

# Fundamentals of Machine Learning: Types, Algorithms, and Applications

Gautam Sharma<sup>1</sup>, Navjot Singh Talwandi<sup>2</sup>

<sup>1</sup>dept. of Apex ,India, Chandigarh University

<sup>2</sup>dept. of Physics ,India, Chandigarh University

**Abstract.** This chapter discusses different flavors of machine learning, including supervised, unsupervised, reinforcement learning, and semisupervised learning. Further, it talks about the characteristics, algorithms, and various applications of each paradigm across sectors like healthcare, finance, marketing, and robotics. Furthermore, it emphasizes how one can properly select a machine learning approach according to the availability of data, problem domain, and computational resources. On top of this, it focuses on trends that are currently surfacing in the near future and could shape the future of machine learning, such as explainable AI and federated learning.

**Keywords:** Machine Learning (ML), Supervised Learning (SL), Unsupervised Learning (UL), Reinforcement Learning (RL), Applications.

## 1 Introduction

In recent years, machine learning (ML) has emerged as a transformative force across various industries, fundamentally altering how organizations analyze data, make decisions, and optimize processes. By enabling systems to learn from data rather than relying solely on pre-programmed rules, ML has paved the way for innovations that were previously thought to be the realm of science fiction. As businesses and researchers grapple with an ever-increasing volume of data, the ability to extract meaningful insights from this data is paramount.

Machine learning encompasses a variety of techniques and approaches, each suited to different types of problems and data structures. This chapter aims to provide a comprehensive overview of the different types of machine learning, including supervised, unsupervised, reinforcement, and semi-supervised learning. Additionally, it will delve into the algorithms associated with each type, their unique characteristics, and the myriad applications that demonstrate their effectiveness in solving real-world problems. By examining the current state and future trends of machine learning, we can better understand how this technology will continue to shape our world.

### 1.1 Understanding Machine Learning

Machine learning can be defined as a subset of artificial intelligence (AI) that focuses on developing algorithms that enable computers to learn from and make predictions or decisions based on data. Unlike traditional programming, where a programmer writes explicit instructions, machine learning algorithms identify patterns in data and improve their performance over time through experience.

The concept of machine learning is rooted in statistics and computational theory. It leverages techniques from both fields to build models that can predict outcomes, classify data, or identify anomalies. This ability to learn from data has led to its adoption across numerous fields, including healthcare, finance, marketing, and robotics.

#### Types of Learning

At its core, machine learning is often categorized into three primary types: supervised learning, unsupervised learning, and reinforcement learning. Each type serves distinct purposes and is suitable for different scenarios:

1. **Supervised Learning:** In this approach, models are trained on labeled datasets, meaning that each input is paired with a corresponding output. The objective is to learn a mapping from inputs to outputs so that the model can make accurate predictions on new, unseen data. Common applications include classification tasks (e.g., spam detection) and regression tasks (e.g., predicting house prices).
2. **Unsupervised Learning:** Unlike supervised learning, unsupervised learning deals with unlabeled data. Here, the model attempts to identify inherent structures or patterns within the data. This can involve

clustering similar data points or reducing dimensionality for visualization. Applications include market segmentation and anomaly detection.

3. **Reinforcement Learning:** This type of learning is inspired by behavioral psychology and involves training agents to make sequences of decisions. Agents learn through trial and error, receiving rewards or penalties based on their actions. This paradigm is commonly used in robotics and game AI, where agents must navigate complex environments.

The choice of machine learning approach largely depends on the nature of the problem at hand, the data available, and the desired outcomes.

## 1.2 Applications of Machine Learning

The versatility of machine learning allows it to be applied in a wide array of domains. Below are some notable applications that illustrate its transformative impact:

### 1.2.1 Healthcare

In healthcare, machine learning is revolutionizing diagnostics, treatment planning, and patient care. Algorithms analyze patient data to predict disease outbreaks, identify high-risk patients, and suggest personalized treatment plans. For instance, ML models have shown promising results in early detection of diseases like diabetes and cancer, enabling timely interventions that can significantly improve patient outcomes.

### 1.2.2 Finance

The financial sector has been quick to adopt machine learning techniques for risk assessment, fraud detection, and algorithmic trading. ML algorithms analyze vast amounts of transaction data to identify suspicious patterns indicative of fraud. Additionally, credit scoring models leverage ML to assess the likelihood of a borrower defaulting on a loan, allowing lenders to make informed decisions while minimizing risk.

### 1.2.3 Marketing

In marketing, machine learning enhances customer targeting and engagement strategies. By analyzing customer behavior and preferences, companies can create personalized marketing campaigns that resonate with individual consumers. Recommendation systems, like those used by Netflix and Amazon, utilize machine learning to suggest products or content based on users' past interactions.

### 1.2.4 Robotics

Machine learning plays a crucial role in the development of autonomous systems. In robotics, ML algorithms enable machines to learn from their environments, improving their ability to perform complex tasks. For example, self-driving cars utilize a combination of supervised and reinforcement learning to navigate streets, avoid obstacles, and make real-time decisions based on sensor data.

These applications underscore the diverse capabilities of machine learning and its potential to drive innovation across various fields.

## 1.3 Challenges and Future Directions

Despite the remarkable advancements in machine learning, several challenges remain that must be addressed to fully harness its potential.

### 1.3.1 Data Quality and Availability

One of the most significant challenges in machine learning is the reliance on high-quality data. In many cases, the data collected may be incomplete, noisy, or biased, leading to suboptimal model performance. Ensuring data quality requires robust data collection, cleaning, and preprocessing techniques. Moreover, the scarcity of labeled data can hinder the effectiveness of supervised learning algorithms.

### 1.3.2 Interpretability and Transparency

As machine learning models become increasingly complex, understanding how they arrive at specific decisions becomes challenging. The "black box" nature of many advanced models, particularly deep learning networks, raises concerns about accountability and trust, especially in high-stakes domains like

healthcare and finance. Developing explainable AI techniques that provide insights into model decision-making processes is essential for fostering user trust and facilitating regulatory compliance.

### 1.3.3 Ethical Considerations

The ethical implications of machine learning cannot be overlooked. Issues such as data privacy, bias in algorithmic decision-making, and the potential for automation to displace jobs require careful consideration. As organizations integrate ML into their operations, they must establish ethical guidelines and frameworks to ensure responsible usage of technology.

### Future Directions

Looking ahead, several trends are likely to shape the future of machine learning:

1. **Explainable AI:** As previously mentioned, the demand for transparency will drive research into explainable AI methods, enabling stakeholders to understand and trust ML models.
2. **Transfer Learning:** This technique allows models trained in one domain to be applied to another, significantly reducing the need for large datasets and accelerating the deployment of machine learning solutions.
3. **Federated Learning:** This decentralized approach to machine learning enables models to be trained across multiple devices without sharing raw data, enhancing privacy while still benefiting from collaborative learning.
4. **Integration with IoT:** As the Internet of Things (IoT) continues to grow, the integration of machine learning with IoT devices will enable real-time data analysis and decision-making, unlocking new opportunities across industries.

In summary, machine learning represents a powerful tool that can transform how we understand and interact with data. By exploring its various types, applications, and challenges, we can better prepare for the future of this rapidly evolving field.

## 2 Types of Machine Learning

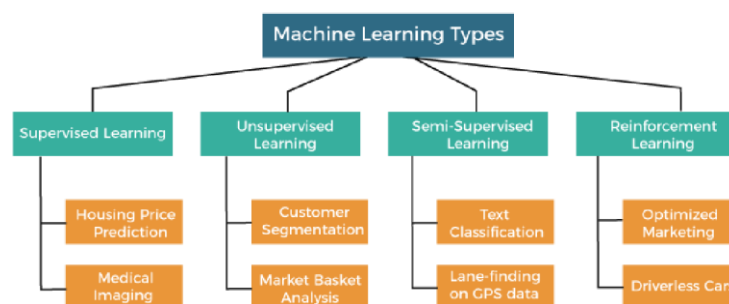


Figure 1 Types of Machine Learning

### 2.1 Supervised Learning

Supervised learning is one of the most widely used paradigms in machine learning. In this approach, models are trained using labeled data, which consists of input-output pairs. The algorithm learns the relationship between the inputs (features) and the outputs (labels) so that it can predict the outcomes for new, unseen data. The essence of supervised learning is to guide the algorithm through examples, allowing it to refine its predictions based on feedback.




Figure 2 Supervised Learning Types

### 2.1.1 Characteristics

Supervised learning is characterized by several key features:

1. **Labeled Datasets:** This is perhaps the most critical aspect of supervised learning. The presence of labeled data allows the algorithm to learn from explicit examples. Each training example consists of input features and a corresponding output label, which the model uses to understand the mapping between the two.
2. **Training and Testing Phases:** The supervised learning process typically involves splitting the dataset into two parts: a training set and a testing set. The model learns from the training set, while its performance is evaluated on the testing set. This separation ensures that the model's ability to generalize to unseen data can be accurately assessed.
3. **Goal of Minimizing Error:** The primary objective in supervised learning is to minimize the error between predicted outcomes and actual labels. This is often achieved through various loss functions, which quantify the difference between predicted and actual values. The model iteratively adjusts its parameters to minimize this error, thus improving its accuracy over time.
4. **Types of Problems:** Supervised learning is typically applied to two main types of problems: classification and regression. Classification involves predicting discrete labels (e.g., identifying whether an email is spam or not), while regression deals with continuous values (e.g., predicting house prices).

### 2.1.2 Algorithms

Several algorithms are commonly used in supervised learning, each with its strengths and weaknesses:

1. **Linear Regression:** This algorithm is used for regression tasks where the relationship between the input features and the output variable is assumed to be linear. It finds the best-fitting line (or hyperplane in higher dimensions) that minimizes the difference between predicted and actual values. Linear regression is straightforward and interpretable, making it a popular choice for many applications.
2. **Decision Trees:** Decision trees are versatile models that can be used for both classification and regression tasks. They work by splitting the data into subsets based on feature values, forming a tree-like structure. Each internal node represents a decision based on a feature, and each leaf node represents an outcome. Decision trees are easy to interpret and visualize but can be prone to overfitting if not properly pruned.
3. **Support Vector Machines (SVM):** SVM is a powerful algorithm used primarily for classification tasks. It works by finding the optimal hyperplane that separates different classes in the feature space. SVMs can handle both linear and non-linear data by using kernel functions, which allow them to operate in higher-dimensional spaces. This flexibility makes SVMs effective in various applications, though they can be computationally intensive.
4. **Neural Networks:** Inspired by the human brain, neural networks consist of interconnected layers of nodes (neurons). They are particularly suited for complex tasks involving large datasets, such as image and speech recognition. Deep learning, a subset of neural networks with multiple layers, has revolutionized fields such as computer vision and natural language processing. However, neural networks require substantial computational resources and can be challenging to interpret.

### 2.1.3 Applications

Supervised learning has a wide range of applications across various domains:


1. **Healthcare:** In the medical field, supervised learning is extensively used for disease prediction and diagnosis. For example, algorithms can analyze patient data (such as age, blood pressure, and glucose levels) to predict the likelihood of developing diabetes. Machine learning models can also assist in diagnosing conditions based on medical imaging, enabling early detection and treatment.
2. **Finance:** The financial industry relies heavily on supervised learning for credit scoring and risk assessment. Algorithms analyze historical data on borrowers, such as income, credit history, and loan repayment

patterns, to predict the likelihood of default. This information helps lenders make informed decisions about extending credit and managing risk.

3. **Marketing:** Supervised learning is instrumental in customer segmentation and targeting. By analyzing customer behavior and demographic data, businesses can classify customers into distinct groups based on their preferences and buying habits. This segmentation allows for personalized marketing strategies, improving engagement and conversion rates. For instance, targeted advertising campaigns can be tailored to specific customer segments, enhancing their effectiveness.
4. **Manufacturing:** In manufacturing, supervised learning can optimize processes by predicting equipment failures and maintenance needs. By analyzing historical sensor data and maintenance records, algorithms can identify patterns that indicate potential breakdowns. This predictive maintenance approach minimizes downtime and reduces costs associated with unplanned repairs.
5. **Natural Language Processing (NLP):** Supervised learning techniques are widely used in NLP tasks, such as sentiment analysis and text classification. For instance, algorithms can analyze social media posts or product reviews to determine the sentiment expressed (positive, negative, or neutral). This information is valuable for businesses seeking to understand customer feedback and improve their products or services.
6. **Image and Speech Recognition:** Supervised learning has revolutionized image and speech recognition technologies. In image classification, models are trained on labeled datasets containing images and corresponding categories (e.g., identifying objects in pictures). Similarly, speech recognition systems use supervised learning to convert spoken language into text, powering virtual assistants and transcription services.

## 2.2 Unsupervised Learning

Unsupervised learning is a key machine learning paradigm that focuses on analyzing and interpreting data without any labeled outputs. This approach allows algorithms to uncover hidden patterns, relationships, and structures within the dataset, enabling a deeper understanding of the underlying phenomena.



*Figure 3 Types of Unsupervised Learning*

### 2.2.1 Characteristics

1. **No Labeled Output:** In unsupervised learning, the algorithm operates on unlabeled data. Unlike supervised learning, where models are trained on input-output pairs, unsupervised learning relies solely on the input data. The absence of labels means that the model must autonomously identify patterns and group similar data points.
2. **Discovery of Inherent Structures:** The primary objective is to explore the data and detect its underlying structure. This can involve clustering similar data points together or finding associations between different variables. As a result, unsupervised learning is particularly useful for exploratory data analysis, where the goal is to gain insights without predefined hypotheses.
3. **Flexibility in Data Exploration:** Unsupervised learning can adapt to various types of data, including numerical, categorical, or text data. Its ability to handle diverse datasets makes it a versatile tool in data science, enabling analysts to extract valuable insights from complex data environments.

### 2.2.2 Algorithms

Several algorithms are commonly employed in unsupervised learning, each with its unique methodology and applications:

- **K-Means Clustering:** This algorithm partitions the dataset into  $k$  distinct clusters based on feature similarity. It iteratively assigns data points to the nearest cluster centroid, adjusting the centroids to minimize the overall distance between data points and their respective centroids.

- **Hierarchical Clustering:** This method creates a hierarchy of clusters through either agglomerative or divisive approaches. Agglomerative clustering starts with each data point as its own cluster and merges them based on similarity, while divisive clustering begins with one cluster and splits it into smaller clusters.
- **Principal Component Analysis (PCA):** PCA is a dimensionality reduction technique that transforms the data into a lower-dimensional space while preserving the most significant variance. It is widely used to simplify datasets, making them easier to visualize and analyze while retaining essential information.
- **Autoencoders:** These are a type of neural network designed to learn efficient representations of data. Autoencoders compress input data into a lower-dimensional space and then reconstruct it, making them valuable for tasks such as denoising or anomaly detection.

### 2.2.3 Applications

Unsupervised learning has a broad range of applications across various domains:

- **Customer Insights:** In retail, unsupervised learning techniques such as market basket analysis can identify purchasing patterns and associations between products. This information helps businesses optimize their inventory and marketing strategies, ultimately enhancing customer satisfaction and increasing sales.
- **Anomaly Detection:** Unsupervised learning plays a crucial role in identifying unusual patterns in data, such as fraudulent transactions in financial systems. By detecting anomalies, organizations can implement preventive measures and enhance security.
- **Image Processing:** Unsupervised learning techniques are commonly applied in image processing for tasks like compressing images and reducing dimensionality. Techniques such as PCA can simplify image data, enabling better visualization and analysis.
- **Document Clustering:** In natural language processing, unsupervised learning can group similar documents or texts based on their content. This application is beneficial for organizing large datasets, enhancing information retrieval, and improving search engines.

## 2.3 Reinforcement Learning

Reinforcement Learning (RL) is a distinct area of machine learning that emphasizes the training of agents to make a sequence of decisions by interacting with their environment. Unlike supervised learning, where models learn from labeled data, reinforcement learning agents learn through trial and error, receiving feedback in the form of rewards or penalties based on their actions. This feedback loop enables the agent to refine its decision-making policies over time. **2.3.1 Characteristics**

1. **Learning Optimal Policies:** The core objective of reinforcement learning is to develop policies that maximize cumulative rewards over time. The agent learns to choose actions that lead to the highest expected reward, continuously adapting its strategy based on the feedback received.
2. **Dynamic Environments:** RL is particularly suited for environments that are dynamic and can change over time. Agents need to be able to adapt to new situations and learn from their interactions, making RL applicable in scenarios where traditional models may struggle.
3. **Exploration vs. Exploitation:** A fundamental concept in reinforcement learning is the balance between exploration (trying new actions to discover their effects) and exploitation (choosing the best-known action). Striking this balance is critical for the agent's success in optimizing its performance. **2.3.2 Algorithms**

Several algorithms are employed in reinforcement learning, each with specific strengths and applications:

- **Q-Learning:** This off-policy algorithm enables agents to learn the value of actions in different states. By using a Q-table to store value estimates, Q-Learning updates these values based on the rewards received, enabling agents to make informed decisions.
- **Deep Q-Networks (DQN):** DQN combines Q-learning with deep neural networks, allowing agents to handle complex, high-dimensional state spaces. By utilizing experience replay and target networks, DQN improves stability and convergence during training.
- **Policy Gradients:** This approach directly optimizes the policy by maximizing the expected reward. Instead of estimating value functions, policy gradient methods learn the policy parameters that yield the highest rewards, making them effective for continuous action spaces. **2.3.3 Applications**

Reinforcement learning has gained traction in various domains due to its ability to handle complex decision-making tasks:

- **Robotics:** In robotics, RL is utilized to teach robots how to navigate environments and perform tasks autonomously. Through continuous interaction and feedback, robots can learn to adapt to dynamic conditions and improve their performance over time.

- **Gaming:** Reinforcement learning has revolutionized the field of AI in gaming. Notable achievements include training AI agents to play games like Go and chess at superhuman levels, demonstrating the potential of RL algorithms to master complex strategies.
- **Autonomous Vehicles:** RL plays a crucial role in the decision-making processes of autonomous vehicles. By learning from real-time data and feedback, these vehicles can optimize navigation routes, enhance safety protocols, and adapt to unpredictable driving conditions.

## 2.4 Semi-Supervised Learning

Semi-supervised learning is an innovative approach that bridges the gap between supervised and unsupervised learning. It leverages a small amount of labeled data alongside a larger set of unlabeled data, making it a practical solution in scenarios where obtaining labeled data is costly or time-consuming. **2.4.1 Characteristics**

1. **Reduced Need for Extensive Labeled Datasets:** By combining labeled and unlabeled data, semi-supervised learning minimizes the reliance on large labeled datasets. This characteristic is particularly beneficial in fields like healthcare, where expert labeling can be expensive and labor-intensive.
2. **Utilization of Unlabeled Data:** The approach capitalizes on the wealth of unlabeled data that is often readily available. By extracting useful information from unlabeled data, models can improve their learning and generalization capabilities.
3. **Improved Model Performance:** Semi-supervised learning can significantly enhance model accuracy, especially when labeled data is scarce. By effectively utilizing both labeled and unlabeled data, models can achieve better performance than those trained solely on labeled datasets. **2.4.2 Applications**

Semi-supervised learning finds applications across various domains, including:

- **Text Classification:** In natural language processing, semi-supervised learning is utilized to improve text classification tasks such as spam detection. By training models on a small set of labeled emails alongside a larger pool of unlabeled emails, systems can enhance their accuracy and reduce false positives.
- **Image Recognition:** In computer vision, semi-supervised learning can enhance classification accuracy in tasks like image recognition. By incorporating a few labeled images with a vast collection of unlabeled images, models can learn more robust features and improve their predictive capabilities.
- **Healthcare:** In medical imaging, where annotated datasets may be limited, semi-supervised learning can help leverage unlabeled patient images to enhance diagnostic models, leading to better patient outcomes.

## 3. Choosing the Right Machine Learning Approach

The selection of an appropriate machine learning (ML) approach is crucial for the success of any project. This decision hinges on various factors, including data availability, the nature of the problem at hand, and the computational resources that can be allocated. **3.1 Data Availability**

Data availability is a primary consideration in determining the suitable ML approach:

- **Labeled Data:** When sufficient labeled data is available, supervised learning methods are the most effective. These methods rely on labeled examples to learn the mapping between input features and target outcomes. The more diverse and representative the labeled dataset, the better the model can generalize to unseen data.
- **Unlabeled Data:** In cases where data is unlabeled or when obtaining labels is impractical, unsupervised or semi-supervised learning approaches should be considered. Unsupervised learning algorithms can identify patterns or groupings in the data without predefined labels, making them ideal for exploratory analysis. Semi-supervised learning, on the other hand, utilizes both labeled and unlabeled data to enhance model performance while minimizing the need for extensive labeling.

### 3.2 Problem Domain

The specific problem domain plays a critical role in determining the appropriate ML approach:

- **Predictive Tasks:** For tasks focused on prediction, such as regression or classification, supervised learning is typically the go-to method. These tasks require a model that can learn from historical data to predict future outcomes accurately.
- **Clustering or Grouping Tasks:** When the goal is to discover inherent structures or groupings within the data, unsupervised learning methods are more suitable. Applications like customer segmentation or anomaly detection fall under this category, where the objective is to identify distinct clusters or patterns without prior knowledge of the class labels.

### 3.3 Computational Resources

The computational resources available can significantly influence the choice of machine learning algorithms:

- **Resource-Intensive Algorithms:** Some ML algorithms, especially those involving deep learning, require substantial computational power and memory. These algorithms may not be feasible for projects with limited

resources. For instance, training deep neural networks often necessitates powerful GPUs and large datasets to achieve optimal performance.

- **Efficiency Considerations:** In contrast, simpler algorithms like linear regression or decision trees may be more suitable for environments with constrained resources. These algorithms can deliver adequate performance with lower computational demands, making them ideal for smaller datasets or real-time applications where speed is critical.

## 4. Future Trends in Machine Learning

The landscape of machine learning is continuously evolving, driven by advancements in technology and increasing demand for intelligent systems.

Several emerging trends are likely to shape the future of ML:

### 4.1 Explainable AI

As ML models become increasingly complex, the need for explainable AI (XAI) grows. Enhancing the interpretability of models is essential for building trust and understanding their decision-making processes. Researchers are developing methods to provide insights into how models arrive at their predictions, allowing users to grasp the rationale behind automated decisions. This trend is particularly important in high-stakes domains like healthcare and finance, where understanding model behavior can have significant implications.

### 4.2 Transfer Learning

Transfer learning involves leveraging knowledge gained from one domain to improve performance in another. This approach can significantly reduce the need for large datasets in target domains, allowing models to adapt quickly to new tasks with minimal training. For instance, a model trained on a large dataset for image recognition can be fine-tuned on a smaller dataset for a specific application, enhancing performance and efficiency. As the demand for specialized models grows, transfer learning is likely to become a standard practice across various industries.

### 4.3 Federated Learning

Federated learning represents a paradigm shift in how machine learning models are trained, focusing on decentralizing data processing. Instead of centralizing data in a single location, federated learning allows models to learn from data distributed across multiple devices or organizations while maintaining data privacy. This approach is particularly advantageous in industries like healthcare, where sensitive data cannot be shared freely. By enabling collaborative learning without compromising privacy, federated learning has the potential to enhance model accuracy and reliability while safeguarding user data.

## 5. Conclusion

Machine learning (ML) presents a diverse array of techniques and applications that hold the potential to revolutionize various industries, from healthcare to finance, marketing, and beyond. As organizations increasingly seek to leverage data for decision-making, understanding the different paradigms of ML becomes essential for practitioners aiming to harness its full power.

By exploring the core types of machine learning—supervised, unsupervised, reinforcement, and semi-supervised—this chapter has illuminated the strengths and weaknesses of each approach. Supervised learning shines in scenarios with abundant labeled data, enabling precise predictions and classifications.

Conversely, unsupervised learning offers insights into unstructured data, allowing for the discovery of hidden patterns and groupings without predefined labels. Reinforcement learning stands out in dynamic environments, enabling systems to learn optimal decision-making strategies through trial and error, while semisupervised learning bridges the gap between labeled and unlabeled data, enhancing model performance with limited resources.

Choosing the right machine learning approach is a multifaceted process, influenced by data availability, problem domain, and computational resources. The ongoing advancements in ML will continue to drive innovation, fostering new applications and enhancing existing systems. The trends toward explainable AI, transfer learning, and federated learning are indicative of a broader movement toward more interpretable, efficient, and privacy-conscious AI systems.

As we move forward, embracing the transformative potential of machine learning will require a commitment to ongoing learning, adaptation, and ethical considerations.

## References

1. S Bishop, C.M.: Pattern Recognition and Machine Learning. Springer, New York (2006)
2. Russell, S., Norvig, P.: Artificial Intelligence: A Modern Approach. Prentice Hall, Upper Saddle River (2010)
3. Goodfellow, I., Bengio, Y., Courville, A.: Deep Learning. MIT Press, Cambridge (2016)



4. Murphy, K.P.: *Machine Learning: A Probabilistic Perspective*. MIT Press, Cambridge (2012)
5. Zhang, Y., Wang, S.: A survey on machine learning in medical imaging. *J. Med. Imaging* 5(3), 031307 (2018)
6. Li, X., Wang, X., Chen, S.: An overview of deep learning in medical imaging. *J. Med. Syst.* 42(11), 204 (2018)
7. Alpaydin, E.: *Introduction to Machine Learning*. MIT Press, Cambridge (2020)
8. Xu, Y., Wu, L., Hu, B.: A survey on machine learning algorithms for social network analysis. *IEEE Access* 8, 148586–148606 (2020)
9. Choudhary, A., Sharma, S., Singh, H.: Reinforcement Learning: A Review. In: *2020 IEEE 7th International Conference on Engineering Technologies and Applied Sciences (ICETAS)*, pp. 1–6. IEEE, Dubai (2020)
10. Kelleher, J.D., Tierney, B.: *Data Science for Professionals: A Comprehensive Guide*. Kogan Page, London (2018)
11. Kearns, M., Neumann, J.: An Introduction to Reinforcement Learning. In: *Advances in Neural Information Processing Systems* 34, pp. 1759–1770. Curran Associates, Inc., New York (2021)
12. Bishop, C.M.: *Neural Networks for Pattern Recognition*. Oxford University Press, Oxford (1995)
13. He, K., Zhang, X., Ren, S., Sun, J.: Deep Residual Learning for Image Recognition. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778. IEEE, Las Vegas (2016)
14. Hinton, G.E., Vinyals, O., Dean, J.: Distilling the Knowledge in a Neural Network. In: *Advances in Neural Information Processing Systems* 27, pp. 1–9. Curran Associates, Inc., New York (2015)
15. Ribeiro, M.T., Singh, S., Guestrin, C.: "Why should I trust you?" Explaining the predictions of any classifier. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 1135–1144. ACM, San Francisco (2016)
16. Yang, Q., Wu, X.: 10 Challenging Problems in Data Mining Research. *International Journal of Information Technology & Decision Making* 5(4), 597–604 (2006)
17. Zhang, J., Wang, W., Wang, X.: Transfer Learning in Data Mining. In: *Proceedings of the 2015 International Conference on Artificial Intelligence and Computer Science*, pp. 35–39. IEEEExplore, New York (2015)
18. McMahan, H.B., Moore, E.J., Ramage, D.: Federated Learning of Deep Networks using Model Averaging. In: *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, pp. 1–8. PMLR, Fort Lauderdale (2017)
19. Chen, J., Xu, B., Li, Z.: A comprehensive survey on explainable artificial intelligence (XAI): Algorithms, applications, and future directions. *IEEE Transactions on Neural Networks and Learning Systems* (2021)
20. Dey, L., Egan, C., Ahsan, A.: A survey of machine learning algorithms in healthcare. *Journal of Biomedical Informatics* 112, 103611 (2020)

