

LEARNING MADE EASY



2nd Edition

Artificial Intelligence

for
dummies[®]
A Wiley Brand



Examine AI's security measures and the ethics of employing it

Discover how robots, drones, and self-driving cars use AI

Explore AI's potential in future human endeavors

John Paul Mueller

Author of *MATLAB For Dummies*

Luca Massaron

Google Developer Expert in machine learning

Artificial Intelligence

for
dummies[®]
A Wiley Brand



Artificial Intelligence

2nd Edition

by John Paul Mueller and
Luca Massaron

for
dummies[®]
A Wiley Brand

Artificial Intelligence For Dummies®, 2nd Edition

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

Copyright © 2022 by John Wiley & Sons, Inc., Hoboken, New Jersey

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: WHILE THE PUBLISHER AND AUTHORS HAVE USED THEIR BEST EFFORTS IN PREPARING THIS WORK, THEY 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 ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES REPRESENTATIVES, WRITTEN SALES MATERIALS OR PROMOTIONAL STATEMENTS FOR THIS WORK. THE FACT THAT AN ORGANIZATION, WEBSITE, OR PRODUCT IS REFERRED TO IN THIS WORK AS A CITATION AND/OR POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE PUBLISHER AND AUTHORS ENDORSE THE INFORMATION OR SERVICES THE ORGANIZATION, WEBSITE, OR PRODUCT MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING PROFESSIONAL SERVICES. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR YOUR SITUATION. YOU SHOULD CONSULT WITH A SPECIALIST WHERE APPROPRIATE. FURTHER, READERS SHOULD BE AWARE THAT WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ. NEITHER THE PUBLISHER NOR AUTHORS SHALL BE LIABLE FOR ANY LOSS OF PROFIT OR ANY OTHER COMMERCIAL DAMAGES, INCLUDING BUT NOT LIMITED TO SPECIAL, INCIDENTAL, CONSEQUENTIAL, OR OTHER DAMAGES.

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 such as a CD or DVD 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: 2021948054

ISBN: 978-1-119-79676-3

ISBN 978-1-119-79677-0 (ebk); ISBN ePDF 978-1-119-79678-7 (ebk)

Contents at a Glance

Introduction	1
Part 1: Introducing AI	7
CHAPTER 1: Introducing AI	9
CHAPTER 2: Defining the Role of Data	25
CHAPTER 3: Considering the Use of Algorithms	49
CHAPTER 4: Pioneering Specialized Hardware	67
Part 2: Considering the Uses of AI in Society	83
CHAPTER 5: Seeing AI Uses in Computer Applications	85
CHAPTER 6: Automating Common Processes	97
CHAPTER 7: Using AI to Address Medical Needs	107
CHAPTER 8: Relying on AI to Improve Human Interaction	127
Part 3: Working with Software-Based AI Applications	139
CHAPTER 9: Performing Data Analysis for AI	141
CHAPTER 10: Employing Machine Learning in AI	155
CHAPTER 11: Improving AI with Deep Learning	175
Part 4: Working with AI in Hardware Applications	207
CHAPTER 12: Developing Robots	209
CHAPTER 13: Flying with Drones	223
CHAPTER 14: Utilizing the AI-Driven Car	237
Part 5: Considering the Future of AI	257
CHAPTER 15: Understanding the Nonstarter Application	259
CHAPTER 16: Seeing AI in Space	273
CHAPTER 17: Engaging in Human Endeavors	287
Part 6: The Part of Tens	301
CHAPTER 18: Ten Occupational Categories that AI Can't Replace	303
CHAPTER 19: Ten Substantial Contributions of AI to Society	311
CHAPTER 20: Ten Ways in Which AI Has Failed	319
Index	327

Table of Contents

INTRODUCTION	1
About This Book	2
Icons Used in This Book	3
Beyond the Book	4
Where to Go from Here	5
PART 1: INTRODUCING AI	7
CHAPTER 1: Introducing AI	9
Defining the Term AI	10
Discerning intelligence	10
Discovering four ways to define AI	11
Understanding the History of AI	17
Starting with symbolic logic at Dartmouth	17
Continuing with expert systems	18
Overcoming the AI winters	19
Considering AI Uses	20
Avoiding AI Hype and Overestimation	21
Defining the five tribes and the master algorithm	21
Considering sources of hype	22
Understanding user overestimation	23
Connecting AI to the Underlying Computer	23
CHAPTER 2: Defining the Role of Data	25
Finding Data Ubiquitous in This Age	26
Understanding Moore's implications	27
Using data everywhere	28
Putting algorithms into action	30
Using Data Successfully	32
Considering the data sources	32
Obtaining reliable data	33
Making human input more reliable	33
Using automated data collection	35
Collecting personal data ethically	35
Manicuring the Data	37
Dealing with missing data	37
Considering data misalignments	38
Separating useful data from other data	39
Considering the Five Mistruths in Data	39
Commission	40
Omission	40
Perspective	41

	Bias	42
	Frame of reference	43
	Defining the Limits of Data Acquisition	43
	Considering Data Security Issues	45
	Understanding purposefully biased data	45
	Dealing with data-source corruption	47
	Cancelling botnets with sinkholing	48
CHAPTER 3:	Considering the Use of Algorithms	49
	Understanding the Role of Algorithms	50
	Understanding what algorithm means	50
	Planning and branching: Trees and nodes	52
	Extending the tree using graph nodes	53
	Traversing the graph	54
	Playing adversarial games	56
	Using local search and heuristics	57
	Discovering the Learning Machine	60
	Leveraging expert systems	61
	Introducing machine learning	64
	Touching new heights	64
CHAPTER 4:	Pioneering Specialized Hardware	67
	Relying on Standard Hardware	68
	Understanding the standard hardware	68
	Describing standard hardware deficiencies	69
	Relying on new computational techniques	71
	Using GPUs	73
	Considering the von Neumann bottleneck	73
	Defining the GPU	74
	Considering why GPUs work well	75
	Working with Deep Learning Processors (DLPs)	76
	Defining the DLP	76
	Using the mobile Neural Processing Unit (NPU)	77
	Accessing the cloud-based Tensor Processing Unit (TPU)	78
	Creating a Specialized Processing Environment	78
	Increasing Hardware Capabilities	79
	Adding Specialized Sensors	80
	Devising Methods to Interact with the Environment	81
	PART 2: CONSIDERING THE USES OF AI IN SOCIETY	83
CHAPTER 5:	Seeing AI Uses in Computer Applications	85
	Introducing Common Application Types	86
	Using AI in typical applications	86
	Realizing AI's wide range of fields	88
	Considering the Chinese Room argument	88

Seeing How AI Makes Applications Friendlier	89
Performing Corrections Automatically.	91
Considering the kinds of corrections	91
Seeing the benefits of automatic corrections	92
Understanding why automated corrections don't work	92
Making Suggestions	93
Getting suggestions based on past actions	93
Getting suggestions based on groups	93
Obtaining the wrong suggestions	94
Considering AI-based Errors	95
CHAPTER 6: Automating Common Processes	97
Developing Solutions for Boredom	98
Making tasks more interesting	98
Helping humans work more efficiently	99
Understanding how AI reduces boredom	100
Considering how AI can't reduce boredom	101
Working in Industrial Settings	101
Developing various levels of automation.	102
Using more than just robots	103
Relying on automation alone.	104
Creating a Safe Environment.	104
Considering the role of boredom in accidents	104
Using AI to avoid safety issues.	105
Understanding that AI can't eliminate safety issues	105
CHAPTER 7: Using AI to Address Medical Needs	107
Implementing Portable Patient Monitoring.	108
Wearing helpful monitors	109
Relying on critical wearable monitors	109
Using movable monitors	110
Making Humans More Capable.	111
Using games for therapy	111
Considering the use of exoskeletons	113
Addressing a Range of Physical Abilities	114
Considering the software-based solutions	115
Relying on hardware augmentation	116
Seeing AI in prosthetics	116
Completing Analysis in New Ways	117
Relying on Telepresence.	118
Defining telepresence.	118
Considering examples of telepresence	118
Understanding telepresence limitations	119

Devising New Surgical Techniques	120
Making surgical suggestions	120
Assisting a surgeon	121
Replacing the surgeon with monitoring.	122
Performing Tasks Using Automation	122
Working with medical records.	123
Predicting the future.	123
Making procedures safer	124
Creating better medications	124
Combining Robots and Medical Professionals	125
CHAPTER 8: Relying on AI to Improve Human Interaction	127
Developing New Ways to Communicate	128
Creating new alphabets	129
Working with emoji and other meaningful graphics	129
Automating language translation	130
Incorporating body language.	131
Exchanging Ideas.	133
Creating connections	133
Augmenting communication	133
Defining trends	134
Using Multimedia.	134
Embellishing Human Sensory Perception	135
Shifting data spectrum	135
Augmenting human senses	136
PART 3: WORKING WITH SOFTWARE-BASED AI APPLICATIONS	139
CHAPTER 9: Performing Data Analysis for AI	141
Defining Data Analysis	142
Understanding why analysis is important	144
Reconsidering the value of data	145
Defining Machine Learning	147
Understanding how machine learning works	148
Understanding the benefits of machine learning	149
Being useful; being mundane	150
Specifying the limits of machine learning	150
Considering How to Learn from Data.	152
Supervised learning.	153
Unsupervised learning	154
Reinforcement learning	154

CHAPTER 10: Employing Machine Learning in AI	155
Taking Many Different Roads to Learning	156
Discovering five main approaches to AI learning	156
Delving into the three most promising AI learning approaches.	159
Awaiting the next breakthrough	160
Exploring the Truth in Probabilities	160
Determining what probabilities can do	162
Considering prior knowledge.	163
Envisioning the world as a graph	166
Growing Trees that Can Classify	170
Predicting outcomes by splitting data	170
Making decisions based on trees	172
Pruning overgrown trees	174
CHAPTER 11: Improving AI with Deep Learning	175
Shaping Neural Networks Similar to the Human Brain	176
Introducing the neuron	176
Starting with the miraculous perceptron.	176
Mimicking the Learning Brain	179
Considering simple neural networks	179
Figuring out the secret is in the weights	180
Understanding the role of backpropagation.	182
Introducing Deep Learning	182
Explaining the differences between deep learning and other forms of neural networks	185
Finding even smarter solutions.	186
Detecting Edges and Shapes from Images	188
Starting with character recognition	189
Explaining how convolutions work.	190
Advancing using image challenges.	191
Learning to Imitate Art and Life.	193
Memorizing sequences that matter	193
Discovering the magic of AI conversations	194
Going for the state of the pretrained art.	196
Making one AI compete against another AI.	198
Pondering reinforcement learning.	201
PART 4: WORKING WITH AI IN HARDWARE APPLICATIONS	207
CHAPTER 12: Developing Robots	209
Defining Robot Roles.	210
Overcoming the sci-fi view of robots	211
Being humanoid can be hard	214
Working with robots	217

Assembling a Basic Robot	220
Considering the components	220
Sensing the world	221
Controlling a robot	222
CHAPTER 13: Flying with Drones	223
Acknowledging the State of the Art	224
Flying unmanned to missions	224
Meeting the quadcopter	226
Defining Uses for Drones	227
Seeing drones in nonmilitary roles	229
Powering up drones using AI	233
Understanding regulatory issues	234
CHAPTER 14: Utilizing the AI-Driven Car	237
Getting a Short History	238
Understanding the Future of Mobility	239
Climbing the six levels of autonomy	239
Rethinking the role of cars in our lives	241
Taking a step back from unmet expectations	244
Getting into a Self-Driving Car	246
Putting all the tech together	247
Letting AI into the scene	249
Understanding that it's not just AI	249
Overcoming Uncertainty of Perceptions	251
Introducing the car's senses	252
Putting together what you perceive	254
PART 5: CONSIDERING THE FUTURE OF AI	257
CHAPTER 15: Understanding the Nonstarter Application	259
Using AI Where It Won't Work	260
Defining the limits of AI	260
Applying AI incorrectly	263
Entering a world of unrealistic expectations	264
Considering the Effects of AI Winters	265
Understanding the AI winter	265
Defining the causes of the AI winter	265
Rebuilding expectations with new goals	267
Creating Solutions in Search of a Problem	269
Defining a gizmo	269
Avoiding the infomercial	270
Understanding when humans do it better	271
Looking for the simple solution	272

CHAPTER 16: Seeing AI in Space	273
Observing the Universe	274
Seeing clearly for the first time	274
Finding new places to go	275
Considering the evolution of the universe	276
Creating new scientific principles	276
Performing Space Mining	277
Harvesting water	277
Obtaining rare earths and other metals	278
Finding new elements	279
Enhancing communication	280
Exploring New Places	280
Starting with the probe	281
Relying on robotic missions	282
Adding the human element	283
Building Structures in Space	284
Taking your first space vacation	285
Performing scientific investigation	285
Industrializing space	286
Using space for storage	286
CHAPTER 17: Engaging in Human Endeavors	287
Keeping Human Beings Popular	288
Living and Working in Space	289
Creating Cities in Hostile Environments	290
Building cities in the ocean	290
Creating space-based habitats	291
Constructing moon-based resources	293
Making Humans More Efficient	294
Fixing Problems on a Planetary Scale	296
Contemplating how the world works	296
Locating potential sources of problems	297
Defining potential solutions	298
Seeing the effects of the solutions	299
Trying again	299
PART 6: THE PART OF TENS	301
CHAPTER 18: Ten Occupational Categories that AI Can't Replace	303
Performing Human Interaction	304
Teaching children	304
Nursing	304
Addressing personal needs	305
Solving developmental issues	305

Creating New Things	306
Inventing	306
Being artistic	307
Imagining the unreal	307
Making Intuitive Decisions	308
Investigating crime	308
Monitoring situations in real time	308
Separating fact from fiction	309
CHAPTER 19: Ten Substantial Contributions of AI to Society	311
Considering Human-Specific Interactions	312
Devising the active human foot	312
Performing constant monitoring	313
Administering medications	313
Developing Industrial Solutions	314
Using AI with 3-D printing	314
Advancing robot technologies	314
Creating New Technology Environments	315
Developing rare new resources	315
Seeing what can't be seen	316
Working with AI in Space	316
Delivering goods to space stations	316
Mining extraplanetary resources	317
Exploring other planets	318
CHAPTER 20: Ten Ways in Which AI Has Failed	319
Understanding	320
Interpreting, not analyzing	321
Going beyond pure numbers	321
Considering consequences	322
Discovering	322
Devising new data from old	323
Seeing beyond the patterns	323
Implementing new senses	324
Empathizing	324
Walking in someone's shoes	325
Developing true relationships	325
Changing perspective	326
Making leaps of faith	326
INDEX	327

Introduction

You can hardly avoid hearing about AI today. 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 so many ways. Some people have come to trust AIs so much, that they fall asleep while their SD cars take them to their destination — illegally, of course (see “Tesla driver found asleep at wheel of self-driving car doing 150km/h” at [The Guardian.com](#).)

Many pundits are burying you in information (and disinformation) about AI, too. Some see AI as cute and fuzzy; others see it as a potential mass murderer of the human race. The problem with being so loaded down with information in so many ways is that you struggle to separate what’s real from what is simply the product of an overactive imagination. Just how far can you trust your AI, anyway? Much of the hype about AI originates from the excessive and unrealistic expectations of scientists, entrepreneurs, and businesspersons. *Artificial Intelligence For Dummies*, 2nd Edition is the book you need if you feel as if you really 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 very 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 the really cool uses for AI. For example, you may not know there is a medical monitoring device that can actually 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 today 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*, 2nd 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 won’t be

able to do other ones until far into the future. Some people try to tell you that these activities are possible for AI, but this book tells you why they can't work, clearing away all the hype that has kept you in the dark about AI. You also discover potential security issues in using AI and the kinds of hardware that work best for implementing it. 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, 2nd 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 some cases. Of course, along with the issues, you also discover the fixes for some 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 an SD 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.

For a technology to survive, it must have a group of solid applications that actually work. It also must provide a payback to investors with the foresight to invest in the technology. In the past, AI failed to achieve critical success because it lacked some of these features. AI also suffered from being ahead of its time: True AI needed to wait for the current hardware to actually succeed. Today, you can find AI used in various computer applications and to automate processes. It's also relied on heavily in the medical field and to help improve human interaction. AI is also related to data analysis, machine learning, and deep learning. Sometimes these terms can prove confusing, so one of the reasons to read this book is to discover how these technologies interconnect.

This book also contains an extraordinary number of links to external information (hundreds, in fact) because AI has become such a huge and complex topic. These links provide you with additional information that just won't fit in the book but that you really do need to know to gain a full appreciation of just how astounding the impact of AI is on your daily life. Many of the links are embedded in the name of the article, and print readers can search for them using your favorite search engine; e-book readers can simply click the links. Many other links use what is called a TinyURL (<https://tinyurl.com/>). The reason for using a TinyURL is that the original link is too long and confusing to type into a search engine without the risk of errors. If you want to check a TinyURL to make sure it's real, you

can use the preview feature by adding the word *preview* as part of the link, like this: <https://preview.tinyurl.com/pd88943u>.

AI has a truly bright future today 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 monospace. 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 that 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 nice because they help you save time or perform some task without 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 get the maximum benefit from learning about AI. Just think of them as extras that we're paying to reward you for reading our book.



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 could find that you engage in the sort of disinformation that has people terrified of AI today.



TECHNICAL
STUFF

Whenever you see this icon, think advanced tip or technique. You could fall asleep 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 to create or use an AI solution. Skip these bits of information whenever you like.



REMEMBER

If you don't get anything 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.

Beyond the Book

This book isn't the end of your AI discovery experience; it's really just the beginning. We provide online content to make this book more flexible and better able to meet your needs. That way, as John receives email from you, we can address questions and tell you how updates to AI or its associated technologies affect book content. In fact, you gain access to all these cool additions:

» **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. You can find the cheat sheet for this book by going to www.dummies.com and searching for *Artificial Intelligence For Dummies*. Under the title, click Cheat Sheet and look for the one for this book. The cheat sheet contains really neat information, such as the meaning of all those strange acronyms and abbreviations associated with AI, machine learning, and deep learning.

» **Updates:** Sometimes changes happen. For example, we might not have seen an upcoming change when we looked into our crystal balls during the writing of this book. In the past, that simply meant that the book would become outdated and less useful, but you can now find updates to the book by going to www.dummies.com and searching this book's title.

In addition to these updates, check out the blog posts with answers to readers' questions and for demonstrations of useful book-related techniques at <http://blog.johnmuelเลอร์books.com/>. In addition, you will find blog posts providing information updates as we become aware of them.

Where to Go from Here

It's time to start discovering AI and see what it can do for you. If you don't know anything 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 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 help you see where AI is used today.

Readers who have a bit more advanced knowledge of AI 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

Introducing AI

Artificial Intelligence (AI) has had several false starts and stops over the years, partly because people don't really understand what AI is all about, or even what it should accomplish. A major part of the problem is that movies, television shows, and books have all conspired to give false hopes as to what AI will accomplish. In addition, the human tendency to *anthropomorphize* (give human characteristics to) technology makes it seem as if AI must do more than it can hope to accomplish. So, the best way to start this book is to define what AI actually 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. It doesn't buy into the hype offered by proponents, nor does it indulge in the negativity espoused by detractors. Instead, it strives to give you the best possible view of AI as a technology. As a result, you may find that you have somewhat different expectations than 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 really tell you anything meaningful, which is why there are 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 rest of 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 get 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 based on needs or wants.
2. Assess the value of any currently 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 information that 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 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, but rather rely on multiple intelligences to perform tasks. Howard Gardner of Harvard has defined a number of these types of intelligence (see the article "Multiple Intelligences" from Project Zero at Harvard University for details), and knowing them helps you to relate them to the kinds of tasks that a computer can simulate as intelligence (see Table 1-1 for a modified version of these intelligences with additional description).

Discovering four ways to define AI

As described in the previous section, the first concept that's important to understand is that AI doesn't really have anything 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

TABLE 1-1

The Kinds of Human Intelligence and How AIs Simulate Them

Type	Simulation Potential	Human Tools	Description
Visual-spatial	Moderate	Models, graphics, charts, photographs, drawings, 3-D modeling, video, television, and multimedia	Physical-environment intelligence used by people like sailors and architects (among many others). To move at all, 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).
Bodily-kinesthetic	Moderate to High	Specialized equipment and real 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, and 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 really 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, audio conferencing, video conferencing, writing, computer conferencing, email	Interacting with others occurs at several levels. The goal of this form of intelligence is to obtain, exchange, give, and 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.

Type	Simulation Potential	Human Tools	Description
Intrapersonal	None	Books, creative materials, diaries, privacy, and time	Looking inward to understand one's own interests and then setting goals based on those interests is currently 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 that it does, nor does it understand anything that it does.
Linguistic (often divided into oral, aural, and written)	Low for oral and aural None for written	Games, multimedia, books, voice recorders, and 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. In many cases, computers can barely parse input into keywords, can't actually understand the request at all, and output responses that may not be understandable at all. In humans, oral, aural, and written linguistic intelligence come from different areas of the brain (see "Say What? How the Brain Separates Our Ability to Talk and Write" from John Hopkins University), which means that even with humans, someone who has high written linguistic intelligence may not have similarly high oral linguistic intelligence. Computers don't currently separate aural and oral linguistic ability — one is simply input and the other output. A computer can't simulate written linguistic capability because this ability requires creativity.
Logical-mathematical	High (potentially higher than humans)	Logic games, investigations, mysteries, and brain teasers	Calculating a result, performing comparisons, exploring patterns, and considering relationships are all areas in which computers currently excel. When you see a computer beat a human on a game show, this is the only form of intelligence that you're actually seeing, out of seven 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.

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:** 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 (see "The Turing Test" at the Alan Turing Internet Scrapbook for details). This category also reflects what the 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). Mitsuku (found at <https://chat.kuki.ai/> and <http://www.square-bear.co.uk/mitsuku/home.htm>), a chatbot that won the Loebner Prize five times for the most human-like artificial intelligence, is an example of such integration.

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 convince a computer that that the human is not a computer (for example, the Completely Automatic 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 through examining its ability to create art.
- **Winograd Schema Challenge:** This test asks multiple-choice questions in a specific format.

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 aerodynamics, which eventually led to human flight. The goal is to fly. Both birds and humans achieve this goal, but they use different approaches.



TECHNICAL
STUFF