



# MACHINES LIKE US

TOWARD AI  
WITH COMMON SENSE

RONALD J. BRACHMAN AND  
HECTOR J. LEVESQUE



## Machines like Us



# **Machines like Us**

**Toward AI with Common Sense**

**Ronald J. Brachman and Hector J. Levesque**

**The MIT Press  
Cambridge, Massachusetts  
London, England**

© 2022 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.

The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in ITC Stone Serif Std and ITC Stone Sans Std by New Best-set Typesetters Ltd.

Library of Congress Cataloging-in-Publication Data is available.

Names: Brachman, Ronald J., 1949- author. | Levesque, Hector J., 1951-author.

Title: Machines like us : toward AI with common sense / Ronald J. Brachman and Hector J. Levesque.

Description: Cambridge, Massachusetts : The MIT Press, [2022] | Includes bibliographical references and index.

Identifiers: LCCN 2021031316 | ISBN 9780262046794 (hardcover)

Subjects: LCSH: Artificial intelligence. | Knowledge representation (Information theory)

Classification: LCC Q335 .B695 2022 | DDC 006.3—dc23

LC record available at <https://lcn.loc.gov/2021031316>

10 9 8 7 6 5 4 3 2 1

# Contents

Preview ix

<b>1</b>	<b>The Road to Common Sense</b>	1
<b>2</b>	<b>Common Sense in Humans</b>	9
<b>3</b>	<b>Expertise in AI Systems</b>	25
<b>4</b>	<b>Knowledge and Its Representation</b>	51
<b>5</b>	<b>A Commonsense Understanding of the World</b>	67
<b>6</b>	<b>Commonsense Knowledge</b>	83
<b>7</b>	<b>Representation and Reasoning, Part I</b>	113
<b>8</b>	<b>Representation and Reasoning, Part II</b>	137
<b>9</b>	<b>Common Sense in Action</b>	159
<b>10</b>	<b>Steps toward Implementation</b>	183
<b>11</b>	<b>Building Trust</b>	197
	<b>Epilogue</b>	215

Appendix 219

Bonus Chapter: The Logic of Common Sense 227

Acknowledgments 247

Notes 249

Bibliography 279

Index 297

About the Authors 305

# 1 The Road to Common Sense

We can only see a short distance ahead, but we can see plenty there that needs to be done.

—Alan Turing, “Computing Machinery and Intelligence”

This is a book about common sense from the standpoint of artificial intelligence (AI). While only a small number of AI researchers actually work on common sense directly, the topic has been around since the beginning of the field. AI was given its name in the 1950s by US computer scientist John McCarthy (1927–2011), and one of the first papers ever written about AI is one he presented in 1958 called “Programs with Common Sense.”

With a history like this, one might think that the common sense part of AI would be well in hand by now, more than sixty years later. But in a 2019 book called *Rebooting AI*, Gary Marcus and Ernest Davis make this observation: “Unfortunately acquiring common sense is much harder than one might think. And as we will see, the need for getting machines to acquire common sense is also far more pervasive than one might have imagined.”

The implication here is clear: AI systems could really use some sort of common sense, but nobody has yet figured out how to provide them with it. And it’s not as if anyone thinks that common sense is some minor aspect of intelligence; arguably, it’s at the very core.

What we propose to do in this book is to take a hard look at this notion of common sense to get a better sense of what it is, why it is important, and what it might take to make it work in an AI system. But before setting out on this long and winding road, let us have a look at where we will be going.



## A Guide to Commonsense Guides

A quick search on Amazon will reveal that there are over seven thousand books with “common sense” (or “common-sense” or “commonsense”) in the title. About a tenth of them also include the word “guide.” To get an idea of their range, here are a dozen of those guides:

*The Common Sense Guide to Real Estate Investing*

*Nanny in a Book: The Common-Sense Guide to Childcare*

*Surviving the Ebola Virus: A Tropical Doctor's Commonsense Guide*

*A Commonsense Guide to Fasting*

*Freedom's Last Stand: A Common-Sense Guide to Understanding the Tyranny of Collectivist Ideology*

*Uncommon Fruits & Vegetables: A Commonsense Guide*

*Technical Writing A-Z: A Commonsense Guide to Engineering Reports and Theses*

*The Common Sense Guide to Handfeeding Baby Birds*

*Your First Gun: A Common Sense Guide to Selection and Usage for Novices and Dummies*

*Dude! What Were You Thinking? A Common Sense Guide to Dating & Relationships*

*Answers to Your Mule Questions: A Common Sense Guide to Understanding the Mule's Point of View*

*Gardening Nude: A Common Sense Guide to Improving Your Health and Lifestyle by Increasing Exposure to Nature, Cultivating a Green Mindset, and Building a Strong Community*

Other than the “common sense” part, it is hard to see what these books could possibly have in common!

One thing to think about in the context of *this* book is what the term “common sense” in these guides is intended to convey. It is clearly not just an exhortation to use common sense; a commonsense guide to fasting should be a bit more than “use common sense when fasting.” What the title *A Commonsense Guide to Fasting* suggests is perhaps something like this: “Here are some things to know about fasting, not to become an expert, but to guide your actions in a sensible way in ordinary situations.” The same applies to investing in real estate, understanding mules, and presumably, gardening in the nude.

The premise here is that while it might be possible to become an *expert* on any of these topics, there are also some basic, commonsense things that even nonexperts will find useful. Here's the pitch at Amazon for the book on hand-feeding baby birds:

Breaking the myth that handfeeding baby birds requires years and years of training, this guide offers a common sense solution. Detailing a basic and workable knowledge of a baby bird's requirements for warmth, cleanliness, frequency of feeding, proper formulas, correct temperature and correct technique, this book will have you feeding baby birds like a pro in no time.

The goal seems reasonable enough. There is clearly a wide range of things that humans deal with every day not as experts, but in simple, ordinary terms. Indeed, how many of us truly understand how an LED TV displays a movie or Alexa tells us the weather? For most of us, a commonsense understanding of the technology is all we have. In some cases, we get our understanding through repeated experience (like how to eat a melting ice cream cone on a hot summer day); in some cases, we get advice from friends (such as what to wear—and not wear—for a party); and in other cases, our understanding may well come from something like a guide or manual (like how to set up a home printer).

What we intend to do in this book is to step back from individual topics like the Ebola virus and hand-feeding baby birds, and look at the underlying idea of common sense itself. We want an understanding of common sense that makes it clear how it leans on the knowledge gained from things such as prior experience, advice from friends, and commonsense guides. As computer scientists, we want to aim for an account of common sense that is specific and detailed enough that we could at least imagine building a computational machine that has it.

### **Artificial Intelligence**

This takes us back to the topic of AI. In general terms, AI is the study of how intelligent behavior can be produced through computational means. At one extreme, we imagine sophisticated behaviors like playing chess, interpreting poetry, and classifying tumors; at the other extreme, we envision more commonplace activities like babysitting a toddler, preparing a meal, and driving a car.

What do these have to do with common sense? This will take some sorting out in the next two chapters, but the basic idea is that an AI system may be able to do remarkably well on certain specific tasks that demand intelligence in humans, but still be unable to behave in an intelligent way more broadly. What we will end up saying is that these AI systems display expertise, not common sense.

Consider the well-known AI system called AlphaGo developed by DeepMind Technologies. This system is unquestionably an expert when it comes to playing the game of Go and likely the top Go player in the world. But it does nothing else. It doesn't have even a rudimentary ability to deal with the rest of the world. We would never expect it to be able to order a pizza, recommend a movie, or diagnose a blood infection. The system can't really be criticized for this; it was never intended to do anything other than play Go.

To someone outside the field of AI, the fact that different AI systems perform different tasks might seem like a small detail. Can't you just take something like AlphaGo, marry it with Roomba and Siri, and end up with a system that plays Go, vacuums floors, and responds to voice commands? Throw in Google Translate, and will the system not respond to voice commands in any language? Keep on doing this long enough and will we not end up with a system with enough versatility to do a good job out there in the real world?

Things have not turned out this way. However impressive each of these AI systems is individually, they remain isolated from each other. They are like small islands of high achievement in a large sea of possible behaviors. There is still no clear sense within the AI community of how to get a system to behave in a reasonable way outside these localized areas of expertise. Even if we were to imagine an AI system armed with the hundreds of commonsense guides mentioned above, it would still be at a loss outside those specific topics.

To give a small example, the guide to uncommon fruits and vegetables listed above might well talk about pomegranates, but it almost certainly does not mention laundry baskets. Nor would we expect a guide for laundry baskets (or any direct experience with laundry baskets) to ever deal with pomegranates. So if we now imagine an AI system that has to decide whether a pomegranate would fit inside a laundry basket (for some reason), it will need something beyond what it has on each individual topic. It will be unable to make a sensible decision if it has no way of bringing these two separate topics together in a meaningful way.

This is of course not how things are in humans. We fully expect that a person who gets to be good at playing Go will still do well on a broad range of non-Go activities including, but not limited to (as the lawyers say), vacuuming floors and handling questions about pomegranates in laundry baskets. This ability to deal with what comes up in a way that is not tied to prowess in any particular domain is precisely what we expect of common sense.

We are not the first to observe that current AI systems and humans are different in this regard. It is one thing that virtually all AI scientists and technologists agree on. Where they do disagree, however—and sometimes quite loudly—is on what the implications are.

For (what might be called) AI boosters, the picture is somewhat like this:

We are in the early stages of AI, and the best of the systems built so far are already intelligent, but clearly not in the same way humans are. “Intelligent” is not a useful label in the end. The systems do better than humans in some ways, and worse in others. The goal is not to emulate humans to the letter, with all their limitations and failings, but rather to behave in a usefully intelligent way on a wide variety of tasks. There will always be a powerful synergy between AI systems and the humans they work with; each will make up for weaknesses in the other.

For (what might be called) AI skeptics, including many of the AI researchers themselves, the picture is quite different:

The AI systems built to date are nothing like intelligent. If real human-level intelligence is what you are after, the whole current AI approach is wrongheaded. Put aside proficiency in ultraspecialized areas like playing Go. You can’t do better than humans in general terms if you can’t do as well as a six-year-old child. And there is no evidence to date that an AI system will ever have the general intelligence of a six-year-old.

It is not our goal in this book to continue this debate; others have written about it. What we aim to do is to focus on just the common sense part: what common sense means, how it works, what it might take to build computational systems that have it, and why we should bother.

## **The Road Ahead**

At the risk of spoiling all the suspense, here is where the book is going:

### Chapter 2: Common Sense in Humans

We look at the idea of common sense as it exists in humans. We make the case that it is tied to knowing certain ordinary things. We argue that

common sense is the ability to make effective use of this knowledge in deciding how to behave and plays a critical role in the spectrum of human cognitive capabilities.

### Chapter 3: Expertise in AI Systems

We turn our attention to AI systems past and present. We make the case that no matter how these systems are built, they deliver expertise on certain tasks, not common sense. We argue that common sense is not a matter of additional expertise in yet another area. If we want AI systems that can deal with the wide range of unanticipated things that come up in real life, a different approach is called for.

### Chapter 4: Knowledge and Its Representation

We consider the issue of knowledge and what it might mean to make use of it. We contend that an item of knowledge is somewhat like a number in that we can make use of it by representing it symbolically and working on the symbols. In the case of numbers, what the symbols should be and how to operate on them is well understood; it's what we call arithmetic. How this should play out for knowledge is much less clear and is the subject of the middle part of the book.

### Chapter 5: A Commonsense Understanding of the World

We begin to consider commonsense knowledge by examining the high-level picture of the world it presupposes: there are things in the world with properties that change over time as a result of events that occur; among these things, there are physical objects with properties like size and location in space; and among these physical objects, there are animate agents with beliefs and goals who can cause events to occur.

### Chapter 6: Commonsense Knowledge

We consider how specific items of commonsense knowledge might be formulated. We make the case that in addition to dealing with things and their properties, we need to deal with concepts—that is, ideas about the kinds of things there can be. We maintain that these concepts are understood in terms of typical, atypical, and borderline cases.

### Chapter 7: Representation and Reasoning, Part I

We take the first step toward representing commonsense knowledge in symbolic form. We argue that we want to represent things and their

properties separately from the concepts involved. We also talk about how a computational system would reason with these representations.

#### Chapter 8: Representation and Reasoning, Part II

We take a second step in the use of symbolic representations. We contend that we need to be able to represent propositions—sentences that can be either true or false—without believing them to be true. This will be essential for reasoning about what it would take to make a proposition true—the basis for planning. And again, we talk about how a computational system would operate on these representations.

#### Chapter 9: Common Sense in Action

We return to common sense at a high level and show how the pieces fit together. We consider a scenario where an agent, while engaging in a routine task, encounters something totally unexpected and needs to use common sense to figure out what to do next. We show how this application of common sense can be understood as computing a certain result from symbolic representations of the sort seen in earlier chapters.

#### Chapter 10: Steps toward Implementation

We look in more practical terms at what it would take to build an AI system with common sense. Besides considering commonsense reasoners, we spend some time thinking about where all the commonsense knowledge will come from, and looking at some interesting ways in which AI systems of the future will need to learn from experience and written texts.

#### Chapter 11: Building Trust

We begin to wrap up by considering whether society should even be thinking about building AI systems with common sense. We make the case that if we are ever going to build systems that make decisions for themselves, we will want them to have reasons for what they do that humans can understand, and then agree or disagree with. In particular, we will want them to have commonsense knowledge and goals, and for their decisions about what to do to be the result of applying this knowledge in pursuit of those goals.

The book also includes a bonus chapter for aficionados on the connection between logic and commonsense reasoning as well as chapter-by-chapter notes that add some color to the text and suggest further resources for the interested reader.

## Setting Out

So this is our journey in a nutshell: from common sense, to commonsense knowledge, to commonsense knowledge representation and reasoning. It's a lot to take in. As the song says, the road is long with many a winding turn. This is a book about common sense, but not really much of a commonsense guide.

There are those who want to believe in a much simpler story for AI, where we would take some sort of learning system, hook it up to the internet (or maybe push a robot out the door), and have it learn to figure things out for itself. Why can't some version of common sense just emerge spontaneously from being embedded in the world over a period of time? All of this talk about computation over symbolic representations of knowledge might feel too mannered, too regimented, too cerebral for common sense.

While we understand the intuitive appeal of this simpler story, there is as yet no good specification of how it could possibly work. That said, we need to be clear that the more involved story we are about to tell here is only a hypothesis—one possible road to common sense. In practical terms, it could well turn out to be a dead end, or eventually be supplanted by more direct highways. But as we hope the reader will see, in intellectual terms, it is still a rewarding road to follow, with wonderful sites to visit along the way in philosophy, psychology, linguistics, logic, and computer science.

The road we are setting out on is one that we ourselves have explored for four decades in our own work. We have built AI systems, proved AI theorems, and led AI projects. We have been inspired by all the terrific people we have worked with and our involvement with the major international AI organizations. As your guides on this journey, we will do our best to make things informative, interesting, and fun. So buckle your seat belt and enjoy the trip!