

---

# Getting Started with Artificial Intelligence

*A Practical Guide to Building Enterprise Applications*

*Tom Markiewicz and Josh Zheng*

Beijing • Boston • Farnham • Sebastopol • Tokyo

**O'REILLY®**



**SHROFF PUBLISHERS & DISTRIBUTORS PVT. LTD.**

*Mumbai*

*Bangalore*

*Kolkata*

*New Delhi*

# Getting Started with Artificial Intelligence

by Tom Markiewicz and Josh Zheng

Copyright © 2018 International Business Machines Corporation. All rights reserved.

ISBN: 978-1-492-02779-9

Originally printed in the United States of America.

Published by O'Reilly Media Inc., 1005 Gravenstein Highway North, Sebastopol, CA 95472.

O'Reilly books may be purchased for educational, business, or sales promotional use. Online editions are also available for most titles (*safari.oreilly.com*). For more information, contact our corporate/institutional sales department: (800) 998-9938 or *corporate@oreilly.com*.

**Editor:** Nicole Tache

**Production Editor:** Justin Billing

**Copyeditor:** Rachel Monaghan

**Proofreader:** Charles Roumeliotis

**Interior Designer:** David Futato

**Cover Designer:** Karen Montgomery

**Illustrator:** Rebecca Demarest

## Printing History:

December 2017: First Edition

**Revision History for the First Edition:** 2017-12-15: First Release,  
2018-04-13: Second Release

**First Indian Reprint:** May 2019

ISBN: 978-93-5213-838-8

The O'Reilly logo is a registered trademark of O'Reilly Media, Inc. *Getting Started with Artificial Intelligence*, the cover image, and related trade dress are trademarks of O'Reilly Media, Inc.

While the publisher and the authors have used good faith efforts to ensure that the information and instructions contained in this work are accurate, the publisher and the authors disclaim all responsibility for errors or omissions, including without limitation responsibility for damages resulting from the use of or reliance on this work. Use of the information and instructions contained in this work is at your own risk. If any code samples or other technology this work contains or describes is subject to open source licenses or the intellectual property rights of others, it is your responsibility to ensure that your use thereof complies with such licenses and/or rights.

This work is part of a collaboration between O'Reilly and IBM. See our *statement of editorial independence*.

**For sale in the Indian Subcontinent (India, Pakistan, Bangladesh, Sri Lanka, Nepal, Bhutan) and African Continent (excluding Morocco, Algeria, Tunisia, Libya, Egypt, and the Republic of South Africa) only. Illegal for sale outside of these countries**

Authorized reprint of the original work published by O'Reilly Media, Inc. All rights reserved. No part of the material protected by this copyright notice may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, nor exported to any countries other than ones mentioned above without the written permission of the copyright owner.

---

Published by **Shroff Publishers and Distributors Pvt. Ltd.** B-103, Railway Commercial Complex, Sector 3, Sanpada (E), Navi Mumbai 400705 • TEL: (91 22) 4158 4158 • FAX: (91 22) 4158 4141  
E-mail: [spdorders@shroffpublishers.com](mailto:spdorders@shroffpublishers.com) • Web: [www.shroffpublishers.com](http://www.shroffpublishers.com)  
CIN : U22200MH1992PTC067760 Printed at Decora Book Prints Pvt. Ltd., Mumbai.

---

# Table of Contents

|  |           |
|--|-----------|
| <b>1. Introduction to Artificial Intelligence.....</b> | <b>1</b>  |
| The Market for Artificial Intelligence                 | 2         |
| Avoiding an AI Winter                                  | 2         |
| Artificial Intelligence, Defined?                      | 3         |
| Applications in the Enterprise                         | 6         |
| Next Steps   | 7         |
| <b>2. Natural Language Processing.....</b>             | <b>9</b>  |
| Overview of NLP  | 10        |
| The Components of NLP                                  | 11        |
| Enterprise Applications of NLP                         | 12        |
| How to Use NLP   | 15        |
| Challenges of NLP                                      | 17        |
| Summary  | 17        |
| <b>3. Chatbots.....</b>                                | <b>19</b> |
| What Is a Chatbot?                                     | 19        |
| The Rise of Chatbots                                   | 20        |
| How to Build a Chatbot                                 | 21        |
| Challenges of Building a Successful Chatbot            | 28        |
| Best Practices   | 28        |
| Industry Case Studies                                  | 34        |
| Summary  | 35        |
| <b>4. Computer Vision.....</b>                         | <b>37</b> |
| Capabilities of Computer Vision for the Enterprise     | 38        |
| How to Use Computer Vision                             | 41        |
| Computer Vision on Mobile Devices                      | 42        |
| Best Practices   | 43        |

|   |           |
|---|-----------|
| Use Cases                                     | 45        |
| Existing Challenges in Computer Vision        | 47        |
| Implementing a Computer Vision Solution       | 48        |
| Summary                                       | 48        |
| <b>5. AI Data Pipeline.....</b>               | <b>49</b> |
| Preparing for a Data Pipeline                 | 51        |
| Sourcing Big Data                             | 51        |
| Storage: Apache Hadoop                        | 52        |
| Hadoop as a Data Lake                         | 52        |
| Discovery: Apache Spark                       | 53        |
| Summary                                       | 54        |
| <b>6. Looking Forward.....</b>                | <b>55</b> |
| What Makes Enterprises Unique?                | 56        |
| Current Challenges, Trends, and Opportunities | 57        |
| Scalability                                   | 60        |
| Social Implications                           | 60        |
| Summary                                       | 61        |

---

# Introduction to Artificial Intelligence

*In the future AI will be diffused into every aspect of the economy.*

—Nils J. Nilsson, Founding researcher, Artificial Intelligence & Computer Science, Stanford University

Beyond the buzzwords, media coverage, and hype, artificial intelligence techniques are becoming a fundamental component of business growth across a wide range of industries. And while the various terms (algorithms, transfer learning, deep learning, neural networks, NLP, etc.) associated with AI are thrown around in meetings and product planning sessions, it's easy to be skeptical of the potential impact of these technologies.

Today's media represents AI in many ways, both good and bad—from the fear of machines taking over all human jobs and portrayals of evil AIs via Hollywood to the much-lauded potential of curing cancer and making our lives easier. Of course, the truth is somewhere in between.

While there are obviously valid concerns about how the future of artificial intelligence will play out (and the social implications), the reality is that the technology is currently used in companies across all industries.

AI is used everywhere—IoT (Internet of Things) and home devices, commercial and industrial robots, autonomous vehicles, drones, digital assistants, and even wearables. And that's just the start. **AI will drive future user experiences that are immersive, continuous, ambient, and conversational.** These conversational services (e.g., chatbots and virtual agents) are currently exploding, while AI will continue to improve these contextual experiences.

Despite several stumbles over the past 60–70 years of effort on developing artificial intelligence, the future is here. If your business is not incorporating at least some AI, you'll quickly find you're at a competitive disadvantage in today's rapidly evolving market.

Just what are these enterprises doing with AI? How are they ensuring an AI implementation is successful (and provides a positive return on investment)? These are only a few of the questions we'll address in this book. From natural language understanding to computer vision, this book will provide you a high-level introduction to the tools and techniques to better understand the AI landscape in the enterprise and initial steps on how to get started with AI in your own company.

It used to be that AI was quite expensive, and a luxury reserved for specific industries. Now it's accessible for everyone in every industry, including you. We'll cover the modern enterprise-level AI techniques available that allow you to both create models efficiently and implement them into your environment. While not meant as an in-depth technical guide, the book is intended as a starting point for your journey into learning more and building applications that implement AI.

## The Market for Artificial Intelligence

The market for artificial intelligence is already large and growing rapidly, with numerous research reports indicating a growing demand for tools that automate, predict, and quickly analyze. **Estimates from IDC** predict revenue from artificial intelligence will top \$47 billion by the year 2020 with a compound annual growth rate (CAGR) of 55.1% over the forecast period, with nearly half of that going to software. Additionally, investment in AI and machine learning companies has increased dramatically—**AI startups have raised close to \$10 billion in funding**. Clearly, the future of artificial intelligence appears healthy.

## Avoiding an AI Winter

Modern AI as we know it started in earnest in the 1950s. While it's not necessary to understand the detailed history of AI, it is helpful to understand one particular concept—the *AI winter*—as it shapes the current environment.

There were two primary eras of artificial intelligence research where high levels of excitement and enthusiasm for the technology never lived up to expectations, causing funding, interest, and continued development to dry up. The buildup of hype followed by disappointment is the **definition of an AI winter**.

So why are we now seeing a resurgence in AI interest? What's the difference today that's making AI so popular in the enterprise, and should we fear another AI winter? The short answer is likely no—we expect to avoid another AI winter this time around due primarily to much more (or big) data and the advent of better processing power and GPUs. From the tiny supercomputers we all carry in our pockets to the ever-expanding role of IoT, we're generating more data now, at an ever-increasing rate. For example, **IDC estimates** that 180 zettabytes of data will be created globally in 2025, up from less than 10 zettabytes in 2015.

Andrew Ng, cofounder of Coursera and Stanford adjunct professor, often presents Figure 1-1 in his courses on machine learning and deep learning.

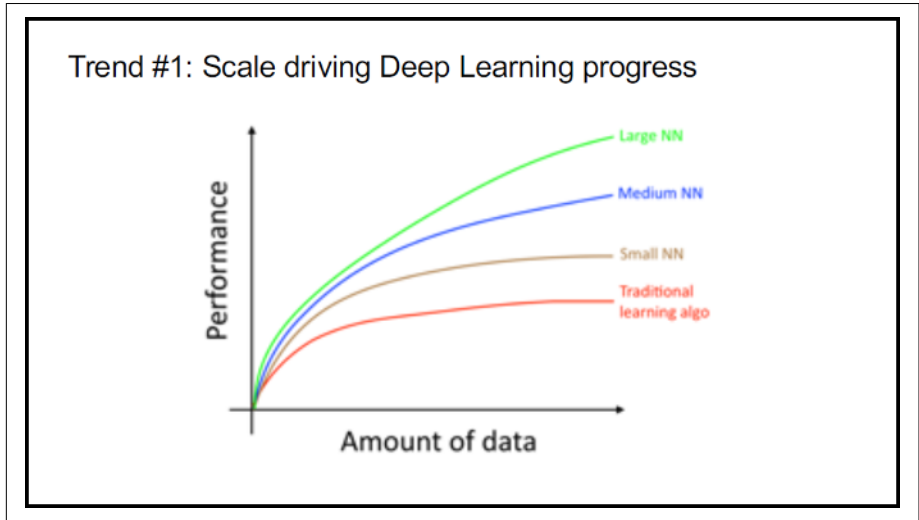


Figure 1-1. Deep learning performance (image courtesy of Andrew Ng, *Deeplearning.ai* course on Coursera)

Conceptually, this chart illustrates how the performance of deep learning algorithms improves with an increasing amount of data—data that we now have in abundance and that is growing at an exponential rate.

So why does it matter whether we understand the AI winter? Well, if companies are going to invest in AI, it's essential to avoid the hyperbole of the past and keep expectations based in reality. While the current situation is much more likely to justify the enthusiasm, that excitement still needs to be tempered with real-world results to avoid another AI winter. As (future) AI practitioners, that's something we can all agree would challenge your businesses.

## Artificial Intelligence, Defined?

The market for artificial intelligence is immense, but what are we truly discussing? While it sounds great to say you're going to implement AI in your business, just what does that mean in practical terms? Artificial intelligence is quite a broad term and is, in reality, an umbrella over a few different concepts.

So, to get started and keep everyone on the same page, let's briefly discuss some of the terms associated with AI that are often confused or interchanged: artificial intelligence, machine learning, and deep learning.

# Artificial Intelligence

Over the years, there have been numerous attempts at precisely defining AI. While there's never really been an official, accepted definition, there are a few high-level concepts that shape and define what we mean. The descriptions range from artificial intelligence merely aiming to augment human intelligence to thinking machines and true computer intelligence. For our purposes, we can think of AI as being any technique that allows computers to bring meaning to data in similar ways to a human.

While most AI is focused on specific problem domains (natural language processing or computer vision, for example), the idea of **artificial general intelligence** or having a machine perform any task a human could (also commonly referred to as “strong AI”) is still more than 10 years out **according to Gartner**.

## Machine Learning

In 1959 Arthur L. Samuel, an IBM researcher and Stanford professor, is said to have stated that “machine learning is the field of study that gives computers the ability to learn without being explicitly programmed,” thereby **becoming the originator of the term**.

Essentially, machine learning is a subset of AI focused on having computers provide insights into problems without explicitly programming them to do so. Most of the tools and techniques that today refer to AI are representative of machine learning. There are three main types of machine learning—supervised learning, unsupervised learning, and reinforcement learning.

By looking at many labeled data points and examples of historical problems, supervised learning algorithms can help solve similar problems under new circumstances. Supervised learning trains on large volumes of historical data and then builds general rules to be applied to future problems. The better the training set data, the better the output. In supervised learning, the system learns from these human-labeled examples.

While supervised learning relies on labeled or structured data (think rows in a database), unsupervised learning trains on unlabeled or unstructured data (the text of a book). These algorithms explore the data and try to find structure. Here, widely used unsupervised learning algorithms are cluster analysis and market basket analysis. Naturally, this tends to be more difficult, as the data has no pre-existing labels to assist the algorithms in understanding the data. While more challenging to process, as we'll discuss later, unstructured data makes up the vast majority of data that enterprises need to process today.

Finally, there's reinforcement learning as a machine learning technique. Reinforcement learning takes an approach similar to behavioral psychology. Instead of training a model with predefined training sets (i.e., where you know the pre-



scribed answers in advance) as in supervised learning, **reinforcement learning rewards the algorithm** when it performs the correct action (behavior). Reinforcement learning resists providing too much training and allows the algorithm to optimize itself for performance-based rewards.

Reinforcement learning was initially conceived in 1951 by Marvin Minsky, but as was the case with many AI implementations, the algorithms were held back by both the scale of data and computer processing power needed for effectiveness. Today, we see many more successful examples of reinforcement learning in the field, with Alphabet subsidiary DeepMind's **AlphaGo** one of the more prominent. Another notable application of reinforcement learning is the development of self-driving cars.

## Deep Learning

No book on developing AI applications in the enterprise would be complete without a discussion of deep learning. One way to understand deep learning is to think of it as a subset of AI that attempts to develop computer systems that learn using neural networks like those in the human brain. While machine learning is mostly about optimization, deep learning is more focused on creating algorithms to simulate how the human brain's neurons work.

Deep learning algorithms are composed of an interconnected web of nodes called neurons and the edges that join them together. **Neural nets receive inputs, perform calculations, and then use this output to solve given problems.**

One of the ultimate goals of AI is to have computers think and learn like human beings. Using neural networks that are based on the human brain's decision-making process, deep learning is a set of computational techniques that move us closest to that goal.

According to Stanford professor Chris Manning, around the year 2010 **deep learning techniques started to outperform other machine learning techniques.** Shortly after that, in 2012, Geoffrey Hinton at the University of Toronto led a team creating a neural network that learned to recognize images. Additionally, that year Andrew Ng led a team at Google that created a system to **recognize cats in YouTube videos** without explicitly training it on cats.

By building multiple layers of abstraction, deep learning technology can also solve complex semantic problems. Traditional AI follows a linear programming structure, thus limiting the ability for a computer to learn without human intervention. However, **deep learning replaces this existing approach** with self-learning and self-teaching algorithms enabling more advanced interpretation of data. For example, traditional AI could read a zip code and mailing address off of an envelope, but deep learning could also infer from the envelope's color and date sent that the letter contains a holiday card.

Deep learning's closer approximation to how the human brain processes information via neural networks offers numerous benefits. First, the approach delivers a significantly improved ability to process large volumes of unstructured data, finding meaningful patterns and insights. Next, deep learning offers the ability to develop models across a wide variety of data, highlighted by advanced approaches toward text, images, audio, and video content. Finally, deep learning's massive parallelism allows for the use of multiple processing cores/computers to maximize overall performance as well as faster training with the advent of GPUs.

As with all advances in AI, deep learning benefits from much more training data than ever before, faster machines and GPUs, and improved algorithms. These neural networks are driving rapid advances in speech recognition, image recognition, and machine translation, with some systems performing as well or better than humans.

## Applications in the Enterprise

From finance to cybersecurity to manufacturing, there isn't an industry that will not be affected by AI. But before we can discuss and examine building applications in the enterprise with AI, we first need to define what we mean by enterprise applications. There are two common definitions of "the enterprise":

- A company of significant size and budget
- Any business-to-business commerce (i.e., not a consumer)

So Geico, Procter & Gamble, IBM, and Sprint would all be enterprise companies. And by this definition, any software designed and built internally would be considered enterprise software. On the other hand, a small company or startup could develop applications to be used by businesses (whether large or small), and this would still be considered enterprise software. However, a photo sharing app for the average consumer would not be considered enterprise software.

This is probably obvious, but since we're discussing enterprise applications with AI in this book, it's important to be explicit about just what we mean when talking about the enterprise. For the rest of the book, the context of an enterprise will be that the end user is a business or business employee.

Does that mean if you are building the next great consumer photo sharing application that this book will be of no use? Absolutely not! Many (if not most) of the concepts discussed throughout the book can be directly applied to consumer-facing applications as well. However, the use cases and discussion will be centered on the enterprise.

# Next Steps

This book is intended to provide a broad overview of artificial intelligence, the various technologies involved, relevant case studies and examples of implementation, and the current landscape. Also, we'll discuss the future of AI and where we see it moving in the enterprise from a current practitioner's viewpoint.

While the book focuses on a more technical audience, specifically application developers looking to apply AI in their business, the overall topics and discussion will be useful for anyone interested in how artificial intelligence will impact the enterprise.

Note that this book is not a substitute for a course or deep study on these topics, but we hope the material covered here jump-starts the process of your becoming a proficient practitioner of AI technologies.

While there are many options for implementing the various AI techniques we'll discuss in this book—coding in-house from scratch, outsourcing, the use of open source libraries, software-as-a-service APIs from leading vendors—our familiarity and deep experience revolves around IBM Watson, so we'll be using those code examples to illustrate the individual topics as necessary.

It's important to note that, while we believe IBM Watson provides a great end-to-end solution for applying AI in the enterprise, it's not the only option. All the major cloud computing providers offer similar solutions. Also, there are numerous open source tools that we'll discuss later.

In the next few chapters, we'll go into more depth on some of the more common uses for AI in the enterprise (natural language processing, chatbots, and computer vision). Then we'll discuss the importance of a solid data pipeline, followed by a look forward at the challenges and trends for AI in the enterprise.

