



# UNDERSTANDING BELIEFS

NILS J. NILSSON



THE MIT PRESS ESSENTIAL KNOWLEDGE SERIES

# UNDERSTANDING BELIEFS



---

**The MIT Press Essential Knowledge Series**

*Understanding Beliefs*, Nils J. Nilsson

*Computing: A Concise History*, Paul Ceruzzi

*Crowdsourcing*, Daren C. Brabham

*Free Will*, Mark Balaguer

*Information and the Modern Corporation*, James Cortada

*Intellectual Property Strategy*, John Palfrey

*Memes in Digital Culture*, Limor Shifman

*Open Access*, Peter Suber

*Paradox*, Margaret Cuonzo

*Waves*, Fred Raichlen

---

# UNDERSTANDING BELIEFS

NILS J. NILSSON

The MIT Press | Cambridge, Massachusetts | London, England

---

© 2014 Massachusetts Institute of Technology

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

MIT Press books may be purchased at special quantity discounts for business or sales promotional use. For information, please email [special\\_sales@mitpress.mit](mailto:special_sales@mitpress.mit).

This book was set in Chaparral Pro by the MIT Press. Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Nilsson, Nils J., 1933-

Understanding beliefs / Nils J. Nilsson.

pages cm

Includes bibliographical references and index.

ISBN 978-0-262-52643-2 (pbk. : alk. paper) 1. Belief and doubt. I. Title.

BD215.N55 2014

121'.6--dc23

2013047479

10 9 8 7 6 5 4 3 2 1

---

[Belief] ... that upon which a man is prepared to act.

—Alexander Bain, psychologist, quoted in Louis Menand,  
*The Metaphysical Club*<sup>1</sup>

Man's most valuable trait is a judicious sense of what not  
to believe.

—Euripides, circa 480–406 BCE

We may not know very much, but we do know something,  
and while we must always be prepared to change our  
minds, we must act as best we can in the light of what we  
do know.

—W. H. Auden, quoted in Adam Gopnik, "The Double Man"<sup>2</sup>



---

# CONTENTS

Series Foreword	ix
Preface	xi
Acknowledgments	xv

1	Beliefs, Knowledge, and Models	1
2	What Do Beliefs Do for Us?	11
3	Where Do Beliefs Come From?	21
4	Evaluating Beliefs	35
5	In All Probability	51
6	Reality and Truth	65
7	The Scientific Method	75
8	Robot Beliefs	105
9	Belief Traps	117

Glossary	131
Notes	137
Further Readings	143
Index	145





---

## SERIES FOREWORD

The MIT Press Essential Knowledge series offers accessible, concise, beautifully produced pocket-size books on topics of current interest. Written by leading thinkers, the books in this series deliver expert overviews of subjects that range from the cultural and the historical to the scientific and the technical.

In today's era of instant information gratification, we have ready access to opinions, rationalizations, and superficial descriptions. Much harder to come by is the foundational knowledge that informs a principled understanding of the world. Essential Knowledge books fill that need. Synthesizing specialized subject matter for nonspecialists and engaging critical topics through fundamentals, each of these compact volumes offers readers a point of access to complex ideas.

*Bruce Tidor*

*Professor of Biological Engineering and Computer Science  
Massachusetts Institute of Technology*



---

## PREFACE

I became interested in the subject of how people come to know things through my work in artificial intelligence and robotics. In order for robots to perform acceptably, they must know something about the worlds they inhabit. We know how robots “know” because we build them. Some of the knowledge that robots possess comes pre-installed by their designers and builders. Some comes more-or-less directly from their perceptual apparatus—what they see, read, touch, and hear. But robots can also “manufacture” additional knowledge in the form of explanations and conclusions for knowledge they already have. It’s the same for us humans. Perhaps our DNA encodes some knowledge as predispositions shaped by our evolution. But the rest is built on our perceptions and on our abilities to reason and to construct theories.

Brain science hasn’t progressed far enough yet for us to have a detailed picture of how we humans know things. Even so, cognitive psychologists and philosophers have had a lot to say about knowledge. Psychologists talk about different kinds of knowledge, two kinds in particular: knowing *how*, which they call “procedural,” and knowing *that*, which they call “declarative.” Knowing how to ride a bicycle is procedural—it’s built into a program in the brain that enables bicycle riding. Knowing that a bicycle has

---

two wheels is declarative—it can be stated as a declarative sentence.

The branch of philosophy concerned with the study of knowledge is called “epistemology.” Philosophers distinguish *how* knowledge from *that* knowledge also. Epistemology is mainly concerned with *that* knowledge—the kind that can be represented by declarative sentences (which philosophers often call “propositions”). For example, the sentence (or proposition), “The sun is fueled by thermonuclear reactions,” constitutes a piece of scientific knowledge.

But what about “beliefs”? Do our beliefs, expressed as declarative sentences, constitute “knowledge”? Even though some of our beliefs are less strongly held than others, I think the sum total of them does comprise a person’s “knowledge” about the world. It’s all he or she has. (In artificial intelligence, it’s common to refer to a set of propositions as a “knowledge base,” even when some of the propositions are uncertain.)

Some epistemologists attempt to distinguish belief from knowledge. After all, they claim, beliefs might not represent “reality” faithfully, whereas knowledge must. Because I don’t think it’s possible to decide whether or not a sentence (or a set of them) represents reality faithfully (whatever that might mean), I don’t think it’s possible to distinguish knowledge from belief in a qualitatively meaningful way. Nevertheless, many people speak as if

---

there were a difference between believing something and knowing something. For example, I've occasionally had arguments with one of my colleagues—I'll call him Charlie—about something or other. These often ended by my saying, "I understand that you *believe* that, Charlie." Charlie would respond, poking me in the sternum, by saying "I don't *believe* it, I *know* it." Whereas Charlie thought there was a real difference between believing and knowing, to me the sternum poke simply indicated the strength of his belief.

When someone says he *knows* some proposition, I interpret that to mean that he *believes* it very, very strongly—even when I might not believe it at all. Such a person may equivalently say that a strongly held proposition is *true*. The same applies for me. I will tend to say that I *know* things that I believe quite strongly, and I will label them *true*. Because people use the word "know" for their strong beliefs, they may think (like Charlie) that the difference between knowing something and believing something involves more than just the strength of a belief. I don't think there is a way to describe what that "more" might be. Chapter 6 of this book will explore knowing and truth more thoroughly.

Many of our beliefs fall in between strong belief and disbelief. Because our beliefs influence our actions, and because some of our actions might have profound consequences, I think it is important to evaluate beliefs carefully.

---

Chapter 4 is devoted to methods for evaluating beliefs. I think the set of practices that has come to be called the scientific method (explored in chapter 7) offers the most helpful way to evaluate beliefs of all kinds. Like scientific theories, all of our beliefs are (or should be) subject to change.

This book describes my beliefs about beliefs. It's written for those who, like myself, are interested in forming their own beliefs about beliefs. Many of my beliefs are quite controversial, and you may disagree with them. But, after all, you can take them simply as *my* beliefs!

### **About the Notes**

So as not to distract the general reader unnecessarily, numbered notes containing citations to source materials appear by chapter at the end of the book. (Notes for the epigraphs at the beginning of the book are listed first under the heading, "Opening Epigraphs.")

---

## ACKNOWLEDGMENTS

Several people read many drafts of this book (including an unpublished early predecessor titled *How Are We to Know?*). All have made helpful comments, suggestions, or corrections. I'm sure that few agreed with everything or even most of what they read in the drafts. I hope I haven't left too many people out of this list: Dave Berwyn, Ike Burke, Daniel Curtis, Oscar Firschein, Michael Genesereth, Ned Hall, Peter Hart, Hugh Haskell, John Iwuc, Mykel and Mary Anne Kochenderfer, Sidney Liebes, Alan Marer, Andy Neher, Walter Nilsson, Bill Rowe, George Slinn, David Stork, Robert Voss, Yin Wang, and Andrew Waterman. Mary Bagg did an excellent job of copyediting, and I also am grateful for the suggestions of six anonymous reviewers. My thanks to all! Most especially, I want to thank my wife, Grace McConnell Abbott, who read through innumerable drafts and suggested many important improvements.





---

## BELIEFS, KNOWLEDGE, AND MODELS

Our beliefs constitute a large part of our knowledge of the world. For example, I believe I exist on a planet that we call Earth and that I share it with billions of other people. I have beliefs about objects, such as automobiles, airplanes, computers, and various tools, and (to various degrees of detail) how they all operate. I have beliefs about the twenty-first-century culture in which I live: about democracy and the rule of law, about the Internet, and about science and the humanities, among other things. I have beliefs about many other people, including family, friends, associates, and even others yet unmet. And, I believe that they have beliefs also. To make an explicit list of all of my beliefs would be impossible. Yet, all of these beliefs are there, somewhere and somehow represented in my brain—changing, growing, shrinking, and mostly ready for use when I need them.

---

If I were to try to list my beliefs, I would do so using sentences in English, such as, “The universe is around 14 billion years old,” “Salem is the capital of Oregon,” “John Jones usually does what he says he is going to do,” and so on. I can also state what I do *not* believe. For example, “I do not believe in extrasensory perception.” And, I can state that I don’t know something. For example, “I don’t know the population of Sri Lanka.”

We often refer to our beliefs (or to sets of them) as “theories.” We construct theories about everyday experience—both social and personal. Why are crime rates falling in New York City? Why did Booth assassinate Lincoln? Why is my child falling behind in school? Why is unemployment so high (or so low)?

Scientific theories are proposed and argued over by scientists. There are theories to account for fossils found in rocks, for the sun’s almost limitless energy, for earthquakes and volcanoes, for the diversity of life forms, for mental behavior, for the birth and death of stars, and for essentially everything else we can perceive about the universe. Usually scientific theories are described by many sentences—replete with mathematics. They are written down in articles and books to supplement what’s stored in the brains of scientists. When scientists say that they “believe” in quantum mechanics, for example, they are assenting to theories contained in certain articles and books about quantum mechanics. Scientific theories usually have

---

to pass more stringent tests than do the personal theories we all have about many things.

Besides the literature of science, nonfiction books such as histories, political analyses, biographies, and narratives purport to set forth their authors' beliefs about something. You and I might adopt some of these beliefs as our own. For example, you might say that you believe Stephen Ambrose's story about the Lewis and Clark expedition as told in his book *Undaunted Courage*. Even books of fiction contain descriptions of the world that we might incorporate into our beliefs.

One of the most important things to say about beliefs is that they are (or at least should be) tentative and changeable. For example, my belief that the weather will be sunny tomorrow (based on a weather forecast I consulted) may change as new weather data arrives. I may change some of my more fundamental beliefs also, such as my beliefs about early childhood education. Science and medicine advance by new experiments and new theoretical explanations, and these entail new or changed beliefs.

Cognitive scientists distinguish various kinds of knowledge. Knowledge represented by beliefs is called "declarative" because beliefs are stated as declarative sentences. No one really knows how beliefs are represented in our brains. The philosopher and cognitive scientist Jerry Fodor proposes that they are represented as sentence-like forms in a "language of thought" he calls "mentalese."

---

One of the most important things to say about beliefs is that they are (or at least should be) tentative and changeable.

---

Neuroscientists, psychologists, and philosophers continue to argue about whether there are any such sentence-like representations in our brains at all. For our purposes we won't worry about how beliefs are actually represented in the brain. Because we state them using sentences, it seems reasonable to think of them as sentences—constrained by the languages we use to construct sentences.

Cognitive scientists also talk about other kinds of knowledge also. One of these is called “procedural.”<sup>1</sup> Procedural knowledge is built into our practiced actions, such as swinging a golf club or riding a bicycle. For tasks that require real-time coordination between sensing and acting, procedural knowledge is more effective than declarative knowledge would be. (After memorizing some sentences about how to do a cartwheel, could you do one?) Analogously, the knowledge that a computer system uses when it parks a car or lands a plane automatically is of this procedural sort. Most likely much of the knowledge that animals have about their world, such as how to construct spider webs, how to migrate, how to chase prey, and so on, is procedural.

Procedural knowledge is important, but it is limited to those specific actions that it enables. The main reason that we humans are so much more versatile than other animals is that a good deal of our knowledge about the world is declarative and thus can be used to guide many different actions. To use a rather mundane example, our belief that

---

Before we trust a belief sufficiently to act on it, we can analyze it and perhaps modify it—taking into account our own experiences, reasoning, and the opinions and criticisms of others.

---

exercise promotes good health can encourage us to swim, to cycle, or to jog.

Of equal importance to its utility in many different situations, declarative knowledge can be discussed and debated. Before we trust a belief sufficiently to act on it, we can analyze it and perhaps modify it—taking into account our own experiences, reasoning, and the opinions and criticisms of others. As the philosopher Karl Popper put it, “By criticizing our theories we can let our theories die in our stead.”<sup>2</sup>

Beliefs constitute one of the ways we describe the world we live in. We also use mathematical equations (such as  $E = mc^2$ ), computer simulations of various phenomena (such as the weather), maps, and stories. The sum total of all of these constitutes a *model* of reality—an accessible substitute for reality itself. We must make do with this substitute because we can’t apprehend reality directly; it’s on the other side of a “sensory curtain.” Even though it’s easy to imagine that the objects, properties, and relations mentioned in our beliefs actually exist as part of reality, they are only components of the declarative part of our *model* of reality—a kind of “virtual reality.” As the physicist David Deutsch wrote:<sup>3</sup>

Reality is out there: objective, physical and independent of what we believe about it. But we never experience that reality directly. Every last



- [click The Story of Kullervo \(US Edition\)](#)
- [click 1983 pdf](#)
- [click Network and System Security \(2nd Edition\) here](#)
- [download online The Massachusetts Eye and Ear Infirmary Review Manual for Ophthalmology \(3rd Edition\) here](#)
- [read online The Mote in God's Eye](#)
- [download Startup Life: Surviving and Thriving in a Relationship with an Entrepreneur](#)
  
- <http://aneventshop.com/ebooks/The-Story-of-Kullervo--US-Edition-.pdf>
- <http://www.shreesaiexport.com/library/Transcritique--On-Kant-and-Marx.pdf>
- <http://thewun.org/?library/The-Fake-Friend---Daphne-s-Diary-of-Daily-Disasters--Book-3-.pdf>
- <http://studystrategically.com/freebooks/Sexuality-and-Being-in-the-Poststructuralist-Universe-of-Clarice-Lispector--The-Diff--rance-of-Desire--Texas->
- <http://serazard.com/lib/The-Mote-in-God-s-Eye.pdf>
- <http://diy-chirol.com/lib/Startup-Life--Surviving-and-Thriving-in-a-Relationship-with-an-Entrepreneur.pdf>