

Artificial Intelligence for Students

*A comprehensive overview of AI's
foundation, applicability, and innovation*

Vibha Pandey



www.bpbonline.com

Copyright © 2023 BPB Online

All rights reserved. No part of this book may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, without the prior written permission of the publisher, except in the case of brief quotations embedded in critical articles or reviews.

Every effort has been made in the preparation of this book to ensure the accuracy of the information presented. However, the information contained in this book is sold without warranty, either express or implied. Neither the author, nor BPB Online or its dealers and distributors, will be held liable for any damages caused or alleged to have been caused directly or indirectly by this book.

BPB Online has endeavored to provide trademark information about all of the companies and products mentioned in this book by the appropriate use of capitals. However, BPB Online cannot guarantee the accuracy of this information.

First published: 2023

Published by BPB Online

WeWork

119 Marylebone Road

London NW1 5PU

UK | UAE | INDIA | SINGAPORE

ISBN 978-93-55517-968

www.bpbonline.com

Dedicated to

My parents

Ms. Lalita Pandey and Dr. R.K. Pandey

&

My mentor

Mr. Dalbir Dhankar

&

My Nephew

Master Abhiram Chaturvedi

Foreword

In today's world, it's essential for everyone to be able to encode and decode data, which has become an integral part of our lives rapidly. This book is for everyone, from beginners to experts; the author proves that 'One size may sometimes fit all.' In this case, the basics of data science are explained in such a way that the often-dreaded subject becomes a lot more accessible to everyone. As a data scientist, I highly recommend this book to anyone interested in gaining a deeper understanding of data science and machine learning. It is a comprehensive resource that covers a range of topics, from data structures and data representation to clustering, classification, and storytelling.

The section on data structures is highly informative and crisp, providing a clear breakdown of the different types of data structures and how they can be used to organize and manage data efficiently. It will prove to be an elixir of knowledge, especially helpful to beginners who may not be familiar with the concept.

The chapter on data representation is also well-written, providing a detailed overview of the different ways data can be presented and how those presentations can impact the insights and conclusions drawn from the data. The section on finding the best-suited graphical representation for the data provides practical insights into how data can be effectively communicated to others.

Concepts of clustering and classification are both covered extensively, with clear explanations of their differences and how they can be used for pattern identification in machine learning. The chapter on classification provides an in-depth explanation of how data labeling works and the importance of understanding true positives, true negatives, false positives, and false negatives. The introduction of logistic regression as a tool for binary classification is also a great addition to the chapter that enhances its practicality.

An often untouched aspect of data science, storytelling, is a unique and valuable addition to the book. It emphasizes the importance of storytelling in conveying insights and making data more accessible. The practical tips and best practices shared in this chapter are sure to be helpful in any storytelling endeavor.

One of the standout features is the section on ethical considerations in machine learning. It highlights the biases that exist in human society and how they can influence the accuracy of machine learning algorithms. The section on principles

for ethical AI is particularly insightful, providing a framework for creating machine learning models that are fair, transparent, and accountable.

Overall, this book is a must-read for anyone interested in data science or machine learning. The chapters are well-structured and informative and provide practical insights that can be applied in real-world situations. A beginner in the field can easily get the gist of what is being shared in the group due to the author's ability to explain heavier concepts in a fluid manner. The author's writing style is warm and professional, making it an enjoyable and engaging read. I would highly recommend this book to anyone looking to expand their knowledge in the field of data science.

-Santanu Bhattacharya

Santanu Bhattacharya currently holds the position of Chief Technologist at NatWest Bank. He also holds a Ph.D. from NASA's Goddard Space Flight Center/UMD and is an alumnus of MIT and IIT-Bombay. As a speaker at the World Economic Forum (WEF) Davos 2020 and a writer, he has been covered in publications such as Mashable, TechCrunch, Forbes, Le Monde, and Economic Times, among others.

About the Author

Vibha Pandey has rich corporate experience of more than 25 years with MNCs like PwC, Oracle, Nortel, Siemens, and Samsung, to name a few. She played central roles in numerous telecom, and IT projects in different profiles, such as a software engineer, product life cycle management, test engineer, presales, sales, business development, and key account management, including project implementation in the US, Indonesia, Japan, and India. Currently, she runs a firm focused on telecom and security projects as well as is a National President of AI at WICCI. She also takes up consulting projects to write research papers. She has taken up various other professions, including being a visiting faculty at NIFT, Delhi, where she introduced Systems Thinking as a subject that was made mandatory.

Furthermore, the author has been associated with Smart India Hackathon right from its inception years for both software and hardware divisions as a mentor and evaluator.

About the Reviewers

- ❖ **Vibhu** is a software professional passionate about Machine Learning, Cloud Computing, and Agile Development. With over 25 years of experience in the software industry, he has worked on a wide range of projects, from small startups to large enterprises, delivering high-quality software demanding 99.999% availability.

Currently, he leads engineering efforts for observability solutions (both on-premise and in the cloud) that touch around 300,000 customers. He is a hands-on leader, unafraid to roll up his sleeves and get involved in the code when needed.

He profoundly understands Cloud Computing (with multiple certifications with different cloud providers) and constantly explores new technologies and techniques to help deliver better software faster. In his spare time, he teaches Machine Learning to professionals trying to break into new areas.

He is passionate about his work and continually seeks new challenges and opportunities to learn and grow as an engineering leader.

- ❖ **Phani** is an experienced Technical Lead with a background in the health care industry. His skillset includes expertise in Predictive modeling, Artificial Intelligence, and Machine Learning using programming languages such as Python, R, and SAS. They also have a solid understanding of Big data, data visualizations, and Storytelling with Data. Throughout their career, they have worked with various relational and non-relational databases as well as data visualization tools such as Tableau and PowerBi.

In terms of education, this individual holds a Bachelor of Technology degree focused in Electrical and Electronics Engineering from Jawaharlal Nehru Technological University. They also pursued a Masters in Business Analytics from the University of Louisville in the USA, where they achieved a remarkable 3.95/4 GPA. Additionally, they are currently pursuing a Ph.D. from the University of Louisville.

Overall, this individual's experience and education demonstrate their expertise in technical fields related to data analysis and their dedication to continuing to expand their knowledge and skillset in this area.

Acknowledgement

I want to express my deepest gratitude to my family and friends for their unwavering support and encouragement throughout this book's writing, especially my parents, my mentor, and my nephew, who has authored books from the very young age of 7 years.

I am also grateful to BPB Publications for their guidance and expertise in bringing this book to fruition. It was a long journey of revising this book, with valuable participation and collaboration of reviewers, technical experts, and editors.

I would also like to acknowledge the valuable contributions of my ex-colleagues and industry alliances during many years working in the tech industry, who have taught me so much and provided valuable feedback on my work.

Finally, I would like to thank all the readers who have taken an interest in my book and for their support in making it a reality. Your encouragement has been invaluable

Preface

AI is a discipline in computer science that focuses on developing intelligent machines, machines that can learn and then teach themselves. These machines, then, can process vast amounts of data than humans can and several times faster. However, AI can go across all disciplines to change the world for the better– from creating new healthcare solutions to designing hospitals of the future, improving farming and our food supply, helping refugees acclimatize to new environments, improving educational resources and access, and even cleaning our oceans, air, and water supply. The potential for humans to improve the world through AI is endless as long as we know how to use it.

This book is designed to provide a comprehensive guide to a planned sequence of instructions consisting of units meant for developing the employability and vocational competencies of students opting for skill subjects along with other educational subjects.

Throughout the book, you will learn about the key aspects of Artificial Intelligence with real-world examples for readers to relate.

This book is intended for readers who are new to Artificial Intelligence and want to explore and experiment in this field.

With this book, you will gain the knowledge and skills to start developing applications using existing frameworks of interest, such as chatbots. I hope you will find this book informative and helpful.

Chapter 1: Introduction: AI for Everyone - explains artificial intelligence and machine learning, terminologies, and related concepts. It also describes the AI products/ applications in society and their being different from non-AI products/ applications. Readers also learn about jobs that may appear in the future.

Chapter 2: AI Applications and Methodologies - presents an overview of areas where artificial intelligence can be applied (like in the field of computer vision, speech, text, etc.). Readers also get an overview of deep learning. This chapter covers the impact of AI on our society and how we can get ready for the future, that is, the AI age.

Chapter 3: Mathematics in Artificial Intelligence – revisit the mathematics involved in artificial intelligence, such as linear algebra, statistics, and set theory. This also covers the basics of graphs and describes the application of math in AI. Readers will learn about representing data in terms of mathematical formulas.

Chapter 4: AI Values (Ethical Decision-Making) - covers the ethics, bias, and impacts of bias on society. It gives special attention to issues and concerns around AI. It helps readers learn to spot issues in the data, make arguments, and apply rules.

Chapter 5: Introduction to Storytelling – this chapter is all about storytelling, the need, storytelling with data, insights from storytelling, and more. It allows the reader to learn to apply imagination, mapping the plot into key events and increasing memory retention. It also guides the creation of blogs, videos, and other content as per the audience and about the conflict and resolution.

Chapter 6: Critical and Creative Thinking – explains the design thinking framework, that is, understanding the problem and being able to express the same. Readers learn to develop/innovate from the design of a solution.

Chapter 7: Data Analysis - explains types of structured data and statistical principles such as frequency tables, mean median, mode, range, and more. Readers learn to represent data in terms of graphs and statistical models. By the end of this chapter, the reader is able to represent a simple problem in terms of numbers.

Chapter 8: Regression – explains mathematical concepts such as correlations, regression, and other related terms. At the end of the chapter, readers learn to relate data with regression and correlation. Readers also get to know about everyday applications of these mathematical concepts.

Chapter 9: Classification and Clustering - explains in detail classification and its types, the kind of problems that may be placed under the category of a classification problem, and where to apply classification principles. Readers are also made aware of the impact of the application of incorrect algorithms on society. In the remaining half of this chapter, readers learn about clustering problems and their application, and why it is called clustering. Readers get an overview of the application of clustering problems using standard models.

Chapter 10: AI Values (Bias Awareness) - explains what ethics are, the impact of ethics on society, as well as the impact of bias on AI functioning. Readers are also able to learn about the impact of biases in data and how to de-bias or neutralize

biased data. By the end of the chapter, readers are able to easily find bias in the acquired dataset

Chapter 11: Capstone Project - introduces readers to commonly used algorithms and the science behind them. This chapter also engages the readers in understanding and decomposing a problem, the analytical approach, and data requirements and collection. This chapter also introduces the validation of the model quality and metrics of model quality. The chapter ends with showcasing a compelling story through all the methodologies and learnings that readers are exposed to in the chapter.

Chapter 12: Model Lifecycle (Knowledge) - this chapter explains different aspects of the model as well as the lifecycle of an AI model.

Chapter 13: Storytelling Through Data - explains the need for storytelling and various related topics, such as the creation and ethics of stories, expressing the related data with suitable charts. This chapter also captures the stories during the step of predictive modeling and ends with the best practices of storytelling.

Chapter 14: AI Applications in Use in Real-World— mentions different fields and the associated real-world AI applications.

Coloured Images

Please follow the link to download the
Coloured Images of the book:

<https://rebrand.ly/skz5nh8>

We have code bundles from our rich catalogue of books and videos available at **<https://github.com/bpbpublications>**. Check them out!

Errata

We take immense pride in our work at BPB Publications and follow best practices to ensure the accuracy of our content to provide with an indulging reading experience to our subscribers. Our readers are our mirrors, and we use their inputs to reflect and improve upon human errors, if any, that may have occurred during the publishing processes involved. To let us maintain the quality and help us reach out to any readers who might be having difficulties due to any unforeseen errors, please write to us at :

errata@bpbonline.com

Your support, suggestions and feedbacks are highly appreciated by the BPB Publications' Family.

Did you know that BPB offers eBook versions of every book published, with PDF and ePub files available? You can upgrade to the eBook version at www.bpbonline.com and as a print book customer, you are entitled to a discount on the eBook copy. Get in touch with us at :

business@bpbonline.com for more details.

At www.bpbonline.com, you can also read a collection of free technical articles, sign up for a range of free newsletters, and receive exclusive discounts and offers on BPB books and eBooks.

Piracy

If you come across any illegal copies of our works in any form on the internet, we would be grateful if you would provide us with the location address or website name. Please contact us at **business@bpbonline.com** with a link to the material.

If you are interested in becoming an author

If there is a topic that you have expertise in, and you are interested in either writing or contributing to a book, please visit **www.bpbonline.com**. We have worked with thousands of developers and tech professionals, just like you, to help them share their insights with the global tech community. You can make a general application, apply for a specific hot topic that we are recruiting an author for, or submit your own idea.

Reviews

Please leave a review. Once you have read and used this book, why not leave a review on the site that you purchased it from? Potential readers can then see and use your unbiased opinion to make purchase decisions. We at BPB can understand what you think about our products, and our authors can see your feedback on their book. Thank you!

For more information about BPB, please visit **www.bpbonline.com**.

Join our book's Discord space

Join the book's Discord Workspace for Latest updates, Offers, Tech happenings around the world, New Release and Sessions with the Authors:

<https://discord.bpbonline.com>



Table of Contents

1. Introduction: AI for Everyone.....	1
Introduction.....	1
Structure.....	1
Data explosion.....	2
What is Artificial Intelligence.....	2
Artificial Intelligence: History and evolution.....	3
The father of AI.....	4
Types of Artificial Intelligence	5
<i>Based on the capabilities of AI.....</i>	<i>5</i>
<i>Based on the functionality of AI.....</i>	<i>6</i>
What is machine learning.....	9
<i>Examples.....</i>	<i>9</i>
What is data.....	10
What is Deep learning.....	11
<i>Examples.....</i>	<i>11</i>
Machine learning techniques and training.....	12
Neural networks.....	15
What machine learning can and cannot do	16
Key differences between artificial intelligence (AI) and machine learning (ML).....	17
Artificial Intelligence project life cycle	18
Career opportunities in artificial intelligence.....	21
Conclusion.....	24
Multiple choice questions.....	24
<i>Answers</i>	<i>24</i>
2. AI Applications and Methodologies	27
Introduction.....	27

Structure.....	27
Key fields of application in AI	28
<i>Chatbots (Natural Language Processing, speech)</i>	<i>28</i>
<i>What are chatbots.....</i>	<i>28</i>
<i>Benefits of chatbots.....</i>	<i>29</i>
<i>Chatbots in key industries.....</i>	<i>31</i>
<i>Machine learning in AI-powered chatbots</i>	<i>33</i>
<i>Generative chatbots – Deep learning.....</i>	<i>33</i>
<i>Artificial neural networks to replicate a human brain – Intelligent chatbot</i>	<i>34</i>
<i>Natural Language Processing (NLP) – Natural conversation</i>	<i>34</i>
<i>Natural language understanding (NLU) – Complex questions.....</i>	<i>34</i>
<i>Voice assistants (Alexa, Siri, and others)</i>	<i>35</i>
<i>Benefits</i>	<i>35</i>
<i>Pitfalls</i>	<i>36</i>
<i>Examples</i>	<i>36</i>
<i>Computer vision</i>	<i>36</i>
<i>Weather predictions</i>	<i>37</i>
<i>Price forecast for commodities</i>	<i>37</i>
<i>Self-driving cars.....</i>	<i>38</i>
Characteristics and types of AI.....	40
<i>Data-driven.....</i>	<i>41</i>
<i>Autonomous systems.....</i>	<i>42</i>
<i>Recommender systems.....</i>	<i>42</i>
<i>Human-like.....</i>	<i>42</i>
Cognitive computing	43
Deep-dive into NLP, CV, and much more.....	43
AI and society.....	44
The future with AI and AI in action.....	47
Non-technical explanation of deep learning	49
Commonly used deep learning algorithm explained	50
Conclusion.....	53

Multiple choice questions.....	53
<i>Answers</i>	54
Questions	54
3. Mathematics in Artificial Intelligence.....	55
Introduction.....	55
Structure.....	55
Overview of matrices.....	56
<i>Definition of matrices</i>	56
<i>Representation</i>	56
<i>Order of a matrix</i>	59
<i>Type of matrices</i>	59
<i>Equality of matrices</i>	60
<i>Operations on matrices</i>	60
<i>Addition</i>	60
<i>Subtraction</i>	60
<i>Multiplication</i>	60
<i>Transpose</i>	61
<i>Symmetric and skew-symmetric</i>	61
<i>Invertible</i>	61
<i>Use of matrix in AI</i>	61
Overview of set theory.....	61
<i>What are sets</i>	62
<i>Empty sets</i>	63
<i>Finite and infinite sets</i>	63
<i>Equal sets</i>	63
<i>Subsets</i>	64
<i>Power set</i>	64
<i>Universal set</i>	65
<i>Venn diagrams</i>	65
<i>Operations on sets</i>	66

<i>Union of sets</i>	66
<i>Intersection of sets</i>	67
<i>Difference of sets</i>	68
<i>Complement of a set</i>	68
<i>Some properties of complement sets</i>	68
<i>Introduction to data table joins</i>	69
<i>Inner joins</i>	69
<i>Left inner joins</i>	70
<i>Right inner joins</i>	70
<i>Left outer joins</i>	71
<i>Right outer joins</i>	71
<i>Outer joins</i>	71
Simple statistical concepts.....	72
<i>Descriptive statistics</i>	72
<i>Correlation</i>	73
<i>Probability distribution</i>	74
<i>Normal distribution</i>	74
<i>Bias</i>	74
<i>Covariance</i>	75
Visual representation of data	75
<i>Bar chart</i>	75
<i>Line chart</i>	76
<i>Scatter plot</i>	77
<i>Box plot</i>	77
<i>Common variations</i>	78
<i>Histogram</i>	78
<i>Stacked bar chart</i>	78
<i>Area chart</i>	79
<i>Dual-axis chart</i>	79
<i>Bubble chart</i>	80
<i>Pie chart</i>	80
<i>Heat map</i>	81

<i>Map-based plots</i>	81
With coordinates and graphs- introduction to the dimensionality of data.....	82
<i>Data dimensionality</i>	82
<i>High and low dimensional data</i>	82
<i>Reduction of dimensionality</i>	82
<i>Curse of dimensionality</i>	82
<i>Data representation on graph</i>	82
Simple linear equation.....	84
<i>Linear regression</i>	84
<i>Least square method of regression</i>	84
Conclusion.....	85
Multiple choice questions.....	85
<i>Answers</i>	86
Questions.....	86
4. AI Values (Ethical Decision-Making)	87
Introduction.....	87
Structure.....	87
Issues and concerns around AI.....	88
AI and ethical concerns.....	89
<i>Privacy and Surveillance</i>	89
<i>Manipulation</i>	89
<i>Opacity</i>	89
<i>Bias in Decision</i>	90
<i>Human-Robot Interaction</i>	90
<i>Automation and labor market</i>	90
<i>Autonomous systems</i>	90
<i>Ethics for ethical machines</i>	91
AI and bias.....	92
<i>Self-driving cars</i>	92
<i>Model disaster</i>	92

<i>Chatbots' inappropriate responses</i>	92
What are the principles of ethical AI?	93
<i>Model interpretability</i>	93
<i>Reliability of the system</i>	94
<i>End-to-end security</i>	94
<i>Who is accountable</i>	94
<i>How is it beneficial</i>	94
<i>Privacy and surveillance</i>	95
<i>Human involvement</i>	95
<i>Policies and regulations</i>	95
<i>Unbias</i>	95
<i>Safety and security</i>	95
Enterprises and use of AI technology	95
<i>Debugging</i>	95
<i>Responsibility</i>	95
<i>No-bias</i>	96
<i>Purpose and misuse</i>	96
<i>What is the AI code of ethics in enterprises</i>	96
AI: Ethics, Bias, and Trust	97
<i>The 23 Asilomar AI principles</i>	97
<i>Research issues</i>	98
<i>Ethics and values</i>	98
<i>Longer-term issues</i>	100
Employment and AI	100
<i>Bias in recruitment</i>	100
<i>AI applications replacing humans</i>	101
Conclusion	101
Multiple choice questions	101
<i>Answers</i>	102
Questions	102

5. Introduction to Storytelling	103
Introduction.....	103
Structure.....	103
Storytelling: communication across the ages.....	104
<i>Significance of storytelling in data storytelling</i>	<i>104</i>
The need for storytelling	105
Storytelling with data.....	107
<i>By the numbers: how to tell a great story with your data</i>	<i>108</i>
<i>Understand data context.....</i>	<i>109</i>
<i>Audience is the key.....</i>	<i>109</i>
<i>Data that matters</i>	<i>109</i>
<i>Data analysis and insights.....</i>	<i>109</i>
<i>Effective data visualizations.....</i>	<i>110</i>
<i>Set the context.....</i>	<i>110</i>
<i>About the story.....</i>	<i>110</i>
<i>Clear and concise.....</i>	<i>111</i>
Conflict and resolution	111
<i>Everyone wants to resolve conflict, and a good data</i> <i>storyteller is there to help!</i>	<i>112</i>
Example of a story	112
Storytelling for audience	114
<i>Your data storytelling depends on the background knowledge of</i> <i>your audience.....</i>	<i>114</i>
Insights from storytelling.....	115
<i>Make the audience care about the data</i>	<i>116</i>
<i>Keep the audience engaged.....</i>	<i>116</i>
<i>Create from the end; present from the beginning.....</i>	<i>116</i>
<i>Start with an anecdote, end with the data</i>	<i>116</i>
<i>Build suspense, not surprise.....</i>	<i>116</i>
Conclusion.....	117
Multiple choice questions.....	117

<i>Answers</i>	118
Questions	118
6. Critical and Creative Thinking.....	119
Introduction.....	119
Structure.....	119
Design thinking framework.....	120
<i>Right questioning (5W and 1H)</i>	124
<i>Identifying the problem to solve</i>	127
<i>Ideate</i>	129
<i>Why is ideation significant in design thinking</i>	129
<i>Ideation techniques and tools</i>	129
<i>Key challenges for ideation</i>	131
<i>Key success factors for ideation</i>	131
Conclusion.....	131
Multiple choice questions.....	132
<i>Answers</i>	133
Questions	133
7. Data Analysis	135
Introduction.....	135
Structure.....	136
AI and data.....	136
<i>Structured data</i>	137
<i>Characteristics</i>	137
<i>Sources</i>	138
<i>Advantages</i>	138
<i>Unstructured data</i>	138
<i>Characteristics</i>	139
<i>Sources</i>	139
<i>Advantages</i>	139
<i>Disadvantages</i>	140

<i>Difficulties faced in storing</i>	140
<i>Solutions for storing</i>	140
<i>Metadata– data about data</i>	140
<i>Semi-structured data</i>	141
<i>Characteristics</i>	141
<i>Source</i>	142
<i>Advantages</i>	142
<i>Disadvantages</i>	142
<i>Difficulties in storing</i>	143
<i>Solutions for storing</i>	143
Types of structured data	144
<i>Date and time</i>	144
<i>String</i>	146
<i>Categorical data</i>	148
<i>Nominal data</i>	148
<i>Ordinal data</i>	148
<i>Binary data</i>	149
Representation of data	149
Exploring data (Pattern recognition).....	150
<i>Cases, variables, and levels of measurement</i>	151
<i>Quantitative or numerical</i>	151
<i>Categorical data</i>	151
<i>Data matrix and frequency table</i>	155
<i>Data matrix</i>	155
<i>Frequency table</i>	155
<i>Graphs and shapes of distributions</i>	157
<i>Categorical variable</i>	157
<i>Quantitative variable</i>	158
<i>Mode, median, and mean</i>	163
<i>Mode</i>	163
<i>Median</i>	164

<i>Mean</i>	164
<i>Range, interquartile range, and box plot</i>	164
<i>Range</i>	165
<i>Interquartile range</i>	165
<i>Variance and standard deviation</i>	168
<i>Z-scores</i>	169
Conclusion.....	170
Multiple choice questions.....	170
<i>Answers</i>	171
Questions.....	171
Practice exercise.....	171
8. Regression	173
Introduction.....	173
Structure.....	173
Correlation and regression.....	174
<i>Correlation</i>	174
<i>Regression</i>	177
<i>Crosstabs and scatterplots</i>	178
<i>Crosstabs</i>	178
<i>Scatterplots</i>	179
<i>Pearson's r</i>	180
<i>Regression - Finding the line</i>	183
<i>Regression - Describing the line</i>	184
<i>Regression - How good is the line</i>	187
<i>R-Squared: goodness of fit</i>	187
<i>Adjusted R Squared</i>	188
<i>Correlation is not causation</i>	188
<i>Examples of correlation and regression</i>	189
<i>Example 1</i>	189
<i>Example 2</i>	190

<i>Caveats and examples</i>	192
Conclusion.....	194
Multiple choice questions.....	194
<i>Answers</i>	194
Questions	195
9. Classification and Clustering	197
Introduction.....	197
Structure.....	198
What is a classification problem	198
<i>Algorithms for classification problems</i>	199
<i>Label encoding</i>	199
<i>Classification terminologies used in machine learning</i>	199
<i>Types of learners in classification</i>	200
<i>Binary classification</i>	200
<i>Multi-class classification</i>	201
<i>Multi-label classification</i>	202
<i>Imbalanced classification</i>	202
Examples.....	204
<i>Simple binary classification</i>	204
<i>Sentiment analysis</i>	204
Introduction to binary classification with logistic regression	207
<i>Logistic regression equation</i>	208
<i>Type of logistic regression</i>	209
True positives, true negatives, false positives, and false negatives.....	209
<i>Where we should care more</i>	212
<i>Criminal courts</i>	212
<i>Metal detectors</i>	213
<i>Medical tests</i>	213
What is a clustering problem	214
<i>Partitioning based algorithm</i>	215

<i>K-Means</i>	215
<i>Hierarchical clustering</i>	216
Why is it unsupervised	219
Examples.....	220
<i>Biology</i>	220
<i>Hierarchical clustering in astronomy</i>	221
<i>Classification and clustering comparison</i>	222
Conclusion.....	223
Multiples choice questions	223
<i>Answers</i>	224
Questions	225
10. AI Values (Bias Awareness)	227
Introduction.....	227
Structure.....	227
AI working for good.....	228
<i>Healthcare</i>	228
<i>Environment</i>	228
<i>Education</i>	229
<i>Public safety</i>	230
<i>Focussed deterrence</i>	230
<i>Investigating crimes</i>	231
<i>Counter-terrorist threats</i>	231
<i>Mitigating damage during natural disasters</i>	231
<i>Crowd management</i>	231
<i>Vehicles: Find optimal routes for an efficient and safe commute</i>	231
<i>Agriculture: Maximizing food production and decreasing waste</i>	231
<i>Save the Bees</i>	232
<i>Accessibility for people with disabilities</i>	232
<i>Conservation of wildlife</i>	232
<i>Fight the World Hunger</i>	232

Principles for ethical AI	232
<i>Respect for freedom, autonomy and dignity of human beings</i>	233
<i>Non-discrimination and fairness</i>	234
<i>Transparency and explainability</i>	234
<i>Human oversight and accountability</i>	235
<i>Privacy, data protection and data security</i>	235
<i>Reliability and safety</i>	235
<i>Inclusiveness and diversity</i>	235
<i>Sustainability and environmental protection</i>	235
<i>Human values</i>	236
<i>Contestability</i>	236
<i>Wellbeing of humans, society, and environment</i>	236
Types of bias (personal / cultural / societal)	236
<i>Personal</i>	236
<i>Social and cultural bias</i>	237
How bias influences AI-based decisions	238
<i>Data gathering bias</i>	239
<i>Data analysis bias</i>	240
<i>Data application bias</i>	240
How data-driven decisions can be debiased	240
Conclusion.....	241
Multiple choice questions.....	241
<i>Answers</i>	243
Questions	243
11. Capstone Project	245
Introduction	245
Structure.....	245
Understanding the problem.....	246
<i>Project 11.1</i>	247
Decomposing the problem through the DT framework.....	248

<i>Project 11.2</i>	250
Analytic approach	250
<i>Project 11.3</i>	251
Data requirements	251
<i>Project 11.4</i>	252
Data collection.....	252
<i>Project 11.5</i>	253
Modeling approach	253
How to validate the model quality	254
<i>By test-train split</i>	255
<i>Introduce concept of cross-validation</i>	257
<i>The cross-validation procedure</i>	257
Metrics of model quality by simple Maths	258
<i>RMSE- Root Mean Squared Error</i>	260
<i>MSE – Mean squared error</i>	262
<i>Project 11.6</i>	263
<i>MAPE – Mean Absolute Percent Error</i>	263
Introduction to commonly used algorithms and the science behind them	264
<i>Supervised learning algorithm</i>	264
<i>Unsupervised learning algorithm</i>	264
<i>Semi-supervised algorithm</i>	265
<i>Reinforcement learning algorithm</i>	265
Showcase through a compelling story.....	265
<i>AlphaGo</i>	265
<i>Deep Patient</i>	266
<i>Project 11.7</i>	266
Conclusion.....	267
Multiple choice questions.....	267
<i>Answers</i>	268
Questions	268

12. Model Lifecycle (Knowledge)	269
Introduction.....	269
Structure.....	269
Different aspects of model.....	270
<i>Train, test, and validate</i>	270
<i>Train with training data set</i>	271
<i>Validation dataset</i>	271
<i>Test set</i>	272
<i>What are hyperparameters?</i>	273
<i>Hyperparameters</i>	273
<i>Commonly used platforms to build and run models (Introduction)</i>	275
<i>Amazon SageMaker</i>	275
<i>Google Cloud AI Platform</i>	276
<i>Microsoft Azure</i>	277
<i>Recommended tools</i>	278
<i>Links to different platforms</i>	279
Lifecycle of an AI model	280
<i>Build</i>	281
<i>Data collection</i>	281
<i>Prepare the data</i>	281
<i>Choosing the algorithm</i>	281
<i>Train the model</i>	282
<i>Test and validate model</i>	282
<i>Deploy</i>	282
<i>Retrain</i>	282
Conclusion.....	283
Multiple choice questions.....	283
<i>Answers</i>	284
Questions	284

13. Storytelling Through Data	285
Introduction.....	285
Simple statistics.....	285
Structure.....	286
The need for storytelling	287
<i>Information processing and recalling stories</i>	289
<i>Why is storytelling important</i>	290
<i>Structure that story</i>	290
How to create stories.....	291
<i>Begin with a pen-paper approach</i>	292
<i>Dig deeper to identify the sole purpose of your story</i>	292
<i>Use powerful headings</i>	292
<i>Design a roadmap</i>	292
<i>Conclude with brevity</i>	292
Ethics of storytelling.....	293
<i>Queries</i>	293
<i>Practices</i>	293
Types of data and suitable charts	293
<i>Text [Wordclouds]</i>	294
<i>Mixed [Facet Grids]</i>	294
<i>Numeric [Line Charts/ Bar Charts]</i>	296
<i>Stocks [Candlestick Charts]</i>	297
<i>Geographic [Maps]</i>	298
Stories during the steps of predictive modeling	299
<i>Data exploration</i>	299
<i>Feature visualizing</i>	300
<i>Model creation</i>	302
<i>Model comparisons</i>	302
<i>Benefits of model visualization</i>	304
Best practices of storytelling	304
Conclusion.....	304

Multiple choice questions.....	305
<i>Answers</i>	305
Questions	305
14. AI Applications in Use in Real-World	307
Introduction.....	307
Structure.....	307
Self-driving cars.....	308
<i>Automobile companies</i>	308
<i>Games based on self-driving cars</i>	310
Industrial AI	310
AI monitored artificial womb	310
Chatbots	310
<i>Travel chatbots</i>	311
<i>Multilingual chatbots</i>	311
AI-enabled voices	311
AI-enabled video creator	312
Audio transcripts.....	312
Content creators	312
Storytelling	313
Meeting assistant	314
Schedule assistant.....	314
Automated website browsing.....	314
Text to image	315
Personal lawyer.....	316
Other AI tools.....	317
AI tools database.....	322
Conclusion.....	322
Index	323-335

CHAPTER 1

Introduction: AI for Everyone

Introduction

The objective of this chapter is to learn the Artificial Intelligence concept, classification, and its components. We will start the chapter with a discussion on what is the need for Artificial Intelligence and end with the career opportunities which are available in this space.

Structure

In this chapter, we will be discussing:

- Data explosion
- What is Artificial Intelligence
- Artificial Intelligence: History and evolution
 - The father of AI
- Types of Artificial Intelligence
 - Based on the capabilities of AI
 - Based on the functionality of AI
- What is machine learning
- What is data

- What is deep learning
- Machine learning techniques and training
- Neural networks
- What machine learning can do and cannot do
- Key differences between artificial intelligence and machine learning
- Artificial Intelligence project life cycle
- Career opportunities in artificial intelligence

Data explosion

We live in a world with an ever increasing amount of data that both humans and machines generate. It far outpaces humans' ability to extract meaningful information and make informed and complex decisions based on the extensive data to process.

Every day, we create roughly 2.5 quintillion bytes of data (that's 2.5, followed by a staggering 18 zeros!)

We may not be aware, but we have been using Artificial Intelligence based technologies in our daily routine. Scientists found that an average person today can process as much as 74 gigabytes (GB) of data a day.

Artificial Intelligence is a technology that is transforming every walk of life with its five basic components include learning, reasoning, problem-solving, perception, and language understanding.

This book is written with the goal of explaining the technology with examples. Let us start with brushing some basic definitions and visiting the history of Artificial intelligence to set the context.

What is a machine?

A machine is a piece of equipment with moving parts that humans design to do a particular job. A machine usually needs electricity, gas, steam, and so on to work.

What is a computer?

A computer is an electronic machine that can store, find and arrange information, calculate amounts, and control other machines.

What is Artificial Intelligence

The human brain has the ability to think, read, learn, remember, reason, and pay attention. Such capabilities are termed cognitive skills. The term “**Intelligence**”

is used for cognitive (connected with the processes of understanding) skills and thinking ability of humans and animals. We may also call it “natural intelligence.”

Then what is *Artificial Intelligence* (referred to as AI in the remaining book)?

The terminology comprises of two words “*Artificial*” and “*Intelligence*.” **Artificial** refers to something that is not natural or is made by humans. AI is, then, intelligence demonstrated by a computer (an electronic machine), hence, it can also be referred to as “machine intelligence.”

In other words, AI is best described as machines having human-like cognitive skills of learning and problem solving by making decisions in such a way that they can be associated with human minds.

To summarize, AI is a field of computer science (not science fiction) combining robust datasets with the aim of having computers simulate intelligent processes. Here the computer needs AI implemented in its system to demonstrate AI capabilities.

Today AI contributes much to our human lives. Industries, including retail, healthcare, manufacturing, agriculture, insurance, and finance, are already harnessing the many benefits of AI. There are companies that provide AI solutions, while others use AI within their organization to manage internal business operations or business growth. A few real world companies in the preceding categories will be described by the end of this book.

Artificial Intelligence: History and evolution

Artificial Intelligence (AI) has been studied for decades and is still one of the most elusive subjects in Computer Science.

The year 1943: Warren McCulloch and Walter pits 1943 proposed a model of artificial neurons.

The year 1949: Donald Hebb demonstrated modifying the connection strength between neurons. His rule is now called Hebbian learning.

The year 1950: Alan Turing, an English mathematician, pioneered Machine learning in 1950. Alan Turing proposed a test in his "Computing Machinery and Intelligence" publication. The test, called a Turing test, can check the machine's ability to exhibit intelligent behavior equivalent to human intelligence.

The period between the 1950s and the 1970s revolved around the research on neural networks; the following three decades (1980s to 2010s) were the development of the applications of Machine Learning.

In *Figure 1.1*, a brief timeline of the past six decades of how AI evolved from its inception has been depicted:

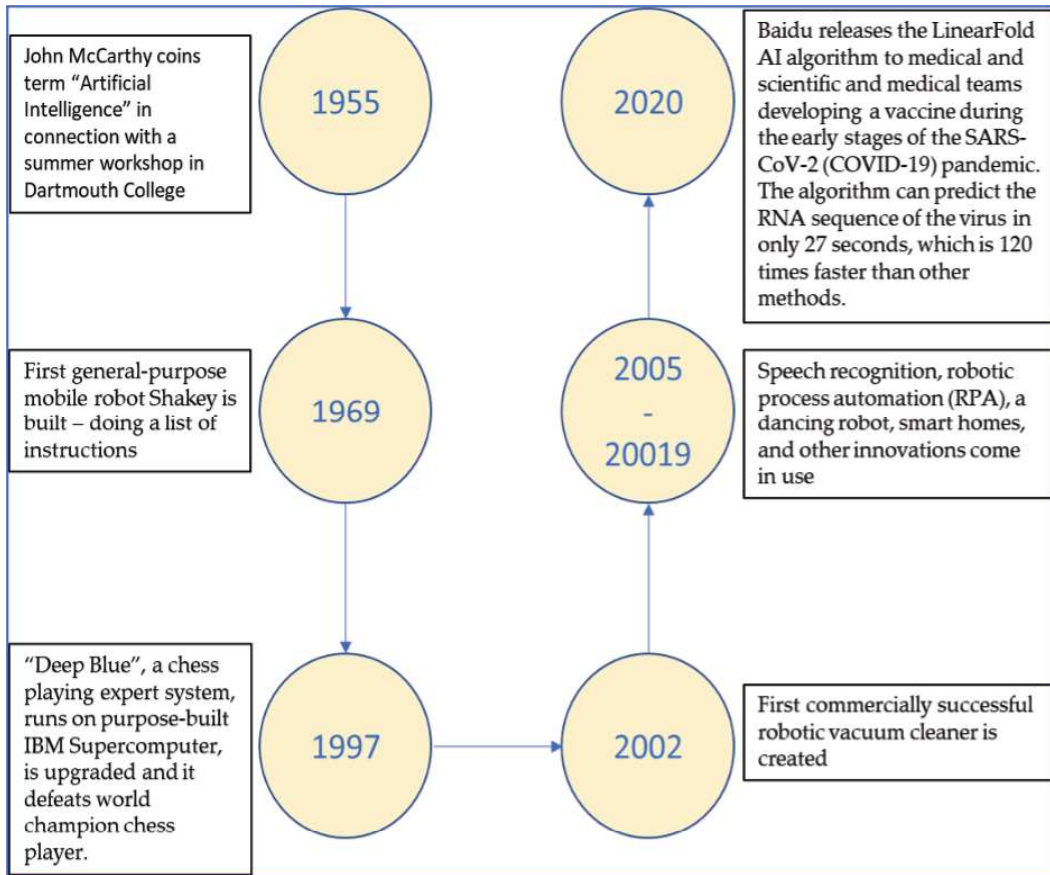


Figure 1.1: The evolution of AI during the last six decades

The father of AI

John McCarthy is widely recognized as the "Father of Artificial Intelligence" due to his astounding contribution and innovations in the field of Computer Science and AI. John McCarthy coined the term "Artificial Intelligence" in his 1955 proposal for the 1956 *Dartmouth Summer Research Project*, the first artificial intelligence conference, which was a seminal event for artificial intelligence as a field. Refer to *Figure 1.2* which depicts the proposal where the term Artificial Intelligence was coined:

A PROPOSAL FOR THE
DARTMOUTH SUMMER RESEARCH PROJECT
ON ARTIFICIAL INTELLIGENCE

J. McCarthy, Dartmouth College
M. L. Minsky, Harvard University
N. Rochester, I. B. M. Corporation
C. E. Shannon, Bell Telephone Laboratories

Figure 1.2: Proposal where the term Artificial Intelligence was coined

In his proposal, he stated that the conference was "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."

In 1956, for the first time, Artificial Intelligence was coined as an academic field. The researchers thought about ways to make machines more cognizant, and they wanted to lay out a framework to better understand human intelligence.

John also paved the way for a few of the world's technological innovations like programming languages, the Internet, the web, and robots, to name just a few

He invented the first programming language for symbolic computation, LISP, and invented and established time-sharing. Human-level Artificial Intelligence and common-sense reasoning were two of his major contributions.

Types of Artificial Intelligence

Artificial Intelligence can be classified into two types:

Based on the capabilities of AI

- **Artificial narrow intelligence**

Artificial narrow intelligence, ANI or Narrow AI, also called "Weak" AI, is goal oriented and is designed to perform singular tasks intelligently and extremely well without any human intervention.

Language translation and image recognition are examples of common uses for narrow AI. Siri is capable of processing human language and submitting a request to a search engine for retrieval. It explains why Siri is unable to

answer abstract and complex queries that require emotional intelligence. It's mere digital assistance to perform basic inquiries and tasks.

Even if Narrow AI appears to be considerably more sophisticated, it operates within a pre-determined, predefined scope. It can attend to a task in real-time, but they pull information from a specific dataset. In fact, what may appear as a complicated AI as a self-driving automobile is labeled Weak AI.

Narrow AI is unable to think. They lack the capability for autonomous reasoning, self-awareness, consciousness, and genuine intelligence.

- **Artificial general intelligence**

Artificial general intelligence (AGI), also called "Strong" AI, is an intelligent system with comprehensive or complete knowledge and cognitive computing capabilities.

In today's world, no true AGI systems exist and remain the stuff of science fiction. Sci-fi movies like "Her," where a human interacts with a machine displaying broad intellectual capabilities to learn, reason, and make own decisions and judgments, while understanding belief systems. True AGI intellectual capacities would exceed human capacities because of its systems' ability to process huge data sets at incredible speeds.

Hence, no real-world systems as examples here.

- **Artificial super intelligence**

Artificial superintelligence, or ASI, will be human intelligence in all aspects. ASI is a futuristic notion and idea about AI capabilities to supersede human intelligence. It will be self-aware and intelligent enough to surpass the cognitive abilities of humans.

Many are concerned about ASI and its impact on humankind. Individuals like Tesla CEO Elon Musk warned about the dangers of ASI-powered robots, even predicting "scary outcomes" like in <the movie> "*The Terminator*."

Based on the functionality of AI

AI can primarily be divided into four different categories based on functionality. Let us have a look at each:

- **Reactive AI**

These machines are the most basic type of AI system and perform best when all parameters are known. These machines do not have any memory or understanding of historical data and will not perform desirably in case of imperfect information input. Refer to *Figure 1.3*: