds. Voltaire was mistaken when he stated, elegantly but erroneously, that “writing is the painting of the voice: the more it bears a resemblance, the better.” A written text is not a high-fidelity recording. Its goal is not to reproduce speech as we pronounce it, but rather to code it at a level abstract enough to allow the reader to quickly retrieve its meaning.

For the sake of argument, we can try to imagine what a purely phonetic writing system would look like—one that Voltaire might

have considered ideal. When we speak, we alter the pronunciation of words as a function of the sounds that surround them. It would be disastrous if spelling were to reflect the obtuse linguistic phenomena of so-called coarticulation, assimilation, and resyllabification, of which most speakers are usually unaware. A matter of context would end up having the same word spelled differently. Should we, for instance, use distinct marks for the various pronunciations of plurals? Should we spell “cap driver,” under the pretext that the sound *b*, when followed by a *d*, tends to be pronounced like a *p*? At one extreme, should we factor in the speaker’s accent (“Do you take me vor a shicken?”). This would be *apsurd* (yes, we do pronounce this word with a *p* sound). The prime goal of writing is to transmit meaning as efficiently as possible. Any servile transcription of sounds would detract from this aim.

English spelling often privileges the transparency of word roots at the expense of the regularity of sounds. The words “insane” and “insanity,” for instance, are so deeply related to their meaning that it would be silly to spell them differently because of their slightly different pronunciation. Similarly, it is logical to maintain the silent *n* at the end of “column,” “autumn,” or “condemn,” given that these words give rise to “columnist,” “autumnal,” or “condemnation.”

Transcription of meaning also explains, at least in part, why English spells the same sounds in many different ways. English words tend to be compact and monosyllabic, and as a result, homophony is very frequent (for instance “eye” and “I,” “you” and “ewe”). If these words were transcribed phonetically, they could not be distinguished from each other. Spelling conventions have evolved with this constraint in mind. Distinctive spelling for the same sounds complicates dictation, but simplifies the task for the reader, who can quickly grasp the intended meaning. Students who complain about the countless forms of spelling the sound *u* as in “two,” “too,” “to,” or “stew” should understand that these embellishments are essential to the speed at which we read. Without them, any written text would become an opaque rebus. Thanks to spelling conventions, written English points straight at meaning. Any spelling reform would have to maintain this subtle equilibrium between sound and meaning, because this balance

reflects a much deeper and more rigid phenomenon: our brain's two reading routes.

The Impossible Dream of Transparent Spelling

Rivalry between reading for sound and reading for meaning is true the world over. All writing systems must somehow manage to address this problem. Which compromise is best depends on the language to be transcribed. Life would certainly be easier if English spelling were as easy to learn as Italian or German. These languages, however, benefit from a number of characteristics that make them easy to transcribe into writing. In Italian as in German, words tend to be long, often made up of several syllables. Grammatical agreement is well marked by resonant vowels. As a result, homonyms are quite rare. Thus, a purely regular transcription of sounds is feasible. Italian and German can afford a fairly transparent spelling system, where almost every letter corresponds to a unique sound.

At the other end of the continuum, there is the case of Mandarin Chinese. The vast majority of Chinese words consist of only one or two syllables, and because there are only 1,239 syllables (410 if one discounts tonal changes), each one can refer to dozens of different concepts ([figure 1.3](#)). Thus a purely phonetic writing system would be useless in Chinese—each of the rebuses could be understood in hundreds of different ways! This is why the thousands of characters in Mandarin script mostly transcribe words, or rather their morphemes—the basic elements of word meaning. Chinese writing also relies on several hundred phonetic markers that further specify how a given root should be pronounced and make it easier for the reader to figure out which word is intended. The character 媽, for instance, which means “mother” and is pronounced *ma*, consists of the morpheme 女 = “woman,” plus a phonetic marker 馬 = *mǎ*. Thus, contrary to common belief, even Chinese is not a purely ideographic script (whose symbols represent concepts), nor a logographic one

Figure 1.3 Spelling irregularities are not as irrational as they seem. For instance, although Chinese writing uses up to twenty or thirty different characters for the same syllable, the redundancy is far from pointless. On the contrary, it is very helpful to Chinese readers because the Chinese language includes a great many homophones—words that sound alike but have different meanings, like the English “one” and “won.” Here, an entire Chinese story was written with the sound “shi”! Any Chinese reader can understand this text, which would clearly be impossible if it were transcribed phonetically as “shi shi shi . . .” Chinese characters disambiguate sounds by using distinct characters for distinct meanings. Similarly, homophony explains why English sticks to so many different spellings for the same sounds (“I scream for ice cream”).

The European Union commissioners have announced that agreement has been reached to adopt English as the preferred language for European communications, rather than German, which was the other possibility. As part of the negotiations, the British government conceded that English spelling had some room for improvement and has accepted a five-year phased plan for what will be known as Euro English (Euro for short).

In the first year, “s” will be used instead of the soft “c.” Certainly, sivil servants will resieve this news with joy. Also, the hard “c” will be replaced with “k.” Not only will this klear up konfusion, but typewriters kan have one less letter. There will be growing publik enthusiasm in the sekond year, when the troublesome “ph” will be replaced by “f.” This will make words like “fotograf” 20 per sent shorter.

In the third year, publik akseptanse of the new spelling kan be expekted to reach the stage where more komplikated changes are possible. Governments will enkorage the removal of double letters, which have always ben a deterent to akurate speling. Also, al wil agre that the horrible mes of silent “e”s in the languag is disgrasful, and they would go.

By the fourth year, peopl wil be reseptiv to steps such as replasing “th” by “z” and “w” by “v.” During ze fifz year, ze unesesary “o” kan be dropd from vords kontaining “ou,” and similar changes vud of kors be aplid to ozer kombinations of leters.

After zis fifz yer, ve vil hav a reli sensibl riten styl. Zer vil be no mor trubls or difikultis and evrivun vil find it ezi tu understand ech ozer.

Ze drem vil finali kum tru!

Two Routes for Reading

Before I go any further, I would like to summarize what has been

covered so far. All writing systems oscillate between an accurate representation of sound and the fast transmission of meaning. This dilemma is directly reflected in the reader's brain. Two information processing pathways coexist and supplement each other while we read. When words are very regular, rare, or novel, we preferentially process them using a "phonological route," in which we first decipher the letter string, then convert it into pronunciation, and finally attempt to access the meaning of the sound pattern (if any). Conversely, when we are confronted with words that are frequent or whose pronunciation is exceptional, our reading takes a direct route that first recovers the identity and meaning of the word and then uses the lexical information to recover its pronunciation (figure 1.4).

The best support for the existence of these two pathways comes from the study of brain injuries and their psychological consequences. Some patients, after a stroke or a brain lesion, lose the ability to quickly compute the pronunciation of written words.³⁰ Clearly, their spelling-to-sound conversion route has been severely damaged. Although they read normally before the lesion, their reading now presents all the characteristics of a syndrome called deep dyslexia or phonological dyslexia. They no longer manage to read infrequent words like "sextant" out loud, even when their spelling is very regular. Furthermore, they are helpless when faced with reading neologisms or invented words such as "departition" or "calbonter." Surprisingly, however, they may still be able to understand words used frequently, and generally manage to read aloud even irregular but common words like "eyes," "door," or "women." Occasionally, they confuse one word with another. A deep dyslexic may, for instance, read the written word "ham" as *meat* or the word "painter" as *artist*. The very nature of these errors demonstrates that access to word meaning is largely preserved. If the patient had not understood, at least in part, the word that he attempted to read, he would not even have been able to approach the right meaning. Patients with deep dyslexia appear to recognize written words even though the computation of their pronunciation has become essentially impossible. It is as if one of the reading routes (from vision to sound) was blocked, while information still circulated along the

other route (from vision to meaning).

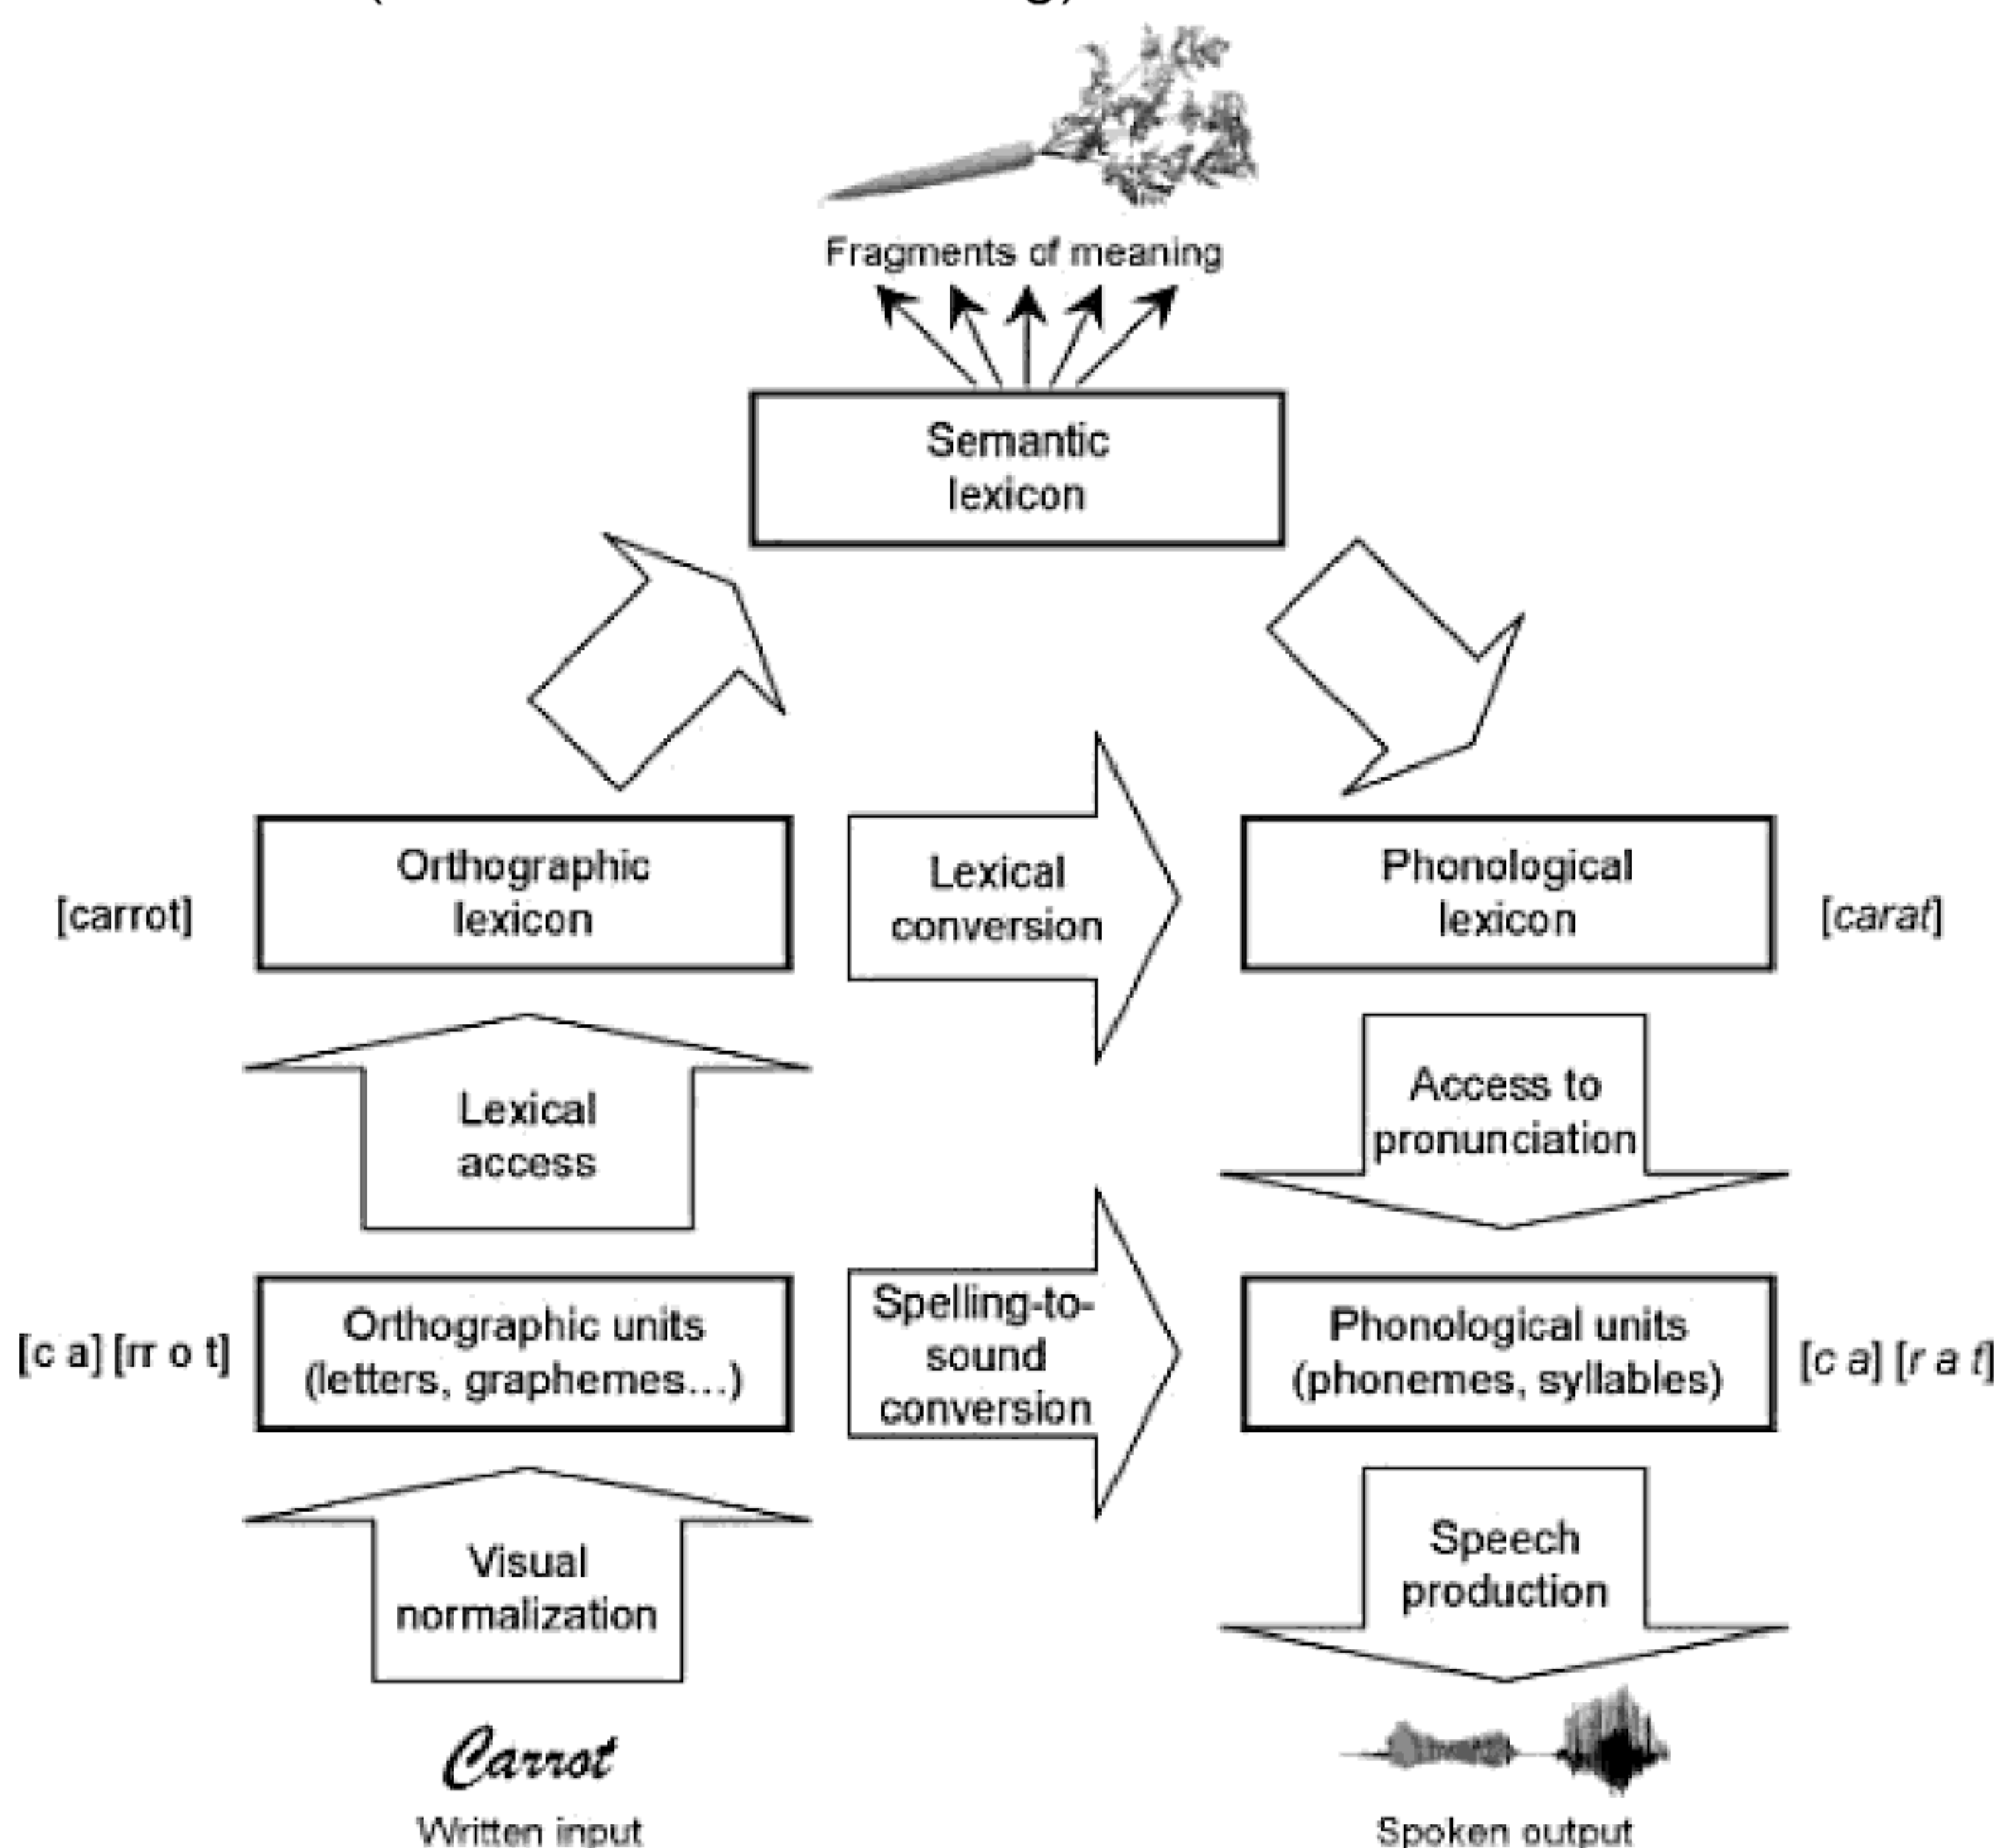


Figure 1.4 Word reading proceeds along several parallel processing routes. To move from the written word (bottom left) to its pronunciation (bottom right), our brain relies on several pathways, indicated here with boxes and arrows. When the word is regular, a surface route directly converts its letters into speech sounds. When the word is irregular, such as “carrot,” deeper representations are involved. They can be compared to mental lexicons that attach word form to meaning.

The opposite situation is also on record. Here, patients who suffer from another syndrome, called surface dyslexia, no longer enjoy direct access to word meaning, but have to slowly work their way through a text and sound out the words. In this situation, the limits of the “silent voice” become obvious. Surface dyslexia patients can still read regular words such as “banana,” or even neologisms such as “chicopar,” but they are virtually helpless with

irregular words. They typically standardize them through blind conversion into sounds. One patient, for instance, read “enough” as *inog* and then swore that he had never heard of this bizarre word. Manifestly, the direct route from vision to the mental lexicon had been blocked, while the conversion from print to sound remained in operation.

While the contrast between these two types of patients proves that there are two quite distinct pathways for reading, it also shows that neither of them, alone, allows us to read all words. The direct route, from letters to words and their meaning, can be used to read most words, as long as they are sufficiently frequent, but it does not work for rare or novel words that are not stored in our mental lexicon. Conversely, the indirect route, from letters to sounds and from sounds to meanings, fails with irregular words like “women” and with homophones such as “too,” but it plays an irreplaceable role when we learn new words. When we read, both routes are in constant collaboration and each contributes to the specification of word pronunciation. Most of a word’s phonemes can be deduced from a letter string through the application of simple spelling-to-sound conversion rules, while occasional ambiguities are resolved by a little nudge from higher lexical and semantic levels. In children, the two routes are often poorly coordinated. Some children rely largely on the direct route. They attempt to second-guess words and often pick a loose synonym for the target (for instance, reading “house” as *home*). Others mumble their way through a sentence and painfully put together an approximate pronunciation based on the words’ letters, but cannot go from the loosely assembled sounds to their meaning. Years of practice are needed before the two routes become close enough to give the impression, in expert adults, of a single integrated reading system.

Most of the current models of reading postulate that fluent reading relies on the close-knit coordination between the two reading routes, where each weighs in to a different degree depending on the word to be read (known or unknown, frequent or rare, regular or irregular) as well as the task (reading aloud or text comprehension). In the 1980s and 1990s, some researchers did try to account for these observations with a single-route system.

At the time, the emergence of neural network models generated a huge wave of interest. Some researchers saw them as universal learning machines that could acquire any skill without the possession of predefined cognitive architecture. They figured that reading acquisition could be modeled by the connection of letter inputs to phonological outputs, while the intermediate connections were tuned by a powerful learning rule. They hoped to arrive at a single network that could simulate normal reading and its pathologies, but did not have to postulate multiple cortical processing pathways. While these networks, at the time, represented a remarkable advance, particularly in the modeling of spelling-to-sound conversion,³¹ most researchers today believe that this approach is inadequate. My own impression is that it is impossible to model reading without a thorough analysis of the brain's architecture, which relies on multiple parallel and partially redundant pathways. Practically all recent models, even though they rely on neural network simulations, incorporate the essential idea of multiple routes for reading.³² When we tackle the cerebral mechanisms for reading later in this book, we will see that an essential feature of cortical architecture is its organization into multiple parallel paths. Thus, even the two-route model probably underestimates the true complexity of neural reading systems. The separation into two pathways, a spelling-to-sound conversion route and a semantic route, is simply a useful approximation.

Mental Dictionaries

As long as we were dealing with the surface route that converts graphemes into phonemes, it was realistic to imagine that mental reading followed a short list of simple procedures. To store a map of the few hundred English graphemes and their pronunciation would seem to suffice. However, when one considers how the deep route, which recognizes thousands of familiar words, may work, a much more massive store is needed. Cognitive psychologists compare it to a dictionary or "mental lexicon." No doubt we should speak of mental dictionaries in the plural, because our heads, in fact, carry around a number of different

succumb to the pitfall of postulating a homunculus, or the little man who according to folk psychology holds the reins of our brain. (Who controls his brain? Another even tinier homunculus?) In this respect, the pandemonium model can be compared to the philosopher Dan Dennett's motto: "One discharges fancy homunculi from one's scheme by organizing armies of such idiots to do the work."³⁴

- **COMPETITION AND ROBUSTNESS:** Daemons fight for the right to represent the correct word. This competition process yields both flexibility and robustness. The pandemonium automatically adapts to the complexity of the task at hand. When there are no other competitors around, even a rare and misspelled word like "astrqlabe" can be recognized very quickly—the daemon that represents it, even if it initially shouts softly, always ends up beating all the others by a comfortable margin. If, however, the stimulus is a word such as "lead," many daemons will activate (those for "bead," "head," "read," "lean," "leaf," "lend" . . .) and there will be a fierce argument before the "lead" daemon manages to take over.

All of these properties, in simplified form, fit with the main characteristics of our nervous system. Composed of close to one hundred billion (10^{11}) cells, the human brain is the archetype of a massively parallel system where all neurons compute simultaneously. The connections that link them, called synapses, bring them evidence from the external sensory stimulus. Furthermore, some of these synapses are inhibitory, which means that when the source neuron fires, the firing of other neurons is suppressed. The result has been likened by the Canadian neurophysiologist Donald Hebb to a network of "cell assemblies," coalitions of neurons that constantly compete. It is therefore no surprise that Selfridge's pandemonium has been a source of inspiration for many theoretical models of the nervous system, including the first neural network models of reading. [Figure 1.5](#) shows one of the very first such models, which was introduced by Jay McClelland and David Rumelhart in 1981.³⁵ The model includes three hierarchical layers of neuron-like units.




Figure 1.5 The word identification process is similar to a vast assembly where thousands of letter and word units conspire to provide the best possible interpretation of the input string. In McClelland and Rumelhart's model, of which only a fragment is shown here, basic features of the input string activate letter detectors, which in turn preferentially connect to detectors of the words that contain them. The links can be excitatory (arrows) or inhibitory (lines ending with discs). A fierce competition between lexical units finally identifies a dominant word that represents the network's preferred hypothesis about the incoming string.

- At the bottom, input units are sensitive to line segments presented on the retina.
- In the middle lie letter-detector units that fire whenever a given letter is present.
- At the top, units code for entire words.

All of these units are tightly linked by a swarm of connections. This enormous connectivity turns the network dynamics into a complex political game in which letters and words support, censor,

or eliminate each other. If you study the graph carefully, you will see excitatory connections, represented by small arrows, as well as inhibitory connections, represented by small circles. Their role is to propagate the votes of each of the daemons. Each input detector, which codes for a specific feature such as a vertical bar, sends stimulation to all the letters that contain this particular feature—one might say, for the sake of simplicity, that each visual neuron “votes” for the presence of these letters. At the next level, similarly, letter detectors conspire to elect specific words through stimulation of their corresponding units. The presence of letters “A” and “N,” for instance, supports the words “RAIN” and “TANK,” but only partially argues for the word “RAIL” and not at all for “PEST.”

Inhibition also contributes toward selection of the best candidate. Thanks to inhibitory connections, letters can vote against the words that do not contain them. For instance, the unit that codes for letter “N” votes against the word “RAIL” by inhibiting it. Furthermore, words that compete inhibit each other. Thus identification of the word “RAIN” is incompatible with the presence of the word “RAIL,” and vice versa.

Finally, it is useful to incorporate top-down connections, from words to their component letters. This process can be compared to a senate where letters are represented by words that, in return, support the letters that voted for them. Reciprocal connections allow for the creation of stable coalitions that can resist an occasional missing letter: if one letter “o” is lacking in the word “crocodile,” for instance, its neighbors will still conspire to elect the word “crocodile,” and in turn the latter will vote for the presence of a middle letter “o” that is not physically present. In the end, millions of connections are needed to incorporate the numerous statistical constraints that link the levels of words, letters, and features.

Other subtleties also allow for the whole network to operate smoothly. For instance, word units can have different thresholds for firing. A word encountered frequently has a lower threshold than a rare word, and with an equal amount of bottom-up support has a better chance to win the race. The most recent models also incorporate additions such as a fine-grained coding of letter

position. The resulting network has such a complex set of dynamics that it is impossible to fully describe it mathematically. One has to resort to computer simulations in order to determine how long the system takes to converge to the correct word, and how often it misidentifies it.

Parallel Reading

If cognitive scientists burden themselves with all these complex models of reading, it is because their predictions fit remarkably well with empirical measurements. Not only do the models inspired by Selfridge's pandemonium manage to reproduce the results of classic experiments on reading speed and errors, but they often lead to the discovery of subtle novel phenomena that constitute the essential properties of human reading behavior.

If you were asked to design software for written word recognition, it is probable, regardless of the solution you choose, that your program would be increasingly slow as words get longer. It would be natural, for instance, for the software to process letters consecutively, from left to right. Since instructions are processed serially, one would expect a six-letter word to take about twice as long as a three-letter word. In any serial model, recognition time should increase in direct relation to the number of letters.

In light of this example, it is rather remarkable that the relation between reading time and number of letters does not hold for the human brain. In adult expert readers, the time to read a word is essentially independent of its length. As long as the word does not have more than six or seven letters, its recognition takes an approximately constant amount of time, regardless of length.³⁶ Obviously, this implies that our brain relies on a parallel letter-crunching mechanism that is able to process all letters at once. This empirical result is incompatible with the metaphor of a computer scanner, but it conforms precisely to the pandemonium hypothesis, whereby millions of specialized processors operate simultaneously and in parallel at each of several levels (features, letters, and words).

Active Letter Decoding

Let us pursue this computer metaphor. In a classic computer program, information is typically processed through a series of steps, from the simplest operations to the most abstract ones. It is logical to begin with a subroutine that recognizes individual letters, followed by one that groups them into graphemes, and finally a third that examines potential words. This type of program, however, is typically highly intolerant of errors. Any failure at the first stage usually results in a complete breakdown of the entire recognition process. Indeed, even the best automated character-recognition software, which is delivered with present-day scanners, remains highly sensitive to any decrease in the quality of the image—a few specks of dust on the scanning window are all it takes to turn a page that humans find perfectly readable into garbled text that the computer considers gibberish.

Our visual system, unlike the computer, revels in the resolution of ambiguities. You can experiment on this for yourself by reading the following sentence:

Honey bccs sovovr sweet ncctar

Unbeknownst to you, your eye just skimmed effortlessly across a series of difficulties that would have confounded any classic computer program. Did you notice that in the word “bee” the repeated letter is actually a “c”? In fact, the very same letter shapes are used to spell the word “nectar,” but your visual system treats the first one as an “e” and the second as a “c.” The word “savour” is worse. The letters “a” and “o” also have exactly the same shape—as do the letters “v” and “u”! Such ambiguities are resolved by the context: the string “souour” would not be legal in English, whereas its interpretation as the verb “savour” fits nicely within the rest of the sentence.

In brief, ambiguities that would create a hitch in any current software cannot even be perceived by a human reader. This resistance to degradation, while hardly compatible with a classic computer program, fits nicely into the pandemonium framework,

with its many neighbors “bare,” “care,” “dare,” “fare,” “mare,” “pare,” “rare,” “here,” “hire,” “hale,” “hate,” “have,” “hard,” “harm,” and “hart,” while the word “idea” is a lonely hermit with no neighbors. Experiments show that the number and especially the relative frequency of a word’s neighbors play a crucial role in the time it takes us to recognize it.⁴⁰

In many cases, the presence of a number of neighbors helps. The more neighbors a word has, the faster we can tell that it belongs to the English lexicon. The presence of neighbors makes the word “hare” more typical of English spelling than the word “idea.” Dense neighborhoods also provide better learning opportunities. We have many more occasions to grasp the pronunciation of words that end in “-are” than of the only one that ends with “-dea.” This orphan is therefore less optimally coded at both the visual and phonological levels.

But too many neighbors can also be intrusive and annoying. To understand or name a word often requires unequivocal identification and hence separation from all the neighbors—an operation that can be particularly slow and difficult if the neighbors are more common and thus perform better in the lexical competition. Thus the word “hare” is named relatively slowly because it competes with very frequently used words like “have” and “hard.”⁴¹ Our lexicon is an arena where competition is fierce and where frequent words enjoy a significant advantage over their lower-ranking neighbors.

Competition also occurs within the spelling-to-sound conversion route. For example, it takes longer to read the word “beach” out loud than the word “black.” In “beach,” the input string must be parsed into the complex graphemes “ea” and “ch,” whose pronunciation departs radically from that of the single letters “e,” “a,” “c,” or “h.” The word “black,” on the other hand, is simpler inasmuch as each of its letters maps directly onto its accepted sound. Experimental psychology easily demonstrates this latent conflict between the letter level and the grapheme level. During the naming of words made up of complex graphemes, there is a brief period of unconscious competition that is quickly resolved but leads to a measurable deceleration of responses compared with the naming of more transparent words like “black.”⁴²

It is noteworthy that most of these conflicts are resolved automatically, without our conscious intervention. When our nervous system is confronted with ambiguity, its fundamental strategy is to leave all possibilities open—something that is only feasible in a massively parallel system where multiple interpretations can be simultaneously entertained. Thanks to this open organization, the subsequent levels of analysis can contribute their own pieces of evidence until a globally satisfactory solution is reached. In some cases, only the context in a sentence allows one to understand the meaning of a word or even its pronunciation—think of a sentence like “the road **winds** through a valley battered by fierce **winds**.” In such cases, experiments show that all of the possible interpretations of the word are unconsciously activated, until the context restricts interpretation down to a single meaning.⁴³ Our reading processes are so efficient that we are hardly ever aware of the ambiguities—unless they are very funny, as in the case of Dorothy Parker, who telegraphed her agent because she had failed to meet a deadline: “Tell the editor I’ve been too fucking busy—or vice versa.”

From Behavior to Brain Mechanisms

Throughout this chapter, we examined how reading can be broken down into a sequence of information processing stages. From image processing in the retina to invariant letter recognition, access to pronunciation, morpheme recognition, all the way up to conflict resolution within the lexicon, the efficiency of the human mechanisms for reading is impressive. In a fraction of a second, with no apparent effort, our brain resolves a visual recognition problem that still defies current computer software. The parallel processing of all a word’s characters, the clever resolution of ambiguities, the immediate access to one out of perhaps 50,000 words in the mental lexicon, all point to the remarkable adaptation of our brain to the task of reading.

The amazing efficiency of our reading process only serves to thicken the mystery surrounding its origins. How can our brain be so well adapted to a problem for which it could not possibly have

evolved? How can the brain architecture of a strange bipedal primate turned hunter-gatherer have adjusted so perfectly, and in only a few thousand years, to the challenges of visual word recognition? To clarify this problem, we will now turn to the cerebral circuits for reading. An amazing recent discovery shows that there is a specific cortical area for written words, much like the primary auditory area or the motor cortex that exist in all our brains. Even more surprisingly perhaps, this reading area seems to be identical in readers of English, Japanese, and Italian. Does this mean that there are universal brain mechanisms for reading?

The Brain's Letterbox

In 1892, the French neurologist Joseph-Jules Déjerine discovered that a stroke affecting a small sector of the brain's left visual system led to a complete and selective disruption of reading. Modern brain imaging confirms that this region plays such an essential part in reading that it can aptly be called "the brain's letterbox." Located in the same brain area in readers the world over, it responds automatically to written words. In less than one-fifth of a second, a time span too brief for conscious perception, it extracts the identity of a letter string regardless of superficial changes in letter size, shape, or position. It then transmits this information to two major sets of brain areas, distributed in the temporal and frontal lobes, that respectively encode sound pattern and meaning.

We are absurdly accustomed to the miracle of a few written signs being able to contain immortal imagery, involutions of thought, new worlds with live people, speaking, weeping, laughing. . . . What if we awake one day, all of us, and find ourselves utterly unable to read?

—VLADIMIR NABOKOV, *PALE FIRE*

Joseph-Jules Déjerine's Discovery

On a beautiful Paris morning in October 1887, Mr. C, a well-read retired salesman and music lover, was comfortably ensconced in an armchair reading a good book when he suddenly found, to his dismay, that he could no longer read a word! For several days previously, he had occasionally felt weakness or numbness in his right arm or leg and the odd mild difficulty in talking. These brief spells of discomfort, however, had only been very fleeting and had

not worried him unduly. Now the problem was far more severe: reading was just impossible! Nonetheless, Mr. C could still speak, recognize objects and the people around him, and even jot down a few notes. It was hard to imagine what could have provoked this frustrating situation.

Convinced that a new pair of eyeglasses would solve his problem, Mr. C consulted the well-known ophthalmologist Edmund Landolt. Unfortunately, Dr. Landolt quickly realized that Mr. C's affliction required more than spectacles. He suspected a neurological problem and decided to enlist the help of the eminent French neurologist Joseph-Jules Déjerine at the Bicêtre Hospital. Mr. C's appointment took place on November 15, 1887. After a remarkably thorough psychological and anatomical examination, Dr. Déjerine was able to make a diagnosis and draw the first scientific conclusions ever made about the cerebral bases of reading.⁴⁴ He termed Mr. C's impairment "pure verbal blindness," meaning a selective loss of the visual recognition of letter strings. That such a disease existed implied the presence in the brain of a cortical "visual center for letters" specialized in reading. The first hint of brain specialization for reading had thus been scientifically demonstrated.

In fact, what Déjerine and Landolt established was that their patient no longer recognized individual letters or a written word. Confronted with an array of letters,

he thinks that he has lost his mind, because he understands that the signs which he fails to recognize are letters, he insists that he can see them perfectly and outlines their shape with his hand but is incapable of naming them. If asked to copy what he sees, he succeeds with great effort by slowly copying each letter one stroke at a time, as though he were working on a technical draft, and by checking each and every curve to ensure the exactness of his drawing. In spite of these efforts, he remains unable to name the letters.

This is the paradox of "verbal blindness": the patient is only blind to letters and words. Visual acuity remains excellent, objects