PONDICHERRY UNIVERSITY

(A Central University)

DIRECTORATE OF DISTANCE EDUCATION

ARTIFICIAL INTELLIGENCE

MBA - IV Semester

Paper Code : MBEP 4005

Author

Dr. Sukhvinder Singh
Assistant Professor,
Dept. of Computer Science
Pondicherry University
Puducherry

MBA - IV Semester **Paper Code: MBEP4005**

ARTIFICIAL INTELLIGENCE

Learning Objective

- **To understand the Artificial Intelligence**
- **To understand the Machine Learning**
- **To learn the enterprise AI strategy and planning techniques**
- **To learn about the Business applications and challenges**
- **To know the role of AI for Enterprise functions.**

UNIT - I Introduction about AI

Needs of Business leaders, basic terminologies in AI, Modern AI Techniques, Machine Intelligence Continuum, promises of AI- Challenges of AI and its effects, Designing safe and Ethical AI

UNIT - II Machine Learning Overview

Types of ML - Accuracy of ML models - Specific ML Methods: A Deep Dive - Model Selection and Validation

UNIT - III Developing an enterprise AI strategy

Invest in technical Talent, Plan Implementation, Collet and prepare data, Building Machine Learning Model, Experiment and Iterate

UNIT - IV Business Applications

Recommender Systems - Impact of recommenders on markets - Other forms of personalization on the web - Challenges with personalization - ML in Finance: Fraud Detection - ML in Finance: Additional applications

UNIT - V AI for Enterprise functions Gandhiji and Khad

Obstacles and opportunities, General and administrative, human resources and talent, business intelligence and analytics, Software Development, Marketing, Sales, Customer support- ethics of enterprise AI Generating AI, Data Protection Laws, Regulatory Aspects of AI.

References

Tom Markiewicz & Josh Zheng, getting started with Artificial Intelligence, Second Edition, O'Reilly Media Inc.2020

Mariya Yao, Marlene Jia, Adelyn Zhou, Applied Artificial Intelligence. www.appliedaibook.com

TABLE OF CONTENTS			
UNIT	LESSON	TITLE	PAGE NO.

UNIT – I

Lesson 1.1 - Artificial Intelligence

Introduction about AI: *Needs of Business leaders, basic terminologies in AI, Modern AI Techniques, Machine Intelligence Continuum, promises of AI- Challenges of AI and its effects, designing safe and Ethical AI*

Learning Objectives

1. Understand the fundamental concepts of Artificial Intelligence (AI)
2. To understand the key terms in AI
3. To understand the concept of the Machine Intelligence Continuum and its implications for AI development.
4. To explore the potential benefits and opportunities offered by AI across different sectors
5. To understand and identify the ethical, social, and economic challenges associated with AI adoption and deployment.

1 Introduction about AI

1.1 Needs of Business Leaders

The concept of leadership has changed completely in recent years. Away from the authoritarian boss, towards the coach. The goal to support employees in their personal development and the necessity to tap their full potential.

This not only creates a conducive environment but also helps in strengthening the teamwork the culture of the team and many other factors. In today's world it is not only the employees who take up the pivotal role but also the technology.

1.2 Need of AI in Business

The integration of AI technologies into everyday management opens up a broad spectrum of new possibilities and at the same time puts proven concepts to the test. This is because artificial intelligence is encroaching on the managers' area of responsibility in many areas.

1.2.1 AI as an Assistant

AI is a tireless assistant. It analyses data in real time and provides recommendations for action. This allows one to make better-informed decisions and react faster to changes.

1.2.2 Relief from Routine Works

Routine tasks eat up valuable time, but are often an important part of everyday work. AI can help – and take over unpleasant tasks. Managers can thus concentrate more on strategic planning, innovative idea development and, above all, on the individual support of their employees.

1.2.3 Improved human Resources Development

AI can help identify individual employee strengths and development areas more precisely. As a result, customized training and development plans can be created that specifically promote professional development. What is important, however, is that the human being must remain at the center.

1.2.4 Ethics and Society

AI in the boardroom is revolutionizing corporate management. But it also raises new questions. Questions about data security, privacy and fair use of AI systems. Tomorrow's leadership must ensure responsible use of AI – and ensure that technological advances are in line with corporate values and the needs of employees.

1.3 Basic Terminologies in AI

Artificial Intelligence is a study to make computers, robots, generally, machines think how the intellect of humans works, think, and learn when it solves any problem.

The main objective of Artificial Intelligence is to enhance computer functions that are related to natural knowledge, for illustration, logic, knowledge, and problem-answering. Computer knowledge attracts AI in the field of science, mathematics, psychology, linguistics, philosophy, and so on.

Terminologies in AI

Term	Meaning
Agent	Agents are systems or software programs suitable for self-governed, purposeful, and logic channeled towards one or another objective. They're also called assistants
Automata theory	The study of abstract machines and automata, as well as the computational problems
Autonomous Robot	The robot is independent of external control or influence and capability to govern itself solely.
Autonomous car	A vehicle that is capable of sensing its environment and moving with little or no human input
Backward Chaining	The scheme of working backward for the Reason/ effect of a problem.
Environment	It's the part of the real or computational world inhabited by the agent.
Heuristics	Knowledge is predicated on Trial-and-error, estimations, and experiments.
Pruning	Overriding nonessential and irrelative considerations in AI systems.
Rule	It's a format of representing knowledge base in Expert System. It's in the form of IF-THEN-ELSE.
Task	It's the aim the agent is trying to achieve.

Table 1.1

1.4 Modern AI Techniques

Artificial Intelligence (AI) refers to the simulation of human intelligence processes by computer systems, encompassing tasks such as learning, reasoning, problem-solving, perception, and decision-making. AI technologies enable machines to perform cognitive functions typically associated with human intelligence, allowing them to analyze data, recognize patterns, and make autonomous decisions.

AI is becoming increasingly part of everyday life, allowing machines to match, or even improve upon, the capabilities of the human mind.

AI represents a groundbreaking field of computer science and engineering that aims to create intelligent systems capable of emulating human-like cognitive abilities. With its ability to analyze vast amounts of data, automate tasks, and solve complex problems, AI holds the promise of driving innovation, enhancing productivity, and improving the quality of life for people around the world. However, realizing the full potential of AI requires careful consideration of its ethical implications and responsible deployment to ensure that AI technologies benefit society as a whole.

1. **Machine learning**: Uses AI algorithms to create and analyze situations to predict their consequences
2. **Deep learning**: Emulates the human brain's neural circuits, processing data and working with patterns to mimic human behavior
3. **Generative AI**: Uses deep learning algorithms to create new data from existing data
4. **Reinforcement learning**: An agent learns how to perform a task by interacting with its environment through action-reward
5. **Supervised learning**: An algorithm learns from a labeled dataset with the output already known
6. **Agentic AI**: Combines machine learning algorithms and deep learning systems
7. **Personalized learning**: AI can deliver custom experiences based on a student's unique preferences and skill set
8. **Speech recognition**: AI-based speech recognition has made it possible for computers to understand and recognize human speech

Other modern AI techniques include: Natural language generation, Virtual agents, Decision management, Biometrics, and Robotic process automation.

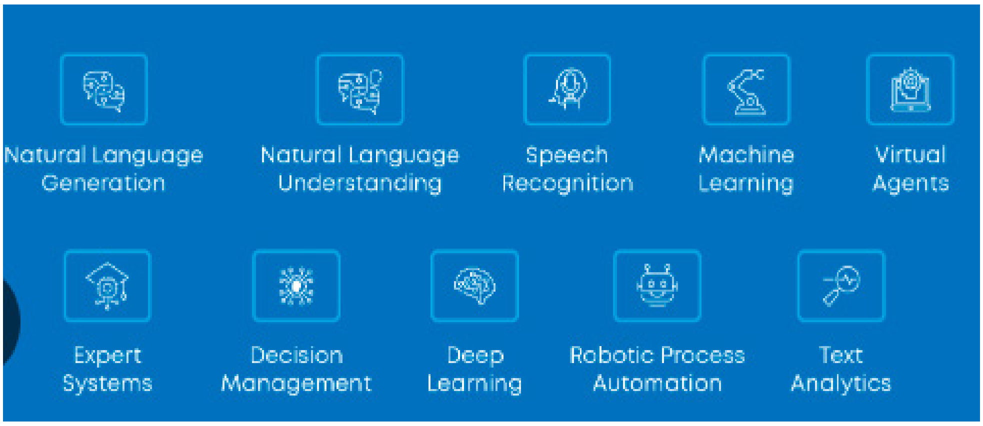

Figure 1.1

1.4.1 Modern AI Techniques in Business

Modern AI techniques are transforming businesses in various ways.

AI Technique	Description
Cybersecurity and Fraud Management	Over half of business owners use AI to protect their systems and manage fraud, making operations more secure.
Customer Service	AI chatbots and virtual assistants are employed by 56% of businesses to enhance customer support and engagement.
Digital Personal Assistants	Nearly half of the businesses use AI for personal assistant functions, streamlining tasks and schedules.
Customer Relationship Management (CRM)	AI is used in CRM to personalize customer interactions and predict customer needs.
Inventory Management	AI helps in forecasting demand and optimizing stock levels, used by 40% of businesses.
Content Production	AI tools like GPT-3 are used for generating content, including multilingual content creation.

AI Technique	Description
Product Recommendations	AI algorithms analyze customer data to provide personalized product recommendations.
Accounting and Finance	AI automates financial processes like bookkeeping and expense categorization.
Supply Chain Operations	AI optimizes logistics, predicting supply chain disruptions and improving efficiency.
Recruitment and Talent Sourcing	AI aids in screening candidates and identifying the best talent for the company.
Audience Segmentation	Businesses use AI to segment their audience for targeted marketing strategies.

Table 1.2

1.5 Machine Intelligence Continuum

To help business executives disentangle the functional differences between different AI approaches, there is a segmented applications as Machine Intelligence Continuum (MIC). The MIC represents a continuum from simple, scripted automation to superhuman intelligence and highlights the functional capabilities of different levels of machine intelligence.

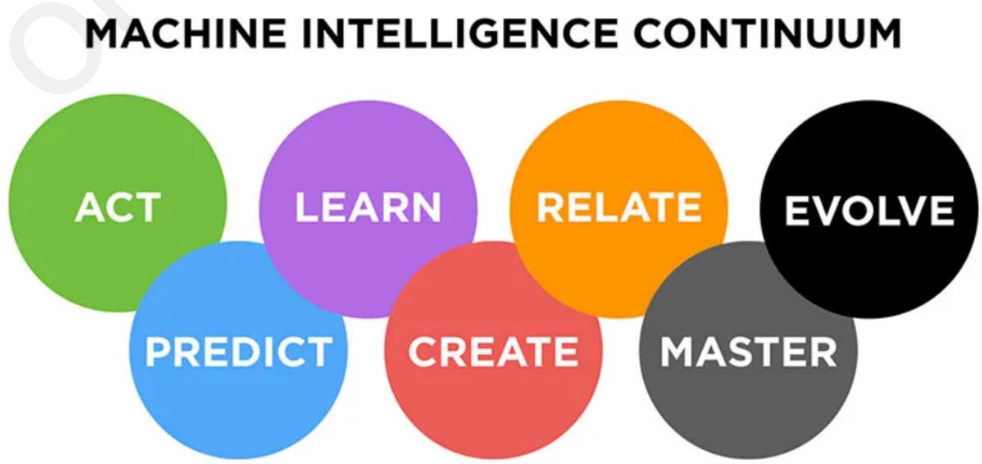

Figure 1.2

1.5.1 Systems that Act

The lowest level of the Machine Intelligence Continuum (MIC) are "Systems That Act" which we define as rule-based automatons. These are systems that are hand-engineered by experts and perform in a scripted fashion, often following if-then types of rules.

Examples include the fire alarm in a house. A fire alarm contains a sensor that detects smoke levels. When smoke levels reach a certain level, the device will play an alarm sound until manually turned off. That would result in very negative outcomes for you. Yet most companies claiming to have "AI" are really just using Systems That Act, or rule-based mechanisms that are incapable of dynamic actions or decisions.

1.5.2 Systems that Predict

"Systems That Predict" are systems that are capable of analyzing data and producing probabilistic predictions based on the data. Note that a "prediction" does not necessarily need to be a future event, but rather a mapping of known information to unknown information.

Automated and computational statistics underlie most "Systems That Predict", but predictions are only as good as the incoming data. If your data is flawed, or you choose a sample set to analyze that does not represent your target population as a whole, you will get erroneous results.

1.5.3 Systems that Learn

Machine learning and deep learning drive most "Systems That Learn". While many learning systems also make predictions like statistical systems do, they differ in that they require less hand-engineering and can learn to perform tasks without being explicitly programmed to do so. For many computational problems, they can function at human or better-than-human levels.

Learning can be automated at different levels of abstraction and for different components of a task. Completing a task requires first acquiring data which is used to generate a prediction about the world. This prediction is combined with higher level judgement and an action to produce a result. Feedback and measurements from the outcome can be fed back to earlier decision points to improve the task performance.

Figure 1.3

Many enterprise applications of statistics and machine learning focus on improving the process of turning data into predictions. In sales, for example, machine learning approaches to lead scoring can perform better than rule-based or statistical methods. Once the machine has produced a prediction of how good a lead is, the salesperson then applied human judgement to take follow up action.

More complex systems, such as self-driving cars and industrial robotics, handle the entire anatomy of a task. An autonomous vehicle must turn video and sensor feeds into accurate predictions of the surrounding world and take the correct action based on the environment. Some complex models can also perform online learning, which entails using real-time data to update machine learning models, versus offline learning, which involves training models on static, pre-existing data.

1.5.4 Systems that Create

Computers have been used for generative design and art for decades. Recent breakthroughs in neural network models have inspired a resurgence of computational creativity, with computers now capable of producing original writing, imagery, music, industrial designs, and even AI software!

Autodesk, the leading producer of CAD software for industrial design, released Dreamcatcher, a program that generates thousands of possible design permutations based on initial constraints set by engineers. Dreamcatcher has produced bizarre yet highly effective designs that challenge traditional manufacturing assumptions and exceed what human designers can manually ideate.

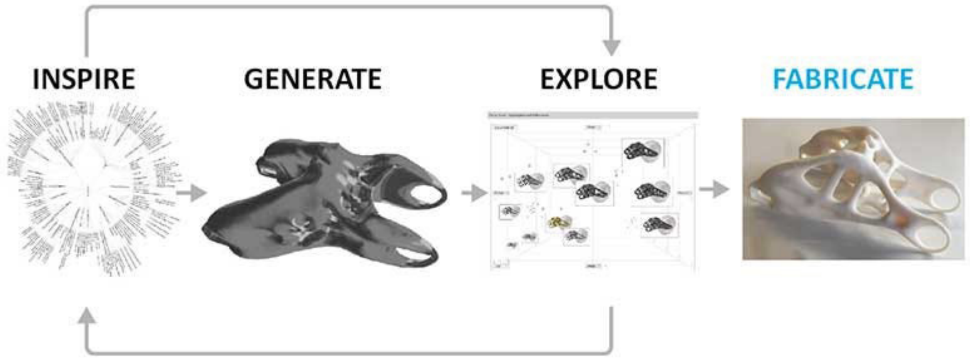

Figure 1.4

AI is even outperforming some artists economically! Google's DeepDream hosted an exhibition and auction of AI-generated art that collectively sold for $97,605.

1.5.5 Systems that Relate

As human employees increasingly collaborate with AI tools at work, and digital assistants like Apple's Siri and Amazon Echo's Alexa infiltrate our personal lives, machines will also need to be emotionally intelligent to succeed in our society.

Sentiment analysis, also known as opinion mining or emotion AI, extracts and quantifies emotional states from our text, voice, facial expressions, and body language. Knowing a user's affective state enables computers to respond empathically and dynamically, as the best humans we know often do. The applications to digital assistants are obvious, and companies like Amazon are already prioritizing emotional recognition for the Echo.

Emotional awareness can also improve interpersonal business functions such as sales, marketing, and communications.

1.5.6 Systems that Master

A human toddler only needs to see a single tiger to develop a mental construct of the animal and recognize other tigers. If humans needed to see thousands of tigers before learning to run away, our species would have died out long ago. By contrast a deep learning algorithm needs to process thousands of tiger images in order to begin recognizing them in images and video. Even then, neural networks trained on tiger photos do not reliably recognize other abstractions and representations of them, such as cartoons or costumes.

Humans have no trouble with this, because we are "Systems That Master". A "System That Masters" is an intelligent agent capable of constructing abstract concepts and strategic plans from sparse data. By creating modular conceptual representations of the world around us, we are able to transfer knowledge from one domain to another, a key feature of general intelligence. But humans are still the Master.

1.5.7 Systems that Evolve

This final category refers to systems that exhibit superhuman intelligence and capabilities. "Systems That Evolve" are entities capable of dynamically changing their own architecture and design to adapt to environmental needs. As humans, we're limited in our intelligence by our biological brains, also known as "wetware".

We evolve through genetic mutations across generations, rather than through re-architecting our own biological infrastructure during our lifetime. We cannot simply insert new RAM if we wish to augment our memory capacity, or buy a new processor if we wish to think faster.

While we continuously search for other intelligent life, we are not aware of any "Systems That Evolve", or superhuman intelligence. Computers are currently constrained by both hardware and software, while humans and other biological organisms are constrained by wetware. Some futurists hypothesize that we may be able to achieve superhuman intelligence by augmenting biological brains with synthesized technologies, but currently this research is more science fiction than science.

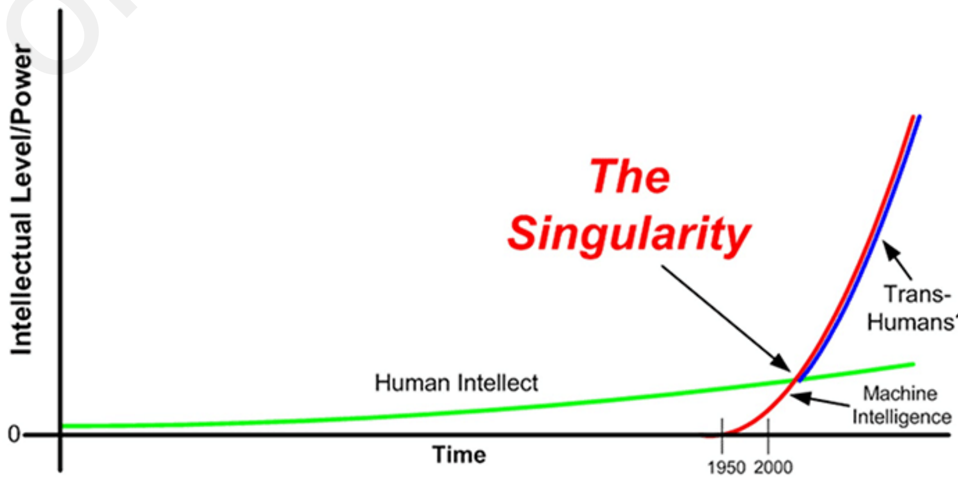

Figure 1.5

Once an upgradable intelligent agent does emerge, we will reach what many experts call the technological "singularity", when machine intelligence surpasses human intelligence. Self-evolving agents will be capable of ever-faster iterations of self-improvements, leading to an intelligence explosion and the emergence of superintelligence.

1.6 Promises of AI

Artificial Intelligence (AI) holds the promise of transforming virtually every aspect of human life. Some of its key promises include:

Increased Efficiency: AI can automate tasks, making processes faster and more efficient. This applies to industries ranging from manufacturing to customer service.

Improved Decision Making: AI algorithms can analyze vast amounts of data quickly and accurately, helping humans make better decisions in fields like healthcare, finance, and logistics.

Enhanced Creativity: AI systems can aid in creative endeavors by generating ideas, assisting in design processes, and even producing art and music.

Personalization: AI enables personalized experiences in fields such as marketing, education, and entertainment, tailoring content and recommendations to individual preferences.

Advanced Healthcare: AI can revolutionize healthcare by assisting in diagnosis, drug discovery, and treatment planning. It can also improve patient outcomes through predictive analytics and personalized medicine.

Automation: AI-powered robots and systems have the potential to automate a wide range of tasks, from simple repetitive actions to complex decision-making processes, leading to increased productivity and reduced labor costs.

Improved Safety and Security: AI can enhance security measures by analyzing patterns to detect and prevent threats in various domains, including cybersecurity, public safety, and transportation.

Environmental Impact: AI can contribute to sustainability efforts by optimizing energy consumption, managing resources more efficiently, and aiding in environmental monitoring and conservation.

Accessible Information: AI-powered systems can help organize and make sense of vast amounts of information, making knowledge more accessible to people worldwide.

Empowering Innovation: AI accelerates innovation by enabling researchers and developers to explore new ideas, test hypotheses, and create groundbreaking technologies across diverse fields.

These promises, however, come with challenges and ethical considerations, such as privacy concerns, job displacement, bias in algorithms, and the need for responsible AI development and deployment. As AI continues to evolve, addressing these challenges will be crucial in realizing its full potential for positive impact on society.

1.7 Challenges of AI and its Effects

While Artificial Intelligence (AI) offers numerous benefits, it also presents several challenges and potential negative effects:

- Job Displacement: Automation driven by AI technologies can lead to job loss in certain sectors, particularly those involving routine tasks. This could exacerbate economic inequality and require retraining programs for affected workers.
- Bias and Fairness: AI algorithms can inherit biases present in training data, leading to discriminatory outcomes in decision-making processes such as hiring, lending, and criminal justice. Ensuring fairness and mitigating bias in AI systems is a significant challenge.
- Privacy Concerns: AI systems often rely on large datasets, raising privacy concerns about the collection, storage, and use of personal information. Unauthorized access to sensitive data poses risks to individuals' privacy and autonomy.
- Security Risks: AI systems can be vulnerable to attacks and manipulation, posing security risks in various domains such as cybersecurity, autonomous vehicles, and critical infrastructure.
- Ethical Dilemmas: AI raises complex ethical questions, such as how to prioritize competing values (e.g., privacy vs. security) and ensure accountability for AI-driven decisions, particularly in autonomous systems where human oversight may be limited.
- Dependency on AI: Overreliance on AI systems without appropriate human oversight and intervention can lead to unforeseen

consequences and diminish human autonomy and decision-making capabilities.

- Algorithmic Transparency: Lack of transparency in AI algorithms and decision-making processes can undermine trust and accountability. Understanding how AI systems arrive at their conclusions is crucial for ensuring their fairness and reliability.
- Social Impact: AI technologies may exacerbate social divides and amplify existing inequalities, particularly if access to AI tools and benefits is not equitable across populations.
- Regulatory Challenges: Developing and implementing effective regulations for AI presents challenges due to the rapid pace of technological advancement, international differences in regulatory approaches, and the complexity of AI systems.
- Existential Risks: Some experts warn of long-term risks associated with advanced AI systems, including the potential for unintended consequences, loss of control, and existential threats to humanity.

Addressing these challenges requires collaboration among policymakers, technologists, ethicists, and other stakeholders to develop responsible AI frameworks that prioritize human well-being, fairness, transparency, and accountability. It also necessitates ongoing research and dialogue to anticipate and mitigate the negative effects of AI while maximizing its potential benefits for society.

1.8 Designing Safe and Ethical AI

Designing safe and ethical AI involves considering various factors throughout the development lifecycle.

By taking a responsible approach, companies will be able to

- create AI systems that are efficient and compatible with regulations;
- ensure that development processes consider all the ethical, legal, and societal implications of AI;
- track and mitigate bias in AI models;
- build trust in AI;
- prevent or minimize negative effects of AI; and
- Get rid of ambiguity about “whose fault it is” if something in AI goes wrong.

Here's a framework for creating AI systems that prioritize safety and ethics:

- Define Ethical Principles: Establish clear ethical principles that guide AI development and deployment. These principles should prioritize values such as fairness, transparency, accountability, privacy, and human well-being.
- Data Collection and Management: Ensure that data collection methods are transparent, lawful, and respectful of privacy rights. Implement robust data governance practices to mitigate bias, protect sensitive information, and maintain data integrity throughout the AI lifecycle.
- Bias Mitigation: Identify and mitigate biases in training data and algorithms to ensure fairness and prevent discriminatory outcomes. Employ techniques such as bias detection, data augmentation, and algorithmic auditing to address bias in AI systems.
- Transparency and Explainability: Design AI systems to be transparent and explainable, enabling users to understand how decisions are made and identify potential biases or errors. Use interpretable machine learning models and provide clear explanations of AI-driven decisions.
- Human Oversight and Control: Incorporate mechanisms for human oversight and control into AI systems, particularly in critical domains where human judgment is essential. Allow users to intervene in AI-driven decisions, provide feedback, and override automated processes when necessary.
- Security and Privacy: Prioritize security and privacy by implementing robust cybersecurity measures, data encryption techniques, and access controls to protect AI systems from unauthorized access, manipulation, and data breaches.
- Accountability and Responsibility: Establish mechanisms for accountability and responsibility in AI development and deployment. Clearly define roles and responsibilities for stakeholders, ensure compliance with ethical guidelines and regulatory requirements, and establish processes for addressing and rectifying AI-related harms.
- Continuous Monitoring and Evaluation: Monitor AI systems continuously to detect and address issues such as bias drift, performance degradation, and unexpected behavior. Conduct

regular evaluations and audits of AI systems to assess their impact on individuals, society, and the environment.

- Ethical Use Cases and Applications: Evaluate the ethical implications of AI use cases and applications before development begins. Consider potential risks and benefits, stakeholder perspectives, and societal impacts to ensure that AI systems are deployed responsibly and ethically.

Collaboration and Engagement: Foster collaboration and engagement with diverse stakeholders, including ethicists, policymakers, regulators, and community members, to promote shared understanding, identify ethical concerns, and co-create solutions for designing safe and ethical AI.

By integrating these principles and practices into the AI development process, organizations can design AI systems that prioritize safety, fairness, transparency, and ethical considerations, ultimately fostering trust and maximizing the positive impact of AI on individuals and society.

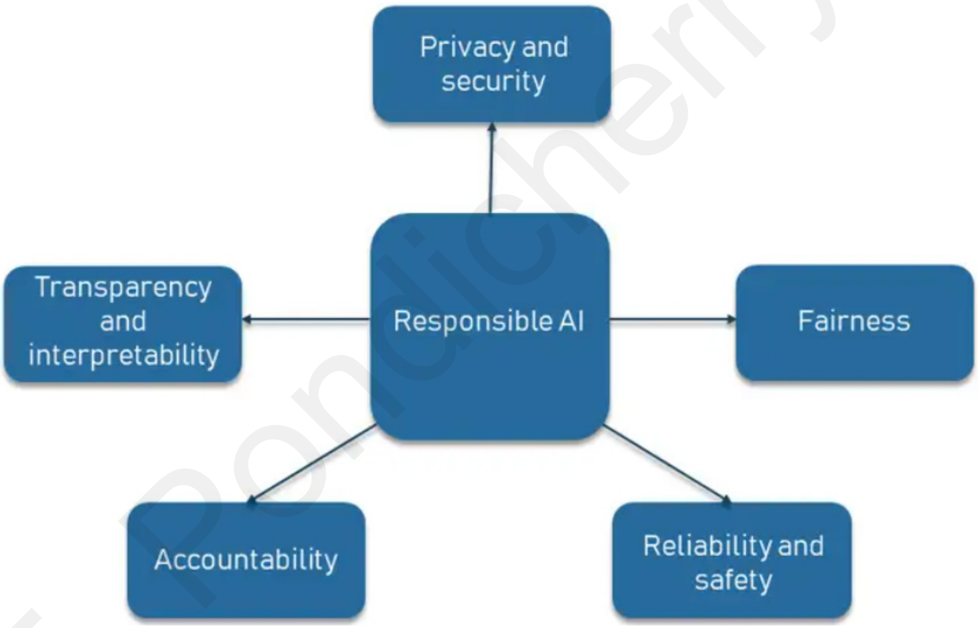

Figure 1.6

1.9 Self-Assessment Question

1. What is the difference between narrow AI and general AI, and can you provide an example of each?
2. What is the difference between supervised learning and unsupervised learning in machine learning?

3. How does the concept of the Machine Intelligence Continuum illustrate the progression from simple automation to advanced AI systems?
4. What are some potential benefits of AI in the healthcare sector?
5. What are some of the ethical concerns associated with the deployment of AI in law enforcement?

UNIT – II

Lesson 2.1 - Machine Learning

Machine Learning Overview: *Types of ML - Accuracy of ML models - Specific ML Methods: A Deep Dive - Model Selection and Validation*

Learning Objectives

1. To understand the concepts of supervised learning, unsupervised learning and reinforcement learning
2. To understand the accuracy of machine learning models
3. To understand linear and logistic regression, decision trees, SVM
4. To understand k-means clustering, Hierarchical clustering and Association rule learning
5. To understand cross-validation, hypermeter tuning, Model evaluation metrics

2.1 Definition of Machine Learning

Machine Learning is the subset of **Artificial Intelligence** which focuses on the creation of **algorithm** that enable computer to **independently** learn from **data** and **past experiences**.

Without being explicitly programmed, machine learning enables a machine to **automatically** learn from data, improve performance from experiences, and predict things. Machine Learning brings together statistics and computer science.

Programming Method

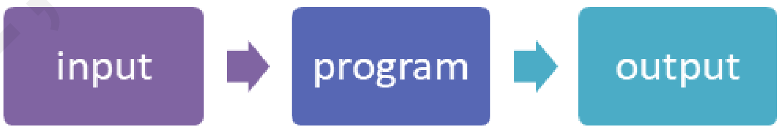

Figure 2.1

Here the user gives the input and writes the program and the computer generates the output.

Machine Learning Method

Figure 2.2

Here the user gives both the input and output and also the machine learning model, the system learns the model and creates the program.

Example: Addition of two numbers in both the cases.

Machine learning allows the user to feed a computer algorithm an immense amount of data and have the computer analyze and make decisions based on only the input data.

If any corrections are identified, the algorithm can incorporate that information to improve its future decision making.

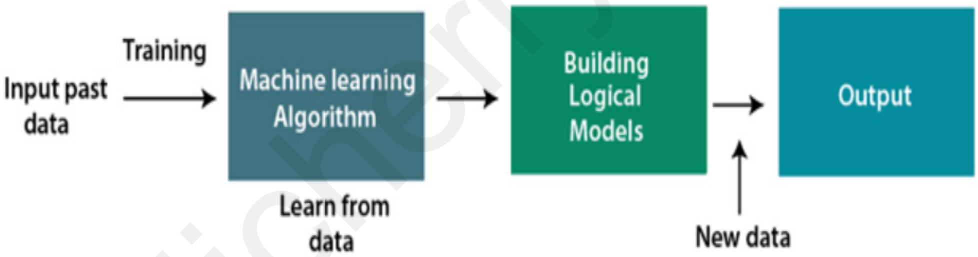

Figure 2.3

2.2 Features of Machine Learning:

1. Uses to detect various patterns in a data
2. Learns from past data and improves automatically
3. Data-driven technology
4. ML is similar to data mining

2.3 Applications of Machine Learning

Machine learning concepts are used almost everywhere, such as

1. Healthcare
2. Finance
3. Infrastructure

4. Marketing
5. Self-driving cars
6. Recommendation systems
7. Chatbots
8. Social sites
9. Gaming
10. Cyber security and many more.

There are certain terms and concepts that are required in Machine Learning. For example: A university has a collection of students and collection of courses. The main problem is to **predict** how much a particular student will like a particular course. With the given historic data or the past data the system may predict what course a student might take. So in simple words, the system predicts using some **examples.** But the system has not be given the idea of other courses that are available. So expecting the system to analyze it would be difficult. So the concept of Generalization is introduced.

Generalization is a term usually refers to a Machine Learning models ability to perform well on the new unseen data. After being trained on a training set, a model can digest new data and can able to make accurate predictions. The main success of the model is the ability of the model to generalize well. If the model has been trained too well on the training data, it will be difficult for the model to generalize.

In order to predict accurately the system has to be trained by giving the **training data** on which the algorithm is expected to learn. Based on this training data, our learning algorithm induces a function **f** that will map a new example to a corresponding prediction. In order to understand whether an algorithm works or predicts correctly, the **test data** (data that is used to test the algorithm) has to be induced in the system. The process of feeding the test data in to the system is called the **induction**.

The goal of inductive machine learning is to take some training data and use it to induce a function f. This function f will be evaluated on the test data. The machine learning algorithm has succeeded if its performance on the test data is high.

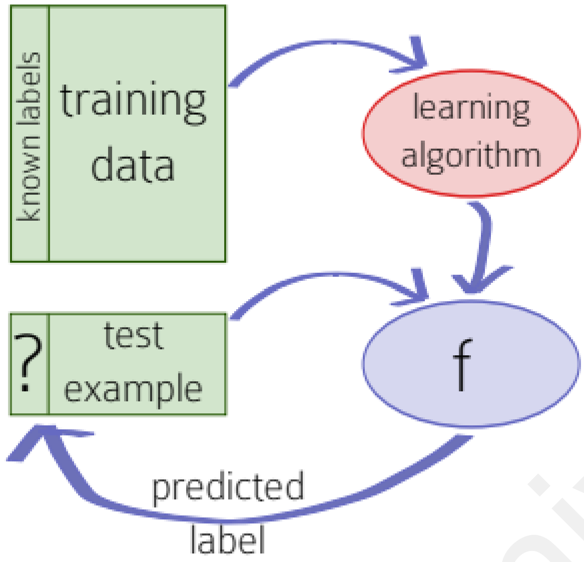

Figure 2.4

2.4 Types of Machine Learning

2.4.1 Supervised Learning

Supervised learning is a type of machine learning that uses labeled data to train machine learning models. In labeled data, the output is already known. The model just needs to map the inputs to the respective outputs.

An example of supervised learning is to train a system that identifies the image of an animal. Few of the top supervised learning applications are weather prediction, sales forecasting, stock price analysis.

The methods in supervised learning are:

Classification

Classification is a supervised machine learning process of categorizing a given set of input data into classes based on one or more variables.

Binary Classification: trying to predict a simple yes/no response that means a problem has only two possible outcomes. For instance, predict whether cat or dog, male or female etc.

Multiclass Classification: trying to put an example into one of a number of classes that means a problem has more than two outcomes. For instance, predict types of music, type of story etc.

Regression: Regression is a supervised machine learning technique which is used to predict continuous values. The goal of the algorithm is to plot a best-fit line or a curve between the data. The three main metrics that are used for evaluating the trained regression model are **variance, bias and error.**

Ranking: trying to put a set of objects in order of relevance for many applications to rank items. For instance, predicting what order to put web pages in, recommendation systems etc.

2.4.2 Unsupervised Learning

Unsupervised learning is a type of machine learning that uses unlabeled data to train machines. Unlabeled data doesn't have a fixed output variable. The model learns from the data, discovers the patterns and features in the data, and returns the output.

An example of an unsupervised learning technique that uses the images of vehicles to classify if it's a bus or a truck. The most popular algorithm used is clustering algorithms.

Few examples are customer segmentation, rate analysis.

The methods used in Unsupervised Learning are:

Clustering

Clustering or cluster analysis is a ML technique, which groups the unlabelled dataset. It can be defined as "A way of grouping the data points into different clusters, consisting of similar data points. The objects with the possible similarities remain in a group that has less or no similarities with another group."

Example: how fruits and vegetables are separated in a store

Association

Association are "if-then" statements that help to show the probability of relationships between data items, within large data sets in various types of databases.

Example a customer who buys bread tends to buy butter and jam too.

2.4.3 Reinforcement Learning

Reinforcement learning (RL) is a machine learning (ML) technique that trains software to make decisions to achieve the most optimal results.

It mimics the trial-and-error learning process that humans use to achieve their goals.

Example: The robotic dog, which automatically learns the movement of his arms

2.5 Accuracy of ML Models

Accuracy is one of the most intuitive performance measures in machine learning. It is a metric that quantifies the number of correct predictions made out of all predictions made. This measure is extremely straightforward in binary and multiclass classification problems, but it's important to understand its nuances and limitations.

Error rate, on the other hand, complements accuracy by quantifying the number of incorrect predictions. It is calculated by subtracting the accuracy from one and often expressed as a percentage. Both accuracy and error rate provide a quick snapshot of model performance, but they may not always give a complete picture, especially in cases where class distributions are imbalanced.

2.5.1 Separation of Training Data and Test Data

Machine Learning enables computers/machines to turn a huge amount of data into predictions. Train and test datasets are the two key concepts of machine learning, where the training dataset is used to fit the model, and the test dataset is used to evaluate the model.

Training Dataset

The training data is the biggest (in -size) subset of the original dataset, which is used to train or fit the machine learning model. Firstly, the training data is fed to the ML algorithms, which lets them learn how to make predictions for the given task.

The training data varies depending on whether we are using Supervised Learning or Unsupervised Learning Algorithms.

For Unsupervised learning, the training data contains unlabeled data points, i.e., inputs are not tagged with the corresponding outputs. For supervised learning, the training data contains labels in order to train the model and make predictions.

The type of training data that we provide to the model is highly responsible for the model's accuracy and prediction ability. It means that the better the quality of the training data, the better will be the performance of the model. Training data is approximately more than or equal to 60% of the total data for an ML project.

Test Dataset

The test dataset is another subset of original data, which is independent of the training dataset. However, it has some similar types of features and class probability distribution and uses it as a benchmark for model evaluation once the model training is completed. Test data is a well-organized dataset that contains data for each type of scenario for a given problem that the model would be facing when used in the real world. Usually, the test dataset is approximately 20-25% of the total original data for an ML project.

Splitting dataset into Train and Test set

Splitting the dataset into train and test sets is one of the important parts of data pre-processing, as by doing so, we can improve the performance of our model and hence give better predictability. It is important to split a dataset into two parts, i.e., train and test set.

In this way, we can easily evaluate the performance of our model. Such as, if it performs well with the training data, but does not perform well with the test dataset, then it is estimated that the model may be overfitted.

Difference between Training Data and Test Dataset

1. The main difference between training data and testing data is that training data is the subset of original data that is used to train the machine learning model, whereas testing data is used to check the accuracy of the model.
2. The training dataset is generally larger in size compared to the testing dataset. The general ratios of splitting train and test datasets are 80:20, 70:30, or 90:10.
3. Training data is well known to the model as it is used to train the model, whereas testing data is like unseen/new data to the model.

Once the model is trained enough with the relevant training data, it is tested with the test data. We can understand the whole process of training and testing in three steps:

1. **Feed**: Firstly, we need to train the model by feeding it with training input data.
2. **Define**: Now, training data is tagged with the corresponding outputs (in Supervised Learning), and the model transforms the training data into text vectors or a number of data features.
3. **Test**: In the last step, we test the model by feeding it with the test data/unseen dataset. This step ensures that the model is trained efficiently and can generalize well.

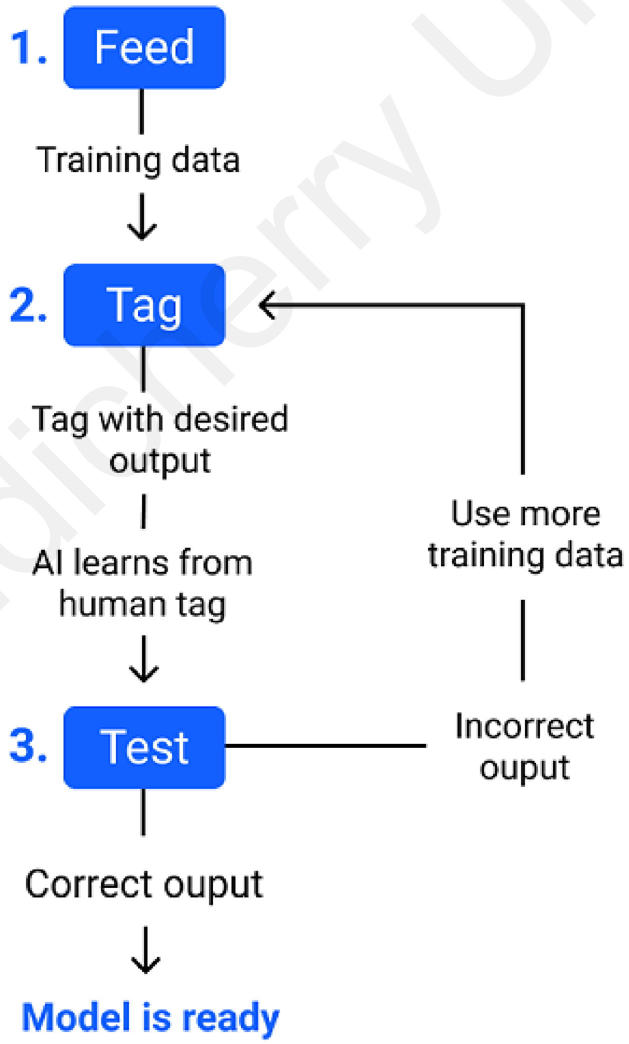

Figure 2.5

The cardinal rule of machine learning is: never touch your test data. Ever. Once you look at the test data, your model's performance on it is no longer indicative of its performance on future unseen data. This is simply because future data is unseen, but your "test" data no longer is.

2.5.2 Underfitting and Overfitting

Overfitting and Underfitting are the two main problems that occur in machine learning and degrade the performance of the machine learning models.

The main goal of each machine learning model is to generalize well. Here generalization defines the ability of an ML model to provide a suitable output by adapting the given set of unknown input.

Important metrics are:

- Signal: It refers to the true underlying pattern of the data that helps the machine learning model to learn from the data. (Opposite of noise)
- Noise: Noise is unnecessary and irrelevant data that reduces the performance of the model.
- Bias: Bias is a prediction error that is introduced in the model due to oversimplifying the machine learning algorithms. Or it is the difference between the predicted values and the actual values.
- Variance: If the machine learning model performs well with the training dataset, but does not perform well with the test dataset, then variance occurs.

Overfitting

Overfitting occurs when our machine learning model tries to cover all the data points or more than the required data points present in the given dataset. Because of this, the model starts caching noise and inaccurate values present in the dataset, and all these factors reduce the efficiency and accuracy of the model. The overfitted model has **low bias** and **high variance.**

The model tries to cover all the data points present in the scatter plot. It may look efficient, but in reality, it is not so. Because the goal of the regression model to find the **best fit line**, but here we have not got any best fit, so, it will generate the prediction errors. Overfitting can be avoided by cross-validation, removing features, regularization etc.

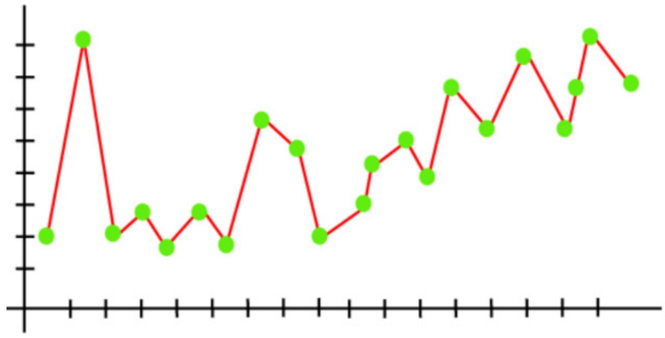

Figure 2.6

Underfitting

Underfitting occurs when our machine learning model is not able to capture the underlying trend of the data. In the case of underfitting, the model is not able to learn enough from the training data, and hence it reduces the accuracy and produces unreliable predictions.

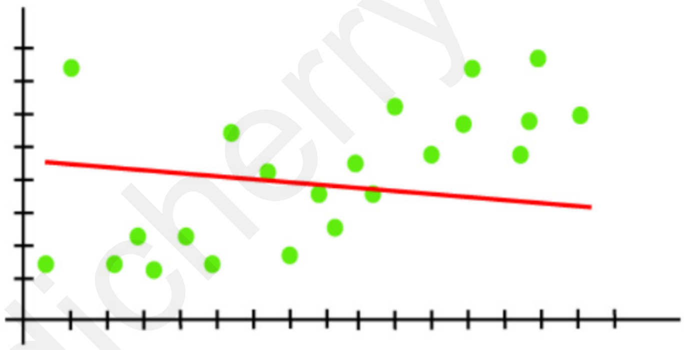

Figure 2.7

The model is unable to capture the data points present in the plot. Underfitting can be avoided by increasing the training time of the model and by increasing the number of features.

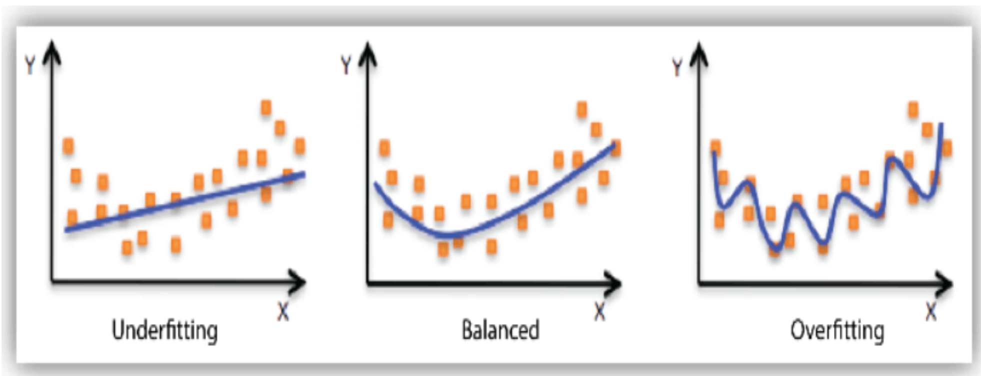

Figure 2.8

The "**Goodness of fit**" term is taken from the statistics, and the goal of the machine learning models to achieve the goodness of fit. The model with a good fit is between the underfitted and overfitted model, and ideally, it makes predictions with 0 errors, but in practice, it is difficult to achieve it.

2.5.3 Formalizing the Learning Problem

As you've seen, there are several issues that we must take into account when formalizing the notion of learning.

- The performance of the learning algorithm should be measured on unseen "test" data.
- The way in which we measure performance should depend on the problem we are trying to solve.
- There should be a strong relationship between the data that our algorithm sees at training time and the data it sees at test time.

In order to accomplish this, let's consider a loss function, $l(\cdot, \cdot)$, of two arguments. The job of l is to tell us how "bad" a system's prediction is in comparison to the truth. In particular, if y is the truth and $\hat{y}$ is the system's prediction, then $l(y, \hat{y})$ is a measure of error.

For three of the canonical tasks discussed, we might use the following loss functions:

Regression: **squared loss** $l(y, \hat{y}) = (y - \hat{y})^2$ or **absolute loss** $l(y, \hat{y}) = |y - \hat{y}|$.

Binary Classification: zero/one loss $l(y, \hat{y}) = 0$ if $y = \hat{y}$ or 1 otherwise.

This notation means that the loss is zero if the prediction is correct and is one otherwise.

Multiclass Classification: also zero/one loss.

2.5.4 Evaluating Model Performance

The main concept in evaluating model performance is, how well the model is performing, is it a useful model?, does it require any features?

There are two major types of binary classification problems. One is **"X versus Y."** For instance, **positive versus negative sentiment.** Another is **"X versus not-X."** This is a subtle and subjective decision.

But “X versus not X” problems often have more of the feel of “X spotting” rather than a true distinction between X and Y.

For spotting problems (X versus not-X), there are often more appropriate success metrics than accuracy. A very popular one from **information retrieval** is the **precision/recall metric**.

Precision means - of all the X’s that you found, how many of them were actually X’s?

Recall means - of all the X’s that were out there, how many of them did you find?

$$P = \frac{I}{S}$$
$$R = \frac{I}{T}$$
$$S = \text{number of Xs that your system found}$$
$$T = \text{number of Xs in the data}$$
$$I = \text{number of correct Xs that your system found}$$

Formally, precision and recall are defined as:

Here, S - “System,” T - “Truth” and I - “Intersection.” It is generally accepted that 0/0 = 1 in these definitions. Thus, if your system found nothing, your precision is always perfect; and if there is nothing to find, your recall is always perfect.

Example

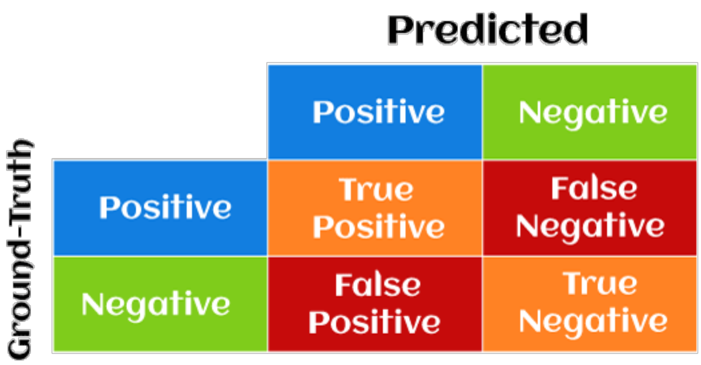

Figure 2.9

The possible combinations are:

- True Positive: This combination tells us how many times a model correctly classifies a positive sample as Positive?
- False Negative: This combination tells us how many times a model incorrectly classifies a positive sample as Negative?
- False Positive: This combination tells us how many times a model incorrectly classifies a negative sample as Positive?
- True Negative: This combination tells us how many times a model correctly classifies a negative sample as Negative?

Once precision and recall is computed, we can produce precision/ recall curves.

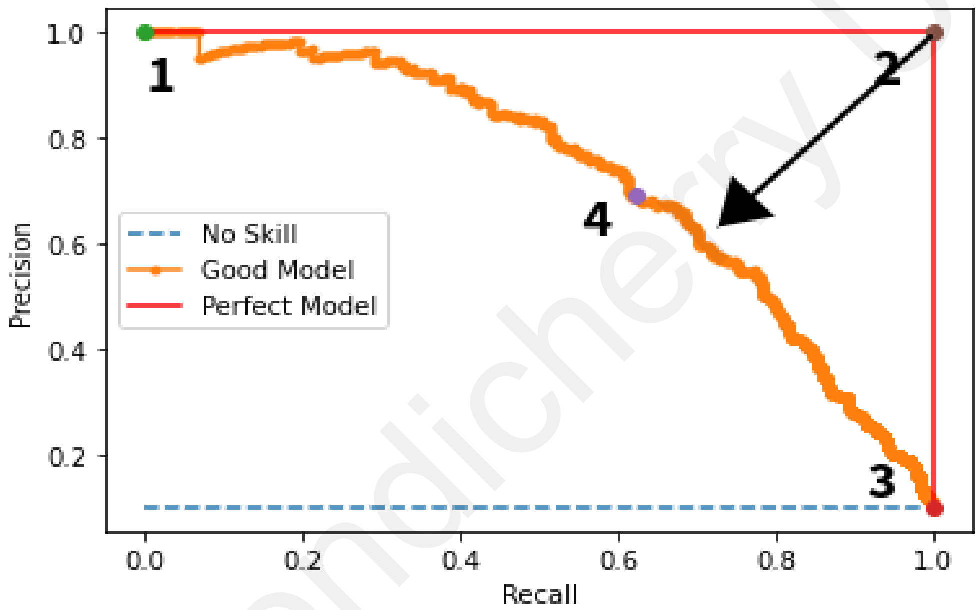

Figure 2.10

2.5.5 ROC Curve and AUC

ROC curve

An **ROC curve** (**receiver operating characteristic curve**) is a graph showing the performance of a classification model at all classification thresholds. This curve plots two parameters:

- True Positive Rate
- False Positive Rate

True Positive Rate (**TPR**) is a synonym for recall and is therefore defined as follows:

$$TPR = TP/TP + FN$$

False Positive Rate (**FPR**) is defined as follows:

$$FPR = FP/FP + TN$$

An ROC curve plots TPR vs. FPR at different classification thresholds. Lowering the classification threshold classifies more items as positive, thus increasing both False Positives and True Positives. The following figure shows a typical ROC curve.

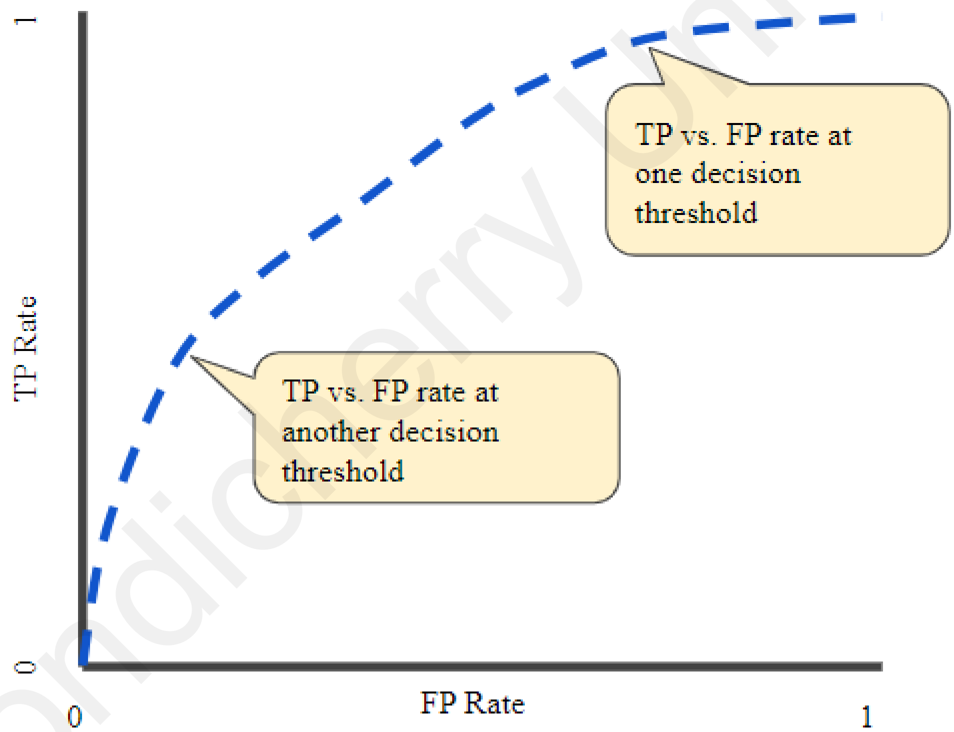

Figure 2.11

To compute the points in an ROC curve, we could evaluate a logistic regression model many times with different classification thresholds, but this would be inefficient. Fortunately, there's an efficient, sorting-based algorithm that can provide this information for us, called AUC.

AUC: Area under the ROC Curve

AUC stands for «Area under the ROC Curve.» That is, AUC measures the entire two-dimensional area underneath the entire ROC curve (think integral calculus) from (0,0) to (1,1).

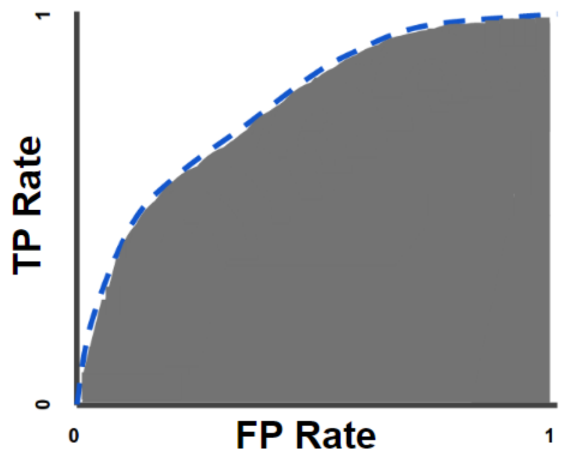

Figure 2.12

AUC provides an aggregate measure of performance across all possible classification thresholds. One way of interpreting AUC is as the probability that the model ranks a random positive example more highly than a random negative example. For example, given the following examples, which are arranged from left to right in ascending order of logistic regression predictions:

2.6 Specific ML Methods: A Deep Dive

2.6.1 Supervised Learning Methods

Decision Tree is a supervised learning technique that can be used for both **classification** and **Regression** problems, but mostly it is preferred for solving Classification problems. It is a tree-structured classifier, where **internal nodes** represent the **features of a dataset**, **branches** represent the **decision rules** and each **leaf node** represents the **outcome**. This algorithm works on the principal of Divide and conquer technique.

Decision Tree Terminologies

- Root Node: The initial node at the beginning of a decision tree, where the entire population or dataset starts dividing based on various features or conditions.
- Decision Node – They are used to make decisions and have multiple branches. This follows the divide technique.
- Leaf Node – They are the output of the decision node and does not have any branches. This follow the conquer technique.

- Sub-Tree – It is referred to the sub section of a decision tree
- Pruning - The process of removing or cutting down specific nodes in a decision tree to prevent overfitting and simplify the model.
- Parent and Child Node - A node that is divided into sub-nodes is known as a parent node, and the sub-nodes emerging from it are referred to as child nodes. The parent node represents a decision or condition, while the child nodes represent the potential outcomes or further decisions based on that condition.

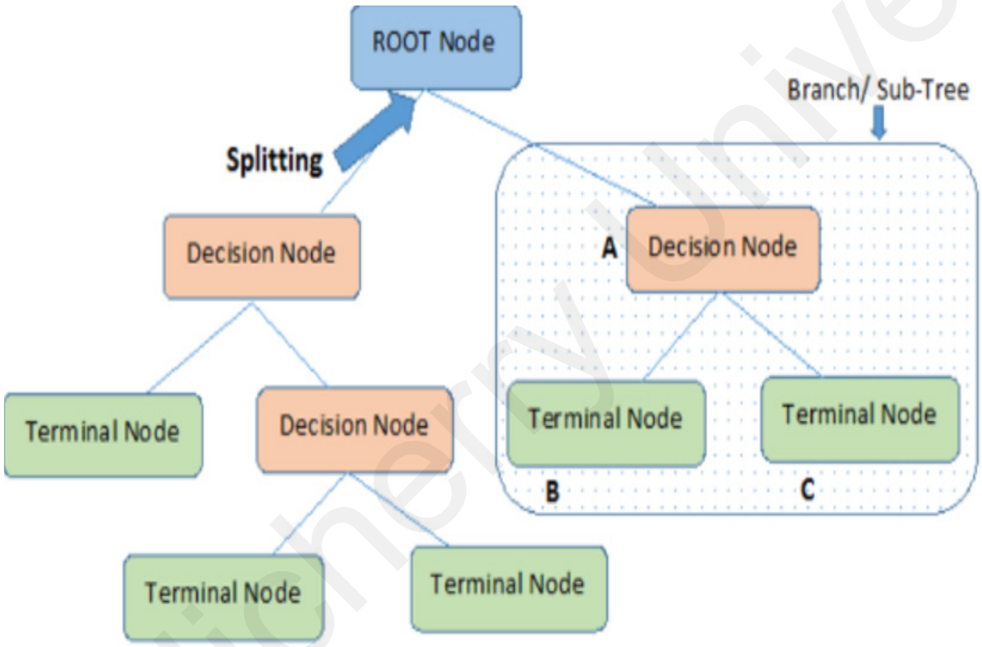

Figure 2.13

The decision tree is so-called because we can write our set of questions and guesses in a tree format. The questions are written in the internal tree nodes (rectangles) and the guesses are written in the leaves (ovals). Each non-terminal node has two children: the left child specifies what to do if the answer to the question is "no" and the right child specifies what to do if it is "yes."

In the decision tree, the **questions** that are asked is referred to as **features** and the **responses** are referred to as **features values**. The **rating** is called as **label**.

Example of a Decision Tree

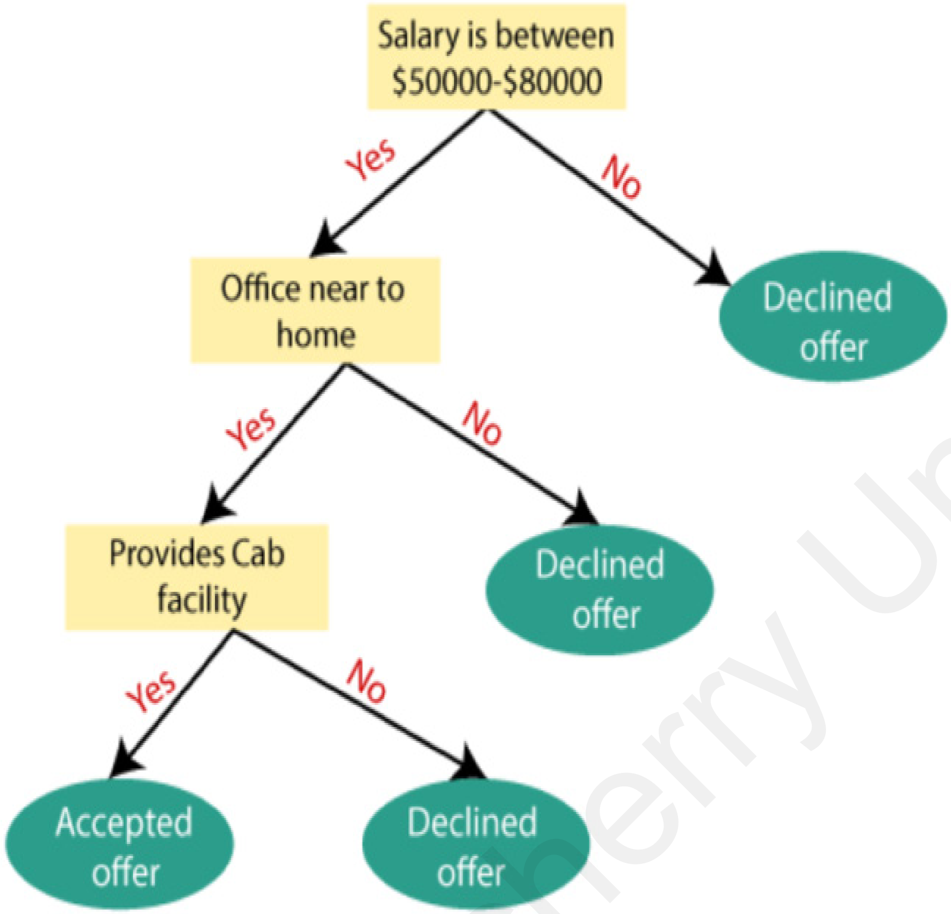

Figure 2.14

K-Nearest Neighbors

- K-Nearest Neighbor is one of the simplest Machine Learning algorithms based on Supervised Learning technique.
- K-NN algorithm assumes the similarity between the new case/data and available cases and put the new case into the category that is most similar to the available categories.
- K-NN algorithm stores all the available data and classifies a new data point based on the similarity. This means when new data appears then it can be easily classified into a well suite category by using K- NN algorithm.
- K-NN algorithm can be used for Regression as well as for Classification but mostly it is used for the Classification problems.

Example

Figure 2.15

Why do we Need a K-NN Algorithm?

Suppose there are two categories, i.e., Category A and Category B, and we have a new data point x1, so this data point will lie in which of these categories. To solve this type of problem, we need a K-NN algorithm. With the help of K-NN, we can easily identify the category or class of a particular dataset. Consider the below diagram:

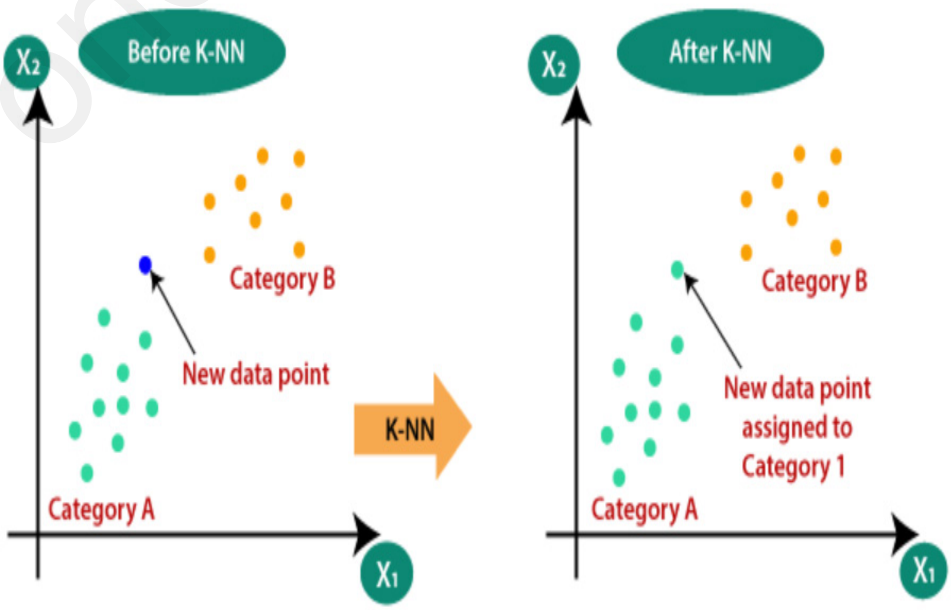

Figure 2.16

How does KNN work?

The K-NN working can be explained on the basis of the below algorithm:

- Step-1: Select the number K of the neighbors
- Step-2: Calculate the Euclidean distance of K number of neighbors
- Step-3: Take the K nearest neighbors as per the calculated Euclidean distance.
- Step-4: Among these k neighbors, count the number of the data points in each category.
- Step-5: Assign the new data points to that category for which the number of the neighbor is maximum.
- Step-6: Our model is ready.

There is no particular way to determine the best value for "K", so we need to try some values to find the best out of them. The most preferred value for K is 5. A very low value for K such as K=1 or K=2, can be noisy and lead to the effects of outliers in the model. Large values for K are good, but it may find some difficulties.

Advantages of KNN Algorithm:

- It is simple to implement.
- It is robust to the noisy training data
- It can be more effective if the training data is large.

Disadvantages of KNN Algorithm:

- Always needs to determine the value of K which may be complex some time.
- The computation cost is high because of calculating the distance between the data points for all the training samples.

2.6.2 Unsupervised Learning Methods

2.6.2.1 K-Means Clustering

K-Means Clustering is an unsupervised learning algorithm that is used to solve the clustering problems in machine learning or data science. It groups the unlabeled dataset into different clusters. Here K defines the number of pre-defined clusters that need to be created in the process, as if K=2, there will be two clusters, and for K=3, there will be three clusters, and so on.

It allows us to cluster the data into different groups and a convenient way to discover the categories of groups in the unlabeled dataset on its own without the need for any training.

It is a centroid-based algorithm, where each cluster is associated with a centroid. The main aim of this algorithm is to minimize the sum of distances between the data point and their corresponding clusters.

The algorithm takes the unlabeled dataset as input, divides the dataset into k-number of clusters, and repeats the process until it does not find the best clusters. The value of k should be predetermined in this algorithm.

The k-means clustering algorithm mainly performs two tasks:

- Determines the best value for K center points or centroids by an iterative process.

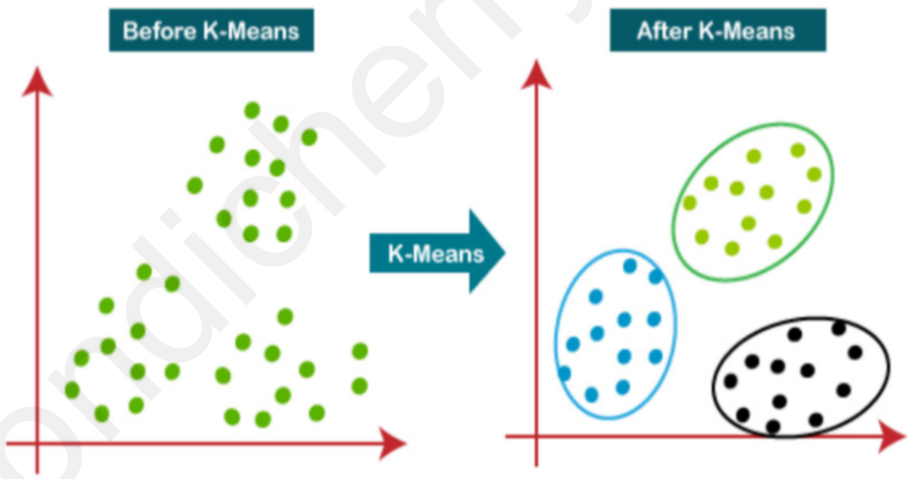

Figure 2.17

- Assigns each data point to its closest k-center. Those data points which are near to the particular k-center, create a cluster.

K-Means Algorithm

Step-1: Select the number K to decide the number of clusters.

Step-2: Select random K points or centroids. (It can be other from the input dataset).

Step-3: Assign each data point to their closest centroid, which will form the predefined K clusters.

Step-4: Calculate the variance and place a new centroid of each cluster.

Step-5: Repeat the third steps, which means reassign each datapoint to the new closest centroid of each cluster.

Step-6: If any reassignment occurs, then go to step-4 else go to FINISH.

Step-7: The model is ready.

Example to be written

Advantages

- Simple and easy to implement
- Fast and efficient
- Scalability – can be used for large datasets
- Flexibility

Disadvantages

- Sensitivity to initial centroids
- Requires specifying the number of clusters
- Sensitive to outliers – data point totally different from others

Choosing the value of K using Elbow Method

The Elbow method is one of the most popular ways to find the optimal number of clusters. This method uses the concept of WCSS value. **WCSS** stands for **Within Cluster Sum of Squares**, which defines the total variations within a cluster.

The formula is for 2 clusters

$$WCSS = \sum_{P_i \text{ in Cluster1}} distance(P_i\ C_1)^2 + \sum_{P_i \text{ in Cluster2}} distance(P_i\ C_2)^2$$

In the above formula of WCSS,

$\sum_{P_i \text{ in Cluster1}} distance(P_i\ C_1)^2$: It is the sum of the square of the distances between each data point and its centroid within a cluster1 and the same for the other term also. To measure the distance between data points and centroid, we can use Euclidean distance.

To find the optimal value of clusters, the elbow method follows the below steps:

- It executes the K-means clustering on a given dataset for different K values (ranges from 1-10).
- For each value of K, calculates the WCSS value.
- Plots a curve between calculated WCSS values and the number of clusters K.
- The sharp point of bend or a point of the plot looks like an arm, then that point is considered as the best value of K.

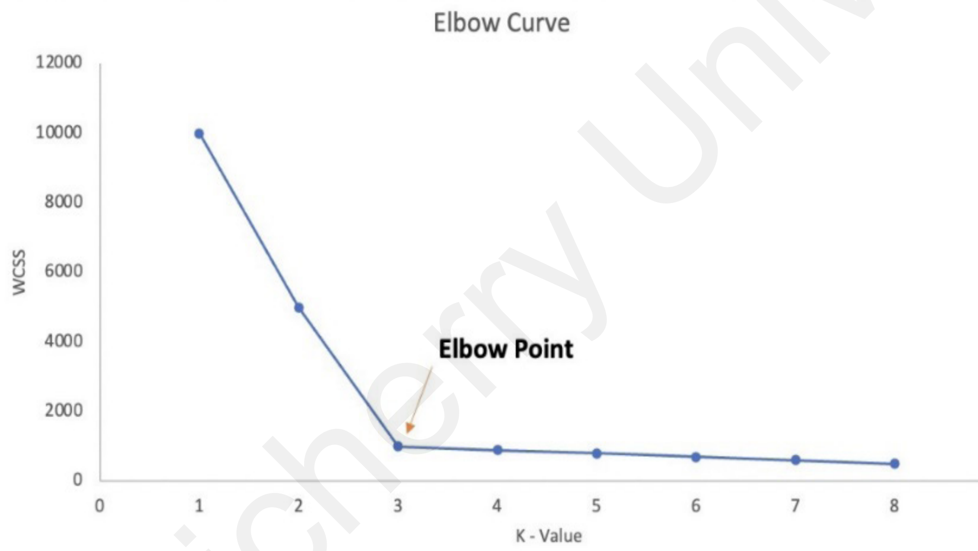

Figure 2.18

2.6.2.2 Association Rules

Association rule learning is a type of unsupervised learning technique that checks for the dependency of one data item on another data item and maps accordingly so that it can be more profitable. It tries to find some interesting relations or associations among the variables of dataset. It is based on different rules to discover the interesting relations between variables in the database.

The association rule learning is one of the very important concepts of machine learning, and it is employed in **Market Basket analysis, Web usage mining, continuous production, etc.** Here market basket analysis is a technique used by the various big retailer to discover the associations between items. We can understand it by taking an example of a supermarket, as in a supermarket, all products that are purchased together are put together.

For example, if a customer buys bread, he most likely can also buy butter, eggs, or milk, so these products are stored within a shelf or mostly nearby.

Association rule learning can be divided into three types of algorithms:

1. Apriori
2. Eclat
3. F-P Growth Algorithm

Working of Association Rule

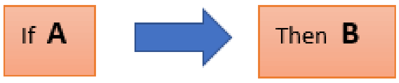

Association rule learning works on the concept of If and Else Statement, such as if A then B.

Here the If element is called **antecedent**, and then statement is called as **Consequent**. These types of relationships where we can find out some association or relation between two items is known *as* **single cardinality**. It is all about creating rules, and if the number of items increases, then cardinality also increases accordingly. So, to measure the associations between thousands of data items, there are several metrics. These metrics are given below:

- Support
- Confidence
- Lift

Support

$$\text{Supp}(X) = \frac{Freq(X)}{T}$$

Support is the frequency of A or how frequently an item appears in the dataset. It is defined as the fraction of the transaction T that contains the itemset X. If there are X datasets, then for transactions T, it can be written as:

Confidence

Confidence indicates how often the rule has been found to be true. Or how often the items X and Y occur together in the dataset when the occurrence of X is already given. It is the ratio of the transaction that contains X and Y to the number of records that contain X.

$$\text{Confidence}=\frac{Freq(X,Y)}{Freq(X)}$$

Lift

It is the ratio of the observed support measure and expected support if X and Y are independent of each other. It has three possible values:

$$\text{Lift}=\frac{Supp(X,Y)}{Supp(X)\times Supp(Y)}$$

- If Lift = 1: The probability of occurrence of antecedent and consequent is independent of each other.
- Lift>1: It determines the degree to which the two itemsets are dependent to each other.
- Lift<1: It tells us that one item is a substitute for other items, which means one item has a negative effect on another.

Applications of Association Rule Learning

It has various applications in machine learning and data mining. Below are some popular applications of association rule learning:

- **Market Basket Analysis:** This technique is commonly used by big retailers to determine the association between items.
- **Medical Diagnosis:** With the help of association rules, it helps in identifying the probability of illness for a particular disease.
- **Protein Sequence:** The association rules help in determining the synthesis of artificial Proteins.
- It is also used for the **Catalog Design** and **Loss-leader Analysis (launching new product with low cost)** and many more other applications.

2.7 Model Selection and Validation

In machine learning, model selection is the process of choosing the best performing model for a specific problem, while model validation is the process of ensuring that the model produces adequate results for its data.

2.7.1 Cross Validation

Cross-validation is a technique for validating the model efficiency by training it on the subset of input data and testing on previously unseen subset of the input data.

In machine learning, there is always the need to test the stability of the model. It means, based only on the training dataset; we can't fit our model on the training dataset. For this purpose, we reserve a particular sample of the dataset, which was not part of the training dataset. After that, we test our model on that sample before deployment, and this complete process comes under cross-validation.

Basic Steps of Cross-Validation

- Reserve a subset of the dataset as a validation set.
- Provide the training to the model using the training dataset.
- Now, evaluate model performance using the validation set. If the model performs well with the validation set, perform the further step, else check for the issues.

Methods used for Cross-Validation

There are some common methods that are used for cross-validation. These methods are given below:

1. **Validation Set Approach**

   In this approach we divide input dataset into a training set and test or validation. Both the subsets are given 50% of the dataset.

   But it has one of the big disadvantages that we are just using a 50% dataset to train our model, so the model may miss out to capture important information of the dataset. It also tends to give the underfitted model.

2. **Leave-P-out cross-validation**

   In this approach, the p datasets are left out of the training data. It means, if there are total n datapoints in the original input dataset, then n-p data points will be used as the training dataset and the p data points as the validation set. This technique is difficult for large set of p.

3. **Leave one out cross-validation**

   This method is similar to the leave-p-out cross-validation, but instead of p, we need to take 1 dataset out of training. It means, for each learning set, only one datapoint is reserved, and the remaining dataset is used to train the model. This process repeats for each datapoint. Hence for n samples, we get n different training set and n test set, so the bias is minimum, process is executed n times and leads to high variation.

4. **K-fold cross-validation**

   K-fold cross-validation approach divides the input dataset into K groups of samples of equal sizes. These samples are called **folds**. For each learning set, the prediction function uses k-1 folds, and the rest of the folds are used for the test set.

   The steps for k-fold cross-validation are:

   - Split the input dataset into K groups
   - For each group:
     - Take one group as the reserve or test data set.
     - Use remaining groups as the training dataset
     - Fit the model on the training set and evaluate the performance of the model using the test set.

5. **Stratified k-fold cross-validation**

   This technique is similar to k-fold cross-validation with some little changes. This approach works on stratification concept, it is a process of rearranging the data to ensure that each fold or group is a good representative of the complete dataset. To deal with the bias and variance, it is one of the best approaches.

2.7.2 Hypothesis Testing and Statistical Significance

The process of hypothesis testing is to draw inferences or some conclusion about the overall data by conducting some statistical tests on a sample. The same inferences are drawn for different machine learning models.

For drawing some inferences, we have to make some assumptions that lead to two terms that are used in the hypothesis testing.

- Null hypothesis: It is regarding the assumption that there is no anomaly pattern or believing according to the assumption made.
- Alternate hypothesis: Contrary to the null hypothesis, it shows that observation is the result of real effect.

The steps involved in the hypothesis testing are as follow:

- Assume a null hypothesis, usually in machine learning algorithms we consider that there is no anomaly between the target and independent variable.
- Collect a sample
- Calculate test statistics
- Decide either to accept or reject the null hypothesis

P-Value

P-value helps us determine how likely it is to get a particular result when the null hypothesis is assumed to be true. The concept of p-value comes from statistics and widely used in machine learning and data science.

The p-value of 0.05 is known as the level of significance (**α**). Usually, it is considered using two suggestions, which are given below:

- If p-value>0.05: The large p-value shows that the null hypothesis needs to be accepted.
- If p-value<0.05: The small p-value shows that the null hypothesis needs to be rejected, and the result is declared as statically significant.

T-test

A t-test is a type of inferential statistic used to determine if there is a significant difference between the means of two groups, which may be related in certain features.

- If t-value is large => the two groups belong to different groups.
- If t-value is small => the two groups belong to same group.

Bootstrapping

Bootstrapping is one of the many methods and techniques that data scientists use. Particularly useful for assessing the quality of a machine learning model, bootstrapping is a method of inferring results for a population from results found on a collection of smaller random samples of the population, using replacement during the sampling process.

2.8 Self-Assessment Question

1. How does reinforcement learning differ from supervised learning and unsupervised learning in terms of learning objectives and feedback mechanisms?
2. What are some common metrics used to assess the accuracy of machine learning models, and how would you choose the appropriate metric for a classification problem?
3. How do decision trees and Support Vector Machines (SVM) differ in their approach to classification, and when might you prefer one over the other?
4. What are the key differences between K-means clustering and hierarchical clustering, and in what scenarios might each method be more advantageous?
5. Why is cross-validation important in machine learning, and how does it help in hyperparameter tuning?

UNIT – III

Lesson 3.1 - Developing an Enterprise AI Strategy

Developing an enterprise AI strategy*: Invest in technical Talent, Plan Implementation, Collet and prepare data, Building Machine Learning Model, Experiment and Iterate*

Learning Objectives

1. To understand identifying key roles and fostering a culture of innovation in enterprise.
2. To understand the developing of a roadmap and establishing governance and oversight
3. To data sources, quality assessment and preprocessing
4. To understand building a Machine Learning Model from selecting appropriate algorithm and feature engineering.
5. To understand embracing agility and conducting experiments

3 Introduction

Artificial Intelligence helps to improve organization and community, but translating those ideas into viable software requires having the right mindset, dedicated leadership, and a diverse support team.

Despite many public claims to innovation, many corporations are still playing catch up on existing technologies such as big data, mobile, and the Internet of Things (IoT). Many brands have built up their social media presence and now offer mobile-friendly apps and websites, but these are merely digital consumer endpoints, not the basis for an enterprise-wide technological transformation.

A key milestone in the corporate digital transformation is the development of a centralized data and technology infrastructure. These two elements connect consumer applications, enterprise systems, and third-party partners and provide access to a single source of truth that contains relevant, up-to date, and accurate information for all parties.

Designing and implementing the infrastructure needed for enterprise-scale AI requires a strong and dedicated technology team that can develop internal application programming interfaces (APIs) to standardize access to both data and company's internal business technology. Doing so will enable the company to streamline enterprise-wide data analysis, accelerate product development, and respond more quickly in evolving markets. Internal APIs will also reduce the communication overhead needed to hunt down specific data, negotiate access, and interpret variations.

3.1 Invest in Technical Talent

The foremost challenge to develop an enterprise AI strategy is to find people who have the necessary technical skills to staff the organization's initiative. Finding the right people can be no less of a challenge.

The high demand for specialized AI talent, coupled with the painfully low supply, means that companies need to adopt new strategies when recruiting for a new AI initiative.

3.1.1 Understanding different job titles.

Many companies struggle to understand what "artificial intelligence" is, much less the myriad of titles, roles, skills, and technologies used to describe a prospective hire. Titles and descriptions vary from company to company, and terms are not well-standardized in the industry.

Few important roles are mentioned here:

3.1.1.1 Data Science Team Manager

A data science team manager understands how best to deploy the expertise of his team in order to maximize their productivity on a project. This manager should have sufficient technical knowledge to understand what his team members are doing and how best to support them; at the same time, this manager must also have good communications skills to liaise with the leadership or non-technical units.

3.1.1.2 Machine Learning Engineers

ML engineers build machine learning solutions to solve business and customer problems. These specialized engineers deploy models, manage infrastructure, and run operations related to machine learning projects. They are assisted by data scientists and data engineers to manage databases and build the data infrastructure necessary to support the products and services used by their customers.

3.1.1.3 Data Scientists

Data scientists typically work in an offline setting and do not deal directly with the production experience, which is what the end user would see. Data scientists collect data, spend most of their time cleaning it, and the rest of their time looking for patterns in the data and building predictive models.

3.1.1.4 Researchers, Research Scientists

Researchers are more focused on driving scientific discovery and less concerned with pursuing industrial applications of their findings. They often build on promising leads uncovered by data scientists and experiment with novel approaches, much of which originates from or is inspired by work done in academic or industry research facilities.

3.1.1.5 Applied Research Scientists, Applied Research Engineers

Applied researchers connect research and engineering. Unlike pure researchers, they are more concerned with practical research, such as identifying and implementing workable solutions to a specific problem or formulating industrial applications for scientific discoveries.

3.1.1.6 Data Engineers, Distributed Systems Engineers

Given the vast amounts of data and computation power required, most ML models face scalability issues. A talented infrastructure engineer can resolve challenges associated with large datasets, allowing researchers and data scientists to focus on their models rather than on data issues.

3.1.2 Characteristics needed for the special talents.

The skills required for successful careers in machine learning are different from those in traditional software development. Software development often has clearly structured tasks with well-defined deadlines for delivery and release.

By contrast, machine learning is highly exploratory and experimental, with less clear timelines and success metrics. Ideal performance targets may not be knowable in advance and may shift during a project. Algorithms require ongoing support, training, and feedback to perform optimally.

3.1.2.1 Mathematical Aptitude

A background in mathematics and statistics is far more valued in machine learning than in traditional software engineering. Training ML models requires a sufficient background to understand which algorithms to apply and how to interpret and improve upon the results. For cutting-edge AI research positions, advanced mathematical intuition is a prerequisite in order to design and develop novel methodologies.

3.1.2.2 Curiosity

Training ML algorithms requires a constant sense of curiosity. The model builder needs to take in abstract information and make sense of it through continuous experimentation.

3.1.2.3 Creativity

As ML tools and methodologies are still relatively new, the ability to think through ideas and to come up with novel ways to tackle a problem is highly valued. There will inevitably be many challenges that require new perspectives and solutions.

3.1.2.4 Perseverance

Artificial intelligence research is an ever-evolving pursuit. There are few simple answers, and it may easily take months to successfully train a viable algorithm. A successful individual will continue to try new techniques in the face of repeated failures until a solution can be found.

3.1.2.5 Rapid Learning

AI is evolving rapidly and keeping up to date with the accomplishments in the field is critical. Successful candidates should be able to stay on top of the latest technical developments and to apply quickly and intelligently what they have learned.

3.2 Plan Implementation

Once it is assessed that an organization has the requisite culture, leadership, and talent to succeed in AI initiatives, the next step is to identify business opportunities with the highest return on investment (ROI).

3.2.1 Ranking Business Goals

Prior to beginning any technology investment, the executive team must be clear on the problems that is to be tackled, the reasons why solving these problems is a priority for the organization, and the metrics for success. One should have clear strategic goals at either the company- or department-level. Common goals include increasing revenue, cutting costs, and entering new business lines. Artificial intelligence can advance many of these goals, but implementation difficulty and impact will vary.

Several frameworks for evaluating current enterprise workflows and technologies, performing opportunity analysis, and clarifying organizational goals and metrics are followed.

3.2.1.1 Gap Analysis

Gap analysis is used to assess where one's business is versus where it would like it to be. The methodology relies on benchmarking, critical analysis, and action planning.

3.2.1.2 Goal and Objectives Setting

The first step to performing a gap analysis is to create clear goals. The objectives selected can vary between companies, organizations, products, processes, etc. The key is to articulate useful goals that have clear objectives with appropriate, well-defined metrics for success.

3.2.1.3 Benchmarking

Benchmarking helps to understand where the organization stands. It is checked by comparing current performance with the performance of organizations that have already implemented automation. If it's difficult to get department-level metrics for competitors, consider hiring experts who have worked extensively with many companies to help establish benchmarks.

3.2.1.4 Gap Identification

After formulating goals and benchmarking current performance against others, list features associated with each objective that has to be achieved. Further break down goals into their constituent parts. Compare where organization stands and identify the gaps between current situation and goals. Where are the biggest gaps? Complete this step for each goal that is analyzed.

3.2.1.5 Action Planning

Once gaps have been identified, create plan of action to address deficient areas. What steps will be needed to take to achieve them? How can one use automation to close this gap? The next step is to create a project plan that identifies how to fill those gaps. SWOT Analysis approach is best suited.

SWOT ANALYSIS

SWOT stands for Strengths, Weaknesses, Opportunities, and Threats. This popular business framework can be used to evaluate AI opportunities. Using this approach uncovers opportunities as well as potential weaknesses that is needed to mitigate.

In SWOT analysis, first evaluate the internal factors affecting the business. What are the strengths? What makes the department so great? Which projects or teams are finding success? Next consider business's weaknesses. Which projects or departments are unprofitable? What resources do one lack? What can be done better?

The second half of SWOT analysis considers the external factors, such as opportunities and threats, in the marketplace. What are the business goals? How can new technologies such as AI drive enterprises forward? Are there new audiences being targeted? Finally, list threats that may derail a company, department, or project. What obstacles does one face? Who is the primary competitor? Can anything be done to prevent or minimize potential threats?

3.3 Collet and Prepare Data

3.3.1 Data is not Reality

Data is a human invention. Humans define the phenomenon that they want to measure, design systems to collect data about it, clean and pre-process it before analysis, and finally choose how to interpret the results. Even with the same dataset, two people can form vastly different conclusions. This is because data alone is not "ground truth," which is defined by machine learning experts as observable, provable, and objective data that reflects reality. If data was inferred from other information, relies on subjective judgment, collected in a slipshod manner, or is of questionable authenticity, then it is not ground truth.

How one choose to conceptualize a phenomenon, determine what to measure, and decide how to take measurements will impact the data that is collected. The ability to solve a problem with artificial intelligence depends heavily on how one frame the problem and also whether one can establish ground truth without ambiguity. Ground truth is used as a benchmark to assess the performance of algorithms.

Unless one is directly involved with defining and monitoring original data collection goals, instruments, and strategy, one is likely missing critical knowledge that may result in incorrect processing, interpretation, and use of that data.

3.3.2 Common Mistakes with Data

What people call "data" can be carefully curated measurements selected purely to support an agenda, haphazard collections of random information with no correspondence to reality, or information that looks reasonable but resulted from unconsciously biased collection efforts. The familiar errors are:

3.3.2.1 Undefined Goals

Failing to pin down the reason for collecting data means that one will miss the opportunity to articulate assumptions and to determine what to collect. The result is likely to collect the wrong data or incomplete data. A common trend in big data is for enterprises to gather heaps of information without any understanding of why they need it and how they want to use it. Gathering huge but messy volumes of data will only impede future analytics.

3.3.2.2 Definition Error

How is a customer defined? Depending on the goals, one might not want to lump everyone into one bucket. One may want to segment customers by their purchasing behavior in order to adjust the marketing efforts or product features more effectively. If that's the case, then it is needed to be sure that useful information is included about the customer, such as demographic information or spending history.

3.3.2.3 Capture Error

Once the type of data that is to be to collected, a mechanism is decided to capture it. Mistakes here can result in capturing incorrect or accidentally biased data. Let's say one want to test whether product A is more compelling than product B, but always product A is displayed first on the website. Because users do not see or purchase product B as frequently, the results of the test will lead to the wrong conclusion.

3.3.2.4 Measurement Error

Measurement errors occur when the software or hardware that is used to capture data goes awry, either failing to capture usable data or producing spurious data. For example, information about user behavior on the mobile app may be lost if usage logs are not synchronized with the servers due to connectivity issues. Similarly, when a microphone is used, the audio recordings may capture background noise or interference from other electrical signals.

3.3.2.5 Processing Error

Many enterprises own data that is decades-old, and the original team capable of explaining important decisions surrounding data collection or data storage may be long gone. Many of their assumptions and issues are likely not documented and will be up to the present team to deduce, which can be a daunting task.

The team may achieve wildly different results by making assumptions that differ from the original ones made during data collection. Common errors include missing a particular filter that may have been used on the data, such as the removal of outliers; using different accounting standards, as in the case with financial reporting; and simply making calculation errors.

3.3.2.6 Coverage Error

Coverage error describes what happens with survey data when there is insufficient opportunity for all targeted respondents to participate. For example, if you are collecting data on the elderly but only offer a website survey, you'll probably miss out on many respondents. In the case of digital products, your marketing teams may be interested in projecting how all mobile smartphone users might behave with a prospective product. However, if you only offer an iOS app and not an Android app, the iOS user data will give you limited insight into how Android users may behave.

3.3.2.7 Sampling Error

Sampling errors occur when data is analyzed from a smaller sample that is not representative of the target population. This is unavoidable when data only exists for some groups within a population. The conclusions that is drawn from the unrepresentative sample will probably not apply to the whole. Asking only few friends for opinions about products and then assuming the user population will feel similarly is a classic sampling error.

3.3.2.8 Inference Error

Inference errors are made by statistical or machine learning models when they make incorrect predictions from the available ground truth. Two types of inference errors can occur: false negatives and false positives. False positives occur when one incorrectly predict that an item belongs in a category when it does not, such as saying that a patient has malaria when he is healthy. False negatives occur when an item is in a category but one predict that it is not, such as when a patient with malaria is predicted to be malaria-free.

Assuming one have a clean record of ground truth, calculating inference errors will help one to assess the performance of the machine learning models. However, the reality is that many real-world datasets are noisy and may be mislabeled, which means that one may not have clarity on the exact inference errors that the AI system is making.

3.3.2.9 Unknown Error

Reality can be elusive, and one cannot always establish ground truth with ease. In many cases, such as with digital products, one can capture tons of data about what a user did on the platform but not his motivation for those actions. One may know that a user clicked on an advertisement, but don't know how annoyed someone may have been with it. In addition to many known types of errors, there are unknown unknowns about the universe that leave a gap between the representation of reality, in the form of data, and reality itself.

3.4 Building Machine Learning Model

Business leaders who want to lead AI initiatives at their companies should develop a high-level understanding of how machine learning models are built.

Machine learning is a powerful tool. Each algorithm has distinct advantages that make it more successful in some scenarios but not in others. AI experts and engineers are well-versed in these details, but most executives who lack technical backgrounds tend to clump all AI technologies together and regard it as the best.

3.4.1 Assessing the performance of the Models

Evaluation metrics are needed to assess the performance of the models. Accuracy, Precision and Recall are the most important concepts and the most common evaluation metrics for classification tasks.

3.4.1.1 Accuracy

Accuracy gives the percentage of classifications that were correctly made. For example, if you are building an email spam filter, the accuracy metric would tell you the number of messages that the filter correctly identified as being spam (true positive) or as being legitimate (true negative) out of all of the messages in your inbox. A perfect model has an accuracy of 1, because it will have correctly classified everything.

3.4.1.2 Precision

Precision measures the percentage of true outcomes that were correctly identified out of all of the true classifications that were made. To put it another way, precision tells us the model's ability to correctly classify the instances that we care about in a dataset. Email spam filters perform a binary classification task, in which it looks at a message and tries to determine whether it is spam (the target category) or not spam. These filters work well most of the time, but they occasionally make a wrong classification, which then either sends a spam message to your inbox or hides a legitimate message in your spam folder. Measuring the precision of a spam filter would tell you the number of messages that the filter correctly identifies as being spam (true positive) out of the total number of messages that it classified as spam, both actual spam messages (true positive) and legitimate messages that were mistakenly labeled (false positive).

3.4.1.3 Recall

Recall measures the percentage of true outcomes that were correctly classified as being true. In other words, recall characterizes a model's ability to identify all of the instances that we should care about in a dataset.

Going back to the spam filter, the true outcomes are all of the spam messages that were received by your email account. However, because spam filters are imperfect, the recall metric would tell you the number of spam messages that were correctly identified as being spam and filtered out (true positive) out of the total number of spam messages, both correctly flagged (true positive) and incorrectly flagged (false negative), that were received.

3.4.2 Machine Learning Workflow

Machine learning projects benefit from following a structured workflow, which starts with clearly defining business goals. Of all the machine learning approaches currently in use, supervised machine learning produces the most business value. The core of supervised machine learning is a mathematical **model** that describes how an algorithm makes predictions after being trained with historical data. The goal of **training** is to develop a model capable of mapping each **input** to a target **output**.

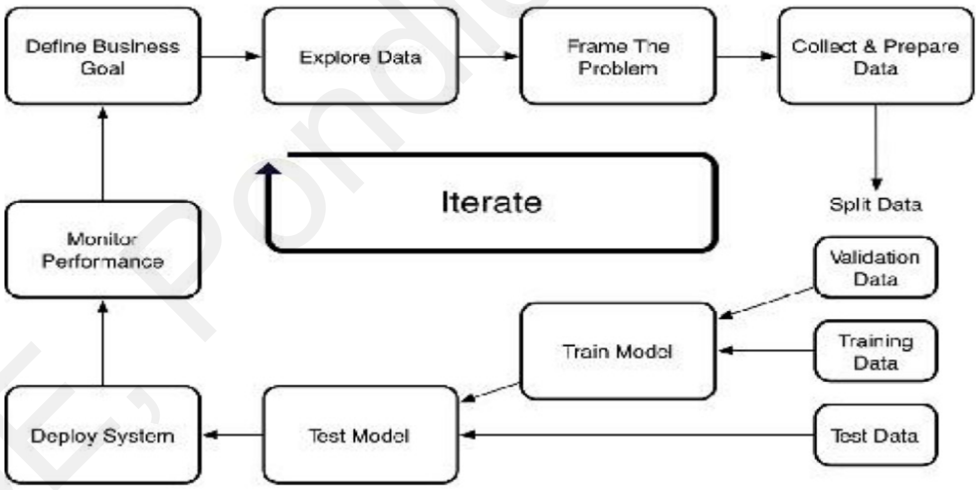

Figure 3.1

3.4.2.1 Define Business Goal

Carefully define the highest priority goal and the key performance indicators (KPI). Keep in mind that optimizing for everything means optimizing for nothing. Choosing too many KPIs will invariably result in conflicts where trying to boost one leads to a performance drop in another.

Try to balance internal business metrics, such as revenue, with metrics related to the customer experience. Avoid vague requests, such as "increase revenue." Increasing revenue could result from entering new markets, cross-promoting products, or reducing customer churn, all of which will require different technical approaches.

3.4.2.2 Examine Existing Data and Processes

If the new model depends on existing data and processes, then it would be need to perform exploratory analysis to understand the nature of the assets, which will inform the machine learning approach in turn. It may find that the data is insufficient or unsuitable for the objectives, requiring the team to collect new data that specifically addresses the problem.

3.4.2.3 Frame the Problem

Once the priority business goal is defined and KPI, and identified data and technology dependencies, then the data scientists and engineers can frame the problem in machine learning terms. They can determine how best to prepare the data, which technical approach to take (e.g., supervised vs. unsupervised learning), and develop a hypothesis about the algorithms that will perform best.

3.4.2.4 Centralize Data

If the desired data resides in different data warehouses or across various departments, it will require a coordinated, cross-functional effort to collate everything into a single training dataset. If data is missing it needs to be coordinated with additional teams to define how it has to capture new information across the products, services, and analytics workflows.

3.4.2.5 Clean Data

Prepare the data for processing by filling in missing values and correcting flaws. Depending on the initial state of the data, one may spend significant amounts of time cleaning and reshaping the raw data into a usable format. For example, locations such as "NYC" may need to be relabeled as "New York City" or with longitude and latitude, and timestamps may need to be converted if they come from different time zones. Many data scientists lament that they spend the majority of their time on data cleaning, but it must be done.

3.4.2.6 Split Data

If one use all the data to train a model, then one cannot easily check to see if that model will perform well on new data. To estimate the generalization error, or the error rate that the model will have while analyzing new data, one will need to randomly split the available data into three sets: training data, validation data, and test data. Training data is the baseline data used to build your model. Validation data is used as an intermediate testing set that will be used to iteratively improve model performance. Once the model is tuned to an acceptable performance level, use test data to estimate the model's generalization ability.

3.4.2.7 Train Model

Model training begins once the data has been split and algorithmic approaches are selected. Experienced data scientists and machine learning engineers may have some sense as to which models work best for specific problems and data types, but machine learning remains as much an art as a science. Finding the best fit will probably be require testing a variety of algorithms, and the team may be surprised by what works best.

3.4.2.8 Validate and Test Model

Measure the model's accuracy by using the validation and test datasets. Metrics for accuracy include recall and precision. Repeat training and testing data until to find the best model that produces the desired performance results.

3.4.2.9 Deploy Model

Finally, deploy the model in the business to reap the benefits of this new technology. Successful models are used to recommend products, customize landing pages, or score new sales leads.

3.4.2.10 Monitor Performance

Machine learning models will decay in performance if they are not regularly retrained on fresh data. One must monitor both a model's performance as well as the integrity of its data inputs. If undetected, corrupt data may not manifest in the predictions until later, which is why it is important to carefully track the data. Changes in data pipelines, data structure, or external conditions all need to be addressed, or they may affect the accuracy of your model.

3.4.2.11 Iterate

Machine learning models are never "done," in the sense that they will need continuous monitoring, iteration, and retraining to maintain required levels of performance over time. One may find that the original business goals and performance targets will shift in response to exogenous events or based on what one learn from previous models.

3.5 Experiment and Iterate

Building artificial intelligence does not have to involve big data, a large team, or expansive changes. Since machine learning processes are iterative and improve over time, one can start small and slowly expand the resources.

3.5.1 Agile Development

Agile software development is a time-boxed, iterative approach to building software incrementally. In a traditional waterfall software development model, a product is delivered in its entirety at the end of a project. By contrast, agile processes break down a project into tiny chunks of user functionality that can be addressed in two- to four-week-long cycles called sprints. This strategy uses continuous review to identify potential for improvement after every sprint. The flexibility of an agile development process works well for developing products that use AI.

3.5.2 Technical Debt

Building a successful machine learning model is just the first step to creating an AI product. Engineers warn that writing the code for a model constitutes only a small fraction of the engineering required for producing a machine learning system. The vast majority of the engineering required to productize a model lies in developing and maintaining the vast and complex infrastructure that surrounds the code.

Poor software design decisions made at the beginning of a project can easily compound into costly problems in the future. Technical debt refers to the cost of additional rework that will be needed in the future when one opt for quick and hacky fixes early on. In machine learning, one can easily incur massive ongoing systems costs by failing to mitigate risks early in the development process.

As more machine learning algorithms are put into production, one will also need to dedicate more resources to model maintenance—monitoring, validating, and updating the model.

Machine learning debt can be divided into three main types: code debt, data debt, and math debt. Code debt arises from the need to revisit and repurpose older code that may no longer suit the project. Data debt focuses on the data that was used to train the algorithm, which may have been incorrect or is no longer relevant. Math debt stems from the complexity of the model's algorithms. Most machine learning algorithms will require ongoing customizations that can make them harder to configure, maintain, and understand.

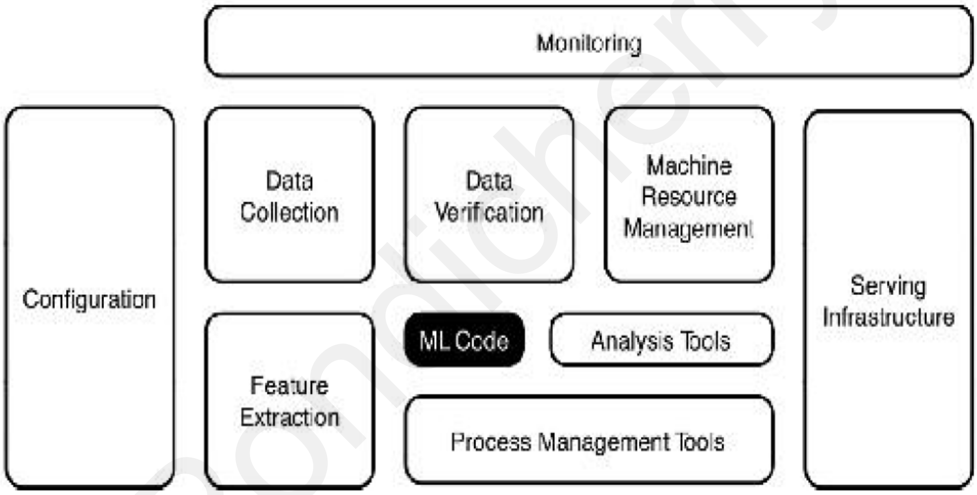

Figuvre 3.2

3.5.3 Deployment and Scaling

In order to support a large number of enterprise-wide machine learning systems, one will need a centralized technology architecture that provides a stable development and deployment environment.

In the beginning, development was decentralized, with data scientists using a variety of tools to create predictive models while engineers built one-off systems to bring these models into production. There were no standardized knowledge databases, no centralized data pipelines, and no streamlined production process.

Companies have created (MLaas) Machine Learning as a service platform to enable their engineering teams to build, deploy, and operate machine learning solutions with ease.

Overall, successful MLaaS systems have the following characteristics:

- Algorithm-agnostic. The platform supports numerous machine learning algorithms and innovative combinations of these algorithms.
- Reusable. Each machine learning algorithm can be reused in other applications.
- Simple. The system is easy for engineers of varying levels of technical experience to understand and use. Over time, the steps should become fully automated.
- Centralized knowledge. Information on past experiments, including results, is easily accessible for future reference.
- Flexible. The platform is capable of handling a variety of data types and learning tasks specific to a company and industry.
- Reliable and scalable. The platform remains resilient and be able to scale with high volumes of data during both the training and production phases.
- Intuitive user interface (UI). The system has a simple user interface to allow engineers and even non-technical domain experts to easily manage experiments as well as visualize and compare outputs.

3.5.4 Iteration and Improvement

Machine learning models are not static. The model will need to be retrained as new data becomes available or as external conditions change. The frequency of updates will depend on the algorithm, the situation, and the availability of data. If one is building a spam email filter, then one may need to retrain constantly if spammers are constantly formulating new attacks. If the modeling user churn in e-commerce, then one may not need to update the model as frequently if customers are slow to change in their turnover behavior.

Models are typically retrained every few weeks or months, or when there is a substantial change in external conditions that fundamentally changes the model trajectory. Some use cases may require daily or even real-time retraining. As a rule of thumb, half of the time should be spent on measurements and maintenance rather than on model creation. While these tasks are not as exciting as building new models, they are just as important when it comes to servicing the machine learning debt.

As these complexities pile up over time, one will likely find it more challenging to conduct root-cause analyses that are vital for maintenance. The black box nature of many deep learning algorithms makes it even more difficult to determine how an algorithm made a particular decision. Therefore, continuous monitoring and servicing of the models will be vital to maintaining the health of your machine learning systems. Tools and techniques for building and deploying enterprise-scale machine learning models are constantly evolving.

3.6 Self-Assessment Questions

1. What key roles should be established in an organization to drive innovation, and how can leaders foster a culture that supports and encourages creative thinking?
2. What are the critical components of an effective innovation roadmap, and how does establishing governance and oversight contribute to successful project execution?
3. Why is it important to assess the quality of data sources before preprocessing, and what are some common techniques for improving data quality?
4. How do you determine the most suitable machine learning algorithm for a given problem, and what is the role of feature engineering in model building?
5. How can adopting an agile approach benefit the development and deployment of machine learning models, and what role do experiments play in this process?

UNIT – IV

Lesson 4.1 - Business Applications

Business Applications: *Recommender Systems - Impact of recommenders on markets – Other forms of personalization on the web - Challenges with personalization - ML in Finance: Fraud Detection -ML in Finance: Additional applications*

Learning Objectives

1. To Understand the Impact of Recommender Systems on Markets
2. Understanding in Exploring Other Forms of Personalization on the Web
3. Understanding and Identifying Challenges with Personalization
4. To Understand Machine Learning in Finance: Fraud Detection and Additional Applications

4 Introduction

An increasing number of online business companies are utilizing recommendation systems to increase user interaction and enrich shopping potential. Use cases of recommendation systems have been expanding rapidly across many aspects of eCommerce and online media over the last 4-5 years.

Recommendation systems (often called "recommendation engines") have the potential to change the way websites communicate with users and to allow companies to maximize their ROI based on the information they can gather on each customer's preferences and purchases.

4.1 What is Recommender System

A recommender system aims to suggest relevant content or products to users that might be liked or purchased by them. It helps to find items that the user is looking for — and they don't even realize it until the recommendation is displayed. Different strategies have to be applied for different clients and they are determined by available data. Since RS has to be a data-driven approach, it can be fueled by machine learning algorithms.

There are two main stages of making recommendations:

- Candidate generation – a creation of a subset of products the user may like.
- Scoring – reduction and sorting a candidate list to the items displayed to a user.

4.1.1 The Power of Product Recommendations

To get the most out of RS and improve user experience, we should understand and dive into relationships between:

- **User and product** – when the user has a preference towards specific products. For example, a Netflix user may like thrillers while another user may prefer comedies.
- **Product and product** – when items are similar. For example, music or movies of the same genre.
- **User and user** – when users have the same or different taste concerning a specific item. For example, teenagers may differ from adults in terms of the content they consume.

Keeping these relationships in mind while designing an RS will lead to a delightful experience for users and consequently boost their engagement with such products.

4.1.2 The Strategies behind Recommender Systems

To select the best strategy for such systems, we must first assess the amount of available user and product data. Below are some popular strategies, sorted by the amount of data required in increasing order:

- Global – serving a user the most frequently purchased, trending, or popular products. They can be relevant for any user.
- Contextual – relying on product attributes and items purchased together, they can be combined with basic user attributes such as geolocation and can be used to target a group.
- Personalised – require not just context data but also user's behaviour, such as purchase history, clicks, etc.

These strategies should be combined on top of each other so as to strengthen the RS performance. For example, online shopping platforms should know the context of the product as well as the history of user purchases. While the "View together" strategy will be only possible for a new user, for old customers the "Purchased together" strategy is a better fit.

4.1.3 Metrics to Measure Recommender System

Candidate generation techniques

The goal of candidate generation is to predict a rating for the products for a certain user and based on that rating select a subset of items they may like.

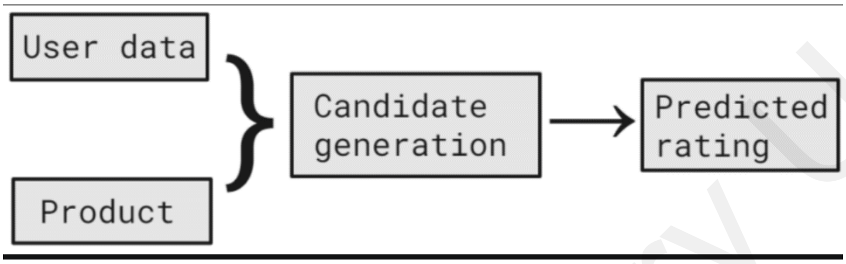

Figure 4.1

There are two main techniques to be described: content-based filtering and collaborative filtering.

4.1.3.1 Content-based filtering

Content-based filtering means that RS will recommend similar items to the liked or purchased ones (contextual strategy). For example, if user A watched two horror movies, another horror movie will be proposed to him. This technique can be user or item-centered.

Item-Centered

Item-centered content-based filtering means that RS recommends new items only based on similarity to the previous items (implicit feedback).

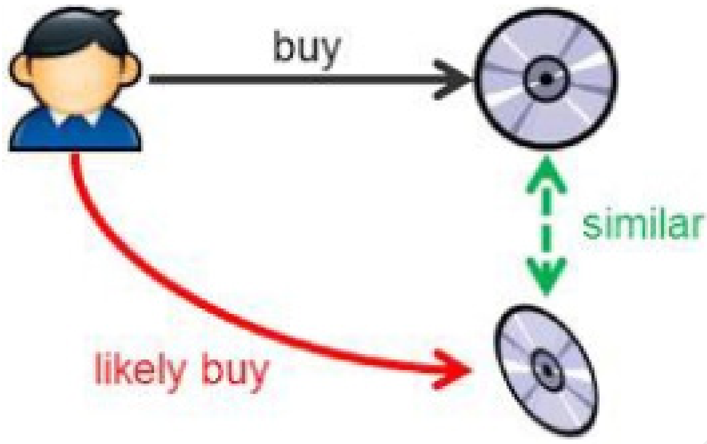

Figure 4.2

User-Centered

In the case of user-centered content-based filtering, information about user preferences is collected, for example via questionnaire form (explicit feedback). Such knowledge leads to recommending items with similar features to the liked one.

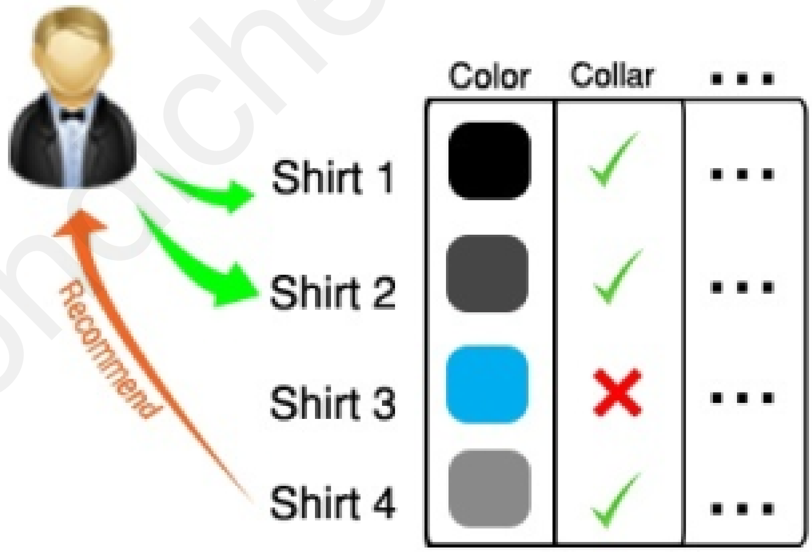

Figure 4.3

4.1.3.2 Collaborative filtering

It addresses some of the limitations of content-based filtering by using similarities between users and items simultaneously. It allows us to recommend an item to user A based on the items purchased by similar user B. Moreover, CF models' main advantage is that they learn users' embeddings automatically, without the need for hand-engineering. That means they are less constrained than content-based methods. Collaborative filtering systems can be split into memory and model-based approaches.

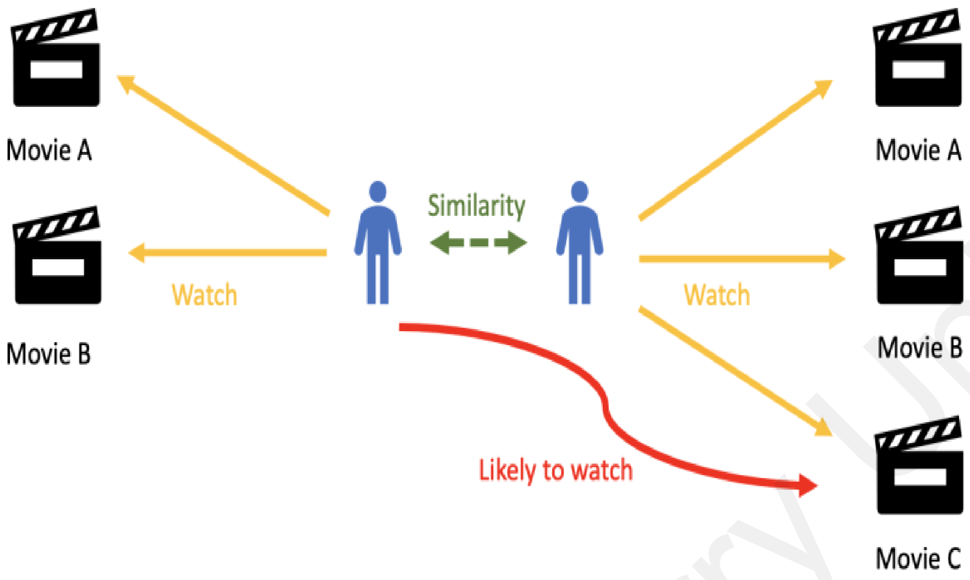

Figure 4.4

Memory-based

Memory-based CF systems work with recorded values from item-item or user-user interactions assuming no model. Search is done based on similarities and nearest neighbors algorithms. For example, find the users that are the closest to user A and suggest items purchased by them.

Model-based

Model-based approaches assume a generative model that explains user-item interactions and makes new predictions. They make use of matrix factorization algorithms that decompose the sparse user-item matrix into a product of two matrices: user-factor and item-factor. Recently, a lot of methods are being researched in the area of model-based RS. For example association rules, clustering algorithms, deep neural networks, etc.

Hybrid

A hybrid recommendation system is a combination of content-based and collaborative filtering methods. These systems help to overcome issues that are faced in those two types of recommenders. It can be implemented in various ways:

The two components can be developed separately and can be combined.

It can also be designed hierarchically based on conditions regarding the amount of user data available. As already mentioned, the "View together" strategy can be applied for new users and content-based item-centered filtering. This helps to overcome the cold-start problem. However, when there is more data available for past purchasers, we can implement collaborative filtering methods for them.

4.2 Impact of Recommendation Systems in Market

The Recommendation System has a positive impact on the market with its approach of Machine Learning at the background, its usage has tremendously increased in a shorter period of time especially for businesses.

4.2.1 Revenue Boost

One of the critical benefits of recommendation systems is their potential to drive revenue growth using data filtering tools. One can increase cross-selling and upselling opportunities by presenting personalized and relevant recommendations. That is why customers are more likely to discover and purchase additional products or services that align with their interests, leading to increased sales and revenue generation.

4.2.2 Enhanced Customer Satisfaction

Recommender systems can elevate customer satisfaction levels, leading to increased customer retention. Recommendations that resonate with particular user preferences may change their perception of the platform as attentive and responsive to their needs. This, in turn, may result in heightened satisfaction and a positive overall user experience.

4.2.3 Personalization at Scale

Recommendation systems excel at delivering personalized experiences. That's because the recommendation engine processes data to create individual profiles and offer tailored recommendations to each customer. This level of personalization increases the likelihood of conversions because customers may appreciate the platform's ability to curate products specifically for them.

4.2.4 Informative Reports

Recommendation engines provide information-filtering research that generates an insightful report about customer behavior, preferences, and trends. These offer a deeper understanding of similar user engagement data, conversion rates, and the effectiveness of different recommendations. Armed with this knowledge, one can make data-driven decisions to drive further growth.

4.2.5 Reduced Workload

By automating the process of suggesting relevant items or content, the system lightens the burden on employees who would otherwise have to curate recommendations manually. This allows employees to focus on other tasks while the recommendation system handles the personalized suggestion process.

4.2.6 Controlled Retailing

Controlled retailing enables a business to showcase new or underexposed items, manage inventory, and drive sales in trending directions. It can influence purchasing behavior and promote specific offerings by strategically guiding customers towards particular products or services.

How can recommendation systems boost the business?

Today, more and more online companies use Recommendation Systems to increase user interaction with the services they provide. Recommendation systems are efficient machine learning solutions that can help increase user retention, and lead to a significant increase in the business revenues.

4.3 Web Personalization and its Forms

At one time, businesses venturing into digital needed only to build a good looking website, publish fresh content, and make smart investments in search, paid media, and email to attract and engage visitors online. As the internet landscape became more saturated, whereby a healthy influx of conversions could be generated by optimizing the site's various elements to determine the best performing experience and serve it across traffic.

But in a world with 1.94 billion websites and an increasing number of online channels being folded into the customer journey, the stakes are even higher. Today, if brands want to influence decision-making, they'll need to invest in personalization as a discipline.

4.3.1 Website Personalization

Website personalization involves creating a unique experience tailored to everyone based on their past interactions, behavior, demographics, and preferences. For example, you can personalize your website by showing product recommendations on the cart page and content-related pop-ups.

4.3.2 Benefits of Website Personalization

- Strengthens customer relationships - Users want to see content that resonates with their preferences, and when you show them exactly that, they feel heard and valued. With personalization, you make the user's experience more humanized. Such experience boosts engagement and conversion metrics.
- **Facilitates repeat purchases** - The number of consumers who will become repeat buyers after a personalized shopping experience drastically improves.
- **Get better-qualified leads - Visitors should instantly know your USP and the benefits they can get from it to get good, qualified leads. Showing personalized content to target different segments will help you attract leads with a high chance of converting. It will also help your sales team eliminate the time and effort they would spend nurturing cold leads.**
- **Increased customer lifetime value - You want your users to buy from you frequently, and personalized content can help you achieve that goal. Using AI and machine learning-based recommendations, you can show the most relevant content based on their browsing data. It will keep users engaged as you show them what they want to see and thus encourage them to convert, increasing customer lifetime value.**

4.3.3 Types of Website Personalization

Site personalization comes in many different shapes and sizes, each use case designed with its own objective in mind. These experiences can be tested, optimized, and personalized for each visitor – and together, are powerful tools for improving the entire funnel.

Type	Use Case
Dynamic Content	Replace on-site hero banners, call-to-action buttons, promotional modules, or any other in-page element with dynamically-generated content variations.
Recommendations	Display content or products according to attributes in the data feed, visitor interactions, and trends in visitor behavior.
Overlays and popups	Highlight a particular offer using a large, prominent popup.
Notifications and widgets	Serve a subtle, unobtrusive element in a corner of the screen or as floating bars and sliding drawers.
Messaging	Tailor promotional, social proof, urgency, or CTA-based copy across the site.
Landing page	Personalize content variations for each visitor instead of a single variation (one-size-fits-all).
Menu personalization	Reorganize or change the order of the navigation bar based on each visitor's affinity or preferred categories.
Search personalization	Populate search results according to visitor preferences and real-time behavior.

Table 4.1

4.3.4 Challenges of Web Personalization

Personalization has become a critical factor in achieving this goal, as it allows one to create unique and one-to-one experiences for each individual.

However, personalization can be tricky – there are many challenges that can prevent one from getting it right. And many businesses face similar personalization challenges.

4.3.4.1 Data Challenges

Customer data is the key to understanding what each individual customer wants and needs. And it's essential for powering marketing personalization engines. That's why marketers are always looking for new and innovative ways to collect it. But what happens when you have too much of it? Or when it's not accurate?

- Here are the primary data challenges:
- Huge Volumes of Customer Data
- Bad Quality Data
- Privacy and Regulations
- Non-Unifies customer Profiles

4.3.4.2 Segmentation Challenges

Segmentation is one of the most important steps in effective personalization. By isolating and targeting specific groups, businesses can improve the customer experience and increase their chances of conversion.

However, this is easier said than done. As businesses strive to provide a more personalized experience for their customers, they are confronted with the challenge of segmenting their customer base.

Creating meaningful segments and accurately targeting customers within each one is critical to the success of any personalization strategy.

4.3.4.3 Talent, Skills and Organization Challenges

Organizational collaboration and capabilities are key differentiators between average and above-average performance in personalization.

However, in many businesses, different teams work with different sets of tech stacks such as CRM, analytics, marketing automation, etc. Therefore, the lack of an aligned tech stack across teams makes it more difficult to create effective personalization efforts.

4.3.4.4 Real-Time Delivery Challenges

Simply put, real-time personalization is delivering a customer experience that is tailored to their specific wants and needs in real-time.

To deliver true personalization, you need to have a system that creates real-time experiences for your customers. This means that as soon as a customer interacts with your brand, their experience should be tailored specifically to them.

Most companies are still struggling to do this effectively because it requires the seamless orchestration of data collection, segmentation, cross-channel profile synchronization, and execution.

There are two major reasons for real-time delivery challenges:

- lack of organizational alignment
- lack of orchestration between tech stack

4.3.4.5 Omnichannel Delivery Challenges

Today, most businesses are looking to improve their customer experience (CX) to stay ahead of the competition. And while many strategies exist for achieving this, omni-channel personalization is one of the most promising approaches.

This involves delivering a personalized experience to customers regardless of their channel of interaction with a company. However, there are several challenges that businesses must overcome to successfully implement omni-channel personalization.

The main challenge for omnichannel personalization is the lack of data synchronization across channels.

Therefore, the omnichannel personalization challenges are very similar to real-time personalization delivery challenges:

- lack of organizational alignment and communication across teams
- lack of orchestration between tech stack
- lack of unified customer profile and data synchronization across channels

4.3.4.6 Technology Challenges

Another major challenge for personalization is selecting the right technology stack.

As companies race to personalize their customer interactions, they realize that this is not as easy as it sounds. Technological challenges abound, from data management to integrating different channels and systems. When you're choosing a personalization tool, be sure that the personalization tool is compatible with your requirements in terms of:

- omni/cross channel personalization capabilities
- customer segmentation
- easy and seamless integrations with other tools
- real-time personalization and optimization
- testing and experimentation
- predictive and rule-based personalization opportunities
- customer data management and unified profile creation

4.3.4.7 Measurement Challenges

This is one of the major challenges for marketers today.

Measuring the results of your personalization efforts can help determine what's working and what needs to be improved. But measuring metrics can also be a challenge, especially when dealing with all the data that's available today.

By taking the time to measure results carefully, one will be able to make better decisions about how to personalize the content and improve the overall business performance.

However, measuring the success of personalization initiatives can be a daunting task. There are a number of factors to consider, including which metrics to use and how to track them.

Because personalization requires a longer time to showcase its effects, and it is difficult to quantify "customer attachment" or "customer trust," marketers struggle to measure the effectiveness of their personalization initiatives.

4.3.4.8 Scalability Challenges

As business grows, the need for more personalized customer experiences also increases. But when it comes to scaling personalized experiences, many businesses hit a wall.

Why is scalability such a challenge for personalization?

Achieving true personalization is not easy – it's a complex process that involves:

- analyzing huge amounts of customer data,
- content design and creation for each customer group, and
- effective usage of communication channels.

Luckily, there are ways to overcome this challenge, and businesses that manage to do so will be well on their way to true personalization nirvana.

Artificial Intelligence (AI) can provide the scalability needed to ensure that every customer receives a unique, personalized experience – even at a massive scale. AI can automate the process of collecting and acting on data, making it possible to deliver a personalized experience at scale.

4.3.4.9 Getting Started Challenges

Many businesses struggle with knowing where to start when it comes to personalizing their interactions with customers.

Even if you have a good understanding of what personalization is and how it works, figuring out how to actually implement it can be a daunting task.

This challenge is only compounded by the fact that there are so many different ways to go about personalizing the customer experience.

To be able to deliver the right personalization, first, you need to understand the fundamentals of personalization:

- First of all, you need to understand that personalization starts with customer segmentation.
- To create customer segments, you need to have customer data. The second backbone of the fundamentals of personalization is having and understanding customer data.

- The third point is understanding your personalization maturity level.
- The next point is understanding and deciding the personalization channels. Usually, when we talk about personalization, we think about web personalization. However, there are different channels for personalization.
- The last point in the fundamentals of personalization is understanding the relationship between personalization and experimentation. You must always keep in mind that personalization is a process. To deliver the best possible customer experience each and every time, you need to create continuous experimentation within the personalization process.

4.4 Machine Learning in Finance

4.4.1 Introduction

Machine learning, or ML, is a branch of computer science and artificial intelligence (AI). It is the design and development of algorithms that are capable of "learning" from data to make predictions. In other words, machine learning models can mimic the cognitive process by acquiring knowledge through data and using it to process and analyze information. It is used to automate cognitive tasks.

Machine learning systems help people understand massive volumes of data and uncover important patterns within them. This information is then used to enhance business processes, make informed decisions, and assist with prediction tasks. Financial services companies use it to offer better pricing, mitigate risks caused by human error, automate repetitive tasks, and understand customer behavior.

4.4.2 ML in Fraud Detection

Machine learning models learn from identifying patterns. These patterns help them understand normal behavior and make it easier to detect suspicious activities, like money laundering or insider trading.

Traditionally businesses relied on rules alone to block fraudulent payments. Today, rules are still an important part of the anti-fraud toolkit but in the past, using them on their own also caused some issues. Rules and machine learning are complementary tools for fraud detection.

Machine learning-based fraud detection systems rely on ML algorithms that can be trained with historical data on past fraudulent or legitimate activities to autonomously identify the characteristic patterns of these events and recognize them once they recur.

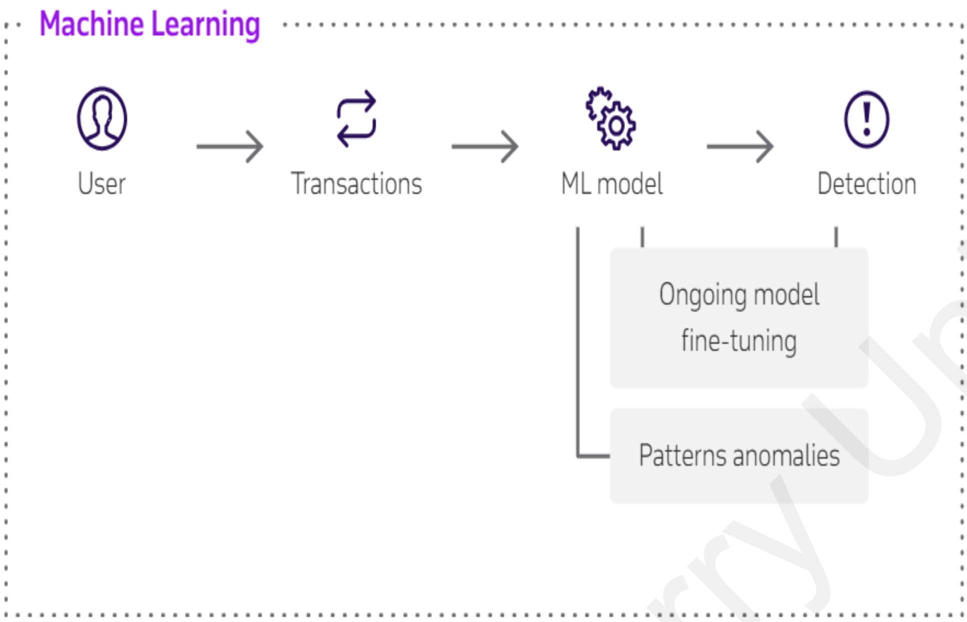

Figure 4.5

4.4.2.1 ML based Fraud Detection

- ML solutions autonomously identify and use more complex and variable rules than traditional systems. To do so, ML algorithms process data on past fraud cases, discover patterns and relationships between data points, and build models trained to identify those patterns once they recur in future datasets.
- ML systems can predict imminent criminal actions by identifying anomalies, namely subtle and unconventional behavioral patterns that humans would probably overlook but that still deviate from the norm, which could be clues to upcoming fraud.
- ML-powered solutions improve with experience, refining their models over time as they process new data, including unmapped data points. So, if they encounter new fraud scenarios, machine learning-based anomaly detection systems will quickly adapt to such threats, automatically integrating and updating the existing rules without human intervention.

4.4.2.2 Technical Overview of ML Fraud Detection

The main approaches to training machine learning algorithms are supervised, unsupervised, and reinforcement learning, depending on the degree of human involvement and control over the ML training process.

Supervised learning

ML-based fraud detection systems are trained with large amounts of labeled data, previously annotated with certain labels describing its key features. This can be data from legitimate and fraudulent transactions described with "fraud" or "non-fraud" labels, respectively. These labeled datasets, which require rather time-consuming manual tagging, provide the system with both the input (transaction data) and the desired output (groups of classified examples), allowing algorithms to identify which patterns and relationships connect them and apply such findings to classify future cases.

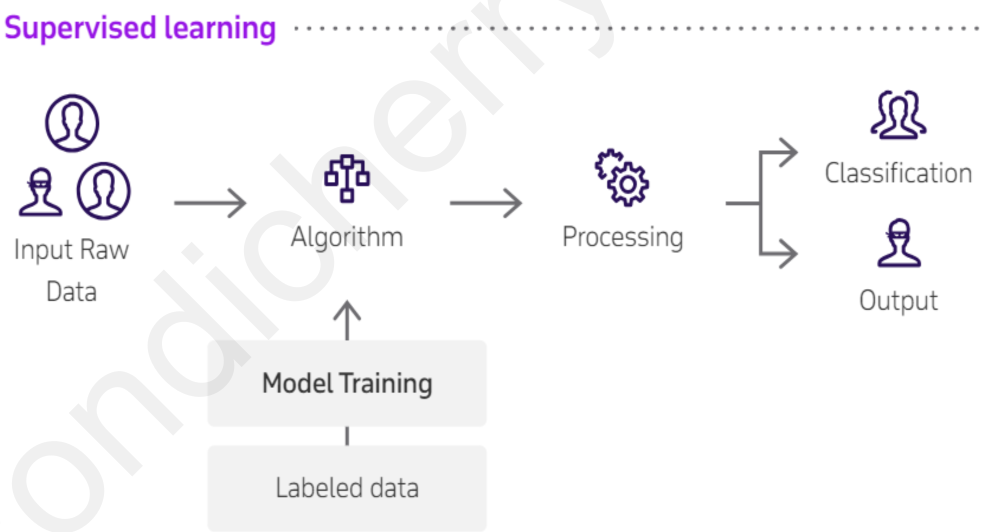

Figure 4.6

Unsupervised learning

These algorithms are fueled with unlabeled transaction data and have to autonomously group these transactions into different clusters based on their similarities (shared behavioral patterns) and differences (typical vs unusual patterns which can correspond to fraudulent activity). This approach, typically associated with deep learning, is computationally demanding but can be the only choice when facing fraud attempts that have never been met before and therefore unlabelled.

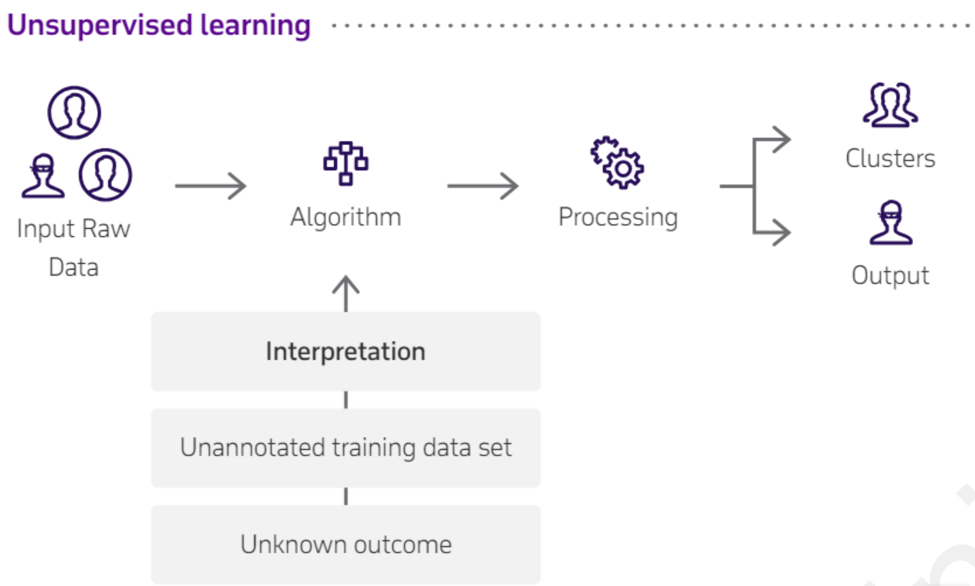

Figure 4.7

Reinforcement learning

This trial-and-error approach involves multiple training iterations in which the algorithm performs a fraud detection task in different ways several times until it can accurately identify fraudulent and non-fraudulent attempts. Since it does not require labeled inputs, reinforcement learning can be applied without prior knowledge of the current fraud scenario. However, it requires considerable computing power.

4.4.2.3 ML in top fraud Scenarios

1. Market Manipulation
2. Money Laundering
3. Credit Card Fraud
4. Identity Theft
5. Fradulent Insurance Claims
6. Tax fraud

4.4.2.4 Setting up an ML System for Fraud Detection

4.4.2.4.1 Business Analysis

- Identify fraud prevention-related needs and challenges
- Evaluate current tech ecosystem

- Understand if it makes sense to opt for ML-based fraud detection software over conventional solutions
- Define the solution's functional and non-functional requirements
- Set up a project roadmap, including scope, deliverables, and timeframes

4.4.2.4.2 Initial Data Analysis

- Perform an exploratory analysis to map available data sources (corporate databases and connected devices such as ATMs or POSs)
- Identify external data sources (public records, law enforcement, or government watch lists)

4.4.2.4.3 Product Design

- Draw up a specification detailing the solution's architecture, modules, core features, UI/UX, and integrations with other software
- Identify a suitable tech stack to build your software
- Optionally, deliver a proof of concept to ensure the project's feasibility and financial viability while pointing out potential adoption challenges

4.4.2.4.4 Building the Solution

- Perform data pre-processing, including data cleansing, annotation, and transformation
- Extract the most relevant features (customer's IP, preferred payment methods, number of failed transactions, average order value, fraud rate of the issuing bank)
- Outline the fraud detection system's evaluation criteria
- Process your datasets with ML algorithms to train the model to recognize patterns and anomalies, or build multiple models until you achieve the desired output

4.4.2.4.5 Model Integration and Deployment

- Integrate the ML model into the solution to power its fraud detection capabilities with the model's output
- Deploy the system to the target environment

- Configure all necessary API integrations with other corporate systems and data sources

4.4.2.4.6 Support

- Closely monitor the system's operation and perform ongoing maintenance with the help of experienced ML engineers
- Provide your staff with user training and support
- Follow best practices, retrain the ML model with new datasets across multiple iterations to fine-tune its output and address model drift issues

4.4.3 ML in Finance – Additional Applications

There are many applications of machine learning in finance, focused on the automation of tasks so that humans can focus on more complex activities. One example is shortening credit timeframes with credit risk prediction models.

Credit scoring prediction models are used to assess potential risks associated with lending decisions based on historical data. By using these models, banks can determine when it would be most profitable to procure loans or when there may be too much risk involved for them.

Another use of machine learning for finance is to recommend the right financial products at the right time, either from financial services companies or robo-advisors. The models can also help banks decide which customers to approach for new services and how best to price their offerings. With these types of predictions, banks can better manage their service portfolio while reducing costs over time (such as by automating repetitive processes).

These models also help with trading decisions and asset management, as AI helps fund managers analyze big data sources, such as stock prices.

4.4.3.1 Applications where ML in finance used.

Chatbots

Chatbots are computer programs that simulate conversation with a human and answer questions. In finance, chatbots can help automate tasks such as answering compliance team inquiries, providing customer service advice, or assisting with financial decisions.

Furthermore, chatbots contribute greatly in lead generation and capture. HubSpot Chatflows are a prime example of this. Integrating it with a reliable email validation tool ensures that only verified and accurate information is captured.

Financial monitoring

Financial monitoring is the process of tracking your financial health over time using tools such as budgeting apps or investor dashboards. In finance, this is commonly known as personal capital management. For example, Cleo is an intelligent, chat-heavy savings app.

Financial advisors also use financial monitoring tools to help their clients track their spending and monitor progress towards achieving their financial goals. These tools can also alert users if they deviate from their budget and provide recommendations for how to adjust accordingly.

Fraud Detection

Fraud prevention is another area where AI can play a role.

Security teams use machine learning algorithms to analyze millions of data points and detect fraud as it's happening, as well as prevent it before funds are released from a client's account. This is possible with large neural networks called deep learning, in this case fueled by massive amounts of financial data.

A fraud prevention system can look at patterns in incoming transactions and compare them to previous data to determine if something looks odd or suspicious, such as a large number of small transactions.

Automation

Many firms use automation to reduce costs associated with manual processes. For example, a bank may have a team responsible for generating new account applications using an application program interface.

Paperwork reduction has been a key goal for many financial firms. Banks spend billions of dollars annually on paperwork and compliance activities like verifying account ownership or monitoring client activity.

This work can be partially or even fully automated using machine learning, freeing up workers to focus on the more complex aspects of clients' accounts, such as helping them make long-term financial decisions or addressing their unique needs.

Automation allows financial advisors to do more with less. That means advisors can spend time advising their clients instead of performing repetitive data entry—or even spending hours each week preparing compliance documents that are then passed along to other teams.

Risk Analysis

Risk analysis is a critical part of any investment strategy. It involves quantifying, aggregating, and understanding risks so that one can better manage them. In finance, this includes identifying potential risks in a transaction using a combination of quantitative and qualitative analysis—such as calculating the expected loss based on historical data or assessing risk based on factors such as industry concentration or macroeconomic conditions.

In addition to providing insight into transaction risks, machine learning algorithms can be used for risk management, by quantifying those risks and giving firms the ability to create policies around them. This helps firms design robust trading strategies, limit their potential losses based on historical patterns, and proactively protect themselves from possible dangers.

Data Management

Data management is the process of gathering, storing, and organizing data so that it can be analyzed. In finance, this often involves monitoring fluctuations in financial markets. For example, a market monitor could look at all trades being conducted by a firm to identify trends or patterns that could indicate potential areas for concern

Using machine learning, the market monitor would then be able to spot these patterns in real time rather than waiting for an analyst to discover them manually. This would free up analysts to focus on more pressing issues—and perhaps alert the business when it's necessary to take action.

Trading

The trading strategy used by a firm also has a big impact on costs and efficiency. A trading strategy could be based on an algorithm that automatically buys and sells based on market conditions. This can help firms avoid placing trades that aren't profitable or aren't needed—saving both time and money in the long run.

Algorithmic trading strategies have become increasingly popular among financial institutions, many of whom find them to be cost-efficient ways of managing risk while generating returns.

Financial advisory

One area where machine learning can have a huge impact on the financial services industry is in the area of financial advisory. If you've ever spoken to an automated phone system, you've interacted with machine learning. This technology uses algorithms to convert speech into text and vice versa—a process using natural language processing (NLP).

Loan Approval

Another area where machine learning can be used in finance is loan approval. In the past, this process has been quite manual, requiring a human to review each loan application and make a decision. This can be time-consuming and costly.

With machine learning, however, firms can develop algorithms that automatically review loan applications and make recommendations about whether to approve or deny them. This not only saves time but also helps to ensure that loans are given to those who are most likely to repay them—reducing risk for the lender.

Job recommendations

In today's job market, it's becoming more and more common for companies to use machine learning to recommend jobs to candidates. This is done by taking into account a variety of factors, such as a candidate's skills, experience, and location.

Machine learning can also be used to assess a candidate's qualifications for a role—saving both the company and the candidate time in the hiring process.

Bankruptcy prediction

Another area where machine learning is being used in finance is bankruptcy prediction. This involves using historical data to build models that can identify which customers are most likely to file for bankruptcy. This information can then be used by banks or other financial institutions to make decisions about lending or credit lines.

Robo-advisors

A robo-advisor is an automated investment management service that provides recommendations about investing based on your goals and risk tolerance. These services have become increasingly popular in recent years as they offer a low-cost alternative to traditional financial advisors.

4.5 Self-Assessment Question

1. How do recommender systems influence consumer behavior and market dynamics, and what are some examples of their impact on major e-commerce platforms?
2. Beyond product recommendations, what are some other forms of web personalization that can enhance user experience, and how can businesses implement these strategies?
3. What are some key challenges associated with implementing personalization strategies on the web, and how can businesses address issues related to privacy and data security?
4. How is machine learning used for fraud detection in the finance sector, and what are some additional applications of machine learning in finance?
5. How can businesses leverage machine learning and personalization to gain a competitive edge, and what are the potential risks associated with these technologies?

UNIT – V

Lesson 5.1 - Artificial Intelligence for Enterprise Functions

AI for Enterprise functions*: Obstacles and opportunities, General and administrative, human resources and talent, business intelligence and analytics, Software Development, Marketing, Sales, Customer support-ethics of enterprise AI Generating AI, Data Protection Laws, Regulatory Aspects of AI.*

Learning Objectives

1. To Understand the key objectives and evaluate opportunities
2. To understand automated decision making and process optimization
3. To understand talent acquisition and employee development and performance
4. To understand data integration and visualization and predictive analysis
5. To understand the ethics of enterprise AI and Generative AI
6. To Understand legal aspects of AI

5 Introduction

Everybody claims to have “AI” in their products today. There are “AIpowered” juicers, “AI-enabled” wifi routers.

Enterprise functions represent one of the easiest entry points for deploying AI within the company. Though not as attention-grabbing as the latest breakthrough research in neural network architectures, finding ways to deploy existing AI techniques to optimize common business functions is usually easier than brainstorming new projects. Multiple products targeted at inefficiencies in virtually every enterprise function are already on the market, promising to revolutionize how we work.

5.1 Obstacles and Opportunities

Companies generate revenue by either cutting costs or finding new ways to make money, with the first being generally more straightforward than the second. Current AI-based solutions are very good at reducing inefficiencies in the workplace. By handing off repetitive tasks to software, employees have more time and energy to spend on high-value tasks.

Unfortunately, for companies that do not have technology as a core competency, even getting started can be a daunting task. AI-based solutions are complex technical products, and one may be asked to choose between vague product descriptions that are heavy on buzzwords but light on exposition or the reverse—convoluted technical descriptions that don't actually tell one what the products do. The resulting confusion increases the risk of buying a product that doesn't actually address the company's needs.

Modern AI solutions are voracious, requiring clean and relevant data by the bucketful. As a rule of thumb, if all data collected so far can be loaded into in to Excel and cleaned by hand, then one probably does not have enough data for robust machine learning applications.

AI solutions are also not cheap. The huge initial outlay of cash can take years to achieve its anticipated ROI, if at all, especially if the company experiences decreased productivity during implementation. The combination of all of these factors can be daunting to interested companies that have no incentive to look beyond their next quarterly earnings report.

Ultimately, integrating an AI-based solution into the workflow of an enterprise function, one of the vital cogs that keep the business running, is a major undertaking. While established technology companies recognize that technical investments take a long time to mature and can afford to wait for success, company may be less happy about weathering the disruptions.

What AI Can Do for Enterprise Functions

Current AI-based solutions are very good at streamlining processes and taking over rote tasks such as triggering a workflow. Automation frees up the cognitive load of the employees so that they can focus on more meaningful aspects of their jobs. For example, mundane tasks can take up to 75 percent of a recruiter's job. Handing off the responsibility of filtering resumés, which consists primarily of matching terms or looking for specific experiential phrases, to an AI-based solution can free up to half of the recruiter's day for other tasks. Recruiters can then use the extra time to get better acquainted with existing candidates, thereby improving the overall hiring process.

5.2 General and Administrative

General and administrative units such as Finance, Legal, and Business Operations are often underappreciated because they do not generate revenue. However, these functions perform some of the most critical jobs within the company, such as keeping track of the money that Sales bring in, using that money to pay for the ads that Marketing will use to attract new leads, and keeping an eye out for regulatory and legal hurdles that Product Management may have to address.

General and administrative roles are riddled with tedious but critical tasks such as manual data entry, which requires extreme precision. The exponentially-increasing volume of data combined with limits on the human capacity for sustained attention to detail is a recipe for corporate disaster. Fortunately, computers do not tire and excel at repetitive tasks where attention to detail is essential.

5.2.1 Finance and Accounting

Instead of forcing the human employees to spend hours over a spreadsheet to audit financial line items for duplicates, errors, expense abuses, and spending anomalies, one can use natural language processing software to automate expense management for accountants and controllers. This cuts down on the possibility of error and removes the most tedious parts of an accounting professional's job. For employees who must regularly submit expense reports, it saves time and cuts down on error by streamlining business travel booking and expense reporting procedures.

Large companies tend to be populated with multiple record-keeping systems, all of them incompatible with each other, and the threat of missing data looms large. Specialist platforms leverage machine learning to better recognize and categorize spending data, even filling in missing information to create a clean, standardized overview of a company's spending patterns.

5.2.2 Records Maintenance

The primary use case in administrative roles is form processing and records maintenance, which involves countless hours of accurate data entry. We've probably all run into organizations that continue to insist on handwritten forms that some poor intern must then painstakingly enter into legacy databases. The need for manual entry creates a bottleneck and increases the risk of error, especially as the prevalence of keyboards has sent handwriting legibility into a steep decline. To deal with this problem, HyperScience utilizes advanced computer vision techniques to scan and process handwritten forms to eliminate the data entry bottleneck. Once a form is scanned, their software cleans the image, matches the format to the correct form, then extracts and stores the relevant information in the correct database.

5.2.3 General Operations

Most companies have tons of repetitive digital workflows. These workflows can be tedious to complete. Employees responsible for these tasks can easily become bored and inattentive, allowing errors to creep into the operations and the data.

Fortunately, these tasks are well-suited for automation by Robotic Process Automation (RPA), which are software robots programmed to perform a specified sequence of actions. Even better, RPA deployment is relatively fast and low risk, so that problematic robots can quickly be removed without detriment to existing systems. Examples of workflows at which RPAs excel include performing regular diagnostics of your software or hardware, creating and updating accounting records (such as payroll), or automatically generating and delivering periodic reports to the relevant stakeholders.

More recent versions of RPAs have self-learning and NLP capabilities can now learn from example, to the point where RPAs can recognize the existence of new input and ask for help. Unfortunately, AI based RPAs need to learn from experience, meaning that this ability to recognize new inputs is still dependent on having access to large amounts of previously generated data.

5.3 Human Resources and Talent

The foundation of every company is its people, and hiring good people who work well together is the most important component to a successful company. Finding and retaining the right people is not easy, and talent development is an ongoing strategic priority for most companies.

5.3.1 Matching Candidates to Positions

Specialists traditionally spend much of their time and energy manually identifying candidates, drawing up compensation packages, designing career plans for new hires, and creating the right corporate culture so that employees can feel comfortable and work productively. Though this model generates a lot of records on prospective hires as well as on both current and former employees, decisions are made primarily by combining intuition with market data.

Fortunately, the nature of HR, which emphasizes matching and planning skills, and the abundance of internal data create opportunities for AI-based optimization throughout the hiring process.

Many platforms use a variety of AI techniques, such as algorithmic matching and predictive scanning, to identify promising candidates. Because these tools are optimized for skills, values, and experience, their use can help reduce unconscious biases that deter from making the best hiring decisions, in the process promoting diversity and a more inclusive corporate culture.

5.3.2 Managing the Interview Process

While the interview process is a slightly different experience with every new candidate, automation can streamline many steps within the process.

Companies use AI to automate candidate communication and assessments, allowing hiring managers to gather more information about each candidate while also scaling up hiring efforts.

5.3.3 Intelligent Scheduling

Scheduling a meeting often requires a lengthy email exchange and the help of one or several human assistants. Products that incorporate natural language processing (NLP) can be used to analyze the contents of email exchanges, extract schedule preferences, and automatically set meetings.

5.3.4 Career Planning and Retention Risk Analysis

Like customer retention, hiring can be a complicated process. HR must match an open position with the right person who has the required skills. Once hired, HR and the employee then face the new challenge of mapping out a career plan and trajectory that is desirable to both parties. This process may become more challenging if an employee expresses a desire to develop in areas where the company has no open positions. An employee may become dissatisfied and look for jobs elsewhere if that person's skill set was badly matched with the demands of the job or finds no opportunities forthcoming.

5.3.5 Administrative Functions

HR specialists must balance strategic assignments with tedious administrative tasks. While they may be given important assignments, such as designing company-wide initiatives about culture or mapping out career trajectories for existing employees, they are also interrupted constantly to service more mundane requests, such as answering questions about workplace policy, filing paperwork, or entering information into databases.

The competing demands can result in a highly inefficient and frustrating work environment. As with the hiring process, the repetitive nature of these administrative tasks is well suited to automation.

5.4 Business Intelligence and Analytics

Business intelligence (BI) creates meaning from data that the company collected. The goal is to leverage that meaning to guide future business decisions. For example, if BI finds that the employees are bored because their skills are being underutilized, HR can use that knowledge to adjust individual advancement plans, increasing employee satisfaction in the process.

5.4.1 Data Wrangling

To create meaning, BI must first convert data into information, then analyze that information to create insights that can then be converted into recommendations for action. Pieces of data can be as simple as the purchasing history of one customer, and they hold very little meaning on their own. Data becomes valuable when all of the pieces have been gathered together, because then technical specialists can structure that collection in ways that aid pattern-matching. Now, a single customer's purchases can be understood in the context of all the purchases that all customers have ever made. Analysts can then draw more general conclusions about product performance and suggest ways to improve sales.

Traditionally, technical specialists had to manually record, prepare, analyze, and interpret the data. This is a labor-intensive process because a company generates a large volume of data, and human involvement incurs the constant risk of error. Data preparation, which cleans, labels, and structures the data for consistency and accuracy, is the most important step in BI.

Without clean data, analysis will not be accurate, and one should not have confidence in the conclusions. However, data preparation is work intensive and tedious, and few data scientists enjoy the drudgery. But advancement in AI had made this work simple and efficient.

5.4.2 Data Architecture

While centralized data allows the company to more efficiently and more intelligently assess its overall state of being, the process of collecting that data is complicated by the existence of data storage. Different business units generate data at different rates and in different, often incompatible, formats.

The existence of data storage are a manifestation of turf wars between business units. Business units have wildly different sets of priorities and goals. In more dysfunctional companies, they may perceive themselves to be in active competition for resources with other units and refuse to sharing data in order to protect their proprietary information.

For example, data-sharing between Sales and Marketing can give Marketing a better understanding of the types of leads that are actually converting to sales, while Sales gets better information about customer segments that salespeople should prioritize.

As a result of the need to manage data that is increasing in scope and complexity and being generated by multiple business units across an organization, new jobs specializing in the care and feeding of shared data have appeared. Chief Data Officer (CDOs) and Chief Data Scientist positions are now becoming common in companies, especially those interested in championing new AI investments.

5.5 Software Development

Software development is no exception to the AI revolution. Not only can machine learning techniques be used to accelerate the traditional software development lifecycle (SDLC), they also present a completely new paradigm for inventing technology.

In traditional development, one have to specify the exact functionality of the computer program before coding it by hand. However, many tasks, such as categorizing objects in a photo, are far too complex to teach to computers in a rigid, rule-based manner.

AI techniques such as machine learning and deep learning instead rely on learning algorithms that are iteratively trained and continuously improved on curated, domain-specific data. This training allows them to deduce which features and patterns are important without being explicitly taught this knowledge. This quality makes them better than the best human engineered code in analyzing image/video, sound/speech, and text.

The most profound impact of AI on computer programming has been the unraveling of how humans perceive, define, and execute software development. The traditional development lifecycle for "Software 1.0" typically starts with a tech spec, which defines functionality requirements for a product. The tech spec is then passed to design and development to guide the development of viable prototypes.

Machine learning-driven development, or "Software 2.0," extrapolates important features and patterns in data and builds mathematical models that leverage these insights. In machine learning, adding functionality can be as simple as re-training your model on new data. While machine learning development has its own debugging and maintenance challenges, it also offers many benefits, including increased homogeneity, ease of management, and high portability. On the flip side, the complexity of these models will make it difficult for humans to fully comprehend how they work, leading them to appear as "black boxes." Worse, the complexity may hide algorithmic biases that can lead to unintended and embarrassing consequences. Software 2.0 will not supplement traditional software development entirely. Training a machine learning model is only a single step in the development process.

5.5.1 Rapid Prototyping

Turning business requirements into actual products typically require months, if not years, of planning. Machine learning has shortened this process by enabling non-specialists to develop technologies using either natural language or visual interfaces.

5.5.2 Intelligent Programming Assistants

Developers spend the vast majority of their time reading documentation or debugging code. Smart programming assistants can reduce the time spent on these tasks by offering just-in-time support and recommendations, such as relevant specifications, best practices, and code examples. Examples of such assistants include Kite for Python and Codota for Java.

5.5.3 Automatic Analytics and Error Handling

Programming assistants can also learn from past experiences to identify common errors and flag them automatically during development. Once a technology has been deployed, machine learning can be used to analyze system logs to flag errors. In the future, it may be possible to allow software to modify itself without human intervention in response to errors.

5.5.4 Automatic Code Refactoring

Clean code is critical for team collaboration and long-term maintenance. Large-scale code refactoring is often an unavoidable necessity as enterprises upgrade to new technologies. Machine learning can be used to analyze code and automatically optimize it for interpretability and performance.

5.5.5 Precise Estimates

Software development has a reputation for exceeding budget and planned timelines. Reliable estimates require deep expertise, understanding of context, and familiarity with the implementation team. Machine learning can be trained on data from past projects—such as user stories, feature definitions, estimates, and actuals—to more accurately predict the effort and budget that will be required.

5.5.6 Strategic Decision-Making

Which products and features should you prioritize, and which ones should you cut? An AI solution trained on past developments and current business priorities can assess the performance of existing applications, helping you and your engineering teams to identify efforts that will maximize impact and minimize risk.

5.6 Marketing

Marketing and Sales are directly responsible for generating revenue. This unique position gives the two functions significantly more power to direct investment into new projects. The need to keep up with competitors in fighting for market share means that both units are, on average, much more willing to try new tools as well.

Marketing and Sales work hand-in-hand to attract and retain customers, which requires understanding what interests customers and motivates them to want to buy products. The goal overlap between them means that they share many tools, though they are applied to different use cases. For example, Marketing may want to use natural language generation products to generate ad copy that increases click rates, while Sales may want to use the same product to create sales pitch email headers that prospective customers will actually want to read.

Marketing aims to educate consumers by communicating about the company's products and services, the brand, and the values that stand for. Catching and retaining the attention of today's fickle customers will require the company to have accurate consumer research to build the branding strategy, create engaging content to excite interest in the audience, and understand how consumers will weigh the message against those of the fiercest competitors.

Marketing depends on extensive digital advertising to generate goodwill towards the brand and the products. This is an especially important operation for B2C companies that don't necessarily control the retail channels through which their products are sold.

Fortunately, digital products can now capture more data about every user interaction, enabling richer analysis and insights to be derived using machine learning approaches that offer superior accuracy to previous methods. The value of new models can also be tested and measured quickly, leading to a direct impact on revenue and costs.

However, the sheer number of individual buyers also means that B2C marketing must weigh individual preferences and the problem of scalability much more heavily.

5.6.1 Digital Ad Optimization

AI research is not yet advanced enough to decipher a person's interests and motivations without human assistance.

However, AI-based content generation, also known as natural language generation (NLG), has been an area of major development in the last few years. Machine learning algorithms can now use content from previous marketing campaigns to create emotional profiles for user groups. Based on these profiles, the NLG solution can create similar content that is tailored to specific platforms and to individual user groups, increasing the likelihood that the targeted audience will engage with the ads. Data about click-through rates can then be fed back to fine-tune the model, building more detailed profiles of how targeted users respond to certain types of messaging.

5.6.2 Recommendations and Personalization

Netflix, Amazon, Yelp, and other sophisticated technology companies all use machine learning to improve product recommendations, such as movies to watch, products to buy, or restaurants to try. Personalizing the product selection that show to customer's results in an improved experience for them and higher revenue for you. Many different machine learning approaches are used to drive recommender system performance. Some use active learning in the form of bandit algorithms, like in digital ad optimizations, while others use ensemble methods that combine the advantages of multiple models.

The next frontier in recommendation systems is the cold-start scenario, in which algorithms must be able to draw good inferences about users or items despite insufficient information.

5.7 Sales

Sales, the other side of the traditional revenue-generating coin, aims to convert the interest that Marketing has whipped up into purchases. The company will need accurate analytics to pinpoint the most likely customers, compelling messages to convert interest into orders, and diligent tending of existing relationships in order to build product and brand loyalty. Though sales is still largely an intuitive process, based on the ability of a salesperson to accurately infer a customer's needs, the Customer Relations Management software (CRM) holds a treasure trove of data about the company's customer relationships that can be used by machine learning algorithms to improve your sales insights and operations.

5.7.1 Customer Segmentation

Customers have different values, preferences, and behaviors. One may be missing opportunities to make meaningful connections if one treat them the same way. The unsupervised machine learning methods can group customers who share common characteristics. Generating clusters of customers in this bottom-up fashion, as compared to defining market segments from the top-down, can detect subtleties of behavior that you may overlook otherwise and helps to identify and qualify new customer segments.

5.7.2 Lead Qualification and Scoring

Lead scoring is typically based on the analysis of static data such as demographics, firmographics, or other behavioral data sources. Because only a few out of thousands of prospective customers ever buy the product, the sales staff need to know how to accurately pinpoint interest and identify the potential for that interest to convert into a sale, or in the best case scenario, repeated sales. Traditionally, qualifying potential leads requires a salesperson to make cold calls or engage in conversations, first to identify the right decision-maker, and then to assess the possibility of a purchase. However, this process is tedious, and a salesperson has to make a lot of unproductive calls before identifying an actual customer. Recent developments in natural language processing, understanding, and generation enable software to automatically handle outbound queries, process prospect replies, and alert salespeople to high-potential opportunities for follow-up.

Applying sentiment analysis to sales correspondence, for example, can filter through replies to predict the interest level in a potential customer and to pinpoint the best leads. A salesperson can use that analysis to gauge whether a lead is worth the effort for further development, thus avoiding wasted efforts.

5.7.3 Sales Development

Once a lead has been qualified, a salesperson takes over to develop a relationship that will hopefully culminate in a sale. Unfortunately, administrative tedium, such as scheduling demos, follow-ups, and a myriad of other social touch points, can consume a large part of a salesperson's day. Artificial intelligence can be used to analyze the contents of email exchanges, extract schedule preferences, and automatically set meetings.

5.7.4 Sales Analytics

By applying machine learning techniques to all of the sales data, including email content, call transcripts, and CRM engagements, one can train a model to predict which specific actions are likely driving her superior conversions and revenue. These insights can be used to improve the sales playbook for training the entire sales force.

5.8 Customer Support

Customer support has become increasingly important, with analysts predicting that it will overtake product and price as the number one way for a business to differentiate itself. However, customer care is expensive.

Though customer service has traditionally relied on human empathy to resolve issues, the pressure to keep costs down has made some degree of automation an imperative.

We believe that customer experience is one of the most fruitful areas for the application of artificial intelligence, and machine intelligence can be used to better understand what customers need and to deliver consistently amazing experiences for them.

5.8.1 Conversational Agents

In the face of skyrocketing costs, consumers also expect an ever-higher quality of customer service. While digital assistants have existed for years as curiosities, recent developments in Natural Language Processing (NLP) and integrated technology ecosystems make them increasingly useful. The most popular examples include Amazon Echo's Alexa, Google Assistant, Apple's Siri, and Microsoft's Cortana. Each is tied to a larger network of native software, native hardware, as well as third-party developer additions to add functionality. Conversational agents, also known as bots or chatbots, also exist on popular social media platforms like Facebook Messenger and Tencent's WeChat, enabling us to execute tasks within our messaging environments rather than external websites or apps.

Currently, companies primarily rely on two models to incorporate conversational agents into customer service: the "bot-only" model and the "bot-assisted agent" (or "cyborg") model. In the former model, a conversational computer program interacts directly with a customer without human intervention, while in the latter model, the bot advises the agent on the best course of action or automates agent functions such as knowledge searches.

At present, conversational agents can only handle the most basic service requests, so human agents are still required for complex or difficult cases.

Machine learning algorithms can be used to assess customer requests as they arrive, routing easier requests to AI agents while prioritizing higher value ones for the human staff. AI solutions also help human agents increase the number of requests that can be addressed.

Meanwhile, startups are competing to provide automated response systems, in which AI agents handles a large portion of requests with minimal human intervention.

5.8.2 Social Listening

Sentiment analysis, either using purely statistical methods (less accurate) or machine learning and deep learning methods (significantly more accurate), can gauge the general mood surrounding conversations that the company care about. Natural language processing (NLP) can predict the topic of conversation, the context, as well as the emotional slant and personality characteristics of the content creator.

The proliferation of conversation agents, discussed above, provide a treasure trove of text data such as reviews, comments, social media shares, or customer support tickets for analysis by companies such as Conversocial and Lexalytics. Additional biometrics such as tone, facial expression, and body language can be analyzed if company possess audio or video data.

5.8.3 Customer Churn

Churn occurs when customers fail to complete a critical task for the business model, such as renewing a subscription, pressing the checkout button in their shopping cart, or sharing content with a friend. Supervised learning methods, such as analyzing the difference between free users and paying customers, can help identify and analyze the factors that contribute to churn, so that one can proactively re-engage lost customers. For example, irate calls can be routed such that agents are given warning and the customers with the highest flight risk are assisted first.

5.8.4 Lifetime Value

Every business wants to maximize the lifetime value (LTV) of their customers, also known as the net profit that a customer brings over their entire expected engagement with the company. In order to pass the best leads to the sales teams, one will need to be able to identify the characteristics and behaviors of the most lucrative customers early. With data from the advertising campaigns, acquisition funnel, and customer engagement history, one can use supervised learning methods to predict the best targets.

5.9 Ethics of Enterprise AI

Ethical AI, also known as Responsible AI, is the practice of using AI to benefit employees, businesses, customers, and society. It's a system of moral principles and techniques that aims to guide the development and use of AI technology.

Some key principles of ethical AI include:

- Transparency: AI should be transparent, from hiring processes to driverless cars.
- Impartiality: AI should be impartial.
- Accountability: AI should be accountable.
- Reliability: AI should be reliable.
- Security and privacy: AI should respect data privacy and keep data secure.
- Social well-being: AI should be available for the benefit of individuals, society, and the environment.
- Avoid unfair bias: AI should not discriminate against individuals or groups, and should provide equitable access and treatment.
- AI inclusivity: AI should be inclusive, minimize harmful bias, and ensure fair and equal treatment and access for individuals.
- AI robust: AI should be engineered to build in quality testing, include safeguards to maintain functionality, and minimize misuse and impact of failure.

5.10 Data Protection Laws in AI

The General Data Protection Regulation (GDPR) and the Data Protection Act 2018 (DPA 2018) regulate the use of personal data by AI systems, including when it's used for training, testing, or deployment. The GDPR requires that AI tools only process the minimum amount of personal data needed for a specific purpose, and that a data protection impact assessment (DPIA) is conducted for processes that could pose a high risk to individuals' rights.

The India Digital Personal Data Protection Act 2023 (DPDPA) came into effect on September 1, 2023, and applies to all organizations that process personal data of individuals in India.

As AI models become more sophisticated, existing privacy laws may need to evolve to account for new ways that personal data is collected and processed. This may include reconsidering definitions of key terms, such as what constitutes processing, transparency, or when data is still considered personal.

5.11 Regulatory Aspects of AI

The regulatory framework for artificial intelligence (AI) covers three main areas:

- Governance: How autonomous intelligence systems are governed
- Responsibility and accountability: How the systems are responsible and accountable
- Privacy and safety: Issues related to privacy and safety

Other regulatory aspects of AI include:

- Data privacy: AI systems use large amounts of data to train and improve their algorithms, so data collection must comply with privacy laws and regulations
- Civil liability: The owner or user of an autonomous AI system is liable for damages caused by the system, unless they can prove they took appropriate measures to prevent the damage
- Bias and discrimination: If AI algorithms are trained using biased data, they can perpetuate social inequalities and discrimination
- Risk assessment: Risk assessment is important in the regulatory framework of AI/ML-based software as a medical device (SaMD)

- Unbiased training: Unbiased training is important in the regulatory framework of AI/ML-based software as a medical device (SaMD)
- Reproducibility: Reproducibility is important in the regulatory framework of AI/ML-based software as a medical device (SaMD)

Policymakers can create regulations and guidelines to balance the benefits of AI with the potential risks, and ensure that AI is used for the benefit of society. For example, the FDA regulates AI in medical devices to ensure patient safety, effectiveness, and transparent AI solutions.

5.12 Self-Assessment Question

1. What are the key objectives for implementing AI in a business strategy, and how can an organization evaluate potential opportunities for AI adoption?
2. How does automated decision-making contribute to process optimization, and what are some examples of its implementation in business operations?
3. How can AI and data analytics improve talent acquisition and employee development, and what role do they play in performance evaluation?
4. What are the benefits of data integration and visualization in business decision-making, and how does predictive analysis enhance strategic planning?
5. What are some ethical considerations for deploying AI and generative AI in a business context, and how can companies address these concerns?
6. What are some key legal issues related to AI that businesses should be aware of, and how can organizations ensure compliance with relevant regulations?