


LEARNING MADE EASY



3rd Edition

Artificial Intelligence

for
dummies[®]
A Wiley Brand



Look under the hood
of AI platforms

—
Connect the dots between
data and AI output

—
Grasp how AI influences
our daily lives

John Paul Mueller
Luca Massaron
Stephanie Diamond



Artificial Intelligence

3rd Edition

**by John Paul Mueller, Luca Massaron,
and Stephanie Diamond**

**for
dummies[®]**
A Wiley Brand

Artificial Intelligence For Dummies®, 3rd Edition

Published by: **John Wiley & Sons, Inc.**, 111 River Street, Hoboken, NJ 07030-5774, www.wiley.com

Copyright © 2025 by John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial technologies or similar technologies.

Media and software compilation copyright © 2025 by John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial technologies or similar technologies.

Published simultaneously in Canada

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the Publisher. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

Trademarks: Wiley, For Dummies, the Dummies Man logo, Dummies.com, Making Everything Easier, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and may not be used without written permission. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc. is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES OR PROMOTIONAL MATERIALS. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING LEGAL, ACCOUNTING, OR OTHER PROFESSIONAL SERVICES. IF PROFESSIONAL ASSISTANCE IS REQUIRED, THE SERVICES OF A COMPETENT PROFESSIONAL PERSON SHOULD BE SOUGHT. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.

For general information on our other products and services, please contact our Customer Care Department within the U.S. at 877-762-2974, outside the U.S. at 317-572-3993, or fax 317-572-4002. For technical support, please visit <https://hub.wiley.com/community/support/dummies>.

Wiley publishes in a variety of print and electronic formats and by print-on-demand. Some material included with standard print versions of this book may not be included in e-books or in print-on-demand. If this book refers to media that is not included in the version you purchased, you may download this material at <http://booksupport.wiley.com>. For more information about Wiley products, visit www.wiley.com.

Library of Congress Control Number is available from the publisher.

ISBN 978-1-394-27071-2 (pbk); ISBN 978-1-394-27073-6 (ebk); ISBN 978-1-394-27072-9 (ebk)

Contents at a Glance

Introduction	1
Part 1: Introducing AI	5
CHAPTER 1: Delving into What AI Means	7
CHAPTER 2: Defining Data's Role In AI.	23
CHAPTER 3: Considering the Use of Algorithms	45
CHAPTER 4: Pioneering Specialized Hardware	63
Part 2: Understanding How AI Works	79
CHAPTER 5: Crafting Intelligence for AI Data Analysis	81
CHAPTER 6: Employing Machine Learning in AI.	97
CHAPTER 7: Improving AI with Deep Learning.	117
Part 3: Recognizing How We Interact with AI Every Day	135
CHAPTER 8: Unleashing Generative AI for Text and Images	137
CHAPTER 9: Seeing AI Uses in Computer Applications.	173
CHAPTER 10: Automating Common Processes	183
CHAPTER 11: Relying on AI to Improve Human Interaction.	193
Part 4: AI Applied in Industries	203
CHAPTER 12: Using AI to Address Medical Needs	205
CHAPTER 13: Developing Robots	225
CHAPTER 14: Flying with Drones	239
CHAPTER 15: Utilizing the AI-Driven Car	257
Part 5: Getting Philosophical About AI	275
CHAPTER 16: Understanding the Nonstarter Application – Why We Still Need Humans	277
CHAPTER 17: Engaging in Human Endeavors	291
CHAPTER 18: Seeing AI in Space.	305
Part 6: The Part of Tens	321
CHAPTER 19: Ten Substantial Contributions of AI to Society	323
CHAPTER 20: Ten Ways in Which AI Has Failed	331
Index	339

Table of Contents

INTRODUCTION	1
About This Book	2
Icons Used in This Book	3
Beyond the Book	3
Where to Go from Here	4
PART 1: INTRODUCING AI	5
CHAPTER 1: Delving into What AI Means	7
Defining the Term AI	8
Discerning intelligence	8
Examining four ways to define AI	11
Reviewing AI categories	15
Understanding the History of AI	16
Considering AI Uses	17
Avoiding AI Hype and Overestimation	18
Defining the five tribes and the master algorithm	18
Considering sources of hype	19
Managing user overestimation	20
Connecting AI to the Underlying Computer	20
CHAPTER 2: Defining Data's Role In AI	23
Finding Data Ubiquitous in This Age	24
Using data everywhere	25
Putting algorithms into action	26
Using Data Successfully	28
Considering the data sources	28
Obtaining reliable data	29
Making human input more reliable	30
Using automated data collection	31
Collecting data ethically	32
Manicuring the Data	33
Dealing with missing data	33
Considering data misalignments	34
Separating useful data from other data	35
Considering the Five Mistruths in Data	36
Commission	36
Omission	37

	Perspective	37
	Bias	38
	Frame of reference	39
	Defining the Limits of Data Acquisition	40
	Considering Data Security Issues	41
	Understanding purposefully biased data	42
	Dealing with data-source corruption	43
	Handling botnets	44
CHAPTER 3:	Considering the Use of Algorithms	45
	Understanding the Role of Algorithms	46
	Examining what an algorithm does	46
	Planning and branching: Trees and nodes	48
	Extending the tree using graph nodes	49
	Traversing the graph	50
	Playing adversarial games	52
	Using local search and heuristics	53
	Discovering the Learning Machine	56
	Leveraging expert systems	57
	Introducing machine learning	60
	Achieving new heights	60
CHAPTER 4:	Pioneering Specialized Hardware	63
	Relying on Standard Hardware	63
	Examining the standard hardware	64
	Describing standard hardware deficiencies	64
	Relying on new computational techniques	65
	Using GPUs	66
	Considering the von Neumann bottleneck	67
	Defining the GPU	68
	Considering why GPUs work well	69
	Working with Deep Learning Processors (DLPs)	69
	Defining the DLP	70
	Using the mobile neural processing unit (NPU)	71
	Accessing the cloud-based tensor processing unit (TPU)	71
	Creating a Specialized Processing Environment	72
	Increasing Hardware Capabilities	73
	Advancing neuromorphic computing	74
	Exploring quantum processors	74
	Adding Specialized Sensors	75
	Integrating AI with Advanced Sensor Technology	76
	Devising Methods to Interact with the Environment	76

PART 2: UNDERSTANDING HOW AI WORKS	79
CHAPTER 5: Crafting Intelligence for AI Data Analysis	81
Defining Data Analysis	82
Understanding why analysis is important	84
Reconsidering the value of data	85
Defining Machine Learning (ML)	87
Understanding how ML works	88
Understanding the benefits of ML	89
Being useful; being mundane	90
Specifying the limits of ML	91
Considering How to Learn from Data	93
Supervised learning	93
Unsupervised learning	94
Reinforcement learning	94
CHAPTER 6: Employing Machine Learning in AI	97
Taking Many Different Roads to Learning	98
Recognizing five main approaches to AI learning	99
Exploring the three most promising AI learning approaches	101
Awaiting the next breakthrough	102
Exploring the Truth in Probabilities	103
Determining what probabilities can do	104
Considering prior knowledge	105
Envisioning the world as a graph	109
Growing Trees That Can Classify	112
Predicting outcomes by splitting data	113
Making decisions based on trees	114
Pruning overgrown trees	116
CHAPTER 7: Improving AI with Deep Learning	117
Shaping Neural Networks Similar to the Human Brain	118
Introducing the neuron	118
Starting with the miraculous perceptron	119
Mimicking the Learning Brain	121
Considering simple neural networks	121
Figuring out the secret is in the weights	123
Understanding the role of backpropagation	124
Introducing Deep Learning	125
Deep learning versus simpler solutions	127
Finding even smarter solutions	128
Detecting Edges and Shapes from Images	130
Starting with character recognition	131
Explaining how convolutions work	132
Advancing using image challenges	133

PART 3: RECOGNIZING HOW WE INTERACT WITH AI EVERY DAY	135
CHAPTER 8: Unleashing Generative AI for Text and Images ...	137
Getting an Overview of Generative AI	138
Memorizing sequences that matter	140
Employing self-attention models	142
Discovering the Magic Smooth Talk of AI	145
Creating generative adversarial networks	148
Revolutionizing art with diffusion models	150
Applying reinforcement learning	151
Recapping the recent history of GenAI	156
Working with Generative AI	157
Creating text using Generative AI	158
Creating images using Generative AI	160
Understanding the Societal Implications of Generative AI	161
Managing media hype	161
Promoting positive uses	162
Addressing security concerns	162
Attempting to mimic human creativity	163
Defining the side effects	164
Deciding What Makes a Good Generative AI App	165
Assessing accuracy and reliability	165
Enhancing the user experience	166
Implementing innovation	166
Monitoring stability and performance	166
Looking at ethical considerations	166
Identifying economic considerations	167
Maximizing ROI potential	167
Commercializing Generative AI	168
Viewing the GPT store	168
Looking at basic app creation techniques	169
Monetizing your Generative AI app	170
Addressing the regulatory environment in commercializing Generative AI	170
CHAPTER 9: Seeing AI Uses in Computer Applications	173
Introducing Common Application Types	174
Using AI in typical applications	174
Realizing AI's wide range of fields	175
Seeing How AI Makes Applications Friendlier	176
Performing Corrections Automatically	178
Considering the kinds of corrections	178
Seeing the benefits of automatic corrections	178
Understanding why automated corrections don't always work ...	179

Making Suggestions	179
Getting suggestions based on past actions	180
Getting suggestions based on groups	180
Obtaining the wrong suggestions	181
Considering AI-Based Errors	182
CHAPTER 10: Automating Common Processes	183
Developing Solutions for Boredom	184
Making tasks more interesting	184
Helping humans work more efficiently	185
Examining how AI reduces boredom	185
Recognizing that AI can't reduce boredom	186
Working in Industrial Settings	187
Developing various levels of automation.	187
Using more than just robots	189
Relying on automation alone.	189
Creating a Safe Environment	190
Considering the role of boredom in accidents	190
Seeing AI to avoid safety issues.	190
Accepting that AI cannot eliminate safety issues	191
CHAPTER 11: Relying on AI to Improve Human Interaction	193
Developing New Ways to Communicate	194
Creating new alphabets	195
Working with emojis and other meaningful graphics	195
Automating language translation	196
Incorporating body language.	196
Exchanging Ideas.	197
Creating connections	198
Augmenting communication	198
Defining trends	199
Using Multimedia	199
Embellishing Human Sensory Perception	200
Shifting data spectrum	200
Augmenting human senses	201
PART 4: AI APPLIED IN INDUSTRIES	203
CHAPTER 12: Using AI to Address Medical Needs	205
Implementing Portable Patient Monitoring.	206
Wearing helpful monitors	206
Relying on critical wearable monitors	207
Using movable monitors	208

Making Humans More Capable	209
Using games for therapy	209
Considering the use of exoskeletons	210
Addressing Special Needs	212
Considering the software-based solutions	213
Relying on hardware augmentation	214
Completing Analysis in New Ways	214
Relying on Telepresence	215
Defining telepresence	215
Considering examples of telepresence	216
Understanding telepresence limitations	216
Devising New Surgical Techniques	217
Making surgical suggestions	217
Assisting a surgeon	218
Replacing the surgeon with monitoring	219
Performing Tasks Using Automation	219
Working with medical records	220
Predicting the future	220
Making procedures safer	221
Creating better medications	221
Combining Robots and Medical Professionals	222
Considering Disruptions AI Causes for Medical Professionals	222
CHAPTER 13: Developing Robots	225
Defining Robot Roles	226
Overcoming the sci-fi view of robots	227
Being humanoid can be hard	230
Working with robots	233
Assembling a Basic Robot	236
Considering the components	236
Sensing the world	237
Controlling a robot	238
CHAPTER 14: Flying with Drones	239
Acknowledging the State of the Art	240
Flying unmanned to missions	240
Meeting the quadcopter	242
Defining Uses for Drones	244
Seeing drones in nonmilitary roles	245
Powering up drones using AI	251
Understanding regulatory issues	252
Reviewing Privacy and Data Protection in Drone Operations	254
Analyzing implications for personal privacy	254
Handling data collected by drones	255
Considering legal considerations and surveillance	255
Developing regulatory frameworks and recommendations	255

CHAPTER 15: Utilizing the AI-Driven Car	257
Examining the Short History of SD Cars	258
Understanding the Future of Mobility	258
Climbing the six levels of autonomy	259
Rethinking the role of cars in our lives	261
Taking a step back from unmet expectations	263
Getting into a Self-Driving Car	265
Putting all the tech together	266
Letting AI into the scene	267
Understanding that it's not just AI	268
Overcoming Uncertainty of Perceptions	270
Introducing the car's senses	270
Putting together what you perceive	273

PART 5: GETTING PHILOSOPHICAL ABOUT AI

CHAPTER 16: Understanding the Nonstarter Application – Why We Still Need Humans	277
Using AI Where It Won't Work	278
Defining the limits of AI	278
Applying AI incorrectly	281
Entering a world of unrealistic expectations	282
Considering the Effects of AI Winters	283
Defining the causes of the AI winter	283
Rebuilding expectations with new goals	285
Creating Solutions in Search of a Problem	287
Defining a gizmo	287
Avoiding the infomercial	288
Understanding when humans do it better	288
Looking for the simple solution	289
CHAPTER 17: Engaging in Human Endeavors	291
Keeping Human Beings Popular	292
Living and Working in Space	293
Creating Cities in Hostile Environments	294
Building cities in the ocean	294
Creating space-based habitats	296
Constructing moon-based resources	297
Making Humans More Efficient	298
Fixing Problems on a Planetary Scale	300
Contemplating how the world works	300
Locating potential sources of problems	301
Defining potential solutions	302
Seeing the effects of the solutions	303
Trying again	303

CHAPTER 18: Seeing AI in Space	305
Integrating AI into Space Operations	305
Seeing clearly with the help of algorithms	306
Finding new frontiers to challenge	308
Unveiling new scientific discoveries	308
Performing Space Mining	310
Harvesting water	310
Obtaining rare earths and other metals	312
Finding new elements	312
Enhancing communication	313
Exploring New Places	313
Starting with the probe	314
Relying on robotic missions	315
Adding the human element	317
Building Structures in Space	317
Taking your first space vacation	318
Industrializing space	318
Using space for storage	319
PART 6: THE PART OF TENS	321
CHAPTER 19: Ten Substantial Contributions of AI to Society	323
Considering Human-Specific Interactions	324
Devising the active human foot	324
Performing constant monitoring	325
Administering medications	325
Developing Industrial Solutions	325
Using AI with 3D printing	326
Advancing robot technologies	326
Creating New Technology Environments	327
Developing rare new resources	327
Seeing what can't be seen	327
Working with AI in Space	328
Delivering goods to space stations	328
Mining extraplanetary resources	329
Exploring other planets	329
CHAPTER 20: Ten Ways in Which AI Has Failed	331
Understanding	332
Interpreting, not analyzing	332
Going beyond pure numbers	333
Considering consequences	334

Discovering	334
Devising new data from old	334
Seeing beyond the patterns	335
Implementing new senses	335
Empathizing	336
Walking in someone's shoes	336
Developing true relationships	337
Changing perspective	337
Making leaps of faith	337
INDEX	339

Introduction

You can hardly avoid hearing about AI these days. You see AI in the movies, in books, in the news, and online. AI is part of robots, self-driving (SD) cars, drones, medical systems, online shopping sites, and all sorts of other technologies that affect your daily life in innumerable ways.

Many pundits are burying you in information (and disinformation) about AI, too. Much of the hype about AI originates from the excessive and unrealistic expectations of scientists, entrepreneurs, and businesspersons. *Artificial Intelligence For Dummies*, 3rd Edition, is the book you need if you feel as though you truly don't know anything about a technology that purports to be an essential element of your life.

Using various media as a starting point, you might notice that most of the useful technologies are almost boring. Certainly, no one gushes over them. AI is like that: so ubiquitous as to be humdrum. You're using AI in some way today; in fact, you probably rely on AI in many different ways — you just don't notice it because it's so mundane. This book makes you aware of these very real and essential uses of AI. A smart thermostat for your home may not sound exciting, but it's an incredibly practical use for a technology that has some people running for the hills in terror.

This book also covers various cool uses of AI. For example, you may not realize that a medical monitoring device can now predict when you might have a heart problem — but such a device exists. AI powers drones, drives cars, and makes all sorts of robots possible. You see AI used now in all sorts of space applications, and AI figures prominently in all the space adventures humans will have tomorrow.

In contrast to many books on the topic, *Artificial Intelligence For Dummies*, 3rd Edition, also tells you the truth about where and how AI *can't* work. In fact, AI will never be able to engage in certain essential activities and tasks, and it will be unable to engage in other ones until far into the future. One takeaway from this book is that humans will always be important. In fact, if anything, AI makes humans even more important because AI helps humans excel in ways that you frankly might not be able to imagine.

About This Book

Artificial Intelligence For Dummies, 3rd Edition starts by helping you understand AI, especially what AI needs to work and why it has failed in the past. You also discover the basis for some of the issues with AI today and how those issues might prove to be nearly impossible to solve in certain cases. Of course, along with the issues, you discover the fixes for various problems and consider where scientists are taking AI in search of answers. Most important, you discover where AI is falling short and where it excels. You likely won't have a self-driving car anytime soon, and that vacation in space will have to wait. On the other hand, you find that telepresence can help people stay in their homes when they might otherwise need to go to a hospital or nursing home.

This book also contains links to external information because AI has become a huge and complex topic. Follow these links to gain additional information that just won't fit in this book — and to gain a full appreciation of just how astounding the impact of AI is on your daily life. If you're reading the print version of this book, you can type the URL provided into your browser; e-book readers can simply click the links. Many other links use what's called a TinyURL (tinyurl.com), in case the original link is too long and confusing to type accurately into a search engine. To check whether a TinyURL is real, you can use the preview feature by adding the word *preview* as part of the link, like this: preview.tinyurl.com/pd88943u.

AI has a truly bright future because it has become an essential technology. This book also shows you the paths that AI is likely to follow in the future. The various trends discussed in this book are based on what people are actually trying to do now. The new technology hasn't succeeded yet, but because people are working on it, it does have a good chance of success at some point.

To make absorbing the concepts even easier, this book uses the following conventions:

- » Web addresses appear in monofont. If you're reading a digital version of this book on a device connected to the Internet, note that you can click the web address to visit that website, like this: `www.dummies.com`. Many article titles of additional resources also appear as clickable links.
- » Words in *italics* are defined inline as special terms you should remember. You see these words used (and sometimes misused) in many different ways in the press and other media, such as movies. Knowing the meaning of these terms can help you clear away some of the hype surrounding AI.

Icons Used in This Book

As you read this book, you see icons in the margins that indicate material of interest (or not, as the case may be). This section briefly describes each icon in this book.



TIP

Tips are gratifying because they help you save time or perform a task without creating a lot of extra work. The tips in this book are time-saving techniques or pointers to resources that you should try in order to gain the maximum benefit from learning about AI. Just think of them as extras that we're paying to reward you for reading our book.



REMEMBER

If you get nothing else out of a particular chapter or section, remember the material marked by this icon. This text usually contains an essential process or a bit of information that you must know to interact with AI successfully.



WARNING

We don't want to sound like angry parents or some kind of maniacs, but you should avoid doing anything marked with a Warning icon. Otherwise, you might find that you engage in the sort of disinformation that now has people terrified of AI.



TECHNICAL
STUFF

Whenever you see this icon, think "advanced tip or technique." You can fall asleep from reading this material, and we don't want to be responsible for that. However, you might find that these tidbits of useful information contain the solution you need in order to create or use an AI solution. Skip these bits of information whenever you like.

Beyond the Book

Every book in the For Dummies series comes supplied with an online Cheat Sheet. You remember using crib notes in school to make a better mark on a test, don't you? You do? Well, a cheat sheet is sort of like that. It provides you with some special notes about tasks that you can do with AI that not everyone else knows about. You can find the cheat sheet for this book by going to www.dummies.com and typing **Artificial Intelligence For Dummies Cheat Sheet** in the search box. The cheat sheet contains neat-to information, such as the meaning of all those strange acronyms and abbreviations associated with AI, machine learning, and deep learning.

Where to Go from Here

It's time to start discovering AI and see what it can do for you. If you know nothing about AI, start with Chapter 1. You may not want to read every chapter in the book, but starting with Chapter 1 helps you understand the AI basics that you need when working through other places in the book.

If your main goal in reading this book is to build knowledge of where AI is used today, start with Chapter 5. The materials in Part 2 can help you see where AI is used today.

If you have a bit more advanced knowledge of AI, you can start with Chapter 9. Part 3 of this book contains the most advanced material that you'll encounter. If you don't want to know how AI works at a low level (not as a developer but simply as someone interested in AI), you might decide to skip this part of the book.

Okay, so you want to know the super fantastic ways in which people are either using AI today or will use AI in the future. If that's the case, start with Chapter 12. All of Parts 4 and 5 show you the incredible ways in which AI is used without forcing you to deal with piles of hype as a result. The information in Part 4 focuses on hardware that relies on AI, and the material in Part 5 focuses more on futuristic uses of AI.

1

Introducing AI

IN THIS PART . . .

Discover what AI can actually do for you.

Consider how data affects the use of AI.

Understand how AI relies on algorithms to perform useful work.

See how using specialized hardware makes AI perform better.

- » Defining AI and its history
- » Using AI for practical tasks
- » Seeing through AI hype
- » Connecting AI with computer technology

Chapter **1**

Delving into What AI Means

Common apps, such as Google Assistant, Alexa, and Siri, have all of us who are online every day, using artificial intelligence (AI) without even thinking about it. Productivity and creative apps such as ChatGPT, Synesthesia, and Gemini help us focus on the content rather than on how to get there. The media floods our entire social environment with so much information and disinformation that many people see AI as a kind of magic (which it most certainly isn't). So the best way to start this book is to define what AI is, what it isn't, and how it connects to computers today.



REMEMBER

Of course, the basis for what you expect from AI is a combination of how you define AI, the technology you have for implementing AI, and the goals you have for AI. Consequently, everyone sees AI differently. This book takes a middle-of-the-road approach by viewing AI from as many different perspectives as possible. We don't buy into the hype offered by proponents, nor do we indulge in the negativity espoused by detractors. Instead, we strive to give you the best possible view of AI as a technology. As a result, you may find that you have expectations somewhat different from those you encounter in this book, which is fine, but it's essential to consider what the technology can actually do for you — rather than expect something it can't.

Defining the Term AI

Before you can use a term in any meaningful and useful way, you must have a definition for it. After all, if nobody agrees on a meaning, the term has none; it's just a collection of characters. Defining the idiom (a term whose meaning isn't clear from the meanings of its constituent elements) is especially important with technical terms that have received more than a little press coverage at various times and in various ways.



REMEMBER

Saying that AI is an artificial intelligence doesn't tell you anything meaningful, which is why people have so many discussions and disagreements over this term. Yes, you can argue that what occurs is artificial, not having come from a natural source. However, the intelligence part is, at best, ambiguous. Even if you don't necessarily agree with the definition of AI as it appears in the sections that follow, this book uses AI according to that definition, and knowing it will help you follow the text more easily.

Discerning intelligence

People define intelligence in many different ways. However, you can say that intelligence involves certain mental activities composed of the following activities:

- » **Learning:** Having the ability to obtain and process new information
- » **Reasoning:** Being able to manipulate information in various ways
- » **Understanding:** Considering the result of information manipulation
- » **Grasping truths:** Determining the validity of the manipulated information
- » **Seeing relationships:** Divining how validated data interacts with other data
- » **Considering meanings:** Applying truths to particular situations in a manner consistent with their relationship
- » **Separating fact from belief:** Determining whether the data is adequately supported by provable sources that can be demonstrated to be consistently valid

The list could easily grow quite long, but even this list is relatively prone to interpretation by anyone who accepts it as viable. As you can see from the list, however, intelligence often follows a process that a computer system can mimic as part of a simulation:

1. Set a goal (the information to process and the desired output) based on needs or wants.
2. Assess the value of any known information in support of the goal.
3. Gather additional information that could support the goal. The emphasis here is on information that *could* support the goal rather than on information you know *will* support the goal.
4. Manipulate the data such that it achieves a form consistent with existing information.
5. Define the relationships and truth values between existing and new information.
6. Determine whether the goal is achieved.
7. Modify the goal in light of the new data and its effect on the probability of success.
8. Repeat Steps 2 through 7 as needed until the goal is achieved (found true) or the possibilities for achieving it are exhausted (found false).



REMEMBER

Even though you can create algorithms and provide access to data in support of this process within a computer, a computer's capability to achieve intelligence is severely limited. For example, a computer is incapable of understanding anything because it relies on machine processes to manipulate data using pure math in a strictly mechanical fashion. Likewise, computers can't easily separate truth from mistruth (as described in Chapter 2). In fact, no computer can fully implement any of the mental activities described in the earlier list that describes intelligence.

As part of deciding what intelligence actually involves, categorizing intelligence is also helpful. Humans don't use just one type of intelligence; rather, they rely on multiple intelligences to perform tasks. Howard Gardner a Harvard psychologist has defined a number of these types of intelligence (for details, see the article "Multiple Intelligences" from Project Zero at Harvard University <https://pz.harvard.edu/resources/the-theory-of-multiple-intelligences>) and knowing them helps you relate them to the kinds of tasks a computer can simulate as intelligence. (See Table 1-1 for a modified version of these intelligences with additional description.)

TABLE 1-1

The Kinds of Human Intelligence and How AIs Simulate Them

Type	Simulation Potential	Human Tools	Description
Bodily kinesthetic	Moderate to High	Specialized equipment and real-life objects	Body movements, such as those used by a surgeon or a dancer, require precision and body awareness. Robots commonly use this kind of intelligence to perform repetitive tasks, often with higher precision than humans, but sometimes with less grace. It's essential to differentiate between human augmentation, such as a surgical device that provides a surgeon with enhanced physical ability, and true independent movement. The former is simply a demonstration of mathematical ability in that it depends on the surgeon for input.
Creative	None	Artistic output, new patterns of thought, inventions, new kinds of musical composition	Creativity is the act of developing a new pattern of thought that results in unique output in the form of art, music, or writing. A truly new kind of product is the result of creativity. An AI can simulate existing patterns of thought and even combine them to create what appears to be a unique presentation but is in reality just a mathematically based version of an existing pattern. In order to create, an AI would need to possess self-awareness, which would require intrapersonal intelligence.
Interpersonal	Low to Moderate	Telephone, audioconferencing, videoconferencing, writing, computer conferencing, email	Interacting with others occurs at several levels. The goal of this form of intelligence is to obtain, exchange, give, or manipulate information based on the experiences of others. Computers can answer basic questions because of keyword input, not because they understand the question. The intelligence occurs while obtaining information, locating suitable keywords, and then giving information based on those keywords. Cross-referencing terms in a lookup table and then acting on the instructions provided by the table demonstrates logical intelligence, not interpersonal intelligence.
Intrapersonal	None	Books, creative materials, diaries, privacy, time	Looking inward to understand one's own interests and then setting goals based on those interests is now a human-only kind of intelligence. As machines, computers have no desires, interests, wants, or creative abilities. An AI processes numeric input using a set of algorithms and provides an output; it isn't aware of anything it does, nor does it understand anything it does.

Type	Simulation Potential	Human Tools	Description
Linguistic (often divided into oral, aural, and written)	Low	Games, multimedia, books, voice recorders, spoken words	Working with words is an essential tool for communication because spoken and written information exchange is far faster than any other form. This form of intelligence includes understanding oral, aural, and written input, managing the input to develop an answer, and providing an understandable answer as output. Discerning just how capable computers are in this form of intelligence is difficult in light of AIs such as ChatGPT because it's all too easy to create tests where the AI produces nonsense answers.
Logical mathematical	High (potentially higher than humans)	Logic games, investigations, mysteries, brainteasers	Calculating results, performing comparisons, exploring patterns, and considering relationships are all areas in which computers now excel. When you see a computer defeat a human on a game show, this is the only form of intelligence you're seeing, out of eight kinds of intelligence. Yes, you might see small bits of other kinds of intelligence, but this is the focus. Basing an assessment of human-versus-computer intelligence on just one area isn't a good idea.
Naturalist	None	Identification, exploration, discovery, new tool creation	Humans rely on the ability to identify, classify, and manipulate their environment to interact with plants, animals, and other objects. This type of intelligence informs you that one piece of fruit is safe to eat though another is not. It also gives you a desire to learn how things work or to explore the universe and all that is in it.
Visual spatial	Moderate	Models, graphics, charts, photographs, drawings, 3D modeling, video, television, multimedia	Physical-environment intelligence is used by people like sailors and architects (among many others). To move around, humans need to understand their physical environment — that is, its dimensions and characteristics. Every robot or portable computer intelligence requires this capability, but the capability is often difficult to simulate (as with self-driving cars) or less than accurate (as with vacuums that rely as much on bumping as they do on moving intelligently).

Examining four ways to define AI

As described in the previous section, the first concept that's important to understand is that AI has little to do with human intelligence. Yes, some AI is modeled to simulate human intelligence, but that's what it is: a simulation. When thinking about AI, notice an interplay between goal seeking, data processing used to achieve

that goal, and data acquisition used to better understand the goal. AI relies on algorithms to achieve a result that may or may not have anything to do with human goals or methods of achieving those goals. With this in mind, you can categorize AI in four ways:

- » Acting humanly
- » Thinking humanly
- » Thinking rationally
- » Acting rationally

Acting humanly

When a computer acts like a human, it best reflects the *Turing test*, in which the computer succeeds when differentiation between the computer and a human isn't possible. (For details, see “The Turing test” at the Alan Turing Internet Scrapbook www.turing.org.uk/scrapbook/test.html). This category also reflects what most media would have you believe AI is all about. You see it employed for technologies such as natural language processing, knowledge representation, automated reasoning, and machine learning (all four of which must be present to pass the test). To pass the Turing test, an AI should have all four previous technologies and, possibly, integrate other solutions (such as expert systems).



TECHNICAL
STUFF

The original Turing test didn't include any physical contact. Harnad's Total Turing Test does include physical contact, in the form of perceptual ability interrogation, which means that the computer must also employ both computer vision and robotics to succeed. Here's a quick overview of other Turing test alternatives:

- » **Reverse Turing test:** A human tries to prove to a computer that the human is not a computer (for example, the Completely Automated Public Turing Test to Tell Computers and Humans Apart, or CAPTCHA).
- » **Minimum intelligent signal test:** Only true/false and yes/no questions are given.
- » **Marcus test:** A computer program simulates watching a television show, and the program is tested with meaningful questions about the show's content.
- » **Lovelace test 2.0:** A test detects AI by examining its ability to create art.
- » **Winograd schema challenge:** This test asks multiple-choice questions in a specific format.

IS THE TURING TEST OUTDATED?

Current discussions about the Turing test have researchers Philip Johnson-Laird, a retired psychology professor from Princeton University, and Marco Ragni, a researcher at the Germany-based Chemnitz University of Technology, asking whether the test is outdated. For example, If AI is making the Turing test obsolete, what might be better? This issue poses several problems with the Turing test and offers a potential solution in the form of a psychological-like evaluation. These tests would use the following three-step process to better test AIs, such as Google's LaMDA and OpenAI's ChatGPT:

- Use tests to check the AI's underlying inferences.
- Verify that the AI understands its own way of reasoning.
- Examine the underlying source code, when possible.

Modern techniques include the idea of achieving the goal rather than mimicking humans completely. For example, the Wright brothers didn't succeed in creating an airplane by precisely copying the flight of birds; rather, the birds provided ideas that led to studying aerodynamics, which eventually led to human flight. The goal is to fly. Both birds and humans achieve this goal, but they use different approaches.

Thinking humanly

A computer that thinks like a human performs tasks that require intelligence (as contrasted with rote procedures) from a human to succeed, such as driving a car. To determine whether a program thinks like a human, you must have some method of determining how humans think, which the cognitive modeling approach defines. This model relies on these three techniques:

- » **Introspection:** Detecting and documenting the techniques used to achieve goals by monitoring one's own thought processes.
- » **Psychological testing:** Observing a person's behavior and adding it to a database of similar behaviors from other persons given a similar set of circumstances, goals, resources, and environmental conditions (among other factors).
- » **Brain imaging:** Monitoring brain activity directly through various mechanical means, such as computerized axial tomography (CAT), positron emission tomography (PET), magnetic resonance imaging (MRI), and magnetoencephalography (MEG).

After creating a model, you can write a program that simulates the model. Given the amount of variability among human thought processes and the difficulty of accurately representing these thought processes as part of a program, the results are experimental at best. This category of thinking humanly is often used in psychology and other fields in which modeling the human thought process to create realistic simulations is essential.

Thinking rationally

Studying how humans think using an established standard enables the creation of guidelines that describe typical human behaviors. A person is considered rational when following these behaviors within certain levels of deviation. A computer that thinks rationally relies on the recorded behaviors to create a guide to how to interact with an environment based on the data at hand.

The goal of this approach is to solve problems logically, when possible. In many cases, this approach would enable the creation of a baseline technique for solving a problem, which would then be modified to actually solve the problem. In other words, the solving of a problem in principle is often different from solving it in practice, but you still need a starting point.

Acting rationally

Studying how humans act in given situations under specific constraints enables you to determine which techniques are both efficient and effective. A computer that acts rationally relies on the recorded actions to interact with an environment based on conditions, environmental factors, and existing data.

As with rational thought, rational acts depend on a solution in principle, which may not prove useful in practice. However, rational acts do provide a baseline on which a computer can begin negotiating the successful completion of a goal.

HUMAN-VERSUS-RATIONAL PROCESSES

Human processes differ from rational processes in their outcome. A process is *rational* if it always does the right thing based on the current information, given an ideal performance measure. In short, rational processes go by the book and assume that the book is correct. Human processes involve instinct, intuition, and other variables that don't necessarily reflect the book and may not even consider the existing data. As an example, the rational way to drive a car is to always follow the law. However, traffic isn't rational. If you follow the law precisely, you end up stuck somewhere because other drivers aren't following the law precisely. To be successful, a self-driving car must therefore act humanly rather than rationally.