

Artificial Intelligence Basics

A Non-Technical Introduction

—

Tom Taulli

Apress®

ARTIFICIAL INTELLIGENCE BASICS

A NON-TECHNICAL INTRODUCTION

Tom Taulli

Apress®

Artificial Intelligence Basics: A Non-Technical Introduction

Tom Taulli
Monrovia, CA, USA

ISBN-13 (pbk): 978-1-4842-5027-3
<https://doi.org/10.1007/978-1-4842-5028-0>

ISBN-13 (electronic): 978-1-4842-5028-0

Copyright © 2019 by Tom Taulli

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

Trademarked names, logos, and images may appear in this book. Rather than use a trademark symbol with every occurrence of a trademarked name, logo, or image we use the names, logos, and images only in an editorial fashion and to the benefit of the trademark owner, with no intention of infringement of the trademark.

The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

While the advice and information in this book are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

Managing Director, Apress Media LLC: Welmoed Spahr
Acquisitions Editor: Shiva Ramachandran
Development Editor: Rita Fernando
Coordinating Editor: Rita Fernando

Cover designed by eStudioCalamar

Distributed to the book trade worldwide by Springer Science+Business Media New York, 233 Spring Street, 6th Floor, New York, NY 10013. Phone 1-800-SPRINGER, fax (201) 348-4505, e-mail orders-ny@springer-sbm.com, or visit www.springeronline.com. Apress Media, LLC is a California LLC and the sole member (owner) is Springer Science + Business Media Finance Inc (SSBM Finance Inc). SSBM Finance Inc is a **Delaware** corporation.

For information on translations, please e-mail rights@apress.com, or visit <http://www.apress.com/rights-permissions>.

Apress titles may be purchased in bulk for academic, corporate, or promotional use. eBook versions and licenses are also available for most titles. For more information, reference our Print and eBook Bulk Sales web page at <http://www.apress.com/bulk-sales>.

Any source code or other supplementary material referenced by the author in this book is available to readers on GitHub via the book's product page, located at www.apress.com/978-1-4842-5027-3. For more detailed information, please visit <http://www.apress.com/source-code>.

Printed on acid-free paper

Contents

About the Author	v
Foreword	vii
Introduction	ix
Chapter 1: AI Foundations	1
Chapter 2: Data	19
Chapter 3: Machine Learning	39
Chapter 4: Deep Learning	69
Chapter 5: Robotic Process Automation (RPA)	91
Chapter 6: Natural Language Processing (NLP)	103
Chapter 7: Physical Robots	125
Chapter 8: Implementation of AI	143
Chapter 9: The Future of AI	161
Appendix: AI Resources	177
Glossary	179
Index	185

About the Author



Tom Taulli has been developing software since the 1980s. In college, he started his first company, which focused on the development of e-learning systems. He created other companies as well, including Hypermart.net that was sold to InfoSpace in 1996. Along the way, Tom has written columns for online publications such as businessweek.com, techweb.com, and Bloomberg.com. He also writes posts on artificial intelligence for Forbes.com and is the advisor to various companies in the AI space. You can reach Tom on Twitter (@ttaulli) or through his web site (www.taulli.com).

Foreword

As this book demonstrates, the adoption of artificial intelligence (AI) will be a major inflection point in human history. Like other similarly groundbreaking technologies, how it's administered and who has access to it will shape society for generations to come. However, AI stands out from the other transformative technologies of the nineteenth and twentieth centuries—think the steam engine, the electrical grid, genomics, computers, and the Internet—because it doesn't depend exclusively on critically expensive physical infrastructure to enable adoption; after all, many of its benefits can be delivered through existing hardware we all carry around in our pockets. Instead, the fundamental limiting factor when it comes to the mass adoption of AI technology is our shared intellectual infrastructure: education, understanding, and vision.

This is a crucial difference because, if handled correctly, AI can act as a sweeping democratizing force. It has and will eliminate from our lives the drudgery of the past and free up a tremendous amount of human energy and capital. But that "if" is far from certain. AI executed irresponsibly has the power to destabilize large parts of the world economy by causing, as many people fear, a shrinking workforce, reduced purchasing power for the middle class, and an economy without a wide and stable base fueled by an endless debt spiral.

However, before we succumb to pessimism on AI, we should take a look back. Historic though AI's transformative capacity may be—and it is historic—these same issues are and have been at play in the economic landscape for decades, even centuries. AI is, after all, an extension of a trend toward automation that has been at play since Henry Ford. In fact, Zoho itself was born from the tension between automation and egalitarian economic principles. Back in the early 2000s, we came to a realization that has shaped our approach to technology: regular people—small business owners, here and abroad—should have access to the same advanced business automations that the Fortune 500 companies have; otherwise, a huge swath of the population will be locked out of the economy.

At the time, powerful digital software was almost unanimously gated behind rigid contracts, exorbitant fee structures, and complicated on-premise implementations. Big companies could shoulder the burden of such systems, while smaller operators were locked out, putting them at a tremendous disadvantage. We sought to disrupt that by opening up the promise of technology to wider and wider audiences. Over the last two decades, we've endeavored to

increase the value of our products without increasing the price by tapping into the scalability of cloud technology. Our goal is to empower people at all levels of society by pushing down the price of business software while expanding the power of the tools. Access to capital shouldn't limit success; businesses should rise or fall based on the strength of their vision for the future.

Viewed this way, AI is the fulfillment of the promise of technology. It frees people from the constraints of time by enabling them to offload tedious or unpleasant rote labor. It helps them identify patterns at microscopic and macroscopic scales, which humans are not naturally well suited to perceive. It can forecast problems, and it can correct errors. It can save money, time, and even lives.

Seeking to democratize these benefits just as we did for general business software, Zoho has threaded AI throughout our suite of apps. We spent the last six years quietly developing our own internal AI technology, built on the bedrock of our own principles. The result is Zia, an AI assistant who is smart, but not clever. This is a crucial distinction. A smart system has the information and functionality to empower the unique vision and intuition of an active operator. A clever system obfuscates the internal workings of the process, reducing the human to a passive user who simply consumes the insights provided by the machine. AI should be a tool to be wielded, not a lens through which we view the world. To steer such a powerful tool, we must be equipped with the knowledge to understand and operate it without eroding the human quality of our human systems.

The need to stay current on this technology is exactly why a book like *Artificial Intelligence Basics* is so important in today's world. It is the intellectual infrastructure that will enable people—regular people—to tap into the power of AI. Without these kinds of initiatives, AI will tip the balance of power in favor of big companies with big budgets. It's crucial that the general population equip themselves with the skills to understand AI systems, because these systems will increasingly define how we interact with and navigate through the world. Soon, the information contained in this book won't be merely a topic of interest; it will be a prerequisite for participation in the modern economy.

This is how the average person can enjoy the fruits of the AI revolution. In the years to come, how we define work and which activities carry economic value will change. We have to embrace the fact that the future of work may be as foreign to us as a desk job would be to our distant ancestors. But we have to—and should—have faith in the human capacity to innovate new forms of work, even if that work doesn't look like the work we're familiar with. But the first step, before everything else, is to learn more about this new, exciting, and fundamentally democratizing technology.

—Sridhar Vembu, co-founder and CEO of Zoho

Introduction

On the face of it, the Uber app is simple. With just a couple clicks, you can hail a driver within a few minutes.

But behind the scenes, there is an advanced technology platform, which relies heavily on artificial intelligence (AI). Here are just some of the capabilities:

- A Natural Language Processing (NLP) system that can understand conversations, allowing for a streamlined experience
- Computer vision software that verifies millions of images and documents like drivers' licenses and restaurant menus
- Sensor processing algorithms that help improve the accuracy in dense urban areas, including automatic crash detection by sensing unexpected movement from the phone of a driver or passenger
- Sophisticated machine learning algorithms that predict driver supply, rider demand, and ETAs

Such technologies are definitely amazing, but they are also required. There is no way that Uber could have scaled its growth—which has involved handling over 10 billion trips—without AI. In light of this, it should be no surprise that the company spends hundreds of millions on the technology and has a large group of AI experts on staff.¹

But AI is not just for fast-charging startups. The technology is also proving a critical priority for traditional companies. Just look at McDonald's. In 2019, the company shelled out \$300 million to acquire a tech startup, Dynamic Yield. It was the company's largest deal since it purchased Boston Market in 1999.²

¹www.sec.gov/Archives/edgar/data/1543151/000119312519120759/d647752ds1a.htm#toc647752_11

²<https://news.mcdonalds.com/news-releases/news-release-details/dynamic-yield-acquisition-release>

Dynamic Yield, which was founded in 2011, is a pioneer in leveraging AI for creating personalized customer interactions across the Web, apps, and email. Some of its customers include the Hallmark Channel, IKEA, and Sephora.

As for McDonald's, it has been undergoing a digital transformation—and AI is a key part of the strategy. With Dynamic Yield, the company plans to use the technology to reimagine its Drive Thru, which accounts for a majority of its revenues. By analyzing data, such as the weather, traffic, and time of day, the digital menus will be dynamically changed to enhance the revenue opportunities. It also looks like McDonald's will use geofencing and even image recognition of license plates to enhance the targeting.

But this will just be the start. McDonald's expects to use AI for in-store kiosks and signage as well as the supply chain.

The company realizes that the future is both promising and dangerous. If companies are not proactive with new technologies, they may ultimately fail. Just look at how Kodak was slow to adapt to digital cameras. Or consider how the taxi industry did not change when faced with the onslaught of Uber and Lyft.

On the other hand, new technologies can be almost an elixir for a company. But there needs to be a solid strategy, a good understanding of what's possible, and a willingness to take risks. So in this book, I'll provide tools to help with all this.

OK then, how big will AI get? According to a study from PWC, it will add a staggering \$15.7 trillion to the global GDP by 2030, which is more than the combined output of China and India. The authors of the report note: "AI touches almost every aspect of our lives. And it's only just getting started."³

True, when it comes to predicting trends, there can be a good deal of hype. However, AI may be different because it has the potential for turning into a general-purpose technology. A parallel to this is what happened in the nineteenth century with the emergence of electricity, which had a transformative impact across the world.

As a sign of the strategic importance of AI, tech companies like Google, Microsoft, Amazon.com, Apple, and Facebook have made substantial investments in this industry. For example, Google calls itself an "AI-first" company and has spent billions buying companies in the space as well as hiring thousands of data scientists.

In other words, more and more jobs will require knowledge of AI. Granted, this does not mean you'll need to learn programming languages or understand advanced statistics. But it will be critical to have a solid foundation of the fundamentals.

³www.pwc.com/gx/en/issues/data-and-analytics/publications/artificial-intelligence-study.html

As for this book, the goal is to provide actionable advice that can make a big difference in your organization and career. Now you will not find deeply technical explanations, code snippets, or equations. Instead, *Artificial Intelligence Basics* is about answering the top-of-mind questions that managers have: Where does AI make sense? What are the gotchas? How do you evaluate the technology? What about starting an AI pilot?

This book also takes a real-world view of the technology. A big advantage I have as a writer for Forbes.com and an advisor in the tech world is that I get to talk to many talented people in the AI field—and this helps me to identify what is really important in the industry. I also get to learn about case studies and examples of what works.

This book is organized in a way to cover the main topics in AI—and you do not have to read each chapter in order. *Artificial Intelligence Basics* is meant to be a handbook.

Here are brief descriptions of the chapters:

- *Chapter 1—AI Foundations*: This is an overview of the rich history of AI, which goes back to the 1950s. You will learn about brilliant researchers and computer scientists like Alan Turing, John McCarthy, Marvin Minsky, and Geoffrey Hinton. There will also be coverage of key concepts like the Turing Test, which gauges if a machine has achieved true AI.
- *Chapter 2—Data*: Data is the lifeblood of AI. It's how algorithms can find patterns and correlations to provide insights. But there are landmines with data, such as quality and bias. This chapter provides a framework to work with data in an AI project.
- *Chapter 3—Machine Learning*: This is a subset of AI and involves traditional statistical techniques like regressions. But in this chapter, we'll also cover the advanced algorithms, such as k-Nearest Neighbor (k-NN) and the Naive Bayes Classifier. Besides this, there will be a look at how to put together a machine learning model.
- *Chapter 4—Deep Learning*: This is another subset of AI and is clearly the one that has seen much of the innovation during the past decade. Deep learning is about using neural networks to find patterns that mimic the brain. In the chapter, we'll take a look at the main algorithms like recurrent neural networks (RNNs), convolutional neural networks (CNNs), and generative adversarial networks (GANs). There will also be explanations of key concepts like backpropagation.

- *Chapter 5—Robotic Process Automation:* This uses systems to automate repetitive processes, such as inputting data in a Customer Relationship Management (CRM) system. Robotic Process Automation (RPA) has seen tremendous growth during the past few years because of the high ROI (Return on Investment). The technology has also been an introductory way for companies to implement AI.
- *Chapter 6—Natural Language Processing (NLP):* This form of AI, which involves understanding conversations, is the most ubiquitous as seen with Siri, Cortana, and Alexa. But NLP systems, such as chatbots, have also become critical in the corporate world. This chapter will show ways to use this technology effectively and how to avoid the tricky issues.
- *Chapter 7—Physical Robots:* AI is starting to have a major impact on this industry. With deep learning, it is getting easier for robots to understand their environments. In this chapter, we'll take a look at both consumer and industrial robots, such as with a myriad of use cases.
- *Chapter 8—Implementation of AI:* We'll take a step-by-step approach to putting together an AI project, from the initial concept to the deployment. This chapter will also cover the various tools like Python, TensorFlow, and PyTorch.
- *Chapter 9—The Future of AI:* This chapter will cover some of the biggest trends in AI like autonomous driving, weaponization of AI, technological unemployment, drug discovery, and regulation.

At the back of the book, you'll also find an appendix of resources for further study and a glossary of common terms related to AI.

Accompanying Material

Any updates will be provided on my site at www.Taulli.com.

AI Foundations

History Lessons

Artificial intelligence would be the ultimate version of Google. The ultimate search engine that would understand everything on the web. It would understand exactly what you wanted, and it would give you the right thing. We're nowhere near doing that now. However, we can get incrementally closer to that, and that is basically what we work on.

—Larry Page, the co-founder of Google Inc. and CEO of Alphabet¹

In Fredric Brown's 1954 short story, "Answer," all of the computers across the 96 billion planets in the universe were connected into one super machine. It was then asked, "Is there a God?" to which it answered, "Yes, now there is a God."

No doubt, Brown's story was certainly clever—as well as a bit comical and chilling! Science fiction has been a way for us to understand the implications of new technologies, and artificial intelligence (AI) has been a major theme. Some of the most memorable characters in science fiction involve androids or computers that become self-aware, such as in *Terminator*, *Blade Runner*, *2001: A Space Odyssey*, and even *Frankenstein*.

But with the relentless pace of new technologies and innovation nowadays, science fiction is starting to become real. We can now talk to our smartphones and get answers; our social media accounts provide us with the content we're

¹Founding CEO of Google Inc. The Academy of Achievement interview, www.achievement.org, October 28, 2000.

interested in; our banking apps provide us with reminders; and on and on. This personalized content creation almost seems magical but is quickly becoming normal in our everyday lives.

To understand AI, it's important to have a grounding in its rich history. You'll see how the development of this industry has been full of breakthroughs and setbacks. There is also a cast of brilliant researchers and academics, like Alan Turing, John McCarthy, Marvin Minsky, and Geoffrey Hinton, who pushed the boundaries of the technology. But through it all, there was constant progress.

Let's get started.

Alan Turing and the Turing Test

Alan Turing is a towering figure in computer science and AI. He is often called the “father of AI.”

In 1936, he wrote a paper called “On Computable Numbers.” In it, he set forth the core concepts of a computer, which became known as the Turing machine. Keep in mind that real computers would not be developed until more than a decade later.

Yet it was his paper, called “Computing Machinery and Intelligence,” that would become historic for AI. He focused on the concept of a machine that was intelligent. But in order to do this, there had to be a way to measure it. What is intelligence—at least for a machine?

This is where he came up with the famous “Turing Test.” It is essentially a game with three players: two that are human and one that is a computer. The evaluator, a human, asks open-ended questions of the other two (one human, one computer) with the goal of determining which one is the human. If the evaluator cannot make a determination, then it is presumed that the computer is intelligent. Figure I-1 shows the basic workflow of the Turing Test.




Figure I-1. The basic workflow of the Turing Test

The genius of this concept is that there is no need to see if the machine actually knows something, is self-aware, or even if it is correct. Rather, the Turing Test indicates that a machine can process large amounts of information, interpret speech, and communicate with humans.

Turing believed that it would actually not be until about the turn of the century that a machine would pass his test. Yes, this was one of many predictions of AI that would come up short.

So how has the Turing Test held up over the years? Well, it has proven to be difficult to crack. Keep in mind that there are contests, such as the Loebner Prize and the Turing Test Competition, to encourage people to create intelligent software systems.

In 2014, there was a case where it did look like the Turing Test was passed. It involved a computer that said it was 13 years old.² Interestingly enough, the human judges likely were fooled because some of the answers had errors.

Then in May 2018 at Google's I/O conference, CEO Sundar Pichai gave a standout demo of Google Assistant.³ Before a live audience, he used the device to call a local hairdresser to make an appointment. The person on the other end of the line acted as if she was talking to a person!

Amazing, right? Definitely. Yet it still probably did not pass the Turing Test. The reason is that the conversation was focused on one topic—not open ended.

As should be no surprise, there has been ongoing controversy with the Turing Test, as some people think it can be manipulated. In 1980, philosopher John Searle wrote a famous paper, entitled “Minds, Brains, and Programs,” where he set up his own thought experiment, called the “Chinese room argument” to highlight the flaws.

Here's how it worked: Let's say John is in a room and does not understand the Chinese language. However, he does have manuals that provide easy-to-use rules to translate it. Outside the room is Jan, who does understand the language and submits characters to John. After some time, she will then get an accurate translation from John. As such, it's reasonable to assume that Jan believes that John can speak Chinese.

² www.theguardian.com/technology/2014/jun/08/super-computer-simulates-13-year-old-boy-passes-turing-test

³ www.theverge.com/2018/5/8/17332070/google-assistant-makes-phone-call-demo-duplex-io-2018

Searle's conclusion:

The point of the argument is this: if the man in the room does not understand Chinese on the basis of implementing the appropriate program for understanding Chinese then neither does any other digital computer solely on that basis because no computer, qua computer, has anything the man does not have.⁴

It was a pretty good argument—and has been a hot topic of debate in AI circles since.

Searle also believed there were two forms of AI:

- *Strong AI*: This is when a machine truly understands what is happening. There may even be emotions and creativity. For the most part, it is what we see in science fiction movies. This type of AI is also known as Artificial General Intelligence (AGI). Note that there are only a handful of companies that focus on this category, such as Google's DeepMind.
- *Weak AI*: With this, a machine is pattern matching and usually focused on narrow tasks. Examples of this include Apple's Siri and Amazon's Alexa.

The reality is that AI is in the early phases of weak AI. Reaching the point of strong AI could easily take decades. Some researchers think it may never happen.

Given the limitations to the Turing Test, there have emerged alternatives, such as the following:

- *Kurzweil-Kapor Test*: This is from futurologist Ray Kurzweil and tech entrepreneur Mitch Kapor. Their test requires that a computer carry on a conversation for two hours and that two of three judges believe it is a human talking. As for Kapor, he does not believe this will be achieved until 2029.
- *Coffee Test*: This is from Apple co-founder Steve Wozniak. According to the coffee test, a robot must be able to go into a stranger's home, locate the kitchen, and brew a cup of coffee.

⁴<https://plato.stanford.edu/entries/chinese-room/>

The Brain Is a...Machine?

In 1943, Warren McCulloch and Walter Pitts met at the University of Chicago, and they became fast friends even though their backgrounds were starkly different as were their ages (McCulloch was 42 and Pitts was 18). McCulloch grew up in a wealthy Eastern Establishment family, having gone to prestigious schools. Pitts, on the other hand, grew up in a low-income neighborhood and was even homeless as a teenager.

Despite all this, the partnership would turn into one of the most consequential in the development of AI. McCulloch and Pitts developed new theories to explain the brain, which often went against the conventional wisdom of Freudian psychology. But both of them thought that logic could explain the power of the brain and also looked at the insights from Alan Turing. From this, they co-wrote a paper in 1943 called “A Logical Calculus of the Ideas Immanent in Nervous Activity,” and it appeared in the *Bulletin of Mathematical Biophysics*. The thesis was that the brain’s core functions like neurons and synapses could be explained by logic and mathematics, say with logical operators like And, Or, and Not. With these, you could construct a complex network that could process information, learn, and think.

Ironically, the paper did not get much traction with neurologists. But it did get the attention with those working on computers and AI.

Cybernetics

While Norbert Wiener created various theories, his most famous one was about cybernetics. It was focused on understanding control and communications with animals, people, and machines—showing the importance of feedback loops.

In 1948, Wiener published *Cybernetics: Or Control and Communication in the Animal and the Machine*. Even though it was a scholarly work—filled with complex equations—the book still became a bestseller, hitting the *New York Times* list.

It was definitely wide ranging. Some of the topics included Newtonian mechanics, meteorology, statistics, astronomy, and thermodynamics. This book would anticipate the development of chaos theory, digital communications, and even computer memory.

But the book would also be influential for AI. Like McCulloch and Pitts, Wiener compared the human brain to the computer. Furthermore, he speculated that a computer would be able to play chess and eventually beat grand masters. The main reason is that he believed that a machine could learn as it played games. He even thought that computers would be able to replicate themselves.

But *Cybernetics* was not utopian either. Wiener was also prescient in understanding the downsides of computers, such as the potential for dehumanization. He even thought that machines would make people unnecessary.

It was definitely a mixed message. But Wiener's ideas were powerful and spurred the development of AI.

The Origin Story

John McCarthy's interest in computers was spurred in 1948, when he attended a seminar, called "Cerebral Mechanisms in Behavior," which covered the topic of how machines would eventually be able to think. Some of the participants included the leading pioneers in the field such as John von Neumann, Alan Turing, and Claude Shannon.

McCarthy continued to immerse himself in the emerging computer industry—including a stint at Bell Labs—and in 1956, he organized a ten-week research project at Dartmouth University. He called it a "study of artificial intelligence." It was the first time the term had been used.

The attendees included academics like Marvin Minsky, Nathaniel Rochester, Allen Newell, O. G. Selfridge, Raymond Solomonoff, and Claude Shannon. All of them would go on to become major players in AI.

The goals for the study were definitely ambitious:

The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.⁵

At the conference, Allen Newell, Cliff Shaw, and Herbert Simon demoed a computer program called the Logic Theorist, which they developed at the Research and Development (RAND) Corporation. The main inspiration came from Simon (who would win the Nobel Prize in Economics in 1978). When he saw how computers printed out words on a map for air defense systems, he realized that these machines could be more than just about processing numbers. It could also help with images, characters, and symbols—all of which could lead to a thinking machine.

Regarding Logic Theorist, the focus was on solving various math theorems from *Principia Mathematica*. One of the solutions from the software turned out to be more elegant—and the co-author of the book, Bertrand Russell, was delighted.

⁵www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html

Creating the Logic Theorist was no easy feat. Newell, Shaw, and Simon used an IBM 701, which used machine language. So they created a high-level language, called IPL (Information Processing Language), that sped up the programming. For several years, it was the language of choice for AI.

The IBM 701 also did not have enough memory for the Logic Theorist. This led to another innovation: list processing. It allowed for dynamically allocating and deallocating memory as the program ran.

Bottom line: The Logic Theorist is considered the first AI program ever developed.

Despite this, it did not garner much interest! The Dartmouth conference was mostly a disappointment. Even the phrase “artificial intelligence” was criticized.

Researchers tried to come up with alternatives, such as “complex information processing.” But they were not catchy like AI was—and the term stuck.

As for McCarthy, he continued on his mission to push innovation in AI. Consider the following:

- During the late 1950s, he developed the Lisp programming language, which was often used for AI projects because of the ease of using nonnumerical data. He also created programming concepts like recursion, dynamic typing, and garbage collection. Lisp continues to be used today, such as with robotics and business applications. While McCarthy was developing the language, he also co-founded the MIT Artificial Intelligence Laboratory.
- In 1961, he formulated the concept of time-sharing of computers, which had a transformative impact on the industry. This also led to the development of the Internet and cloud computing.
- A few years later, he founded Stanford’s Artificial Intelligence Laboratory.
- In 1969, he wrote a paper called “Computer-Controlled Cars,” in which he described how a person could enter directions with a keyboard and a television camera would navigate the vehicle.
- He won the Turing Award in 1971. This prize is considered the Nobel Prize for Computer Science.

In a speech in 2006, McCarthy noted that he was too optimistic about the progress of strong AI. According to him, “we humans are not very good at identifying the heuristics we ourselves use.”⁶

Golden Age of AI

From 1956 to 1974, the AI field was one of the hottest spots in the tech world. A major catalyst was the rapid development in computer technologies. They went from being massive systems—based on vacuum tubes—to smaller systems run on integrated circuits that were much quicker and had more storage capacity.

The federal government was also investing heavily in new technologies. Part of this was due to the ambitious goals of the Apollo space program and the heavy demands of the Cold War.

As for AI, the main funding source was the Advanced Research Projects Agency (ARPA), which was launched in the late 1950s after the shock of Russia’s Sputnik. The spending on projects usually came with few requirements. The goal was to inspire breakthrough innovation. One of the leaders of ARPA, J. C. R. Licklider, had a motto of “fund people, not projects.” For the most part, the majority of the funding was from Stanford, MIT, Lincoln Laboratories, and Carnegie Mellon University.

Other than IBM, the private sector had little involvement in AI development. Keep in mind that—by the mid-1950s—IBM would pull back and focus on the commercialization of its computers. There was actually fear from customers that this technology would lead to significant job losses. So IBM did not want to be blamed.

In other words, much of the innovation in AI spun out from academia. For example, in 1959, Newell, Shaw, and Simon continued to push the boundaries in the AI field with the development of a program called “General Problem Solver.” As the name implied, it was about solving math problems, such as the Tower of Hanoi.

But there were many other programs that attempted to achieve some level of strong AI. Examples included the following:

- *SAINT or Symbolic Automatic INTEgrator (1961)*: This program, created by MIT researcher James Slagle, helped to solve freshman calculus problems. It would be updated into other programs, called SIN and MACSYMA, that did much more advanced math. SAINT was actually the first example of an expert system, a category of AI we’ll cover later in this chapter.

⁶www.technologyreview.com/s/425913/computing-pioneer-dies/

- *ANALOGY (1963)*: This program was the creation of MIT professor Thomas Evans. The application demonstrated that a computer could solve analogy problems of an IQ test.
- *STUDENT (1964)*: Under the supervision of Minsky at MIT, Daniel Bobrow created this AI application for his PhD thesis. The system used Natural Language Processing (NLP) to solve algebra problems for high school students.
- *ELIZA (1965)*: MIT professor Joseph Weizenbaum designed this program, which instantly became a big hit. It even got buzz in the mainstream press. It was named after Eliza (based on George Bernard Shaw's play *Pygmalion*) and served as a psychoanalyst. A user could type in questions, and ELIZA would provide counsel (this was the first example of a chatbot). Some people who used it thought the program was a real person, which deeply concerned Weizenbaum since the underlying technology was fairly basic. You can find examples of ELIZA on the web, such as at <http://psych.fullerton.edu/mbirnbaum/psych101/Eliza.htm>.
- *Computer Vision (1966)*: In a legendary story, MIT's Marvin Minsky said to a student, Gerald Jay Sussman, to spend the summer linking a camera to a computer and getting the computer to describe what it saw. He did just that and built a system that detected basic patterns. It was the first use of computer vision.
- *Mac Hack (1968)*: MIT professor Richard D. Greenblatt created this program that played chess. It was the first to play in real tournaments and got a C-rating.
- *Hearsay I (Late 1960s)*: Professor Raj Reddy developed a continuous speech recognition system. Some of his students would then go on to create Dragon Systems, which became a major tech company.

During this period, there was a proliferation of AI academic papers and books. Some of the topics included Bayesian methods, machine learning, and vision.

But there were generally two major theories about AI. One was led by Minsky, who said that there needed to be symbolic systems. This meant that AI should be based on traditional computer logic or preprogramming—that is, the use of approaches like If-Then-Else statements.