Transactions on Computer Systems and Networks

Govindakumar Madhavan

Mastering Machine Learning: From Basics to Advanced

COURSE INSIDE

MOREMEDIA

Springer

Transactions on Computer Systems and Networks

Transactions on Computer Systems and Networks is a unique series that aims to capture advances in evolution of computer hardware and software systems and progress in computer networks. Computing Systems in present world span from miniature IoT nodes and embedded computing systems to large-scale cloud infrastructures, which necessitates developing systems architecture, storage infrastructure and process management to work at various scales. Present day networking technologies provide pervasive global coverage on a scale and enable multitude of transformative technologies. The new landscape of computing comprises of self-aware autonomous systems, which are built upon a software-hardware collaborative framework. These systems are designed to execute critical and non-critical tasks involving a variety of processing resources like multi-core CPUs, reconfigurable hardware, GPUs and TPUs which are managed through virtualisation, real-time process management and fault-tolerance. While AI, Machine Learning and Deep Learning tasks are predominantly increasing in the application space the computing system research aim towards efficient means of data processing, memory management, real-time task scheduling, scalable, secured and energy aware computing. The paradigm of computer networks also extends it support to this evolving application scenario through various advanced protocols, architectures and services. This series aims to present leading works on advances in theory, design, behaviour and applications in computing systems and networks. The Series accepts research monographs, introductory and advanced textbooks, professional books, reference works, and select conference proceedings.

Govindakumar Madhavan

Mastering Machine Learning: From Basics to Advanced

Springer

Govindakumar Madhavan
SeaportAi
Chennai, India

ISSN 2730-7484 ISSN 2730-7492 (electronic)
Transactions on Computer Systems and Networks
ISBN 978-981-97-9913-8 ISBN 978-981-97-9914-5 (eBook)
https://doi.org/10.1007/978-981-97-9914-5

This Springer imprint is published by the registered company Springer Nature Singapore Pte Ltd.
The registered company address is: 152 Beach Road, #21-01/04 Gateway East, Singapore 189721, Singapore

Contents

About the Author

Govindakumar Madhavan is the founder and CEO of SeaportAi. He completed his B.Tech. and MBA from top institutes in India. He also holds certifications in six sigma and project management. He has over two decades of experience managing technology, operations, and quality in large MNCs and start-ups. He has led high performance multi-national teams and managed businesses across the Asia Pacific region. He has successfully incubated Centers of Excellence for fraud prevention and service analytics. He has significant experience in design thinking-based product development and management. He has also played a critical role in developing products for emerging markets and won global awards in the areas of Customer Experience, Leadership Excellence, Quality, and Technology. He is a prolific author of about 40 video courses and 10 books.

Chapter 1
About the Book

Abstract This chapter introduces the book, what to expect besides links to code and datasets.

Keyword Machine learning

The interest in machine learning (ML) as well as the application of machine learning is growing rapidly. Since machine learning is at the intersection of mathematics, programming, and domain knowledge, machine learning is best learnt from an application perspective rather than from a theoretical perspective. This is especially true for beginners as well as young practitioners.

This book attempts to bring application and industry perspectives to both the mathematics aspects and the programming aspects of machine learning. All the codes associated with algorithms are explained at length (Fig. 1.1).

Unique attempt: video lectures

In a distinctive approach, the book's content is complemented by video lectures that are available on iversity.com. This offers readers a multimedia learning experience, accommodating different learning preferences, and reinforcing the material through visual and auditory means.

The book covers the following areas.

- Machine learning concepts.
- Explanation of mathematical concepts relevant to machine learning, explained from an application perspective.
- Supervised learning algorithms.
- Unsupervised learning algorithms.
- Deep learning algorithms.

Depending on your interest and knowledge level, you can either go through the book in the above order or jump to specific areas.

sn.pub/Lia1RD

G. Madhavan, *Mastering Machine Learning: From Basics to Advanced*, Transactions on Computer Systems and Networks, https://doi.org/10.1007/978-981-97-9914-5_1

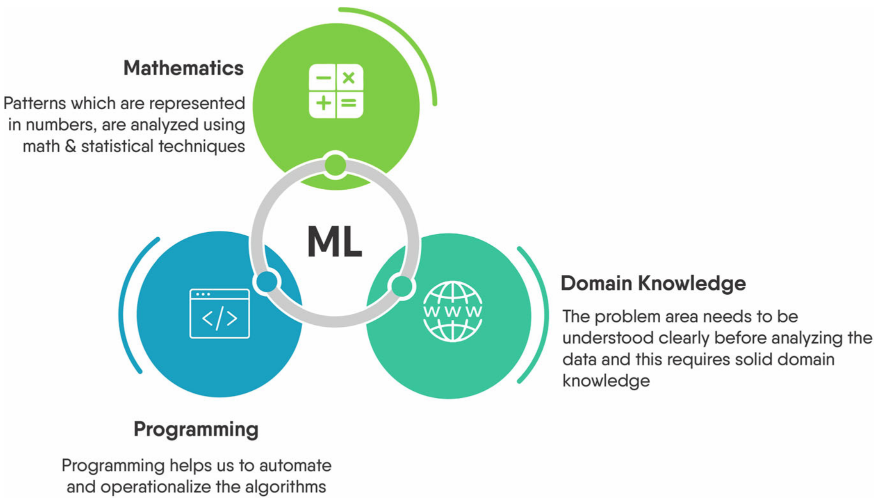

Fig. 1.1 Machine learning at intersection of three disciplines

Who this book is for:

- **Students**: This book is ideal for students at various stages of their education who are interested in entering the field of machine learning. It provides foundational knowledge that is crucial for beginners.
- **Young professionals in machine learning**: Young professionals who have recently started their careers in machine learning will find this book beneficial for consolidating their basic understanding and enhancing their practical skills.
- **Professionals from other disciplines wanting to get into machine learning**: For those looking to pivot into machine learning from different professional backgrounds, this book serves as a comprehensive guide to the basics and practical applications of ML concepts.

Who this book is not for:

- **Experts in machine learning**: Individuals who already have an advanced understanding of machine learning and are proficient in the latest ML techniques and theories might not find this book useful.
- **Those looking for advanced concepts** in machine learning: Readers in search of in-depth exploration of advanced machine learning topics may need to look beyond this book, which is tailored more toward foundational and intermediate knowledge.

Pre-requisites:

- Basic knowledge of Python programming including Pandas.

Datasets and codes

- The datasets used in this book along with the codes will be available at https://github.com/amgovindkumar/springer.

Errata

The book has undergone multiple revisions aimed at accuracy and clarity, thus minimizing the likelihood of errors. However, should readers encounter any mistakes, they are encouraged to report them to info@seaportAi.com so that the book can be further improved in future editions. This open channel for feedback demonstrates the author's commitment to the quality and reliability of the material and openness to receive feedback.

Chapter 2
Introduction to AI

AI Is the New Electricity, Illuminating Paths in Data Darkness

Abstract This chapter introduces readers to Artificial Intelligence (AI), exploring its fundamental nature as systems designed to emulate human cognition, problem-solving, and decision-making capabilities. Readers can look forward to gaining clarity on how Machine Learning enables AI systems to independently recognize patterns and improve performance from data without direct programming. A concise overview of AI history highlights milestones, periods of significant advancement, setbacks, and the influence of evolving technology. Special attention is given to the current "ChatGPT Era," showcasing how sophisticated language models have transformed communication and interactions between humans and machines. Additionally, the chapter explains the significance of the Turing Test, devised by Alan Turing, as a benchmark for evaluating a machine's intelligence based on human-like conversational ability.

Keywords Artificial intelligence · ChatGPT · Turing Test · Traditional programming · Machine learning

2.1 What Is Artificial Intelligence?

Artificial intelligence (AI), in simple terms, refers to the capability of a computer or a machine to imitate intelligent human behavior. It encompasses a broad range of technologies that enable machines to sense, comprehend, act, and learn. AI can be as simple as a chess program or as complex as a self-driving car.

Artificial intelligence (AI) has evolved significantly since its inception. The term "Artificial Intelligence" was first coined in 1956 at a conference at Dartmouth College. Initially, AI research focused on problem-solving and symbolic methods. In the 1960s, the US Department of Defense took interest in this type of work and began training computers to mimic basic human reasoning. However, the limitations of the then computing power slowed progress.

sn.pub/Lia1RD

G. Madhavan, *Mastering Machine Learning: From Basics to Advanced*, Transactions on Computer Systems and Networks, https://doi.org/10.1007/978-981-97-9914-5_2

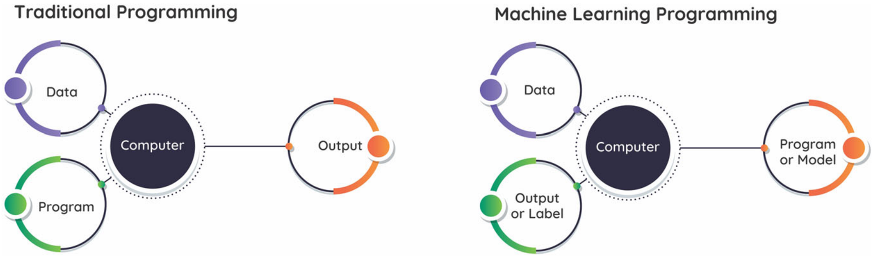

Fig. 2.1 Difference between traditional programming versus machine learning programming

As AI developed, a new branch emerged—machine learning (ML), which is a subset of artificial intelligence. Machine learning involves the development and use of algorithms that allow computers to learn from data. This learning is subsequently used to make predictions or decisions. This approach differs from traditional programming, where humans explicitly code the computer's every action (Fig. 2.1).

2.2 How Does the Machine Learn?

Machine learning algorithms use statistical techniques and equations to enable machines or software to improve at tasks with experience. The process starts with feeding the computer a large set of data. The machine then uses this data to train itself to recognize patterns and make decisions. Over time, as it is exposed to more data, its ability to make decisions and predictions improves.

At the heart of machine learning is pattern recognition. Computers are taught to recognize patterns and make connections between them. This ability is what enables a facial recognition system to identify faces or a recommendation system to suggest products you might like. Pattern recognition involves understanding the linkages in the historical data, which the machine uses to make predictions about new data.

2.3 History of Artificial Intelligence

The history of artificial intelligence (AI) is a tale of dreams, innovation, and milestones that stretch back much further than many people realize. It all began in the realm of science fiction, with stories of intelligent machines that could think and act on their own. But the real journey of AI started in the mid-twentieth century.

In the 1950s, a mathematician named Alan Turing asked a groundbreaking question: "Can machines think?" This question kicked off serious scientific efforts to create intelligent machines. Turing also developed the Turing Test, a method for determining whether a machine is capable of thinking like a human.

The term "artificial intelligence" was first coined by John McCarthy in 1956 at a conference at Dartmouth College. Here, McCarthy and others believed that machines could be made to simulate every aspect of learning or any other feature of intelligence.

The 1960s and 70s saw the first wave of AI, where researchers were optimistic about the future. They built the first AI programs that could solve algebra problems, prove theorems, and even understand natural language to a certain extent.

But then, AI hit its first roadblocks in the 1970s, a period known as "AI winter," and again in the 1980s. Funding dried up, and the high expectations were not met primarily because of the limitations in computer power and complexity of human intelligence.

The 1990s brought a resurgence of interest and progress in AI, thanks to improved computer technology and new approaches like machine learning, where computers could learn from data. The victory of IBM's Deep Blue over chess champion Garry Kasparov in 1997 was a significant milestone that showed the world what AI could do.

The twenty-first century has seen rapid progress in AI. We have voice-activated assistants, cars that can drive themselves, and algorithms that can diagnose diseases. AI has become a part of everyday life, with systems that can understand and process human language, recognize faces, and make decisions.

In recent years, the field of AI has continued to advance, with technologies like deep learning, where neural networks—systems inspired by the human brain—can learn from vast amounts of data. This has led to improvements in image and speech recognition that were hard to imagine just a few decades ago.

As we move forward, AI continues to push the boundaries of what machines are capable of, creating new possibilities and raising important questions about the future of technology and society. The story of AI is still being written, and the next chapters are likely to be even more astonishing as the line between human and machine intelligence becomes increasingly blurred.

2.4 The ChatGPT Era

As the history of artificial intelligence unfolded, a new chapter began with the arrival of ChatGPT, a sophisticated language model developed by OpenAI. ChatGPT, which stands for "Chat Generative Pre-trained Transformer," represents a significant leap forward in the field of natural language processing (NLP).

The roots of ChatGPT lie in the development of machine learning models known as transformers, introduced in a 2017 paper titled "Attention is All You Need." These models revolutionized the way machines understand and generate human language. They could pay "attention" to different parts of a sentence to understand context and meaning more effectively than ever before.

Building on this, ChatGPT was trained on a vast corpus of text data, allowing it to generate human-like text. It can engage in conversations, answer questions, write essays, and even create poetry. This is not just about finding the right words; it's about understanding the nuances of human language—a task that has traditionally been incredibly complex for machines.

The era of ChatGPT began when OpenAI released versions of the model to the public, where it was met with both enthusiasm and astonishment. Users found that they could have meaningful and coherent conversations with an AI. This was a far cry from the rule-based chatbots of the past that could only handle very specific queries.

ChatGPT has found its place in various applications:

- Customer Service: Automating responses to customer inquiries with a level of understanding that standard chatbots can't match.
- Education: Assisting in learning and tutoring by providing explanations to complex concepts and answering student questions.
- Content Creation: Helping writers by suggesting ideas, drafting content, and even writing code.
- Entertainment: Engaging users in interactive storytelling and creating personalized gaming experiences.

However, the ChatGPT era is not without its challenges. Concerns about misinformation, the potential for generating biased content, and the implications for job displacement are part of ongoing discussions. Furthermore, the technology raises important ethical questions about the use and misuse of AI-generated text.

The development of ChatGPT has also sparked a broader conversation about the future of AI and its integration into society. As AI becomes more advanced, issues of governance, security, and the societal impact of AI will need to be addressed more fully.

Despite these challenges, the ChatGPT era is an exciting time in the evolution of artificial intelligence. It marks a moment where AI began to interact with us in a way that feels more human, reshaping our relationship with technology and opening up new possibilities for how we communicate, work, and create. The story of ChatGPT continues to unfold, and its full impact on society remains to be seen.

2.5 The Turing Test

The Turing Test is a famous concept in the world of artificial intelligence, introduced by the British mathematician and computer scientist Alan Turing in 1950. The test was a part of Turing's paper, "Computing Machinery and Intelligence," where he pondered the question: "Can machines think?"

The Turing Test is a simple yet powerful idea. Turing proposed a test to determine whether a machine can exhibit intelligent behavior that is indistinguishable from that of a human. The test involves a human judge who interacts with two participants: one human and one machine. The interactions are typically text-based, so the judge cannot see or hear the participants. The judge asks questions, and both the human and the machine respond (Fig. 2.2).

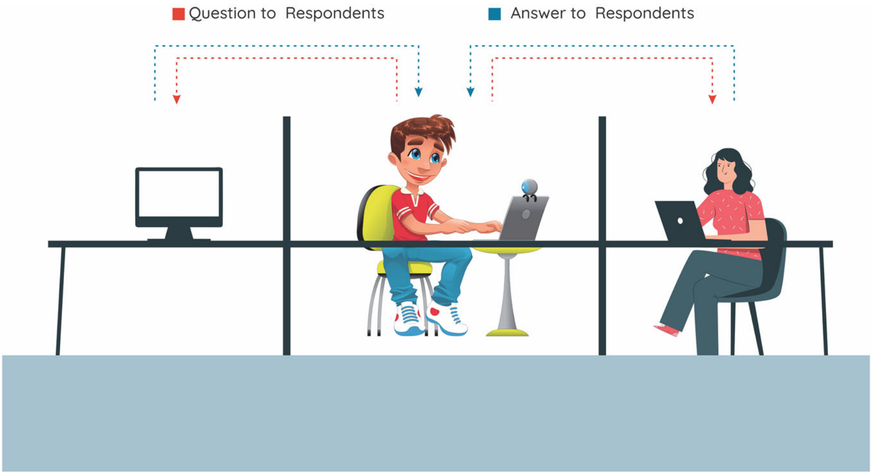

Fig. 2.2 Turing Test

The essence of the Turing Test is deception. If the judge can't reliably tell which participant is the machine and which is the human, the machine is said to have passed the Turing Test. In other words, if the machine can mimic human responses well enough that it fools the judge, it's considered a sign of intelligence.

The Turing Test has been both influential and controversial. It shifted the focus from whether machines can think to whether machines can behave like they think. It's more about the appearance of intelligence than intelligence itself. This has led to debates about what it truly means to be intelligent and whether the Turing Test is a valid measure of AI.

Despite its simplicity, the Turing Test has had a significant impact on the development of AI. It has been used as a benchmark for AI progress, inspiring the creation of more sophisticated and human-like AI systems.

However, the test has its limitations. Critics argue that passing the Turing Test doesn't necessarily mean a machine is truly intelligent. It might just be good at imitating human speech. Also, the test doesn't account for non-verbal intelligence or other forms of intelligence that don't involve human-like conversation.

Do you think ChatGPT will pass the Turing Test with flying colors?

Chapter 3
Pattern Recognition

Patterns Are Nature's Whispers, Revealing Secrets in Data's Chaos

Abstract In the previous chapter, readers explored foundational concepts of AI and learned how machines use data-driven approaches to emulate human intelligence. Building on this understanding, this chapter introduces readers to the core concept of pattern recognition—a vital process underlying machine learning. Through engaging quizzes, including classic problems such as predicting the next number in a series and the Fibonacci sequence, readers actively practice identifying and interpreting patterns. Each exercise encourages learners to follow key steps: understanding the given data, uncovering hidden relationships or patterns, and utilizing those insights to make accurate predictions. This approach mirrors exactly what occurs in machine learning, where algorithms analyse data to detect patterns and predict outcomes. By the chapter's end, readers will recognize how human pattern recognition serves as a blueprint for the techniques employed by intelligent systems.

Keywords Fibonacci series · Pattern recognition

3.1 Identify the Next Number in the Series

Question 1: What is the next number in the series: 2, 4, 8, 16, 32, …?

(A) 40
(B) 48
(C) 64
(D) 56

Question 2: What is the next number in the series: 5, 10, 20, 40, 80, …?

(A) 100
(B) 120
(C) 160
(D) 200

sn.pub/Lia1RD

G. Madhavan, *Mastering Machine Learning: From Basics to Advanced*, Transactions on Computer Systems and Networks, https://doi.org/10.1007/978-981-97-9914-5_3

Question 3: What is the next number in the series: 1, 1, 2, 3, 5, 8, …?

(A) 10
(B) 11
(C) 13
(D) 15

Question 4: What is the next number in the series: 3, 9, 27, 81, …?

(A) 162
(B) 243
(C) 324
(D) 405

Question 5: What is the next number in the series: 2, 3, 5, 8, 12, …?

(A) 15
(B) 17
(C) 18
(D) 20

3.2 Answers and Explanation to the Multiple-Choice Questions

Question 1

Series: 2, 4, 8, 16, 32, …

Answer: (C) 64

Explanation: The pattern in this series is each number being multiplied by 2. So, $32 \times 2 = 64$.

Question 2

Series: 5, 10, 20, 40, 80, …

Answer: (C) 160

Explanation: Similar to the first, this series multiplies each number by 2. Thus, $80 \times 2 = 160$.

Question 3

Series: 1, 1, 2, 3, 5, 8, …

Answer: (C) 13

Explanation: This is the Fibonacci sequence, where each number is the sum of the two preceding ones. Therefore, $8 + 5 = 13$. The golden ratio as shown below gets repeated across numbers (Fig. 3.1).

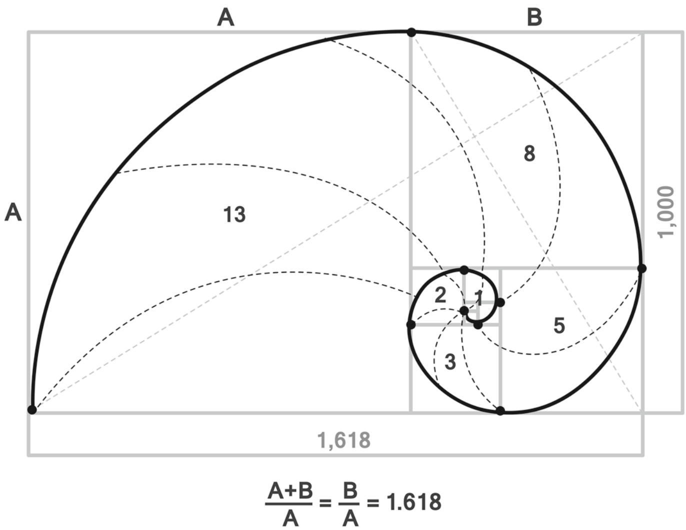

Fig. 3.1 Fibonacci series

Question 4

Series: 3, 9, 27, 81, …

Answer: (B) 243

Explanation: This series involves multiplying each number by 3. Hence,

$81 \times 3 = 243$

Question 5

Series: 2, 3, 5, 8, 12, …

Answer: (B) 17

Explanation: This series adds 1, 2, 3, 4, … to each number. So, the next number is

$12 + 5 = 17$

Each series follows a distinct numerical pattern, and identifying the pattern is key to predicting the next number.

3.3 Pattern Recognition in Machine Learning

Pattern recognition is a crucial part of machine learning, much like solving a puzzle by recognizing shapes and colors. In machine learning, pattern recognition is about teaching a computer to spot patterns and make sense of them.

The multiple-choice questions (MCQs) about predicting the next number in a series are simple examples of pattern recognition. In each series, there's a hidden rule. If you can find the rule, you can predict the next number. This is similar to what we ask computers to do in machine learning, but with more complex patterns.

In machine learning, pattern recognition involves algorithms that look at data and learn to recognize patterns. This can be anything from spotting a cat in a photo to predicting what a customer will buy next. The machine looks at examples (data) and starts to see patterns, like how numbers increase in a series or what features make up a cat's face.

The link between those MCQs and machine learning is the concept of identifying a pattern. In the MCQs, you might recognize that numbers are doubling or following the Fibonacci sequence. In machine learning, an algorithm might learn to recognize that certain email messages are spam or which transactions are likely fraudulent.

But in machine learning, the patterns can be much more complex, and the data can be huge. That's why we need computers and programming. They can analyze vast amounts of data and find patterns that might be too subtle or complicated for a human to see.

Pattern recognition is essential in machine learning because it's the foundation of learning from data. Whether it's a simple number series or a complex problem like speech recognition, the ability to recognize patterns helps machines to learn and make decisions. This capability is at the heart of making machines smarter and more useful in our daily lives.

Chapter 4
Introduction to Machine Learning

Teaching Machines to Learn, One Algorithm at a Time

Abstract In the previous chapter, readers discovered how recognizing patterns underpins the foundation of machine learning. Building upon this, the current chapter provides an introduction to machine learning (ML), clearly differentiating between supervised and unsupervised learning—the primary focus areas of this book. Readers will explore industry specific ML applications spanning diverse sectors such as agriculture, education, energy, entertainment, finance, manufacturing, and healthcare, highlighting ML's wide-ranging impact and utility. Additionally, the chapter briefly introduces reinforcement learning and semi-supervised learning, offering learners a broader perspective of the ML landscape. To simplify complex concepts, everyday analogies and relatable examples illustrate how ML models and algorithms function, making the topic approachable even to beginners. This chapter thus sets the stage for deeper exploration of supervised and unsupervised methods, preparing readers effectively for subsequent chapters dedicated to these key ML approaches. Readers will also understand why choosing the right ML technique is crucial for solving real-world problems. Finally, the chapter encourages readers to appreciate the power and potential of machine learning in driving innovation and improving decision-making across various domains.

Keywords Machine learning · Spam filter · Supervised learning · Unsupervised learning · Semi supervised learning · Reinforcement learning · Application

4.1 An Introduction to Machine Learning

Machine learning is a branch of artificial intelligence (AI) and computer science which focuses on the use of data and algorithms to imitate the way that humans learn. Machine learning models are a class of computer algorithms designed to improve their performance on specific tasks by learning from data. This sets them apart from

sn.pub/Lia1RD

G. Madhavan, *Mastering Machine Learning: From Basics to Advanced*, Transactions on Computer Systems and Networks, https://doi.org/10.1007/978-981-97-9914-5_4