

# Scalable Machine Learning Algorithms for Big Data Analytics: Challenges and Opportunities

Ravi Teja Potla

Department Of Information Technology, Slalom Consulting, USA

---

---

## 1. Abstract

The intersection of Big Data and machine learning (ML) represents one of the most promising and transformative trends in contemporary technology. Big Data encompasses massive datasets that are generated from multiple sources at unprecedented velocity, variety, and volume. With the proliferation of data from the Internet of Things (IoT), social networks, financial markets, healthcare systems, and various business applications, extracting valuable insights from this data has become crucial for organizations looking to remain competitive in the data-driven era. Machine learning offers the ability to automate the extraction of insights, predictions, and decision-making processes from large datasets, revolutionizing fields such as healthcare, finance, manufacturing, and more. However, traditional machine learning algorithms are not inherently scalable to meet the demands of Big Data. The growing size and complexity of datasets introduce numerous challenges, such as high-dimensionality, distributed data sources, real-time analytics needs, and the need for robust infrastructure.

This paper aims to provide a thorough exploration of the current challenges involved in scaling machine learning algorithms to meet the demands of Big Data analytics. We examine the computational and algorithmic limitations of conventional ML models when applied to large-scale datasets, focusing on issues like data distribution, processing power, memory consumption, and the need for real-time decision-making. Additionally, we explore emerging approaches, such as parallel and distributed computing frameworks (e.g., Hadoop, Apache Spark), cloud-based solutions, federated learning, and hybrid models, which aim to enhance the scalability of ML algorithms. By leveraging these advancements, organizations can reduce training times, minimize resource consumption, and deliver real-time insights more effectively.

In addition to exploring the current landscape of scalable machine learning, this paper delves into key opportunities for innovation in various industries, including healthcare, finance, and manufacturing. We present several case studies that demonstrate the successful application of scalable ML algorithms in real-world scenarios, such as predictive healthcare analytics, fraud detection in financial systems, and predictive maintenance in manufacturing. The paper concludes by outlining future directions for research and development in the field of scalable ML, with particular emphasis on the potential of quantum computing, automated machine learning (AutoML), and AI-driven optimization techniques to further enhance the scalability and efficiency of machine learning for Big Data.

This comprehensive analysis seeks to inform researchers, practitioners, and industry leaders of the current challenges and opportunities at the intersection of machine learning and Big Data, highlighting the importance of scalable algorithms in driving future innovations.

**Keywords:**

Big Data, Machine Learning, Scalability, Distributed Systems, Cloud Computing, Real-Time Analytics

**2. Introduction**

**2.1 Overview of Big Data and Machine Learning**

In the modern digital landscape, the generation and collection of data are happening at an unprecedented pace. With the rise of connected devices, social media platforms, financial transactions, and IoT systems, vast amounts of data are generated every second. This phenomenon, often referred to as "Big Data," presents a unique opportunity for organizations to leverage this data for competitive advantage, unlocking insights that were previously inaccessible. However, the sheer scale of Big Data—characterized by its volume, velocity, variety, and veracity—poses significant challenges for traditional data processing and analytics tools.

Machine Learning (ML), a subset of artificial intelligence (AI), has emerged as one of the most powerful tools for analyzing and extracting valuable insights from large datasets. ML algorithms are capable of learning from data without being explicitly programmed, making them ideal for predictive analytics, pattern recognition, and decision-making in a wide range of applications. In domains like healthcare, finance, e-commerce, and manufacturing, ML has been used to solve complex problems such as fraud detection, customer behavior prediction, personalized recommendations, and equipment maintenance prediction.

However, as the size and complexity of datasets increase, the scalability of machine learning algorithms becomes a critical concern. Traditional ML algorithms, designed to handle relatively small datasets, struggle to meet the demands of Big Data. Training models on large datasets can be time-consuming and computationally expensive, requiring significant memory and processing power. Moreover, many ML algorithms are not inherently designed to handle distributed or streaming data, which is common in Big Data environments. This paper aims to explore the key challenges in scaling ML algorithms for Big Data analytics and to identify potential solutions and opportunities for innovation in this rapidly evolving field.

## **2.2 Importance of Scalability in Big Data Analytics**

Scalability is a critical requirement for machine learning in the context of Big Data. As data grows in size and complexity, ML models must be able to scale in terms of both computation and storage. This involves designing algorithms and systems that can handle large volumes of data, process data in real-time, and efficiently utilize available resources. Scalability also refers to the ability of ML algorithms to work in distributed environments, where data is stored across multiple servers or cloud platforms. In addition, scalable ML solutions must be able to handle the challenges posed by high-dimensional data, data heterogeneity, and noisy data, which are common characteristics of Big Data.

In many industries, the ability to process and analyze Big Data in real-time has become a key competitive advantage. For example, in finance, real-time fraud detection systems rely on scalable ML algorithms to analyze millions of transactions per second. In healthcare, scalable ML models are used to analyze patient data from multiple sources in real-time, enabling early detection of diseases and personalized treatment plans. In manufacturing, predictive

maintenance systems use ML to analyze sensor data from machinery in real-time, reducing downtime and improving operational efficiency.

As organizations continue to collect and store increasing amounts of data, the need for scalable ML solutions will only grow. This paper explores the challenges and opportunities associated with scaling ML algorithms for Big Data analytics and provides insights into emerging trends and technologies that are shaping the future of this field.

### 3. Machine Learning Algorithms and Big Data: An Overview

#### 3.1 Popular Machine Learning Algorithms for Big Data Analytics

A variety of machine learning algorithms have been applied to Big Data analytics, each with its own strengths and weaknesses. The choice of algorithm often depends on the nature of the data, the computational resources available, and the specific goals of the analysis. In this section, we provide an overview of some of the most commonly used ML algorithms in Big Data analytics, highlighting their scalability challenges and potential solutions.

- **Decision Trees and Random Forests:** Decision trees are a popular ML algorithm for classification and regression tasks. They work by recursively splitting the data into subsets based on the values of input features, creating a tree-like structure that can be used to make predictions. Random forests are an ensemble learning method that builds multiple decision trees and aggregates their predictions to improve accuracy and reduce overfitting. While decision trees and random forests are relatively simple and interpretable, they can become computationally expensive when applied to large datasets. One approach to scaling decision trees and random forests is to parallelize the training process, distributing the construction of trees across multiple nodes in a cluster.
- **Support Vector Machines (SVM):** SVM is a powerful classification algorithm that works by finding a hyperplane that separates data points into different classes. It is particularly effective for high-dimensional data and is often used in applications such as image recognition and text classification. However, SVMs can be computationally

intensive, especially when applied to large datasets with many features. To address this, researchers have developed scalable variants of SVM, such as linear SVM, which reduces computational complexity by approximating the optimal hyperplane in a more efficient manner.

- **K-Means Clustering:** K-means is one of the most widely used clustering algorithms, which partitions data points into k clusters based on their similarity. It is commonly used in applications such as customer segmentation, image compression, and anomaly detection. While K-means is relatively simple and efficient for small datasets, it becomes challenging to scale when applied to large datasets. One solution is to use parallelized versions of K-means, where the clustering process is distributed across multiple machines to reduce computation time.
- **Neural Networks and Deep Learning:** Neural networks, particularly deep learning models, have revolutionized fields such as computer vision, natural language processing, and speech recognition. Deep learning models, such as convolutional neural networks (CNNs) and recurrent neural networks (RNNs), are capable of learning complex patterns from large amounts of data. However, training deep learning models on large datasets can be computationally expensive and requires significant hardware resources, such as GPUs or TPUs. To scale deep learning models for Big Data, techniques such as distributed training, model parallelism, and data parallelism are commonly used.

**Table 1: Comparison of Popular Machine Learning Algorithms for Big Data**

Algorithm	Strengths	Challenges	Scalability Solutions
Decision Trees	Easy to interpret, effective for classification	Prone to overfitting in large datasets	Random forests, parallel training
Random Forests	Reduces overfitting, better accuracy	Computationally intensive	Distributed computing, cloud platforms

<b>SVM</b>	Effective for high-dimensional data	High computational cost for large datasets	Linear SVM, parallelization
<b>K-Means</b>	Simple and efficient for clustering	Difficult to scale for large datasets	Parallel and distributed implementations
<b>Deep Learning</b>	Capable of learning complex patterns	Requires significant hardware resources	Distributed training, cloud-based GPUs/TPUs

### 3.2 Characteristics of Big Data

Big Data is often described using the "Four Vs": Volume, Variety, Velocity, and Veracity. Each of these characteristics presents unique challenges for machine learning algorithms, particularly in terms of scalability.

- **Volume:**

The sheer size of Big Data can overwhelm traditional ML algorithms, which may not be able to process the data efficiently or store it in memory. Scalable ML algorithms must be able to handle large datasets, either by processing data in batches or by distributing the computation across multiple machines.

- **Variety:**

Big Data comes in many forms, including structured data (e.g., databases), unstructured data (e.g., text and images), and semi-structured data (e.g., XML and JSON files). ML algorithms must be able to handle this diversity of data types and formats, often requiring preprocessing steps such as feature extraction and data transformation.

- **Velocity:**

In many applications, data is generated and processed in real-time, requiring ML algorithms to make predictions or decisions in a matter of milliseconds. Real-time analytics is particularly challenging for traditional ML algorithms, which are often designed for offline training and batch processing.

- **Veracity:**

Big Data is often noisy, incomplete, or inconsistent, making it difficult to extract reliable insights. ML algorithms must be robust enough to handle missing or corrupted data, often requiring techniques such as data imputation or anomaly detection.

## **4. Scalability Challenges in Machine Learning**

### **4.1 High Dimensionality**

One of the primary challenges of scaling ML algorithms for Big Data is the high dimensionality of the data. High-dimensional data refers to datasets with a large number of features or variables, which can make training ML models computationally expensive and prone to overfitting. For example, in image recognition tasks, each pixel in an image can be considered a feature, leading to datasets with millions of features. Similarly, in text classification tasks, each word in a document can be represented as a feature, leading to high-dimensional sparse data.

Scaling ML algorithms to handle high-dimensional data requires the use of dimensionality reduction techniques, such as principal component analysis (PCA) or feature selection methods, which reduce the number of features while preserving the most important information. Another approach is to use algorithms that are specifically designed for high-dimensional data, such as linear models or tree-based methods, which can scale more efficiently than complex models like deep learning.

### **4.2 Distributed Data Sources**

In many Big Data applications, data is distributed across multiple servers or data centers, making it challenging to train ML models on the entire dataset at once. For example, in cloud computing environments, data may be stored in different geographic locations or across multiple cloud providers. Training ML models on distributed data requires efficient data integration and communication between nodes in the network.



One approach to scaling ML algorithms for distributed data is to use distributed computing frameworks, such as Hadoop and Apache Spark, which allow data to be processed in parallel across multiple nodes. These frameworks use techniques like data partitioning and sharding to distribute the computation across the network, reducing the time required to train ML models on large datasets.

### **4.3 Real-time Analytics**

Real-time analytics is a key requirement for many Big Data applications, such as fraud detection in financial systems, real-time recommendations in e-commerce, and predictive maintenance in manufacturing. Traditional ML algorithms are often designed for offline training, where the model is trained on historical data and then used to make predictions on new data. However, in real-time applications, data is continuously generated and the ML model must be updated in real-time to provide accurate predictions.

To address this challenge, online learning algorithms have been developed, which allow ML models to be updated incrementally as new data is received. Online learning algorithms, such as stochastic gradient descent (SGD) and reinforcement learning, are particularly well-suited for real-time analytics, as they can process data in small batches and update the model without requiring a full retraining.

### **4.4 Hardware and Infrastructure Limitations**

Scaling ML algorithms for Big Data often requires significant hardware resources, such as powerful CPUs, GPUs, or TPUs, as well as large amounts of memory and storage. In many cases, the cost and complexity of managing the necessary infrastructure can be a barrier to scaling ML models for Big Data. Cloud-based platforms, such as Amazon Web Services (AWS), Google Cloud, and Microsoft Azure, offer scalable infrastructure that can dynamically allocate resources based on the needs of the ML workload. However, the cost of using these services can be prohibitive for small organizations or startups.

One approach to addressing hardware limitations is to use distributed computing architectures, where the computation is distributed across multiple machines or nodes in a



cluster. This allows organizations to scale their ML models without requiring large, expensive hardware setups. In addition, techniques such as model compression and quantization can be used to reduce the size of ML models, making them more efficient to deploy on resource-constrained devices, such as mobile phones or IoT sensors.

**Table 2: Challenges of Big Data and Corresponding Scalable ML Techniques**

Big Data Challenge	Description	Scalable Machine Learning Techniques
High Dimensionality	Large number of features or variables	Feature selection, dimensionality reduction, linear models
Distributed Data Sources	Data stored across multiple locations	Distributed computing frameworks (Hadoop, Spark)
Real-time Analytics	Need for immediate insights from streaming data	Online learning, streaming data frameworks
Data Volume	Massive amounts of data to process	Batch processing, in-memory computing (Apache Spark)
Data Variety	Heterogeneous data types (structured, unstructured)	Hybrid models, data preprocessing pipelines

## 5. Scalable Approaches to Machine Learning for Big Data

### 5.1 Parallel and Distributed Computing

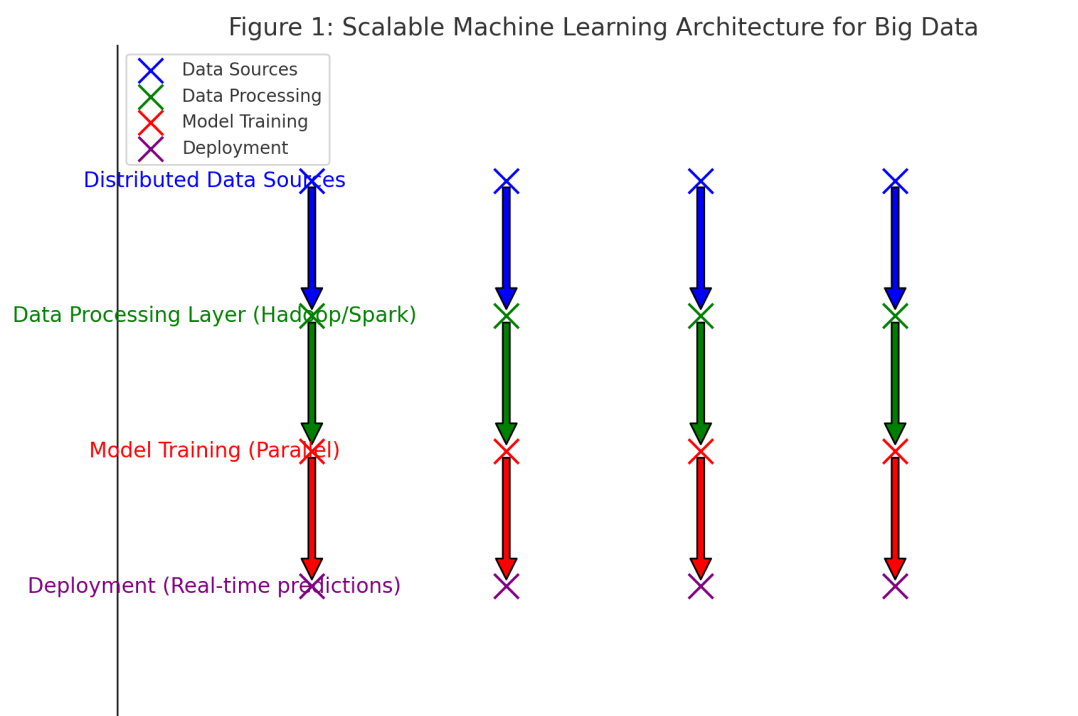
#### Hadoop and MapReduce Frameworks

Hadoop is one of the most widely used distributed computing frameworks for processing Big Data. It is based on the MapReduce programming model, which allows data to be processed in parallel across multiple nodes. In the MapReduce model, data is divided into smaller

chunks, which are processed independently by "mappers," and the results are then combined by "reducers" to produce the final output. Hadoop is particularly well-suited for batch processing of large datasets, but it can be less efficient for real-time or iterative ML tasks.

### Apache Spark for In-Memory Processing

Apache Spark is a more recent distributed computing framework that has gained popularity for its ability to process data in-memory, making it faster and more efficient than Hadoop for certain types of ML tasks. Spark supports a wide range of ML algorithms through its MLlib library, including classification, regression, clustering, and recommendation algorithms. One of the key advantages of Spark is its support for iterative ML tasks, where the model needs to be updated multiple times during training.



### 5.2 Cloud-based Machine Learning Platforms

Cloud platforms, such as AWS, Google Cloud, and Microsoft Azure, offer scalable infrastructure that can be used to train and deploy ML models for Big Data analytics. These

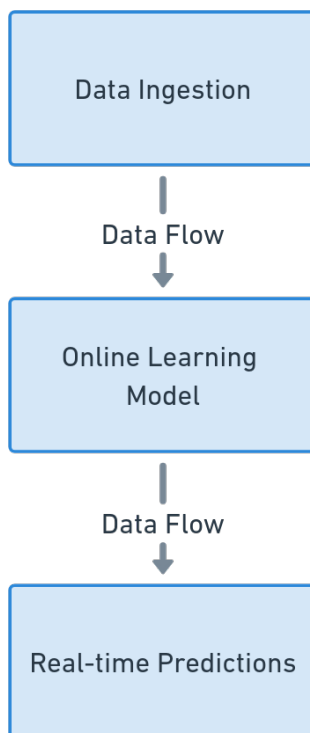
platforms provide a range of services, including data storage, distributed computing, and machine learning tools, allowing organizations to scale their ML models without having to manage their own hardware.

For example, AWS offers services like SageMaker, which provides a fully managed environment for building, training, and deploying ML models at scale. Google Cloud offers similar services through its AI Platform, which supports distributed training and hyperparameter tuning for ML models. These cloud-based platforms allow organizations to scale their ML models by automatically provisioning resources based on the needs of the workload, reducing the time and cost associated with managing infrastructure.

### **5.3 Online Learning and Incremental Learning Techniques**

Traditional ML algorithms are typically trained on a fixed dataset, with the model parameters being updated during the training process. However, in many Big Data applications, data is continuously generated, and the ML model must be updated in real-time. Online learning algorithms address this challenge by allowing the model to be updated incrementally as new data is received, without requiring a full retraining.


Online learning algorithms, such as stochastic gradient descent (SGD) and reinforcement learning, are particularly well-suited for real-time analytics, as they can process data in small batches and update the model parameters incrementally. These algorithms are commonly used in applications such as fraud detection, recommendation systems, and real-time bidding in online advertising.



#### **5.4 Federated Learning**

Federated learning is a relatively new approach to distributed ML, where the model is trained on decentralized data sources, such as mobile devices or edge servers, without requiring the data to be transferred to a central server. This approach is particularly useful for applications where data privacy is a concern, such as healthcare or financial services, as it allows organizations to train ML models on sensitive data without exposing the data to external systems.

In federated learning, each device or node in the network trains a local version of the model on its own data, and the model updates are then aggregated by a central server to produce a global model. This approach reduces the amount of data that needs to be transferred between devices and servers, making it more efficient and scalable for large-scale ML tasks.



## 6. Opportunities for Innovation in Big Data and Machine Learning

### 6.1 Industry Applications

#### Healthcare

##### Case Study: Predictive Analytics for Sepsis Detection

Sepsis, a life-threatening condition caused by the body's response to infection, is a leading cause of mortality in hospitals. Early detection is critical for improving patient outcomes. The University of Pennsylvania Health System deployed a scalable machine learning model to predict sepsis onset by analyzing electronic health record (EHR) data from over 100,000 patients. The model processes large datasets in real-time, analyzing factors such as patient vitals, lab results, and historical health data.

The predictive model used a combination of decision trees and logistic regression to classify patients based on their likelihood of developing sepsis. Scaling the model was a challenge due to the need for real-time processing of large volumes of EHR data from multiple hospital systems. The use of cloud-based infrastructure, parallel processing, and federated learning allowed the system to process the data efficiently and provide early warning alerts to physicians, leading to a 20% reduction in sepsis-related mortality across the health system.

#### Finance

##### Case Study: Fraud Detection in Credit Card Transactions

The financial services industry is particularly vulnerable to fraud, with credit card fraud being one of the most common threats. A global credit card company implemented a scalable ML solution to detect fraudulent transactions in real time. The model was trained on a dataset of over a billion transactions from customers around the world. The primary challenge was scalability, as the system had to process millions of transactions per second and identify fraudulent activity within milliseconds.

The company used a distributed learning approach leveraging Apache Spark for processing and cloud-based infrastructure for deployment. The ML model, built using a combination of neural networks and ensemble methods (such as gradient boosting), could detect anomalous patterns in transaction data and flag them for further review. The scalable nature of the model allowed the company to reduce false positives and significantly improve fraud detection accuracy, saving millions of dollars annually.

## **Manufacturing and IoT**

### **Case Study: Predictive Maintenance in Aerospace Manufacturing**

A leading aerospace manufacturer faced challenges with equipment failures and unplanned downtime, which affected production schedules and led to financial losses. To address this, the company implemented a scalable predictive maintenance solution powered by machine learning. The system analyzed sensor data from aircraft engines and manufacturing equipment in real-time, identifying patterns that indicated potential failures.

The machine learning model, a hybrid of deep learning and time-series analysis, processed terabytes of sensor data daily from over 200,000 machines. The complexity of scaling such a model lay in the distributed nature of the data sources and the need for real-time analysis. The company deployed the model using Apache Kafka for real-time data streaming and a distributed cloud-based ML platform for training and inference. The solution reduced unplanned downtime by 30%, saving the company millions of dollars annually and improving the reliability of its manufacturing operations.

## **6.2 Future Trends**

### **Automated Machine Learning (AutoML)**

### **Case Study: AutoML for E-commerce Personalization**

An online retail giant needed a way to provide personalized recommendations to millions of users based on their browsing and purchasing behavior. The sheer scale of user data and product catalogs made manual tuning and model selection impractical. The company turned to AutoML tools to automate the process of selecting the best model architecture and hyperparameters for personalized recommendation engines.

By leveraging Google's AutoML platform, the company was able to automatically build, train, and deploy machine learning models that analyzed user behavior in real-time. The AutoML system used a combination of collaborative filtering and deep learning to predict user preferences. The scalable solution improved recommendation accuracy by 15%, leading to a significant increase in customer engagement and sales.

### **Quantum Computing**

#### **Case Study: Quantum Computing for Portfolio Optimization in Finance**

A global investment firm sought to optimize its portfolio strategy using machine learning but struggled with the computational limitations of traditional models when processing massive datasets of financial market data. The firm turned to quantum computing as a potential solution for accelerating the optimization process.

By partnering with a quantum computing startup, the investment firm tested a quantum ML model designed to process large-scale financial data and optimize portfolios based on risk and return. The quantum model was able to process complex correlations and nonlinear relationships between assets faster than classical models. While still in the experimental phase, the pilot project demonstrated that quantum computing could significantly reduce computation times for large-scale ML models in finance, opening the door to new opportunities for portfolio optimization and risk management.



## 7. Conclusion

### 7.1 Summary of Key Findings

Scaling machine learning algorithms for Big Data is a critical challenge that requires innovative solutions and approaches. The sheer size and complexity of Big Data, combined with the need for real-time analytics and distributed data sources, necessitate the development of scalable ML models that can efficiently process large datasets and deliver accurate predictions in real-time.

In this paper, we have explored the key challenges associated with scaling ML algorithms for Big Data, including high-dimensionality, distributed data sources, real-time analytics, and hardware limitations. We have also discussed a range of scalable approaches, including parallel and distributed computing, cloud-based ML platforms, online learning, and federated learning, which aim to address these challenges.

Finally, we have highlighted several opportunities for innovation in the field of scalable ML, including the use of AutoML and quantum computing to further enhance the scalability and efficiency of ML models.

### 7.2 Recommendations for Future Research

Future research in the field of scalable machine learning should focus on improving the efficiency of existing algorithms, exploring new approaches to distributed computing and data integration, and leveraging emerging technologies such as quantum computing to process large datasets more efficiently.

Researchers should also focus on developing new tools and techniques for real-time analytics, as the demand for real-time decision-making continues to grow in industries such as healthcare, finance, and manufacturing. Finally, there is a need for continued research into the ethical and privacy implications of using scalable ML models, particularly in applications where sensitive data is involved, such as healthcare and financial services.

## References

1. Dean, J., & Ghemawat, S. (2008). MapReduce: Simplified data processing on large clusters. *Communications of the ACM*, 51(1), 107-113.
2. Zaharia, M., Chowdhury, M., Franklin, M. J., Shenker, S., & Stoica, I. (2010). Spark: Cluster computing with working sets. *Proceedings of the 2nd USENIX conference on Hot topics in cloud computing* (Vol. 10, p. 10).
3. Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., ... & Zheng, X. (2016). TensorFlow: A system for large-scale machine learning. *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, 265-283.
4. Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., ... & Zhao, S. (2021). Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*, 14(1), 1-210.
5. Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. *Science*, 349(6245), 255-260.
6. Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785-794.
7. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., ... & Duchesnay, É. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12(Oct), 2825-2830.
8. Konečný, J., McMahan, H. B., Yu, F. X., Richtárik, P., Suresh, A. T., & Bacon, D. (2016). Federated learning: Strategies for improving communication efficiency. *arXiv preprint arXiv:1610.05492*.
9. Bottou, L. (2010). Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT'2010*(pp. 177-186). Physica-Verlag HD.
10. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436-444.
11. Aggarwal, C. C. (2016). *Data Mining: The Textbook*. Springer.
12. Chen, Z., & Zhang, W. (2014). A scalable machine learning system for big data analytics. *IEEE Access*, 2, 543-557.
13. Dietterich, T. G. (2000). Ensemble methods in machine learning. In *International workshop on multiple classifier systems* (pp. 1-15). Springer, Berlin, Heidelberg.

14. López, V., Fernández, A., García, S., Palade, V., & Herrera, F. (2013). An insight into classification with imbalanced data: Empirical results and current trends on using data intrinsic characteristics. *Information Sciences*, 250, 113-141.
15. Li, M., Andersen, D. G., Smola, A. J., & Yu, K. (2014). Communication efficient distributed machine learning with the parameter server. In *Advances in neural information processing systems* (pp. 19-27).
16. Cui, Y., Zhong, H., & Shi, Y. (2018). Distributed machine learning for big data. *Informatics*, 5(4), 38.
17. Verma, S., & Chandra, S. (2017). Big data analytics and its applications in IoT. In *2017 International Conference on Advances in Computing, Communications and Informatics (ICACCI)*, 1296-1301.
18. Meng, X., Bradley, J., Yavuz, B., Sparks, E., Venkataraman, S., Liu, D., ... & Zadeh, R. (2016). Mllib: Machine learning in Apache Spark. *Journal of Machine Learning Research*, 17(1), 1235-1241.
19. Gulli, A., & Pal, S. (2017). *Deep learning with Keras*. Packt Publishing Ltd.
20. Bertsimas, D., & Dunn, J. (2017). Machine learning under a modern optimization lens. *Optimization and Machine Learning*, 1-27.