

AI and ML for Coders

*Applying Core ML algorithms, deep learning
models, and MLOps best practices*

Suddhasatwa Bhaumik



www.bpbonline.com

First Edition 2025

Copyright © BPB Publications, India


ISBN: 978-93-65897-821

All Rights Reserved. No part of this publication may be reproduced, distributed or transmitted in any form or by any means or stored in a database or retrieval system, without the prior written permission of the publisher with the exception to the program listings which may be entered, stored and executed in a computer system, but they can not be reproduced by the means of publication, photocopy, recording, or by any electronic and mechanical means.

LIMITS OF LIABILITY AND DISCLAIMER OF WARRANTY

The information contained in this book is true and correct to the best of author's and publisher's knowledge. The author has made every effort to ensure the accuracy of these publications, but the publisher cannot be held responsible for any loss or damage arising from any information in this book.

All trademarks referred to in the book are acknowledged as properties of their respective owners but BPB Publications cannot guarantee the accuracy of this information.



Dedicated to

To Atasi and Aadriti

my constant sources of inspiration and joy – this one's for you

About the Author

Suddhasatwa Bhaumik is a seasoned software engineer, passionate about AI/ML. With 18+ years of experience, he has designed data and machine learning systems, led teams, architected solutions, and delivered impactful projects. He has always believed in continuous learning and thoroughly enjoys sharing his limited yet practical knowledge to empower software engineers to embrace the world of data science, machine learning, and AI.

In his current role at Google Cloud India, based out of Pune, he feels extremely proud to be working as a Strategic Cloud Engineer helping customers worldwide in implementing AI and machine learning solutions on Google Cloud.

About the Reviewers

- ❖ **Sai Chaitanya** is a data scientist with 7 years of experience (8.5 years overall) in building machine learning models and solving real-world problems. He is currently pursuing a master's in data science at Scaler Academy, focusing on **data structures and algorithms (DSA)**, machine learning, deep learning, and generative AI.

Sai is an aspiring technical author with plans to write a machine learning book that simplifies complex algorithms using a unique visual approach. His aim is to make learning intuitive and long-lasting for readers. This exciting project is set to begin in 2025, reflecting his passion for demystifying data science concepts.

In addition to his professional pursuits, Sai has recently started prioritizing his overall well-being, balancing physical, mental, and spiritual health while expanding his knowledge in personal finance. He is also exploring the stock market and is dedicated to personal growth. Sai envisions a future where his journey inspires others to pursue excellence in the data science field and personal growth.

- ❖ **Pronnoy Goswami** is a seasoned software engineer and thought leader specializing in distributed systems, cloud infrastructure, and large-scale observability. He has led the architecture and design of next-generation telemetry platforms optimized for high availability, scalability, and operational excellence. With deep expertise in DevOps, AI-driven automation, and performance optimization, Pronnoy builds systems that solve complex problems at scale.

His technical skill set spans Kubernetes, public cloud platforms, ELK, Prometheus, and Terraform, and he is passionate about mentoring, technical writing, and advancing the future of cloud-native technologies. Pronnoy frequently shares insights on cloud computing, infrastructure, and applied AI through articles, tutorials, and expert commentary.

As an active technical reviewer, he contributes to books and publications across top-tier platforms in cloud, AI, and distributed architectures. Pronnoy holds a master's degree in computer engineering from Virginia Tech and spends his free time reading, hiking, and exploring emerging technologies.

Acknowledgement

There are a few people I want to thank for the continued and ongoing support they have given me during the writing of this book. First and foremost, I would like to thank my better half for continuously encouraging me to write the book — I could have never completed this book without her support.

I am grateful to the innumerable amount of courses that I have taken till date from platforms like iNeuron, Udemy and Edureka, my Online Master's programs from Great Learning, and the organisations which gave me support throughout the learning process of Artificial Intelligence and Machine learning, notably my experience at _VOIS and Google Cloud India. I thank all my seniors and colleagues (most of whom are my closest friends!) in supporting me in every possible way to learn and grow throughout my career, especially in the last 12+ years!

My gratitude also goes to the team at BPB Publications for being supportive enough to provide me with enough support and time to finish this book, and in gratefully coordinating in getting the technical reviews and chapter reviews completed on time with the highest levels of dedication and quality.

Preface

This book offers a hands-on, code-centric journey into the world of artificial intelligence and machine learning. Readers will immediately dive into practical coding exercises and real-world examples, mastering fundamental concepts. Starting with machine learning principles, the book progresses through computer vision, teaching image processing techniques like feature detection to enable applications to analyze visual data. It then explores **natural language processing (NLP)**, equipping software with language understanding and manipulation skills, covering techniques like tokenization and sentence sequencing for applications like chatbots and sentiment analysis. The book also delves into sequence modeling, covering RNNs and LSTMs for handling sequential data.

Building upon these core concepts, readers will learn to develop TensorFlow models and deploy AI solutions across various platforms. The book culminates in strategies for confidently serving models online and in the cloud, ensuring scalable and robust AI applications. By the end of this book, readers will emerge as confident AI developers, ready to contribute valuable expertise to any organization and innovate in the field of AI and ML.

This book is divided into 13 chapters, details of which are listed below.

Chapter 1: Introducing Artificial Intelligence and Machine Learning - In this chapter, we introduce the ever-growing world of AI and ML using practical examples and applications where such systems are in use and where such systems help millions of people worldwide to successfully execute and use products in their day-to-day. This chapter also sets the clear boundaries of what an AI can do as of the current day and what it should be expected to do.

Chapter 2: Machine Learning Fundamentals - Before the readers of the book deep dive into variety of applications of AI and ML, and how to design and develop an entire system covering it, it is essential to know the basics of machine learning, its types, the various algorithms, where they fit into the landscape, and how they are beneficial in solutioning of particular use cases across domains. This chapter provides the landing zone to realize the fundamentals of ML and its use cases. Additionally, it also covers the details around hyperparameter tuning and its importance in production AI/ML systems.

Chapter 3: TensorFlow Essentials - Now that the readers of this book have some idea on what machine learning is and what kind of use cases are catered by their use in the

real-world, this chapter sets the stage with the basics of TensorFlow, the most popular and production grade ML engineering library, with Python language as the mode of usage. The readers are assumed to have the very basic hands-on knowledge of the Python programming language to understand all the nuances presented in this chapter.

Chapter 4: Engineering for Machine Learning - It is essential at this stage now for readers of this book to understand and grasp the fundamental ways and tools of programming for machine learning, including how to modularize their software and even design to containerize them later for easy shipments. In this section, we cover the most essential aspects of software engineering required to make the best use of ML driven applications. Additionally, we also introduce the usage and benefits, by some case studies, of RDBMS systems (MySQL and/or Oracle) and EDW (BigQuery, on Google Cloud Platform), for readers to be made aware of how real-world industrial data is stored, managed and processed in both batch and real-time modes.

Chapter 5: Machine Learning Algorithms - Since the basics are now ready with the readers of this book, in this chapter we deep dive into the working of all of the traditional ML algorithms including the most common regression, classification and clustering techniques. Readers are also presented with approaches to use these algorithms given the right use cases, which are suitable for each of these algorithms. We start to implement these algorithms from the next chapter onwards.

Chapter 6: Implementing First ML Models - In this chapter, readers get to learn and focus on the industry recommended and time tested methods of designing and developing all the common ML algorithms using scikit-learn (a starter ML Library, also used in production) and TensorFlow (which is a production grade library for ML). We focus on the nuances of traditional ML algorithms, while the advanced topics like neural networks will be covered starting in the next chapter. From time to time, this chapter refers back to the knowledge gained in Chapters 2,3,4 and 5.

Chapter 7: Computer Vision - We begin this chapter with an introduction on neural networks and the variety of ever increasing use cases where they are finding their place, because of the humongous speed and size at which we (humans) have started to generate and accumulate data. Then, we start to deep dive into one of the most interesting areas in deep learning, a.k.a computer vision, where we understand the basics, the architectures, the use cases, and industries where the methods apply, and the best ways of articulating a successful computer vision project. Readers get to implement the aforementioned computer vision methods using TensorFlow and Keras, another high-level abstraction over TensorFlow library. Readers are also introduced to the challenges in implementing

computer vision algorithms in any given project or IT landscape, including relevant portions of data privacy and AI ethics.

Chapter 8: Natural Language Processing - Readers are now introduced to NLP and its nuances in this chapter, which is by far the most common, the most useful, as well as the most challenging category of applications where AI and ML find their use. Readers are introduced to the basics of how to analyze and process text, followed by a variety of methods and algorithms to engineer and recognize patterns from textual data. Finally, readers are introduced to useful case studies covering the most commonly built applications around NLP.

Chapter 9: Sequence Modelling and Transformers - In this chapter, we introduce the importance of sequence models and use cases like forecasting and language translation where there is an increasing use of sequence models. Readers are introduced to forecasting, one of the most common methods used across industries for planning and demand management, followed by modelling language as a sequence of events. We introduce readers to generative adversarial networks which are broadly used in language and image use cases. Finally, readers are introduced to what is today the heart of generative AI systems, a.k.a Transformers, covering what the attention mechanism is. All of these aforementioned topics are implemented hands-on using TensorFlow and Keras for practical realization.

Chapter 10: MLOps and Deployment - In this chapter, we introduce the basic concepts and importance of MLOps in the real-world. MLOps is one of the most important backbone of AI and machine learning applications in Production, which in fact, contributes to more than 90% of the overall components in a system, apart from what the readers have learnt in the earlier chapters, which merely contributes for 5-10% in the overall scheme of things. Once MLOps and its importance are realized, readers are introduced to deployment patterns on-premise and cloud in this chapter, where we use in-house or cloud based IT infrastructure of any given organization to deploy and operate on AI and ML systems. Readers are also briefly introduced to data and model security, along with applicable design guidelines. In this chapter as well, readers are introduced to hands-on MLOps and deployments using **TensorFlow Extended (TFX)** and Kubeflow. For cloud implementations and references, readers are introduced to Google Cloud Platform, and the state-of-the-art ML and operations suite of offerings in GCP, a.k.a Vertex AI. We also touch base about model monitoring and CI/CD pipelines for ML models in production.

Chapter 11: Model Serving and Scalability - In this chapter, we introduce the nuances of how to serve models in the real-world, both using in-house IT infrastructure and Google Cloud Platform. This is where readers are also introduced to concepts of model scalability and how to design the right infrastructure for a given application which is powered by AI/

ML. We introduce TensorFlow Serving as one prudent option for effective service of ML models, along with Vertex AI on Google Cloud Platform, and how to design scalable ML infrastructure. Although this chapter is more on system design, readers are also introduced to hands-on practical examples using Python and TensorFlow on aforementioned topics of model serving and scalability.

Chapter 12: Model Deployment for Mobile - In the chapter of this book, we provide the readers with detailed knowledge on how to deploy ML models and make them work on mobile (or related remote) devices. This is a crucial topic to be explored as of date, because of the inescapable rise and role of mobile devices like smartphones, tablets, and smart home devices in our lives! In this chapter, hence, we introduce TensorFlow Lite, which is one of the common methods, along with design and development samples and/or guidelines, to develop and package ML model driven applications for mobile devices. Further, we look at the integration of such applications (driven by machine learning) on Android and iOS platforms. Finally, we conclude the chapter with a couple of interesting case studies on two of the most common current world applications.

Chapter 13: Summary, Future, and Resources - In the last chapter of this book, we provide the readers with a detailed summary of what we have learnt and programmed in this book, touching on the most important design and engineering principles and tools, along with best practices. Readers are introduced to the current state of AI and ML implementations, as a revision/summary, as well as a realization, along with a variety of interesting implications and uses, like generative AI, in the current world. Finally, readers are encouraged with the right amount of additional resources for professional growth in the ever-growing field of AI and ML, including additional focus on AI ethics and data privacy in the age of ChatGPT and similar.

Code Bundle and Coloured Images

Please follow the link to download the
Code Bundle and the *Coloured Images* of the book:

<https://rebrand.ly/0184de>

The code bundle for the book is also hosted on GitHub at

<https://github.com/bpbpublications/AI-and-ML-for-Coders>.

In case there's an update to the code, it will be updated on the existing GitHub repository.

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. Introducing Artificial Intelligence and Machine Learning.....	1
Introduction.....	1
Structure.....	2
Objectives	2
Introduction to AI/ML	2
Applications and use cases	4
Responsible AI principles.....	9
Expectations from AI	10
Career opportunities	11
Conclusion.....	11
2. Machine Learning Fundamentals	13
Introduction.....	13
Structure.....	14
Objectives	14
Introduction to machine learning	14
Types of machine learning	15
Evaluating machine learning models.....	17
Use cases of machine learning.....	19
Common algorithms	20
<i>Linear regression</i>	21
<i>Logistic regression</i>	23
<i>Naïve Bayes</i>	26
<i>K-nearest neighbors</i>	28
<i>Decision trees</i>	31
<i>Random forests</i>	33
Ways to make ML models better.....	37
<i>Data matters</i>	37
<i>Model selection and optimization</i>	38
<i>Interpretation and debugging</i>	38

<i>Production readiness</i>	39
Conclusion.....	39
Exercises.....	39
3. TensorFlow Essentials	41
Introduction.....	41
Structure.....	41
Objectives	42
TensorFlow	42
<i>Building blocks</i>	42
<i>Quick hands on</i>	44
<i>Creating tensors in Google Colab</i>	47
<i>Quick summary</i>	50
Scikit-learn.....	51
ML models with TensorFlow and scikit-learn.....	53
Merits and demerits of TensorFlow	56
TensorFlow vs scikit-learn vs PyTorch	58
Keras.....	59
Conclusion.....	62
Exercises.....	63
References.....	63
4. Engineering for Machine Learning	65
Introduction.....	65
Structure.....	66
Objectives	66
Python for data analysis.....	66
Jupyter Notebooks	81
Visual Studio Code.....	84
Version control with Git	89
Containers and Docker	94
Databases.....	96
Conclusion.....	99
Exercises.....	99

5. Machine Learning Algorithms	101
Introduction.....	101
Structure.....	102
Objectives	102
Regression models.....	102
<i>Linear regression</i>	102
<i>Support vector regression</i>	107
<i>Random forest regression</i>	110
<i>Gradient boosting regression</i>	112
Classification models	114
<i>K-nearest neighbors classifier</i>	114
<i>Naïve Bayes</i>	117
<i>Decision trees</i>	119
<i>Random forest algorithm</i>	123
<i>XGBoost</i>	126
Clustering models	128
<i>K-means clustering</i>	128
<i>DBSCAN</i>	131
Neural networks.....	134
Conclusion.....	135
Exercises.....	136
6. Implementing First ML Models	137
Introduction.....	137
Structure.....	137
Objectives	138
Estimating social media interactions.....	138
Prediction of heart disease type	155
Clustering water treatment plants	168
Conclusion.....	179
Exercises.....	179

7. Computer Vision	181
Introduction.....	181
Structure.....	181
Objectives	182
Artificial neural networks	182
<i>Activation functions</i>	183
<i>Optimizers</i>	184
<i>Backpropagation</i>	186
<i>Applications</i>	187
<i>House price prediction</i>	188
<i>Spam emails</i>	191
<i>Playground</i>	194
Convolutional neural networks	194
<i>Filters</i>	196
<i>Pooling</i>	197
<i>Fully connected layers</i>	199
<i>MNIST handwritten digits</i>	200
Real-world object classification	203
Transfer learning.....	215
Practical uses and limitations	219
Conclusion.....	221
Exercises.....	222
8. Natural Language Processing	223
Introduction.....	223
Structure.....	223
Objectives	224
Importance of natural language processing.....	224
Natural language data analysis.....	226
Introduction to sequential models	237
<i>Recurrent neural networks</i>	237
<i>Vanishing gradients</i>	241
<i>Long short-term memory</i>	242
Sentiment analysis.....	244

Case studies.....	249
<i>Chatbots</i>	250
<i>Machine translation</i>	251
Conclusion.....	252
Exercises.....	253
9. Sequence Modelling and Transformers.....	255
Introduction.....	255
Structure.....	255
Objectives	256
Introduction to sequential models	256
Time series forecasting.....	256
<i>Training and tuning sequential models</i>	257
<i>Dataset preparation</i>	258
<i>Loss functions and optimizers</i>	260
<i>Hyperparameter optimization</i>	261
Forecasting	263
Language modelling	269
Generative adversarial networks.....	272
Transformers and attention mechanism	274
<i>Generative artificial intelligence</i>	278
Real-world case studies.....	283
<i>Sequential models</i>	284
<i>Transformers</i>	284
<i>GANs</i>	284
<i>Generative artificial intelligence offerings</i>	285
Conclusion.....	285
Exercises.....	286
10. MLOps and Deployment.....	287
Introduction.....	287
Structure.....	287
Objectives	288
Introduction to MLOps.....	288

Level 0: Manual process	291
Level 1: Pipeline automation	293
Level 2: CI/CD.....	295
AI and ML deployments	298
Model monitoring	300
CI/ CD for ML	304
TensorFlow Extended	305
Kubeflow	315
Containers, Docker, and Kubernetes.....	322
Containers	322
Docker.....	327
Kubernetes.....	330
Conclusion.....	333
Exercises.....	333
11. Model Serving and Scalability	335
Introduction.....	335
Structure.....	336
Objectives	336
Introduction to model serving.....	336
Model serving options	337
Deployment strategies.....	338
TensorFlow Serving.....	339
Vertex AI in Google Cloud.....	340
Deployment strategies.....	342
ML in production using containers	342
ML in production using GCP.....	345
Model service strategies	348
Real-world challenges.....	350
Scalability design	350
Performance considerations	352
Best practices	355
Conclusion.....	357

Exercises.....	357
References.....	357
12. Model Deployment for Mobile	359
Introduction.....	359
Structure.....	360
Objectives	360
Machine learning on mobile	360
Design and implementation	362
<i>TensorFlow Lite</i>	362
<i>Models on Android and iOS</i>	367
<i>Real-world applications</i>	368
Conclusion.....	370
Exercises.....	371
References.....	371
13. Summary, Future, and Resources	373
Introduction.....	373
Structure.....	373
Objectives	374
Overview and insights.....	374
<i>Machine learning</i>	374
<i>Deep learning</i>	375
<i>Computer vision</i>	376
<i>Sequential models</i>	377
<i>Machine learning operations</i>	378
<i>Model deployment and scaling</i>	379
<i>Model deployment on mobile</i>	380
Future of AI and ML	381
Resources	383
Conclusion.....	384
References.....	384
Index	385-391

CHAPTER 1

Introducing Artificial Intelligence and Machine Learning

Introduction

In a shocking outcome of the infamous Chinese board game *Go*, in early March 2016, the *Go* champion *Lee Se-dol* and the observers of a historic game were in awe when *Google DeepMind's* AI, *AlphaGo*, won by 4-1. At least quite a few from the audience had thought that *Mr. Lee's* impressive win in the fourth round of the game not only redefined the final outcome, but also that he was learning how the computer worked and strategized its steps during the game. However, the end result was nothing less than a significant leap forward for **artificial intelligence (AI)** systems of the kind being designed by *Google DeepMind*, in line with many other crucial players around the world. Simultaneously, and in the same game, *Mr. Lee* could be seen to have realized the non-human and eccentric playing style of *AlphaGo*, hence identifying its weak points and proving that the machine was not infallible.

There are many similar stories, competitions, and applications that we all see around us today, which helps us realize the importance of AI and **machine learning (ML)** driven applications and systems, making our lives easy and nothing more. In this chapter, we formally define AI and ML, and look for practical examples all around us. We will finally clear some concepts and boundaries on what an AI or ML system can do as of the current day, along with what it should be expected to do.

Structure

The chapter covers the following topics:

- Introduction to AI/ML
- Applications and use cases
- Responsible AI principles
- Expectations from AI
- Career opportunities

Objectives

By the end of this chapter, you will be able to understand what AI and ML are and how they are used in our day-to-day lives across multiple domains and applications. We will also look at its transformative power (even more with the recent rise in generative AI hype), the career opportunities in this field, and finally, a set of realistic expectations.

Introduction to AI/ML

AI has become an indispensable tool in the professional landscape. From automating repetitive tasks to enabling complex decisions, AI is redefining our working environments in profound ways.

To formally define it:

AI refers to the simulation of human intelligence processes by machines, specifically the ability to acquire, interpret, and apply knowledge and data. It entails developing algorithms and models that can learn and make intelligent decisions without explicit instructions.

As a novice, we can categorize an AI system into one of the following two broad classes:

- **Strong AI (General AI):** This is where AI has a comparable performance to humans. In simple words, this type of AI can learn, reason, and hence solve problems across different areas of work or domains. General or strong AI remains to be a theoretical concept as of today, and reaching a stage of artificial general intelligence is still a work in progress. Although it sounds intriguing and magnetic, and as we will discuss later in this book, it is imperative to perform our research, development, and applications in line with a solid and broad set of responsible AI principles.
- **Weak AI (Narrow AI):** This refers to AI systems designed to perform a specific task or set of tasks with high efficiency. Examples include speech recognition, image classification, and game-playing algorithms. These systems form the backbone of many IT systems and would be our primary point of focus during the subsequent chapters of this book. Even if we call this category of systems as weak or narrow

AI, one must realize that the responsible AI principles are equally applicable in these systems, like the former type.

Irrespective of their type, an AI system is designed to process enormous amounts of data to learn and identify patterns, make predictions (and estimations in many cases, as we will learn further in this book) based on the information they have been trained on, and assist humans in unseen cases. This is an important point to note at this early stage of understanding AI, where one must oblige and acknowledge the fact that an AI system is merely an assistant, and in many given contexts and industries, only fits as one or more helpers to achieve a broader goal.

At this stage, it would be logically important to understand the key components of AI so that we can also understand the areas we will be learning in this book, along with hands-on examples. Let us consider the following points:

- **ML:**
 - ML algorithms enable computers to learn from data without being explicitly programmed.
 - They can identify patterns, make predictions and estimations, and improve their performance over time as they are exposed to more data.
- **Deep learning:**
 - Deep learning is a subset of ML that utilizes artificial neural networks, which are loosely inspired by the structure of the human brain.
 - These networks can learn complex relationships within data and make accurate predictions even when presented with incomplete or noisy information.
 - In many practical applications and the onset of exemplary research seen in the last few decades, deep learning and ML have now become interchangeable terms.
- **Natural language processing:**
 - **Natural language processing (NLP)** allows computers to understand and generate human language.
 - NLP algorithms can perform tasks such as sentiment analysis, machine translation, and speech recognition.
 - Both ML and deep learning methods and algorithms are used day in and day out to approach and solve NLP problems. However, since natural language is the most approached method of human communication, it takes a separate stand for itself.

- **Computer vision:**
 - Computer vision algorithms enable computers to interpret and understand visual information.
 - They can identify objects, detect faces, and track movement in images and videos, to quote a few examples in the real-world.
 - Again, similar to NLP, computer vision is a domain by itself, and to approach relevant use cases of computer vision, engineers like us make use of ML and deep learning methods.
- **Robotics:**
 - Robotics combines AI with physical hardware to create autonomous machines that can perform tasks in the real-world.
 - Robots can navigate their environment, interact with objects, and make decisions based on sensor data.
 - One can argue that robotics does not naturally fall in the AI realm by itself, but in our context, these are intelligent machines that are learning to perform human-like tasks using methods like reinforcement learning.

Let us now get over the basics and try to realize the vast number of applications in the real-world that make use of AI and ML.

Applications and use cases

Firstly, let us understand some of the applications of AI and ML from a generalist's perspective. Let us look at a broader set of applications and uses of AI and ML across domains:

- **Healthcare:**
 - **Disease diagnosis:** AI algorithms analyze medical images, such as X-rays, CT scans, and MRIs, to detect diseases like cancer, pneumonia, and bone fractures with high accuracy.
 - **Drug discovery:** ML models help researchers identify potential drug candidates and optimize their properties, accelerating the drug development process.
 - **Personalized medicine:** AI systems analyze individual patient data to tailor treatment plans, predict disease risks, and provide personalized healthcare recommendations.
- **Finance:**
 - **Fraud detection:** AI algorithms analyze transaction patterns to identify suspicious activities and prevent fraudulent transactions in real time.

- **Credit scoring:** ML models assess an individual's creditworthiness based on various factors, enabling lenders to make informed decisions.
- **Stock market prediction:** AI systems analyze historical data, market trends, and news sentiments to predict stock price movements and assist investors in making informed decisions.
- **Retail:**
 - **Product recommendations:** AI algorithms analyze customer behavior, purchase history, and preferences to recommend personalized products and improve the shopping experience.
 - **Inventory management:** ML models optimize inventory levels by predicting demand, preventing stockouts, and minimizing waste.
 - **Customer service:** AI-powered chatbots and virtual assistants provide 24/7 customer support, answering queries, resolving issues, and improving customer satisfaction.
- **Transportation:**
 - **Self-driving cars:** AI algorithms enable autonomous vehicles to navigate roads, detect obstacles, and make real-time decisions, enhancing safety and reducing accidents.
 - **Traffic management:** ML models analyze traffic patterns and predict congestion, helping authorities optimize traffic flow and reduce travel times.
 - **Fleet management:** AI systems track vehicle locations, monitor fuel consumption, and schedule maintenance, improving fleet efficiency and reducing costs.
- **Manufacturing:**
 - **Quality control:** AI algorithms inspect products for defects, ensuring quality and consistency.
 - **Predictive maintenance:** ML models analyze sensor data to predict equipment failures, enabling proactive maintenance and minimizing downtime.
 - **Process optimization:** AI systems analyze production data to identify inefficiencies, optimize processes, and increase productivity.
- **Agriculture:**
 - **Crop yield prediction:** AI algorithms analyze weather data, soil conditions, and historical yields to predict crop yields, helping farmers make informed decisions about planting and harvesting.