

# Artificial Intelligence for Business

© 2020 Jeffrey L. Coveyduc and Jason L. Anderson

Published by John Wiley & Sons, Inc., Hoboken, New Jersey.  
Published simultaneously in Canada.

The right of Jason L. Anderson and Jeffrey L. Coveyduc to be identified as the author(s) of the editorial material in this work has been asserted in accordance with law.

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 Section 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 646-8600, or on the Web at [www.copyright.com](http://www.copyright.com). 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 [www.wiley.com/go/permissions](http://www.wiley.com/go/permissions).

Limit of Liability/Disclaimer of Warranty: While the publisher and author have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. Neither the publisher nor author 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 or for technical support, please contact our Customer Care Department within the United States at (800) 762-2974, outside the United States at (317) 572-3993, or fax (317) 572-4002.

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](http://www.wiley.com).

#### *Library of Congress Cataloging-in-Publication Data*

Names: Anderson, Jason L, author. | Coveyduc, Jeffrey L, author.

Title: Artificial intelligence for business : a roadmap for getting started with AI / Jason L. Anderson, Jeffrey L. Coveyduc.

Description: First edition. | Hoboken : Wiley, 2020. | Includes index.

Identifiers: LCCN 2020004359 (print) | LCCN 2020004360 (ebook) | ISBN 9781119651734 (hardback) | ISBN 9781119651413 (adobe pdf) | ISBN 9781119651802 (epub)

Subjects: LCSH: Artificial intelligence—Economic aspects. | Business enterprises—Technological innovations. | Artificial intelligence—Data processing.

Classification: LCC HC79.I55 .A527 2020 (print) | LCC HC79.I55 (ebook) | DDC 006.3068—dc23

LC record available at <https://lccn.loc.gov/2020004359>

LC ebook record available at <https://lccn.loc.gov/2020004360>

Cover Design: Wiley

Cover Image: © Yuichiro Chino/Getty Images

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

# Contents

---

<i><a href="#">Preface</a></i>	ix
<i><a href="#">Acknowledgments</a></i>	xi
<b>CHAPTER 1</b> <a href="#">Introduction</a>	1
<a href="#">Case Study #1: FANUC Corporation</a>	2
<a href="#">Case Study #2: H&amp;R Block</a>	4
<a href="#">Case Study #3: BlackRock, Inc.</a>	5
<a href="#">How to Get Started</a>	6
<a href="#">The Road Ahead</a>	10
<a href="#">Notes</a>	11
<b>CHAPTER 2</b> <a href="#">Ideation</a>	13
<a href="#">An Artificial Intelligence Primer</a>	13
<a href="#">Becoming an Innovation-Focused Organization</a>	23
<a href="#">Idea Bank</a>	25
<a href="#">Business Process Mapping</a>	27
<a href="#">Flowcharts, SOPs, and You</a>	28
<a href="#">Information Flows</a>	29
<a href="#">Coming Up with Ideas</a>	31
<a href="#">Value Analysis</a>	31
<a href="#">Sorting and Filtering</a>	34
<a href="#">Ranking, Categorizing, and Classifying</a>	35
<a href="#">Reviewing the Idea Bank</a>	37
<a href="#">Brainstorming and Chance Encounters</a>	38
<a href="#">AI Limitations</a>	41
<a href="#">Pitfalls</a>	44
<a href="#">Action Checklist</a>	45
<a href="#">Notes</a>	46

<b>CHAPTER 3</b>	<u>Defining the Project</u>	47
	<u>The <i>What</i>, <i>Why</i>, and <i>How</i> of a Project Plan</u>	48
	<u>The Components of a Project Plan</u>	49
	<u>Approaches to Break Down a Project</u>	53
	<u>Project Measurability</u>	62
	<u>Balanced Scorecard</u>	63
	<u>Building an AI Project Plan</u>	64
	<u>Pitfalls</u>	66
	<u>Action Checklist</u>	69
<b>CHAPTER 4</b>	<u>Data Curation and Governance</u>	71
	<u>Data Collection</u>	73
	<u>Leveraging the Power of Existing Systems</u>	81
	<u>The Role of a Data Scientist</u>	81
	<u>Feedback Loops</u>	82
	<u>Making Data Accessible</u>	84
	<u>Data Governance</u>	85
	<u>Are You Data Ready?</u>	89
	<u>Pitfalls</u>	90
	<u>Action Checklist</u>	94
	<u>Notes</u>	94
<b>CHAPTER 5</b>	<u>Prototyping</u>	97
	<u>Is There an Existing Solution?</u>	97
	<u>Employing vs. Contracting Talent</u>	99
	<u>Scrum Overview</u>	101
	<u>User Story Prioritization</u>	103
	<u>The Development Feedback Loop</u>	105
	<u>Designing the Prototype</u>	106
	<u>Technology Selection</u>	107
	<u>Cloud APIs and Microservices</u>	110
	<u>Internal APIs</u>	112
	<u>Pitfalls</u>	112
	<u>Action Checklist</u>	114
	<u>Notes</u>	114
<b>CHAPTER 6</b>	<u>Production</u>	117
	<u>Reusing the Prototype vs. Starting from a Clean Slate</u>	117

	<a href="#">Continuous Integration</a>	119
	<a href="#">Automated Testing</a>	124
	<a href="#">Ensuring a Robust AI System</a>	128
	<a href="#">Human Intervention in AI Systems</a>	129
	<a href="#">Ensure Prototype Technology Scales</a>	131
	<a href="#">Cloud Deployment Paradigms</a>	133
	<a href="#">Cloud API's SLA</a>	135
	<a href="#">Continuing the Feedback Loop</a>	135
	<a href="#">Pitfalls</a>	135
	<a href="#">Action Checklist</a>	137
	<a href="#">Notes</a>	137
<b>CHAPTER 7</b>	<b>Thriving with an AI Lifecycle</b>	139
	<a href="#">Incorporate User Feedback</a>	140
	<a href="#">AI Systems Learn</a>	142
	<a href="#">New Technology</a>	144
	<a href="#">Quantifying Model Performance</a>	145
	<a href="#">Updating and Reviewing the Idea Bank</a>	147
	<a href="#">Knowledge Base</a>	148
	<a href="#">Building a Model Library</a>	150
	<a href="#">Contributing to Open Source</a>	155
	<a href="#">Data Improvements</a>	157
	<a href="#">With Great Power Comes Responsibility</a>	158
	<a href="#">Pitfalls</a>	159
	<a href="#">Action Checklist</a>	161
	<a href="#">Notes</a>	161
<b>CHAPTER 8</b>	<b>Conclusion</b>	163
	<a href="#">The Intelligent Business Model</a>	164
	<a href="#">The Recap</a>	164
	<a href="#">So What Are You Waiting For?</a>	168
<b>APPENDIX A</b>	<b>AI Experts</b>	169
	<a href="#">AI Experts</a>	169
	<a href="#">Chris Ackerson</a>	169
	<a href="#">Jeff Bradford</a>	173
	<a href="#">Nathan S. Robinson</a>	175
	<a href="#">Evelyn Duesterwald</a>	177

Jill Nephew	179
Rahul Akolkar	183
Steven Flores	187
<b>APPENDIX B</b> Roadmap Action Checklists	191
Step 1: Ideation	191
Step 2: Defining the Project	191
<a href="#">Step 3: Data Curation and Governance</a>	<a href="#">192</a>
<a href="#">Step 4: Prototyping</a>	<a href="#">192</a>
<a href="#">Step 5: Production</a>	<a href="#">193</a>
<a href="#">Thriving with an AI Lifecycle</a>	<a href="#">193</a>
<b>APPENDIX C</b> <a href="#">Pitfalls to Avoid</a>	<a href="#">195</a>
<a href="#">Step 1: Ideation</a>	<a href="#">195</a>
<a href="#">Step 2: Defining the Project</a>	<a href="#">196</a>
Step 3: Data Curation and Governance	199
<a href="#">Step 4: Prototyping</a>	<a href="#">203</a>
Step 5: Production	204
<a href="#">Thriving with an AI Lifecycle</a>	<a href="#">206</a>
<b><i>Index</i></b>	<a href="#">209</a>

# Preface

---

Artificial intelligence (AI) has become so ingrained in our daily lives that most people knowingly leverage it every day. Whether interacting with an artificial “entity” such as the iPhone assistant Siri, or browsing through Netflix’s recommendations, our functional adoption of machine learning is already well under way. Indirectly, however, AI is even more prevalent. Every credit card purchase made is run through fraud detection AI to help safeguard customers’ money. Advanced logistical scheduling software is used to deliver tens of millions of packages daily, to locales around the world, with minimal disruption. In fact, the e-commerce giant Amazon alone claims to have shipped 5 billion packages with Prime in 2017 (see [businesswire.com/news/home/20180102005390/en/](http://businesswire.com/news/home/20180102005390/en/)). None of this would be possible on such a grand scale without the advances we have seen in AI systems and in machine learning technology over the last few decades.

Historically, these AI systems have been developed in-house by skilled teams of programmers, working around the clock at great expense to employers. This reality is now changing. Companies like IBM, Google, and Microsoft are making AI capabilities available on a pay-as-you-go basis, dramatically lowering the barrier to entry. For example, each of these companies provide speech-to-text and text-to-speech services to easily build voice interfaces for pennies per use. This is opening the door for smaller companies with less disposable capital to introduce AI initiatives that will produce substantial results. With the aforementioned backdrop of consumers interacting with AI on a daily basis, these consumers are becoming increasingly more comfortable and receptive to their brands adopting and incorporating more AI technology. The combination of all of these

components makes it a smart bet for any modern company to start down the road toward AI adoption.

But how do these companies get started? This question is one we have seen time and again working with clients in the AI space. The drive and enthusiasm are there, but what organizational thought leaders are missing is the “how to” and overall direction. In our day jobs working with IBM Watson Client Engagement Centers and clients around the world, we repeatedly saw this pattern play out. Clients were eager to incorporate AI systems into their business models. They understood many of the benefits. They just needed a way in. While attending tech conferences and meetups, we find similar stories as well. Though the technological barriers are lower, with vendors providing accessible AI technology in the cloud, the challenge of coming up with the overall plan was still preventing many businesses from adopting AI. Having a good roadmap is essential to feeling comfortable with starting the journey. It is for this very reason that we wrote this book. Our goal is to empower you with the knowledge to successfully adopt AI technology into your organization. And you’ve already taken the first step by opening this book.

In addition to helping you adopt and understand emerging AI technology, this book will give you the tools to use AI to make a measurable impact in your business. Perhaps you will find some new cost-saving opportunities to unlock. Maybe AI will allow your business to uniquely position itself to enter new markets and take on competitors. Although AI has become more widespread and mainstream in its use in recent years, we are still seeing a tremendous amount of room for disruption in every field. That’s the great thing about AI—it can be applied in an interdisciplinary fashion to all domains, and the more it grows, the more its capabilities grow along with it. All that we ask of you, the reader, is to start with an open mind while we provide that missing roadmap to help you successfully navigate your way to driving value within your organization using AI.



## Introduction

The modern era has embedded code in everything we use. From your washing machine to your car, if it was made any time in the last decade, there is likely code inside it. In fact, the term “Internet of Things (IoT)” has emerged to define all Internet-connected devices that are not strictly computers. Although the code on these IoT devices is becoming smarter with every upgrade, the devices are not exactly learning autonomously. A programmer has to code every new feature or decision into a model. These programs do not learn from their mistakes. Advancement in AI will help solve this problem, and soon we will have devices that *will* learn from the input of their human creators, as well as from their own mistakes. Today we are surrounded by code, and in the near future, we will be surrounded by embedded artificially intelligent agents. This will be a massive opportunity for upgrades and will enable more convenience and efficiency.

Although companies may have implemented software projects on their own or with the help of outside vendors in the past, AI projects have their own set of quirks. If those quirks are not managed properly, they may cause a project to be a failure. A brilliant idea must be paired with brilliant execution in order to succeed. Following the path laid out in this book will put you on a trajectory toward managing AI projects more efficiently, as well as prepare you for the

age of intelligent systems. Artificial intelligence is very likely to be the next frontier of technology, and in order for us to maximize this opportunity, the groundwork must be laid today.

Every organization is different, and it is important to remember not to try to apply techniques like a straitjacket. Doing so will suffocate your organization. This book is written with a mindset of best practices. Although best practices will work in most cases, it is important to remain attentive and flexible when considering your own organization's transformation. Therefore, you must use your best judgment with each recommendation we make. There is no one-size-fits-all solution, especially not in a field like AI that is constantly evolving.

Ahead of the recent boom in AI technologies, many organizations have already successfully implemented intelligent solutions. Most of these organizations followed an adoption roadmap similar to the one we will describe in this book. It is insightful for us to take a look at a few of these organizations, see what they implemented, and take stock of the benefits they are now realizing. As you read through these organizations' stories, keep in mind that we will be diving into aspects of each approach in more detail during the course of this book.

## Case Study #1: FANUC Corporation

---

Science fiction has told of factories that run entirely by themselves, constantly monitoring and adjusting their input and output for maximum efficiency. Factories that can do just-in-time (JIT) ordering based on sales demand, sensors that predict maintenance requirements, the ability to minimize downtime and repair costs—these are no longer concepts of speculative fiction. With modern sensors and AI software, it has become possible to build these efficient, self-bolstering factories. Out-of-the-box IoT equipment can do better monitoring today than industrial sensors from 10 years ago. This leap in accuracy and connectivity has increased production threshold limits, enabling industrial automation on a scale never before imagined.

FANUC Corporation of Japan,<sup>1</sup> a manufacturer of robots for factories, leads by example. Its own factories have robots building other robots with minimal human intervention. Human workers are able to focus on managerial tasks, whereas robots are built in the dark. This gives a whole new meaning to the industry saying “lights-out operations,” which originally meant servers, not robots with moving parts, running independently in a dark data center. FANUC Japan has invested in Preferred Networks Inc. to gather data from their own robots to make them more reliable and efficient than ever before. Picking parts from a bin with an assortment of different-sized parts mixed together has been a hard problem to solve with traditional coding. With AI, however, FANUC has managed to achieve a consistent 90 percent accuracy in part identification and selection, tested over some 5,000 attempts. The fact that minimal code has gone into allowing these robots to achieve their previously unobtainable objective is yet another testament to the robust capabilities of AI in the industrial setting. FANUC and Preferred Networks have leveraged the continuous stream of data available to them from automated plants, underlining the fact that data collection and analysis is critical to the success of their factory project. FANUC Intelligent Edge Link & Drive (FIELD) is the company’s solution for data collection to be later implemented using deep learning models. The AI Bin-Picking product relies on models created via the data collected from the FIELD project. Such data collection procedures form a critical backbone for any industrial process that needs to be automated.

FANUC has also enabled deep learning<sup>2</sup> models for situations where there are too many parameters to be fine-tuned manually. Such models include AI servo-tuning processes that enable high-precision, high-speed machining processes that were not possible until recently. In the near future, your Apple iPhone case will probably be made using a machine similar to the one in Figure 1.1.

Most factories today are capable of utilizing these advancements with minor modifications to their processes. The gains that can be achieved from such changes will be able to exponentially elevate the output of any factory.




FIGURE 1.1 Example of a FANUC Robot<sup>3</sup>

## Case Study #2: H&R Block

H&R Block is a U.S.-based company that specializes in tax preparation services. One of their customer satisfaction guarantees is to find the maximum number of tax deductions for each of their customers. Some deductions are straightforward, such as homeowners being able to deduct the mortgage interest on their primary residence. Other deductions, however, may be dependent on certain client-specific variables, such as the taxpayer's state of residence. Deduction complexity can then be further compounded by requiring multiple client-dependent variables to be considered simultaneously, such as a taxpayer with multiple sources of income who also has multiple personal deductions. The ultimate result is that maximizing deductions for a given customer can be difficult, even for a seasoned tax professional. H&R Block saw an opportunity to leverage AI to help their tax preparers optimize their service. In order to help facilitate the adoption process, H&R Block partnered with IBM to leverage their Watson capabilities.<sup>4</sup>

When a customer comes into H&R Block, the tax preparer engages them in a friendly discussion. “Have you experienced any life-changing events in the last year?” “Have you purchased a home?” and so on. As they talk, the tax preparer types relevant details of the conversation into their computer system to be used as reference later. If the customer mentions that they purchased a house last year, that will be an indicator that they may qualify for a mortgage interest deduction this year.

H&R Block saw the opportunity here to leverage the use of AI to compile, cross-reference, and analyze all of these notes. Natural language processing (NLP) can be applied to identify the core intent of each note, which then can be fed into the AI system to automatically identify possible deductions. The system then presents the tax professionals with any potentially relevant information to ensure that they do not miss any possible deductions. In the end, both tax professionals and their customers can enjoy an increased sense of confidence that every last applicable deduction was found.

### **Case Study #3: BlackRock, Inc.**

Financial markets are a hotbed for data. The data can be collected accurately and in real time for most financial instruments (stocks, options, funds, etc.) listed on stock markets. Metadata (data about data) can also be curated from analytical reports, articles, and the like. The necessity for channeling the sheer amount of information that is generated every day has given rise to professional data stream providers like Bloomberg. The immense quantity of data available, along with the potential for trend prediction, growth estimations, and increasingly accurate risk assessment, makes the financial industry ripe for implementing AI projects.

BlackRock, Inc., one of the world’s largest asset managers, deploys the Aladdin<sup>5</sup> (Asset, Liability, Debt, and Derivative Investment Network) software, which calculates risks, analyzes financial data, supports investment operations, and offers trade executions. Aladdin’s key strength lies in using the vast amount of data to arrive at models of risk that give the user more confidence in deploying investments and hedging. The project was started nearly two decades ago, and it

User stories will also be a large part of the project plan. User stories are an excellent way to break down a project into functional pieces of value. They define a user, the functionality that the system will provide for the user, and the value that the function will provide to the organization. Well-defined user stories also quantify their results to empirically know when success has been achieved. These success criteria make it much easier to see when we have accomplished our user story's goal and communicate a clear course of action for everyone involved. Specificity is the key.

### 3. Data Curation and Governance

Data is paramount to every AI system. A system can only be as good as the data that is used to build it. Therefore, it is important to take stock of all the possible data sources at your disposal. This is true whether it is data being collected and stored internally or data that you externally license.

After you have identified your data, it is time to leverage technology to further improve the data's quality and prepare it to train an AI system. Crowdsourcing can be a valuable tool to enhance existing data, and data platforms such as Apache Hadoop can help consolidate data from multiple sources. Data scientists will be key in orchestrating this process and ensuring success. The quality of your data will determine the success of your project in a huge way. It is therefore essential to choose the best available data on hand. The old saying about "garbage in, garbage out" applies to AI as well.

### 4. Prototyping

With your project plan and data defined, it is time to start building an initial version your system. As with any project, it is best to take an iterative approach. In the prototype step, you will select a subset of your use cases to validate the idea. In this way, you are able to see if the expected value is materializing before you are completely invested. This step also enables you to adjust your approach early if you see any problems arise. Developing a prototype will help you to see, with actual results, whether the ideas and plans you defined in the previous steps have promise. In the event that they do not, you

should be able to recover quickly and adjust them using the knowledge gained from prototyping, without the wasted investment of building a full system.

During the prototyping phase, it is necessary to have realistic expectations. With most AI systems, they improve with more data and parameter tweaking, so you should expect to see increasing improvements over time. Luckily, metrics like precision and recall can be empirically measured and used to track this improvement. We will also cover the cases when more data is not the answer and what other techniques can be pursued to continue improving the system.

## 5. Production

With a successful prototype under your belt, you have been able to see the value of the technology in action. Now it is time to further invest and complete your system. At this point, it is also a good idea to revisit your user stories and plan as a whole to determine if any priorities have changed. You can then proceed with building the production system.

The production step is the process of converting the prototype into a full-fledged system. This includes conducting a technological evaluation, building user security models, and establishing testing frameworks.

### Technological Evaluation

During the prototype phase, developers select technologies appropriate for a prototype, including using technologies and languages that are easy to work with. This mitigates risk by determining the project's feasibility quickly before investing a lot of time and money. That said, during the production step the technology must be evaluated for other factors as well. For instance, will the technology scale to a large number of users or massive amounts of data? Will the technology be supported in the long term and be flexible enough to change as requirements do? If not, pieces of the prototype might have to be rebuilt to accommodate.

### User/Security Model

During the prototype phase, the project is typically only running on locked-down development machines or internal servers.

While they require some security, high levels of security are not typically needed during prototyping and will only slow down the prototyping process. Work, such as integrating an organization's user directory (single sign-on [SSO]) and permission structures, will be part of the production process.

### Testing Frameworks

To ensure code quality, testing frameworks should be built alongside the production code. Testing ensures that the code base does not regress as new code is added. Development teams may even adopt a "test first" approach called test-driven development (TDD) to ensure that all pieces of code have tests written before starting their implementation. If TDD is used, developers repeat very short development cycles, writing only enough code for the tests to pass. In this way, tests reflect the desired functionality and code is written to implement that functionality.

## Thriving with an AI Lifecycle

Once you have adopted AI and your organization is realizing its benefits, it is time to switch into the lifecycle mode. At this point, you will be maintaining your AI systems while consistently looking for ways to improve. This might mean leveraging system usage data to improve your machine learning models or keeping an eye on the latest technology announcements. Perhaps the AI models you have implemented can also be used in another part of your organization. Furthermore, it is important that the knowledge gained during the implementation of your first AI system be saved for future projects. As we will discuss in this book, this can take the shape of either an entry in your organization's model library or a lessons learned document.

## The Road Ahead

---

Adopting artificial intelligence in your organization can feel like a daunting task, especially since the technology is changing so frequently. The main idea is to be aware of all the benefits, as well as the pitfalls, so that you can adequately discern between them and



navigate your way to success. Mistakes are inevitable. Keeping them small and easy to recover from will ensure that your AI transformation has the resilience it needs to prevail. To minimize the likelihood of mistakes, we list the common pitfalls associated with each step at the end of each chapter so you can take notice and avoid them. With sufficient planning and foresight provided by this book, you will be able to acquire the tools necessary to make your organizational adoption of AI a great success.

## Notes

---

1. <https://preferred.jp/en/news/tag/fanuc/>
2. [www.bloomberg.com/news/features/2017-10-18/this-company-s-robots-are-making-everything-and-reshaping-the-world](http://www.bloomberg.com/news/features/2017-10-18/this-company-s-robots-are-making-everything-and-reshaping-the-world)
3. [https://en.wikipedia.org/wiki/File:FANUC\\_R2000iB\\_AtWork.jpg](https://en.wikipedia.org/wiki/File:FANUC_R2000iB_AtWork.jpg)
4. [www.hrblock.com/tax-center/newsroom/around-block/partnership-with-ibm-watson-reinventing-tax-prep/](http://www.hrblock.com/tax-center/newsroom/around-block/partnership-with-ibm-watson-reinventing-tax-prep/)
5. [www.blackrock.com/aladdin/offerings/aladdin-overview](http://www.blackrock.com/aladdin/offerings/aladdin-overview)
6. <https://ir.blackrock.com/Cache/1001247206.PDF?O=PDF&T=&Y=&D=&FID=1001247206&iid=4048287>
7. [www.institutionalinvestor.com/article/b1dn-7pgfhbxpsg/BlackRock-s-Aladdin-Adds-Alts-Power](http://www.institutionalinvestor.com/article/b1dn-7pgfhbxpsg/BlackRock-s-Aladdin-Adds-Alts-Power)
8. [www.ft.com/content/9ab2d910-1816-11e9-9e64-d150b3105d21](http://www.ft.com/content/9ab2d910-1816-11e9-9e64-d150b3105d21)



have claimed to break the test, it has not been defeated without cheating and using tricks and hacks that do not guarantee a long-term correct result.

Modern AI has come a long way from the humble beginnings of the analytical engine and the one simple question of “Can machines think?” Today, we have AI that can understand the sentiment and tone of a text message, identify objects in image, search through thousands of documents quickly, and almost converse with us flawlessly with natural language. Artificial intelligence has become a magic assistant in our phones that awaits our questions in natural language, interprets them, and then returns an answer in the same language, instead of just showing a web result. In the next section, we will have a brief look at the state of modern AI and its current set of capabilities.

## Natural Language Processing

The ability to converse with humans, as humans do with one another, has been one of the most coveted feats of AI ever since a thinking machine has been thought of. The Turing test measures a computer's ability to “speak” with a human and fool that person into thinking they are speaking to another human. This branch of AI, known as *natural language processing* (NLP), deals with the ability of the computer to understand and express itself in a natural language. This has proven to be especially difficult, since human conversations are loaded with context and deep meanings that are not explicitly communicated and are simply “understood.” Computers are bad at dealing with such loosely defined problems, since they work on well-defined programs that are unambiguous and clear. For example, the phrase “it is raining cats and dogs” is difficult for a computer to understand without the entirety of history and literature accessible inside the computer. To us, such a sentence is obvious even if we're previously unaware of the meaning, because we have the entire context of our lives to judge that raining animals is an impossibility.

## Programmatic NLP

The first chatbots and natural language processing programs used tricks and hacks to translate human speech into computer instructions.

ELIZA was one of the first few programs to make people believe with certain limitations that it was capable of intelligent speech. This was accomplished by Joseph Weizenbaum at the Massachusetts Institute of Technology (MIT) Artificial Intelligence Laboratory in the 1960s. ELIZA was designed to mimic psychologists by echoing the user's answer back to them. In this way, the computer seemed to hold intelligent conversation, but it clearly was not. There are other forms of NLP seen in the 1980s; they were text-based adventure games. The games understood a certain set of verbs—such as *go*, *run*, *fight*, and *eat*—and modified their feedback based purely on language parsing. This was accomplished by having a set of words that the game understood mapped to functions that would execute based on the keyword. The limitation of words that could be stored in memory meant that these early natural language parsers could not understand everything and would return errors and thus ruin the illusion very quickly.

A method of using programming techniques to parse natural language, programmatic NLP uses string parsing with regular expressions (regex) along with a dictionary of words the program can execute on. The regular expressions match the specified patterns, and the program adjusts its control flow based on the information gleaned from sentences, discarding everything except the main word. For example, the following is a simple regular expression that could be used to determine possible illness names:

```
diagnosed with \w+
```

This example looks for the phrase “diagnosed with” followed by a single word, which would assumingly be the name of an illness (such as “diagnosed with cancer”). A more complex regular expression is required to identify illnesses with multiple words (such as “diagnosed with scarlet fever”). A full discussion of regular expressions is outside the scope of this book—*Mastering Regular Expressions*<sup>6</sup> by Jeffrey Friedl is a great resource if you want to learn more.

Although the techniques we've discussed can make wondrous leaps in parsing a language, they fall short very quickly when applied more generally, because the dictionary supplied with the program can be exhausted. This is where AI steps in and outperforms these traditional methods by a huge margin. Natural languages, while they

do follow the rules of grammar, follow rules that are not universal and thus something that holds meaning in one region might not hold true for another region. Languages vary so much because they are a fluid concept; new words are being added constantly to the vernacular, old idioms and words are retired, grammar rules change. This makes a language a perfect candidate for stochastic modeling and other statistical analysis, covered under the umbrella term of *machine learning for natural language processing*.

## Statistical NLP

The techniques we've described are limited in scope to what can be achieved via parsing. The results turn out to be unsatisfactory when used for longer conversations or larger bodies of texts like an encyclopedia or the body of literature on even just one illness (per our earlier example). This necessitates a method that can learn new concepts while trying to understand a text, much like a human does. This method must be able to encounter new words in the same way that a human does and ask questions about what they mean given their context. Although a true AI agent that can perform automatic dictionary and other necessary contextual lookups instantly is years away, we can improve on the programmatic parsing of text by performing a statistical analysis.

For almost all the AI-based natural language parsers, there are some key steps in the algorithm: tokenization and keys, term frequency-inverse document frequency (tf-idf), probability, and ranking. The first step in parsing a sentence is *chunking*. Chunking is the process of breaking down a sentence based on a predetermined criterion; for example, a single word or multiple words in the order subject-adverb-verb-object, and so forth. Each such chunk is known as a *token*. The set of tokens is then analyzed and the duplicates are discarded. These unique chunks are the "keys" to the text. The tokens and their unique keys are used as the building blocks for probability distributions and for understanding the text in more detail. The next step is to identify the frequency distribution of the tokens and keys in the training data. A histogram of the number of occurrences of each key in the text can be used to plot the data to help better visualize the data. Using these frequencies, we can arrive

at probabilities that a word will be followed by another in the text. Some words like “a” and “the” will be used the most, whereas others like names, proper nouns, and jargon will be more sparingly used. The frequency by which a given word appears in a document is called a *term frequency* and the frequency of the same word across various documents is called *inverse document frequency*. Inverse document frequency aids in reducing the impact of commonly used words like “a” and “the.”

## Machine Learning

Machine learning is the broad classification of techniques that involve generating new insights by statistical generalization of the earlier data. Machine learning algorithms typically work by minimizing a penalty or maximizing a reward. A combination of functions, parameters, and weights come together to make a machine learning model. The learning techniques can be grouped into three major categories: supervised learning, unsupervised learning, and reinforcement learning. Supervised learning is the most common type of technique used for constructing AI models, whereas unsupervised learning is more useful for identifying patterns in an input. Supervised learning will calculate a cost for each answer in the training dataset that was incorrectly answered during the training phase. Based on this error estimate, the weights and parameters of the function are adjusted recursively so that we end up with a general function that can match the training questions with answers with a high level of confidence. This process is known as *back propagation*. Unsupervised learning is only given a dataset without any corresponding answers. The objective here is to find a general function to better describe the data. Sometimes, unsupervised learning is also used for simplifying the input stream to be fed into a supervised learning model. This simplification reduces the complexity required for supervised learning. Just as cost and error is defined as two functions in supervised learning, under reinforcement learning an arbitrary cost is assigned based on the action taken. Such a cost will need to minimize or maximize a similar arbitrary reward. In this example:

```
Raw Data: [fruit: apple, animal: tiger, flower: rose]
```

supervised learning would be provided with the entirety of this set of data and would use the answers—say, “apple is a fruit”—to test itself. In unsupervised learning, only the following data would be given to the algorithm:

```
[apple, tiger, rose]
```

and the algorithm would then find a pattern among the given data. Reinforcement learning could have the computer guessing the type of noun, and the user would assign a penalty/reward for each correct/incorrect guess accordingly.

## Markov Chains

In mathematics, a stochastic process is a random variable that changes with time. This random variable can be modeled if it has a *Markov property*. A Markov property indicates that the state of the variable is affected only by the present state; if you know the present state, you can predict the future state with reasonably high accuracy. *Markov chains* work on the principle of the Markov property; by incrementally “guessing” the next word, a sentence can be formed, and sentences together make a paragraph, and so on. The Markov chain uses the tf-idf frequencies to assess the probability of each word that can appear next and then chooses the word with the highest probability. Because the reliance on tf-idf is quite high, Markov chains require large amounts of training data to be accurate. It is of vital importance that Markov models are parametrized and developed properly. Such a statistical approach has been proven better than modeling grammar rules. Variations on this model include smoothing the frequency and giving access to more data to the model. Although Markov chains are very fast and can give a better result than traditional approaches, they have their limitations and are unable to produce long forms of coherent text on a particular topic.

## Hidden Markov Models

A *hidden Markov model* is one where the state of the program is “hidden.” Unlike the regular models, which act in a very deterministic

One use case for deep learning is creating neural networks that are capable of creating new sentences from scratch based on a prompt. In this case, neural networks serve as a probability calculator that ranks the probability of a word “making sense” as part of a sentence. Using the simplest form of a neural “network,” a single neuron could be asked what the next word is in the new sentence it is generating, and based on the training data, it would reply with a best guess for the next word.

Words are not stored directly in the neuron’s memory but instead are encoded and decoded as numbers using *word embedding*. The encoding process converts the word to a number, which is more easily manipulated by the computer, whereas the decoding process converts the number back to a word. Such a simplistic sentence-generating model with a single neuron can hardly be used for any serious applications. In practice, the number of neurons would be directly proportional to the complexity of text being analyzed and the quality of expected output. Adding more neurons (either in the same layer or by adding additional layers) does not automatically make a neural network better. The techniques to improve a neural network will vary based on the problem at hand, and the model will need to be adjusted for unexpected outcomes. In a practical implementation, there would be multiple neurons, each having a specific weight; the weight of the neuron would determine the final output of the neural network. This method of adjusting weights by looking at the output during the training stage is known as a *back propagation neural network*.

If we pass the output received from a neuron back through it again, it would give us a better opportunity of analyzing and generating the data. A recurrent neural network (RNN) does exactly that, and multiple passes are made recursively, feeding the output back into the neuron as input. RNNs have proved to be much better at understanding and generating large amounts of text than traditional neural networks. A further improvement over RNNs are long short-term memory (LSTM) neural networks. LSTM networks are able to remember previous states as well, and then output an answer based on these states. LSTM networks can also be finely tuned with gate-like structures that can restrict what information is input and output by a neuron. The gates use pointwise multiplication to determine if all or no information will go in through the gate. The gates are



operated by a sigmoid layer that judges each point of data to determine whether the gate should be opened or closed. Further variations include allowing the gate to view the output of a neuron and then pass a judgment, thus modifying the output on the fly.

Chatbots and text generators are some of the biggest use cases for NLP-based neural networks. Speech recognition is another area where such neural networks are used. Amazon's Alexa uses an LSTM.<sup>7</sup>

### Image Recognition/Classification

Images contain a lot of data, and the permutation and combination of each pixel can change the output drastically. We can preprocess these images to reduce the size of the input by adding a convolutional layer that reduces the size of the input data, thus reducing the computing power required to scan and understand larger images. For image processing, it is critical to not lose sight of the bigger picture—a single pixel cannot tell us if we are looking at an image of a plane or a train. The process of convolution is defined mathematically to mean how the shape of one function affects another. In a convolutional neural network (CNN), maximum pooling and average pooling are also applied after the convolutional process to further reduce the parameters and generalize the data. This processed data from the image is then fed into a fully connected neural network for classification. CNNs have proven their effectiveness in image recognition; they are very efficient at recognizing images, and the reduced parameterization aids in making simpler models.

### Becoming an Innovation-Focused Organization

The world of technology moves at a high speed. Innovation will be the key business differentiator going forward. The advantage of being the company to present a key piece of innovation in the marketplace is game changing. This first-mover advantage can be achieved only via the relentless pursuit of innovation and rigorous experimentation with new ideas. Innovations using AI can also lead to cost-saving practices, offering the competitive edge needed to fiscally outmaneuver the competition. Organizations should develop a culture of innovation by establishing suitable processes and incentives to encourage their workforce.

A culture of innovation is hard to implement correctly, but it yields splendid results when done well. Your organization should have policies that encourage innovation and not limit it. A motivated workforce is the cornerstone of innovative thinking. If the employees are unmotivated, they will think only about getting through the day and not about how to make their systems more efficient.

IBM pioneered the concept of Think Fridays, where employees are encouraged to spend Friday afternoons working on self-development and personal research projects.<sup>8</sup> This in part has led IBM to being called one of the most innovative companies given that 2018 marks their 26th consecutive year of being the entity with the highest number of U.S. patents issued for the year.<sup>9</sup>

Google similarly has a 20 percent rule that allows employees to work on personal projects that are aligned with the company's strategy.<sup>10</sup> This means that Google employees, for the equivalent of one day a week, are allowed to work on a project of their own choosing. This morale-boosting perk enables employees to work on what inspires while Google maintains the intellectual property their employees generate. Famously, Gmail and Google Maps came out of 20 percent projects and are now industry-leading Google products in their respective fields.

Not every organization needs to be as open as Google, but even a 5 percent corporate-supported allowance can have a big impact, because employees will use this time to focus on revolutionizing the business and will feel an increased level of creative validation with their work. Ensuring that employees are adequately enabled and motivated is a task of paramount importance and gives them the tools to modernize and adapt their workflows.

An organization that fosters creativity among its employees is bolstered to succeed. It will be the first among its peers to develop and implement solutions that will directly impact key metrics and result in savings of time and money. When existing business processes can be revamped, major transformations of 80 to 90 percent increased efficiency and productivity are possible. The main idea is to allow employees time to think freely. Every employee is a subject matter expert about their own job. Transferring this personal knowledge into shared knowledge for the organization can lead to a valuable well of ideas.

Innovation should be a top priority for the modern organization. At its core, prioritizing innovation means adapting the business to the constantly changing and evolving technological landscape. A business that refuses to innovate will slowly but surely wither away.

The organization that chooses to follow the mantra “Necessity is the mother of all invention” will continuously lag behind its peers. It will be following a reactive approach instead of a proactive approach. Such an organization will, by definition, be perpetually behind the technological curve. With that said, there are certain upsides to this strategy. If you are merely following the curve, you avoid the mistakes made by early adopters and can learn from those mistakes. Such an organization would save on short-lived, interim technologies that are discarded and projects that are shelved after partial implementation.

At the other end of the spectrum is an organization that yearns for change. This organization will try to implement new technologies like AI and robotics as much as it can. This proactive policy has the potential to yield far greater returns than the organization that follows the curve. This modern business, however, will need to control and manage its change costs carefully. It runs the risk of budget overruns and faces the disastrous possibility of being thrown completely out of the race. Care should be taken to minimize the cost of mistakes by employing strategies like feedback loops, diversification, and change management, all of which we will discuss in more detail in the coming sections. The rewards in this scenario are great and will lead an organization to new heights.

This path of constant innovation and learning is the smartest path to follow in the modern world. Passively reacting to technological changes only once they have become an industry standard will greatly hamper any organization.

## Idea Bank

---

Organizational memory can be fickle. To aid the growth of an innovative organization, an “idea bank” should be maintained. The idea bank stores all ideas that have been received but not yet implemented. An innovation focus group should be given the authority to add and delete ideas from the bank, though the idea bank should be