

40 Algorithms Every Data Scientist Should Know

Navigating through essential AI and ML algorithms

Jürgen Weichenberger

Huw Kwon



www.bpbonline.com

First Edition 2025

Copyright © BPB Publications, India

ISBN: 978-93-55519-832

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 to correct and the best of author's and publisher's knowledge. The author has made every effort to ensure the accuracy of these publications, but 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.

To View Complete
BPB Publications Catalogue
Scan the QR Code:



Dedicated to

My beloved wife: Li and

My Daughter Sophia

– Jürgen

*Ross McDonald, Matthew Jones, Edward Huang, Neil Graham,
Raj Dash, Mark Somers, Philip Treleaven, Mark Rowland, Matt
O’Kane, Andy Huang, Huw Tindall, Luigi Masoero, Michael Glinski,
Gerard Crispie, Rob Handcock, George Marcotte*

and the three professors named Alan at the University

– Huw

About the Authors

- **Jürgen Weichenberger** has been working in Artificial Intelligence and Machine Learning Development for more than 20 years, playing central roles in numerous projects as a technical leader and chief data scientist, delivering projects using numerous different algorithms and even developing entire new algorithms for big companies, including well succeeded projects in Europe, Asia, and North America. Currently, he is a Head of AI Strategy and Innovation at Schneider Electric and a Senior Advisor at P.E.T. Consulting. He is also an accomplished postgraduate completing a degree in Executive General Management and holds three master's degrees focused on Applied Computer Science, Bioinformatics, and Cybernetics. He has many certifications for various types of Artificial Intelligence Technologies. Furthermore, the author participates as a speaker in international AI Conferences and writes technical articles on AI Technologies, Algorithms and related topics. Based on all his contributions to AI communities worldwide, he was awarded over 20 AI patents.
- **Huw Kwon** is a well-known authority in AI and ML within the management consulting world, bringing over two decades of academic rigor and truly extensive practical experience to the field as a first-hand practitioner. With a strong foundation in management science and advanced computational algorithms, his work is unified by a singular focus: leveraging AI to create tangible, sustained, and lasting real-world impact.

Huw's career spans senior roles in banking and consulting at some of the world's most prestigious firms, including Ernst & Young, Accenture, and McKinsey. In these roles, he has been pivotal in transforming businesses by developing AI capabilities and leading complex, large-scale transformation programs across industries such as financial services, automotive, and high-performance manufacturing.

As a member of several advisory boards and a recognized thought leader, Huw frequently contributes to industry forums and publications, shaping the discourse on the future of AI. Beyond consulting, he is a driving force in AI innovation, working closely with academic institutions and startups to push the boundaries of what AI and ML can achieve in practical applications.

About the Reviewer

Dr. Zachary Elewitz is a data scientist with over a decade of experience, currently serving as the Head of AI at Fortune Brands Innovations. He holds two AI-related patents, sits on Texas A&M Commerce's Venture College Board, and serves in several AI-related groups, including the National Institute of Standards and Technology Generative AI Public Working Group. In his spare time, he is pursuing a Masters in Viking studies and enjoys snowboarding, indoor bouldering, and playing the guitar.

Acknowledgements

- **Jürgen:** 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 wife Li and my daughter Sophia.

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 colleagues and co-workers during many years working in the AI 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 this book and for their support in making it a reality. Your encouragement has been invaluable.

- **Huw:** Your guidance and support have been invaluable on my journey to solve real-world problems with Data and AI. Each of you has been statistically significant in this pursuit of mine, and for that, I am deeply grateful.

Preface

Building Artificial Intelligence and Machine Learning Solutions is a complex task that requires a comprehensive understanding of the latest technologies and algorithms available to us. Artificial Intelligence has become an increasingly powerful tool over the last couple of years and as such the amount of algorithm available to us have explode.

This book is designed to provide a comprehensive guide through the world of Artificial Intelligence Algorithms and be a practical and hands-on support to every new data scientist as well as experienced data scientists. It covers a wide range of topics, including the basic definition of Artificial Intelligence and Machine Learning, basic data concepts, and basic and advanced algorithms for supervised, unsupervised, semi-supervised, and reinforcement learning algorithms.

Throughout the book, you will learn about the key features of every algorithm, their mathematical foundation, and how to use them to build Artificial Intelligence solutions that are efficient, reliable, and easy to maintain. You will also learn about best practices and design patterns for Artificial Intelligence solutions and will be provided with numerous practical examples to help you understand the algorithms.

This book is intended for new data scientists who want to learn which algorithms are available and how to build Artificial Intelligence solutions with them. It is also helpful for experienced data scientists who want to expand their knowledge of these algorithms and improve their skills in building robust and reliable Artificial Intelligence solutions.

With this book, you will gain the knowledge and skills to become a proficient data scientist and be able to build Artificial Intelligence solutions we hope you will find this book informative and helpful.

Chapter 1: Fundamentals – Introduction into the world of AI and ML algorithms covering a little historical extract to the origins of AI and how it has developed to what we know today. Every modern AI/ML algorithm follows a basic structure which assures that the training process will converge and the inference will deliver a reasonable result. Furthermore, it will cover the process of retraining an algorithm to refit its parameters and hyperparameters.

Chapter 2: Typical Data Structures – An AI/ML algorithm can neither be trained nor run an inference without being fed with the right data structure. The process of preparing the data is known as feature engineering and requires the right school of thought.

Chapter 3: 40 AI/ML Algorithms Overview – Introduction to 40 AI/ML algorithms, including the classification and structure for the 40 algorithms, in the following chapters.

Chapter 4: Basic Supervised Learning Algorithms – Chapters 4-11 will cover the 40 algorithms which comprise the core of the book.

This chapter covers essential supervised learning algorithms in machine learning. It starts with **Linear Regression**, a widely used method for modeling the linear relationship between input features and continuous target variables. It then introduces **Logistic Regression**, commonly used for binary classification by estimating the probability of class membership. The chapter also explores **Decision Trees**, which partition data based on feature values and are applicable to both regression and classification tasks, often forming the basis of **Random Forests**—an ensemble model combining multiple decision trees for more accurate predictions. Lastly, **Naive Bayes** is discussed, a probabilistic algorithm based on Bayes’ theorem, known for its efficiency in tasks like text classification and spam filtering.

Chapter 5: Advanced Supervised Learning Algorithms – This chapter explores advanced supervised learning algorithms, starting with **k-Nearest Neighbors (k-NN)**, which predicts based on the majority vote of nearest neighbors and is effective for non-linear boundaries. It covers **Support Vector Machines (SVMs)**, which find an optimal hyperplane to separate classes or predict values by maximizing class margins. The chapter also examines **Gradient Boosting Machines (GBM)**, an ensemble method that combines weak models to build a strong predictor by focusing on errors from previous models. **XGBoost**, an optimized gradient boosting method known for its performance and scalability, is also discussed. Finally, it introduces **Neural Networks**, complex models inspired by the brain that learn intricate patterns through layers of artificial neurons, driving advancements in areas like image recognition and natural language processing.

Chapter 6: Basic Unsupervised Learning Algorithms – This chapter explores basic unsupervised learning algorithms, beginning with **K-means Clustering**, which groups data into clusters based on proximity. It also covers **Hierarchical Clustering**, which builds a tree-like structure of clusters by merging or splitting based on similarity. **Principal Component Analysis (PCA)** is introduced as a technique for reducing dimensionality while preserving key features by identifying principal components. **t-Distributed Stochastic Neighbor Embedding (t-SNE)** is discussed for visualizing high-dimensional data in lower dimensions, emphasizing local structures. Finally, **Association Rule Mining** with the **A priori Algorithm** is examined for discovering relationships between items in a dataset by identifying frequent item sets and generating association rules.

Chapter 7: Advanced Unsupervised Learning Algorithms – This chapter covers advanced unsupervised learning algorithms, including **Density-Based Spatial Clustering of Applications with Noise (DBSCAN)**, which identifies clusters of varying shapes and handles noise. It discusses **Gaussian Mixture Models (GMM)**, which model data as a mixture of Gaussian distributions to uncover subpopulations. **Autoencoders** are introduced for unsupervised representation learning and dimensionality reduction. The chapter also explores **Anomaly Detection**, which identifies rare or unusual instances using various techniques. Finally, **Latent Dirichlet Allocation (LDA)** is covered for topic modeling, discovering hidden topics in documents, and assigning topic distributions.

Chapter 8: Basic Reinforcement Learning Algorithms – This chapter covers basic reinforcement learning algorithms, starting with **Q-Learning**, which estimates optimal action values for state-action pairs through iterative updates. It then introduces **Deep Q-Networks (DQN)**, which uses deep neural networks to handle high-dimensional state spaces. **Policy Gradient Methods** optimize policy parameters directly to maximize rewards with algorithms like REINFORCE and **Proximal Policy Optimization (PPO)**. The chapter also explores **Advantage Actor-Critic (A2C)**, which combines policy gradient and value-based methods for stable learning. Finally, **Trust Region Policy Optimization (TRPO)** improves policies iteratively while staying close to the original policy using trust regions.

Chapter 9: Advanced Reinforcement Learning Algorithms – This chapter covers advanced reinforcement learning algorithms, starting with **Asynchronous Advantage Actor-Critic (A3C)**, which uses parallel agents to improve sample efficiency and learning speed. **Proximal Policy Optimization (PPO)** is discussed next, using a trust region approach for stable policy updates. **Deep Deterministic Policy Gradient (DDPG)** combines deep Q-networks with actor-critic methods for continuous action spaces, while **Twin Delayed Deep Deterministic Policy Gradient (TD3)** enhances DDPG by addressing overestimation with multiple critics and delayed updates. Finally, **Soft Actor-Critic (SAC)** is introduced, optimizing both reward and exploration using the maximum entropy framework.

Chapter 10: Basic Semi-Supervised Learning Algorithms – This chapter covers basic semi-supervised learning algorithms, including **Self-training**, where a model iteratively adds high-confidence predictions from unlabeled data to its training set. **Co-training** involves multiple models training on different data views and refining each other's predictions. **Multi-view Learning** enhances learning by using various data representations to ensure prediction agreement. **Expectation-Maximization (EM)** estimates parameters and missing labels in probabilistic models. Finally, **Graph-based Methods** propagate labels from labeled to unlabeled data using the data's structure, with techniques like Label Propagation and Manifold Regularization.

Chapter 11: Advanced Semi-Supervised Learning Algorithms – This chapter covers advanced semi-supervised learning algorithms, including **Transductive Support Vector Machines (TSVM)**, which uses both labeled and unlabeled data to learn decision boundaries. **Co-regularization** combines different regularization strategies to maintain consistency and reduce sensitivity to noisy labels. **Deep Generative Models**, such as VAEs and GANs, learn from both labeled and unlabeled data to generate new samples and representations. **Virtual Adversarial Training (VAT)** adds robustness to models by addressing adversarial perturbations from both data types. **Tri-training** trains three models on different labeled feature subsets, using their consistent predictions on unlabeled data to expand the labeled set.

Chapter 12: Natural Language Processing – Natural Language Processing (NLP) is a subfield of computer science, artificial intelligence, and computational linguistics that focuses on the interaction between computers and humans in natural language. NLP enables computers to process, understand, and generate natural language, which is the language used by humans to communicate with each other.

NLP involves developing algorithms and computational models that can analyze, interpret, and generate human language, including tasks such as language translation, sentiment analysis, text summarization, speech recognition, and language generation.

The goal of NLP is to enable computers to understand and respond to natural language in the same way that humans do, allowing for more natural and intuitive communication between humans and machines. NLP has applications in a wide range of fields, including machine translation, information retrieval, customer service, healthcare, and education.

Chapter 13: Computer Vision – Computer vision is a field of artificial intelligence and computer science that focuses on enabling computers to interpret and understand the visual world around them, similar to how humans perceive and process visual information.

Computer vision involves developing algorithms and computational models that can analyze and interpret images and videos. This includes tasks such as object detection, image classification, facial recognition, scene understanding, and image segmentation.

The field of computer vision has made significant progress in recent years, with the development of deep learning algorithms and convolutional neural networks, which have led to breakthroughs in tasks such as image recognition and object detection.

Computer vision has many applications in various industries, including healthcare, transportation, retail, and entertainment. For example, it can be used for medical image analysis, self-driving cars, visual search in e-commerce, and augmented reality in gaming and entertainment.

Chapter 14: Large-Scale Algorithms – Large-scale algorithms are computational methods designed to handle massive amounts of data, such as those generated by modern digital technologies. These algorithms typically involve processing large datasets in parallel or distributed systems and require specialized hardware and software architectures to achieve high performance.

The development of large-scale algorithms has become increasingly important in recent years due to the exponential growth of data generated by various sources such as social media, scientific simulations, and internet of things (IoT) devices. Large-scale algorithms are needed to handle these large and complex datasets efficiently and effectively.

Examples of large-scale algorithms include distributed machine learning algorithms such as Spark MLlib and TensorFlow, graph processing algorithms such as Apache Giraph and GraphX, and parallel processing algorithms such as Hadoop MapReduce and Apache Flink. These algorithms are widely used in various industries, such as finance, healthcare, and social media to analyze large datasets and make data-driven decisions.

Chapter 15: Outlook into the Future: Quantum Machine Learning – Quantum machine learning is an emerging field that combines quantum computing and machine learning. Quantum computing uses the principles of quantum mechanics to perform certain computations much faster than classical computing. Machine learning, on the other hand, involves developing algorithms that can learn patterns and insights from data.

Quantum machine learning aims to leverage the power of quantum computing to develop more efficient algorithms for machine learning tasks, such as classification, clustering, and regression. These algorithms could potentially provide significant speedup and better accuracy compared to classical machine learning algorithms.

There are various approaches to quantum machine learning, including quantum-inspired classical algorithms, quantum-enhanced classical algorithms, and fully quantum algorithms. Some of the challenges in quantum machine learning include designing quantum algorithms that can take advantage of the unique properties of quantum computing, such as superposition and entanglement, and developing hardware and software infrastructure for quantum computing that can support large-scale machine learning tasks.

Quantum machine learning has the potential to revolutionize many industries, including finance, healthcare, and cybersecurity, by providing faster and more accurate predictions and insights from large datasets.

Code Bundle and Coloured Images

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

<https://rebrand.ly/8qenuj3>

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

<https://github.com/bpbpublications/40-Algorithms-Every-Data-Scientist-Should-Know>.

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. Fundamentals	1
Introduction	1
Structure	1
Objectives	2
Fundamentals of AI and ML.....	2
Defining AI and ML.....	3
<i>Artificial Intelligence</i>	3
<i>Machine learning</i>	4
History of AI and ML	4
<i>Classic examples of AI and ML</i>	6
AI and ML algorithms	9
<i>Examples of AI and ML algorithms</i>	10
<i>Structure of a typical AI and ML algorithm</i>	15
Conclusion	16
Points to remember	16
2. Typical Data Structures	19
Introduction	19
Structure	19
Objectives	19
Introducing data structures	20
<i>Examples of typical data structures</i>	21
<i>Arrays</i>	21
<i>Matrices</i>	22
<i>Tensors</i>	24
<i>Linked lists</i>	26
<i>Graphs</i>	28
<i>Hash tables</i>	30

Queues	32
Trees.....	34
Knowledge graph.....	36
Conclusion	39
Points to remember	40
Exercises	40
<i>Exercise 1: Data preparation for sentiment analysis</i>	40
<i>Exercise 2: Data preparation for image classification</i>	41
3. 40 AI/ML Algorithms Overview	43
Introduction	43
Structure	43
Objectives	44
Classification of AI and ML algorithms	44
<i>Supervised learning algorithms</i>	44
<i>Unsupervised learning algorithms</i>	46
<i>Reinforcement learning algorithms</i>	47
<i>Semi-supervised learning algorithms</i>	48
Overview of the 40 AI/ML algorithms.....	49
<i>Supervised learning algorithms</i>	49
<i>Unsupervised learning algorithms</i>	50
<i>Reinforcement learning algorithms</i>	51
<i>Semi-supervised learning algorithms</i>	52
Conclusion	53
Points to remember	53
4. Basic Supervised Learning Algorithms	55
Introduction	55
Structure	55
Objectives	55
Introduction to supervised learning.....	56
<i>Linear regression</i>	57
<i>Mathematical foundation</i>	59

<i>Advantages of linear regression</i>	60
<i>Disadvantages of linear regression</i>	61
<i>Real-world applications</i>	61
<i>Real-world coding example</i>	62
<i>Logistic regression</i>	64
<i>Mathematical foundation</i>	66
<i>Advantages of logistic regression</i>	67
<i>Disadvantages of logistic regression</i>	68
<i>Real-world applications</i>	69
<i>Real-world coding example</i>	69
<i>Decision trees</i>	71
<i>Mathematical foundation</i>	73
<i>Advantages</i>	74
<i>Disadvantages</i>	75
<i>Real-world applications</i>	75
<i>Real-world coding example</i>	76
<i>Random forests</i>	77
<i>Mathematical foundation</i>	79
<i>Advantages</i>	81
<i>Disadvantages</i>	81
<i>Real-world applications</i>	82
<i>Real-world coding example</i>	83
<i>Support Vector Machines</i>	84
<i>Mathematical foundation</i>	87
<i>Advantages of SVM algorithms</i>	88
<i>Disadvantages of SVM algorithms</i>	89
<i>Real-world applications</i>	90
<i>Real-world coding example</i>	91
<i>Conclusion</i>	92
<i>Points to remember</i>	92

Exercises	93
<i>Linear regression exercise: Predict house prices</i>	93
<i>Logistic regression exercise: Predicting customer churn</i>	95
<i>Decision tree exercise: Diagnosing plant diseases</i>	97
<i>Random forest exercise: Predicting wine quality</i>	98
<i>SVM exercise: Classifying handwritten digits</i>	101
5. Advanced Supervised Learning Algorithms	103
Introduction	103
Structure	103
Objectives	103
Introduction into advanced supervised learning.....	104
Naive Bayes.....	104
<i>Mathematical foundation</i>	106
<i>Bayes' theorem</i>	106
<i>Application to classification</i>	107
<i>Independence assumption</i>	107
<i>Types of Naïve Bayes</i>	107
<i>Parameter estimation</i>	108
<i>Advantages and disadvantages</i>	108
<i>Advantages of Naïve Bayes algorithms</i>	108
<i>Disadvantages of Naïve Bayes algorithms</i>	109
<i>Real-world applications</i>	109
<i>Real-world coding example</i>	110
k-Nearest Neighbors.....	112
<i>Mathematical foundation</i>	114
<i>Basic idea</i>	114
<i>Distance metrics</i>	114
<i>Making decisions</i>	114
<i>Choosing k</i>	116
<i>Complexity of the algorithm</i>	116

<i>Advantages and disadvantages</i>	116
<i>Advantages of k-NN algorithms</i>	116
<i>Disadvantages of k-NN algorithms</i>	117
<i>Real world applications</i>	117
<i>Real-world coding example</i>	118
Neural networks.....	120
<i>Neural network architecture</i>	120
<i>Forward propagation</i>	121
<i>Backpropagation</i>	121
<i>Training process</i>	121
<i>Mathematical foundation</i>	123
<i>Basic neuron model</i>	123
<i>Activation functions</i>	124
<i>Feedforward mechanism:</i>	124
<i>Cost function</i>	124
<i>Backpropagation</i>	125
<i>Optimization</i>	125
<i>Regularization</i>	125
<i>Advantages and disadvantages</i>	125
<i>Advantages of Neural network algorithms</i>	125
<i>Disadvantages of neural network algorithms</i>	126
<i>Real-world applications</i>	127
<i>Real-world coding example</i>	127
Gradient Boosting Machines	130
<i>Mathematical foundation</i>	132
<i>Boosting</i>	132
<i>Objective function</i>	132
<i>Loss function</i>	133
<i>Regularization</i>	133
<i>Gradient boosting</i>	133

<i>Shrinkage</i>	133
<i>Stopping criteria</i>	134
<i>Advantages and disadvantages of GBM algorithms</i>	134
<i>Advantages of GBM algorithms</i>	134
<i>Disadvantages of GBM algorithms</i>	135
<i>Real-world applications for GBM algorithms</i>	135
<i>Real-world GBM coding example</i>	137
XGBoost.....	138
<i>Mathematical foundation</i>	141
<i>Objective function</i>	141
<i>Regularization</i>	141
<i>Taylor expansion for approximation</i>	142
<i>Optimal leaf weights</i>	142
<i>Pruning</i>	142
<i>Handling missing values</i>	142
<i>Column block and parallelization</i>	143
<i>Advantages and disadvantages of XGBoost algorithms</i>	143
<i>Advantages of XGBoost algorithms</i>	143
<i>Disadvantages of XGBoost algorithms</i>	144
<i>Real-world applications for XGBoost algorithms</i>	144
<i>Real-world XGBoost coding example</i>	145
Conclusion	147
Points to remember	147
Exercises and solutions	148
<i>Naïve Bayes exercise: Classifying email messages as spam or not spam</i>	148
<i>k-NN exercise: Classifying types of flowers based on measurements</i>	150
<i>Neural Network exercise: Handwritten digit classification with the MNIST dataset</i> ..	152
<i>GBM exercise: Predicting house prices with the Boston Housing dataset</i>	154
<i>XGBoost exercise: Predicting iabetes Ouotcomes with the</i> <i>Pima Indians Diabetes dataset</i>	156

6. Basic Unsupervised Learning Algorithms	159
Introduction	159
Structure	159
Objectives	159
Introduction to unsupervised learning.....	160
<i>Key concepts in unsupervised learning.....</i>	<i>160</i>
<i>Popular basic unsupervised learning algorithms.....</i>	<i>160</i>
K-means clustering	161
<i>Mathematical foundation</i>	<i>163</i>
<i>Properties and considerations.....</i>	<i>164</i>
<i>Advantages and disadvantages</i>	<i>164</i>
<i>Real world applications</i>	<i>166</i>
<i>Real world coding example.....</i>	<i>166</i>
Hierarchical clustering	168
<i>Mathematical foundation</i>	<i>170</i>
<i>Properties and considerations.....</i>	<i>171</i>
<i>Advantages and disadvantages</i>	<i>172</i>
<i>Real world applications</i>	<i>173</i>
<i>Real world coding example.....</i>	<i>174</i>
Principal Component Analysis	176
<i>Mathematical foundation</i>	<i>178</i>
<i>Properties of PCA.....</i>	<i>178</i>
<i>Limitations</i>	<i>179</i>
<i>Advantages and disadvantages</i>	<i>179</i>
<i>Real world applications</i>	<i>180</i>
<i>Real world principal coding example</i>	<i>181</i>
t-Distributed Stochastic Neighbor Embedding	183
<i>Mathematical foundation</i>	<i>185</i>
<i>Conditional probabilities in high-dimensional space</i>	<i>185</i>
<i>Conditional probabilities in low-dimensional space</i>	<i>185</i>
<i>Symmetrized version.....</i>	<i>185</i>

<i>Cost function</i>	186
<i>Gradient descent</i>	186
<i>t-Distribution</i>	186
<i>Key points</i>	186
<i>Limitations and considerations</i>	186
<i>Advantages and disadvantages</i>	187
<i>Advantages of t-SNE algorithms</i>	187
<i>Disadvantages of t-SNE algorithms</i>	187
<i>Real-world applications</i>	188
<i>Real-world coding example</i>	189
Association Rule Mining: A priori Algorithm.....	191
<i>Mathematical foundation</i>	193
<i>Basic terminologies</i>	193
<i>A priori principle</i>	193
<i>Advantages and disadvantages</i>	194
<i>Advantages</i>	195
<i>Disadvantages</i>	195
<i>Real world applications</i>	196
<i>Real-world coding example</i>	196
Conclusion	198
Points to remember.....	198
Exercises and solutions	199
<i>Using K-means clustering to segment mall customers based on spending habits</i>	199
Categorizing iris flowers based on features using hierarchical clustering.....	201
<i>Dimensionality reduction for breast cancer data visualization using PCA</i>	203
<i>Visualizing MNIST dataset in 2D using t-SNE</i>	205
<i>Analyzing retail transactions using Association Rule Mining</i>	207
7. Advanced Unsupervised Learning Algorithms.....	209
Introduction	209
Structure	209
Objectives	209

Introduction to unsupervised learning.....	210
Density-Based Spatial Clustering of Applications with Noise.....	211
<i>Mathematical foundations</i>	212
<i>Main concepts of DBSCAN</i>	212
<i>DBSCAN algorithm</i>	213
<i>Advantages and disadvantages</i>	213
<i>Advantages of DBSCAN</i>	214
<i>Disadvantages of DBSCAN</i>	214
<i>Real-world applications</i>	215
<i>Real-world coding example</i>	215
Gaussian Mixture Models.....	217
<i>Expectation-Maximization for GMM</i>	218
<i>Mathematical foundations</i>	219
<i>Gaussian distribution</i>	220
<i>Mixture of Gaussians</i>	220
<i>Expectation-Maximization for GMMs</i>	220
<i>Advantages and disadvantages</i>	221
<i>Advantages</i>	221
<i>Disadvantages</i>	222
<i>Real-world applications</i>	222
<i>Real-world coding example</i>	223
Autoencoders.....	224
<i>Defining an autoencoder</i>	224
<i>Loss function</i>	225
<i>Applications</i>	225
<i>Simple Autoencoder with TensorFlow and Keras</i>	225
<i>Mathematical foundations</i>	227
<i>Objective function</i>	227
<i>Encoder and decoder functions</i>	228
<i>Variational Autoencoders</i>	228

<i>Bottleneck and sparsity</i>	228
<i>Regularization and noise</i>	229
<i>Importance of Autoencoders in Generative AI</i>	229
<i>Advantages and disadvantages</i>	230
<i>Advantages</i>	230
<i>Disadvantages</i>	231
<i>Real-world applications</i>	232
<i>Real-world coding example</i>	232
<i>Using Autoencoders for image abstraction</i>	232
<i>Anomaly detection: Outlier detection</i>	234
<i>Coding example using Scikit-learn: Isolation forest</i>	235
<i>Mathematical foundation</i>	236
<i>Advantages and disadvantages</i>	238
<i>Advantages</i>	238
<i>Disadvantages</i>	238
<i>Real world applications</i>	239
<i>Real-world coding example</i>	240
<i>Latent Dirichlet Allocation</i>	242
<i>Mathematical foundation</i>	244
<i>Basics</i>	244
<i>LDA generative process</i>	244
<i>Variable Definitions</i>	244
<i>Goal</i>	244
<i>Gibbs sampling for LDA</i>	245
<i>Application</i>	245
<i>Advantages and disadvantages</i>	245
<i>Advantages of LDA</i>	245
<i>Disadvantages of LDA</i>	246
<i>Real-world applications</i>	247
<i>Real-world coding example</i>	247

Conclusion	249
Points to remember	249
Exercises and solutions	250
<i>DBSCAN Exercise: Clustering geographical data</i>	250
<i>GMM exercise: Clustering customer spending data</i>	252
<i>Autoencoder exercise: Image denoising</i>	254
<i>Anomaly detection exercise: Detecting fraudulent transactions</i>	256
<i>LDA exercise: Topic modelling on news articles</i>	258
8. Basic Reinforcement Learning Algorithms	261
Introduction	261
Structure	261
Objectives	262
Introduction to reinforcement learning	262
<i>Reinforcement learning process</i>	262
<i>Key elements</i>	263
<i>Reinforcement learning algorithms</i>	263
Q-learning	264
Mathematical foundation	267
<i>Q-learning update rule</i>	267
<i>Intuition</i>	268
<i>Convergence</i>	268
<i>Advantages and disadvantages</i>	268
<i>Advantages</i>	268
<i>Disadvantages</i>	269
<i>Real world applications</i>	270
<i>Real-world coding example</i>	270
Deep Q-Networks	273
Mathematical foundation	276
<i>Mathematical relationship with Q-learning</i>	277
<i>Challenges and extensions</i>	277
<i>Advantages and disadvantages</i>	277

<i>Advantages of DQNs</i>	278
<i>Disadvantages of DQNs</i>	278
<i>Real world applications</i>	279
<i>Real-world coding example</i>	279
Policy Gradient Methods.....	282
<i>Reinforce algorithm</i>	282
<i>Mathematical foundation</i>	284
<i>Intuition</i>	285
<i>Objective function</i>	285
<i>Policy Gradient theorem</i>	285
<i>Reinforce</i>	285
<i>Challenges and extensions</i>	286
<i>Advantages and disadvantages</i>	286
<i>Advantages of PGMs</i>	286
<i>Disadvantages of PGMs</i>	287
<i>Real world applications</i>	287
<i>Real world coding example</i>	288
Advantage Actor-Critic.....	290
<i>Mathematical foundation</i>	294
<i>Intuition and advantages</i>	294
<i>Advantages and disadvantages</i>	295
<i>Real-world applications</i>	296
<i>Real-world coding example</i>	296
Trust Region Policy Optimization.....	299
<i>Mathematical foundation</i>	302
<i>Objective function</i>	302
<i>Surrogate objective function</i>	302
<i>Trust region constraint</i>	303
<i>Optimization</i>	303
<i>Intuition</i>	303

<i>Advantages and disadvantages</i>	303
<i>Real-world applications</i>	304
<i>Real world coding example</i>	305
Conclusion	306
Points to remember	307
Exercises and solutions	308
<i>Q-learning exercise: Navigate a grid world</i>	308
<i>DQN exercise: CartPole balancing with DQN</i>	310
<i>Policy gradient exercise: Solve the LunarLander environment</i>	312
<i>A2C exercise: MountainCar continuous control</i>	314
<i>TRPO exercise: Robot locomotion using TRPO</i>	315
9. Advanced Reinforcement Learning Algorithms	319
Introduction	319
Structure	319
Objectives	320
Introduction into Reinforcement learning.....	320
Advanced Reinforcement Learning algorithms	320
<i>Asynchronous Advantage Actor-Critic</i>	321
<i>Mathematical foundation</i>	324
<i>Advantages and disadvantages</i>	325
<i>Real world applications</i>	326
<i>Real-world coding example</i>	326
<i>Proximal Policy Optimization</i>	328
<i>Mathematical foundation</i>	330
<i>Advantages and disadvantages</i>	331
<i>Real-world applications</i>	332
<i>Real-world coding example</i>	332
<i>Deep Deterministic Policy Gradient</i>	334
<i>Mathematical foundation</i>	335
<i>Advantages and disadvantages</i>	336

<i>Real-world applications</i>	337
<i>Real-world coding example</i>	338
<i>Twin Delayed Deep Deterministic Policy Gradient</i>	339
<i>Mathematical foundation</i>	340
<i>Advantages and disadvantages</i>	341
<i>Real-world applications</i>	342
<i>Real-world coding example</i>	343
<i>Soft Actor-Critic</i>	344
<i>Mathematical foundation</i>	345
<i>Advantages and disadvantages</i>	347
<i>Real-world applications</i>	348
<i>Real-world coding example</i>	348
Conclusion	349
Points to remember	350
Exercises and solutions	351
<i>A3C exercise: Implementing asynchronous training for the CartPole game</i>	351
<i>PPO exercise: Balancing the Lunar Lander with PPO</i>	353
<i>DDPG exercise: Navigating a Pendulum using DDPG</i>	355
<i>TD3 exercise: Controlling a Bipedal Robot with TD3</i>	357
<i>Soft Actor-Critic exercise: Training an Agent to Balance a Pendulum</i>	359
10. Basic Semi-Supervised Learning Algorithms	361
Introduction	361
Structure	362
Objectives	362
Introduction to semi-supervised learning.....	362
<i>Need for semi-supervised learning</i>	362
<i>Techniques in semi-supervised learning</i>	363
<i>Challenges</i>	363
Self-training.....	364
<i>Mathematical foundation</i>	366

Intuition	366
Risks	367
Advantages and disadvantages	367
Real-world applications	368
Real-world coding example	369
Scenario	369
Co-training	371
Mathematical foundation	373
Assumptions	373
The Algorithm	374
Mathematical justification	374
Advantages and disadvantages	375
Real-world applications	376
Real-world coding example	376
Multi-view learning	378
Coding example using co-training	379
Mathematical foundation	380
Co-training	380
Multiple kernel learning	380
Canonical correlation analysis based methods	381
Shared and individual feature learning:	381
Joint and individual feature learning	381
Advantages and disadvantages	381
Real-world applications	383
Real-world coding example	383
Scenario	383
Expectation-Maximization	385
Mathematical foundation	387
Advantages and Disadvantages	388
Real-world applications	389

<i>Real-world coding example</i>	390
Graph-based methods	391
<i>Label propagation as a graph-based semi-supervised learning example</i>	392
<i>Mathematical foundation</i>	393
<i>Objective function</i>	393
<i>Laplacian matrix</i>	394
<i>Label propagation</i>	394
<i>Algorithm steps</i>	394
<i>Regularization and modifications</i>	394
<i>Advantages and disadvantages</i>	394
<i>Real world applications</i>	396
<i>Real world coding example</i>	396
Conclusion	398
Points to remember	399
Exercises and solutions	400
<i>Exercise: Implementing a self-training algorithm for text classification</i>	400
<i>Exercise: Implementing a o-training algorithm for sentiment analysis</i>	401
<i>Exercise: Implementing a multi-view learning algorithm for</i> <i>image and text classification</i>	403
<i>Exercise: Implementing the Expectation-Maximization algorithm</i> <i>for Gaussian Mixture Models</i>	404
<i>Exercise: Implementing a graph-based semi-supervised learning algorithm</i>	406
11. Advanced Semi-Supervised Learning Algorithms	409
Introduction	409
Structure	409
Objectives	410
Introduction to semi-supervised learning	410
Transductive Support Vector Machines	410
<i>Mathematical foundation</i>	412
<i>Formulation</i>	412
<i>Transductive inference</i>	412

<i>Mathematical properties</i>	413
<i>Advantages and disadvantages</i>	413
<i>Real-world applications</i>	414
<i>Real-world coding example</i>	415
Co-regularization: Label propagation.....	416
<i>Mathematical foundation</i>	418
<i>Objective function</i>	419
<i>Advantages and disadvantages</i>	420
<i>Real world applications</i>	421
<i>Real-world coding example</i>	422
Deep generative models.....	423
<i>Mathematical foundation</i>	426
<i>Variational Autoencoders in semi-supervised learning</i> :.....	426
<i>Generative Adversarial Networks in semi-supervised learning</i>	426
<i>Unified objective for semi-supervised deep generative models</i>	426
<i>Advantages and disadvantages</i>	427
<i>Real-world applications</i>	428
<i>Real-world coding example</i>	429
Virtual Adversarial Training.....	431
<i>Mathematical foundation</i>	434
<i>Mathematical notation</i>	434
<i>Objective function</i>	434
<i>Loss function</i>	434
<i>Algorithm</i>	434
<i>Advantages and disadvantages</i>	435
<i>Real world applications</i>	436
<i>Real world coding example</i>	437
Tri-training	439
<i>Mathematical foundation</i>	441
<i>Ensemble learning</i>	441
<i>Bootstrapping</i>	442

Voting mechanism.....	442
Iterative refinement	442
Consensus	442
Error rate.....	442
Advantages and disadvantages	442
Real world applications	444
Real-world coding example	444
Conclusion	446
Points to remember	446
Exercises	447
Exercise on Transductive Support Vector Machines.....	447
Exercise on co-regularization for semi-supervised learning	449
Exercise on deep generative models.....	450
Exercise on Virtual Adversarial Training for semi-supervised learning	452
Exercise on tri-training for semi-supervised learning	453
12. Natural Language Processing.....	455
Introduction	455
Structure	455
Objectives	455
Natural Language Processing	456
Natural Language Understanding.....	457
Python coding example for NLU	458
Natural Language Generation	459
Python coding example for NLG	460
Large Language Models	461
History of LLMs.....	461
What are LLMs?	462
Python coding example for LLMs.....	463
Generative AI	464
Python coding example for Generative AI.....	465

<i>Mathematical foundations of Natural Language Processing</i>	466
<i>Mathematical foundation of NLU</i>	466
<i>Mathematical foundation of Natural Language Generation</i>	469
Advantages and disadvantages	470
Real-world applications	472
Coding example	473
Conclusion	474
Points to remember	475
Exercise	476
<i>Exercise: Sentiment analysis on movie reviews</i>	476
13. Computer Vision	479
Introduction	479
Structure	479
Objectives	479
Computer vision.....	480
<i>Mathematical foundation</i>	482
<i>Convolutional Neural Networks</i>	483
<i>Spiking Neural Networks: Outlook into the future of CV</i>	485
Advantages and disadvantages of computer vision.....	486
Real-world applications for computer vision	487
Coding example for computer vision	488
<i>Example for Convolutional Neural Networks</i>	490
<i>Example for Spiking Neural Networks</i>	492
Conclusion	494
Points to remember	494
Exercise	495
<i>Exercise: Image classification with Computer Vision</i>	495
14. Large-Scale Algorithms	497
Introduction	497
Structure	497
Objectives	497

Large-scale algorithms	498
<i>MapReduce</i>	499
<i>What is MapReduce?</i>	501
<i>Distributed machine learning</i>	502
<i>Graph processing algorithms</i>	504
<i>Large-scale optimization</i>	505
<i>Mathematical foundation of large-scale algorithms</i>	507
Advantages and disadvantages	508
Real-world applications	510
Coding example	511
Conclusion	513
Points to remember	513
Exercise	514
<i>Exercise: Word count using MapReduce</i>	514
15. Outlook into the Future: Quantum Machine Learning	517
Introduction	517
Structure	517
Objectives	518
A Short Introduction to Quantum and Quantum Computing	518
Quantum Machine Learning	519
<i>Challenges and limitations</i>	520
<i>Mathematical foundation</i>	521
Quantum Machine Learning algorithms	522
<i>Quantum Support Vector Machine</i>	522
<i>Quantum Neural Networks</i>	524
<i>Quantum Principal Component Analysis</i>	525
<i>Quantum k-Means Clustering</i>	527
<i>Quantum Boltzmann Machine</i>	528
<i>Quantum Genetic Algorithms</i>	530
Advantages of Quantum Machine Learning	531
Disadvantages of Quantum Machine Learning	532

Real-world applications	533
Coding example	533
Conclusion	535
Points to remember	535
Exercise	536
<i>Exercise: Implementing a Simple Quantum Circuit for Quantum Machine Learning</i>	<i>536</i>
Index	539-553

CHAPTER 1

Fundamentals

Introduction

This chapter of the book will cover the fundamentals of **artificial intelligence (AI)** and **machine learning (ML)**. We would start by laying out the fundamentals and their definitions to create a common understanding of the field. We will dive into the world of AI and ML by defining the fields and their impact on the world inside and outside of AI. We will as well include the critical concepts and what kind of industry problems could be solved with AI and ML. We will close out the chapter with simple examples to make a differentiation between an AI/ML application and an AI/ML algorithm.

Structure

The chapter covers the following topics:

- Fundamentals of AI and ML
- Defining AI and ML
- History of AI and ML
 - Classic examples of AI and ML
- AI and ML algorithms
 - Examples of AI and ML algorithms
 - Structure of a typical AI and ML algorithm

Objectives

By the end of this chapter, you will gain a comprehensive understanding of AI and ML as general concepts and their underlying fundamentals. Additionally, you will learn about the origins of AI and ML and be exposed to some basic examples. Furthermore, you will grasp the concept of basic data structures associated with these fields.

Fundamentals of AI and ML

The fundamentals of AI and ML encompass a wide range of concepts and techniques. Here are some key fundamentals of AI and ML:

- **Data:** High-quality data is essential for AI and ML. It serves as the foundation for training and evaluating models. Understanding the data, its quality, structure, and representation is crucial for successful AI and ML applications.
- **Algorithms:** Algorithms are mathematical and computational procedures used to solve specific problems or perform tasks. In AI and ML, algorithms are used to train models, make predictions, and make decisions based on data. Examples include decision trees, neural networks, **Support Vector Machines (SVM)**, and clustering algorithms.
- **Feature engineering:** Feature engineering involves selecting, transforming, and creating relevant features from raw data to improve the performance of ML models. This process helps extract meaningful information and patterns from the data, making it easier for models to learn and make accurate predictions.
- **Model training:** Model training is the process of feeding labeled data into an algorithm or model to learn patterns and relationships. During training, the model adjusts its internal parameters to minimize the difference between predicted and actual outputs. This process often involves optimization techniques, such as gradient descent, to find the best parameter values.
- **Model evaluation:** Evaluating the performance of ML models is crucial to ensure their effectiveness and generalization. Various metrics, such as accuracy, precision, recall, and F1-score, are used to assess the model's predictive capabilities. Cross-validation techniques, such as k-fold cross-validation, help estimate the model's performance on unseen data.
- **Generalization and overfitting:** Generalization refers to a model's ability to perform well on unseen data. Overfitting occurs when a model becomes overly complex and performs well on the training data but fails to generalize to new data. Techniques such as regularization and early stopping are employed to prevent overfitting and promote better generalization.
- **Model deployment:** Deploying ML models involves making them available for use in real-world applications. This includes optimizing the model for efficiency, scalability, and compatibility with the target environment. The model deployment

also involves monitoring the model's performance and retraining or updating it when necessary.

- **Ethics and bias:** As AI and ML systems have a societal impact, understanding the ethical implications and addressing potential biases is crucial. Ensuring fairness, transparency, and accountability in AI systems is an essential consideration. Ethical considerations involve issues such as data privacy, algorithmic bias, and the potential impact of AI on various stakeholders.

These fundamentals provide a solid foundation for understanding and developing AI and ML applications. Mastering these concepts allows practitioners to build robust and effective AI systems. This book will focus on the algorithms that are forming the core of every modern AI and ML application.

Defining AI and ML

Let us now discuss AI and ML in detail.

Artificial Intelligence

Artificial Intelligence is a broad field in computer science that focuses on the creation of systems capable of performing tasks that would typically require human intelligence. This includes tasks like understanding natural language, recognizing patterns, solving problems, learning from experience, and making decisions.

Examples of AI in use today by data scientists include:

- **Natural Language Processing (NLP):** NLP algorithms are used to create systems like *Siri*, *Google Assistant*, and *ChatGPT* (which you currently interact with) that can understand and generate human language.
- **Computer vision:** Algorithms in this domain are designed to interpret and understand the visual world. For instance, Facebook uses computer vision AI to recognize and tag faces in images.
- **Recommendation systems:** Websites like *Amazon* and *Netflix* use AI to recommend products or movies based on a user's past behavior and the behavior of similar users.
- **Predictive analytics:** Many industries use AI to predict future outcomes, like predicting stock prices in finance or disease outbreaks in healthcare.
- **Autonomous vehicles:** Companies like Tesla use AI to enable cars to navigate and understand the world around them.

These examples only scratch the surface of AI's potential. Its reach is continually expanding, making it a crucial tool in a modern data scientist's arsenal.

Machine learning

Machine learning is a subset of AI that gives computers the ability to learn from data and make decisions or predictions without being explicitly programmed to do so. This process involves the development of algorithms that can process large amounts of data, learn patterns within that data, and use this learned information to predict future outcomes or behavior. This *learning* is accomplished by improving the performance of the system over time as it is exposed to more data.

There are three main types of ML, supervised learning, unsupervised learning, and reinforcement learning, which are discussed below:

- **Supervised learning** involves training a model on a labeled dataset, that is, a dataset where the outcome or target variable is known. The model learns the relationship between the features and the target and can then predict the outcome for new, unseen data. For example, a bank might use supervised learning to predict whether a loan applicant will default based on their previous loan history and financial profile.
- **Unsupervised learning** involves training a model on an unlabeled dataset, that is, a dataset where the outcome or target variable is not known. The goal is to discover hidden patterns or intrinsic structures within the data. Common uses include clustering and dimensionality reduction. For example, a retail company might use unsupervised learning to segment its customers into different groups based on their buying behavior.
- **Reinforcement learning** involves training a model to make a series of decisions by rewarding or punishing the model (the *agent*) based on the actions it takes in an environment to reach a goal. The model learns to perform actions that maximize some reward over time. This is often used in robotics, gaming, and navigation. For example, reinforcement learning has been used to train AI to play and win complex games like *Go* and *Chess*.

Modern data scientists need to understand these concepts and techniques to build and deploy effective ML models. Moreover, they often need to use different types of machine learning in concert, depending on the task at hand. They should also be aware of new trends in ML, such as Deep Learning, transfer learning, and active learning, which have led to significant advancements in fields like computer vision, natural language processing, and recommender systems.

History of AI and ML

The development of AI and ML has been an incremental journey spanning several decades. The evolution of these fields has been influenced by various domains like mathematics, statistics, computer science, cognitive psychology, and neuroscience, as discussed in the following points:

- **The 1950s - Birth of AI and ML:** The birth of AI as a distinct field happened during a summer conference at Dartmouth College in 1956, which was attended by pioneers like *John McCarthy*, *Marvin Minsky*, *Allen Newell*, and *Herbert Simon*. Here, they proposed that *every feature of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it.*

Even before this, in 1950, *Alan Turing* introduced the concept of machine intelligence with the *Turing Test*, a measure of a machine's ability to exhibit intelligent behavior equivalent to, or indistinguishable from, that of a human.

In 1959, *Arthur Samuel* developed a program that could play checkers and learn from its mistakes, marking one of the first self-learning programs and a seminal moment in ML.

- **The 1960s - Growth and consolidation:** In the 1960s, AI research focused on problem-solving and symbolic methods. AI programs like **DENDRAL** and **ELIZA** were developed during this time.

ML saw a significant development in 1967 with the creation of the **Nearest Neighbor algorithm**, which started basic pattern recognition.

- **The 1970s - AI Winter and rule-based systems:** The mid-1970s marked the beginning of the first *AI Winter*, a period of disappointment resulting from the overhyping of AI capabilities and subsequent cuts in funding. The focus shifted towards *expert systems* – rule-based systems that tried to mimic the decision-making of human experts.
- **The 1980s – Revival and ML expansion:** In the 1980s, AI saw a revival with the rise of ML. The development of the backpropagation algorithm enabled more efficient training of neural networks, and the advent of SVM led to significant progress in ML.
- **The 1990s – AI and ML Maturity:** The 1990s saw ML mature into a field of its own, with the growth of decision tree algorithms, reinforcement learning, and Bayesian networks. AI and ML began to be used in practical applications, from data mining to industrial robotics.
- **The 2000s - The data boom and rise of deep learning:** The explosion of data in the 2000s, due to the rise of the internet and, later, social media, alongside advancements in computational power and storage, created the perfect conditions for AI and ML to flourish. Deep learning, a subset of ML, started to become feasible, driven by the development of new neural network architectures.
- **The 2010s - AI and ML breakthroughs:** This decade witnessed rapid progress in AI and ML. The development of advanced neural network architectures, like **Convolutional Neural Networks (CNNs)** and **Recurrent Neural Networks (RNNs)**, led to breakthroughs in image and speech recognition and natural language processing.