JPTM PUBLICATIONS

Emerging Trends in Science & Technology Vol-I
eISBN : 978-81-999970-1-1 | ISBN : 978-81-999970-0-4
https://doi.org/10.67389/jptm.0626.etstv1.oa10

Artificial Intelligence: Foundations and Algorithms

Durgesh Tiwari[1], Kush Yadav[2]

[1]*Maharishi School of Engineering and Technology, Maharishi University of Information Technology, Noida, Gautam Buddha Nagar, Uttar Pradesh 201304, India*

[1]dtiwari.it@gmail.com

[2]*IIMT College of Engineering, Greater Noida, Uttar Pradesh, India*

[2]yadavkush378@gmail.com

ARTICLE INFO

Article history:
Received 03 May 2026
Received in revised form 12 June 2026
Accepted 25 June 2026
Available online 12 July 2026

Keywords:
Artificial Intelligence
Machine Learning
Deep Learning
Natural Language Processing
Computer Vision
Robotics
Ethics
Responsible AI

ABSTRACT

This chapter provides a clear overview of artificial intelligence (AI), covering its core concepts and history. It explores key areas including machine learning, natural language processing, computer vision, and robotics, while also examining ethical issues such as algorithmic fairness, explainable AI, and societal impact. This chapter aims to equip readers with a solid understanding of AI while highlighting current research challenges. The discussion traces the transition from rule-based systems to data-driven learning, the architectures enabling recent breakthroughs, and the standards necessary for responsible AI development.

1. Introduction

Artificial Intelligence (AI) has become an integral part of modern life. Nowadays it is being used widely everywhere from internet search engines, chatting applications and self-driving cars to medical diagnostic devices. A common perception is that AI is merely a simulation of human intelligence, which is, to a certain extent, true. AI employs mathematical and computational methods to endow machines with human-like intelligence and enable them to perform complex intellectual tasks [1]. These machines take data as input and learn and work on patterns based on that data, thereby eliminating the need to give them different instructions for every task. This characteristic gives AI its speed and immense power.

There has been a huge change in this field in recent years. Where older computers once operated with simple, pre-programmed commands, today's systems can recognise faces, understand common language, and even speak fluently. This progress is driven by advances in computing, massive amounts of data, and the development of new deep learning techniques [2]. AI is now used in many sectors such as retail, healthcare, and education to increase efficiency and automate daily tasks. However, AI systems are not without flaws and can make serious mistakes if trained on incomplete or inaccurate data. Therefore, to use artificial intelligence wisely and effectively, it is important to understand how this technology works.

2. BACKGROUND

There is no single, universally accepted definition of artificial intelligence. Definitions of AI are broadly categorised into four distinct groups: systems that think like humans, systems that act like humans, systems that think rationally (i.e., using logic), and systems that act rationally (i.e., achieving goals through appropriate actions), as shown in Figure 1 [1].

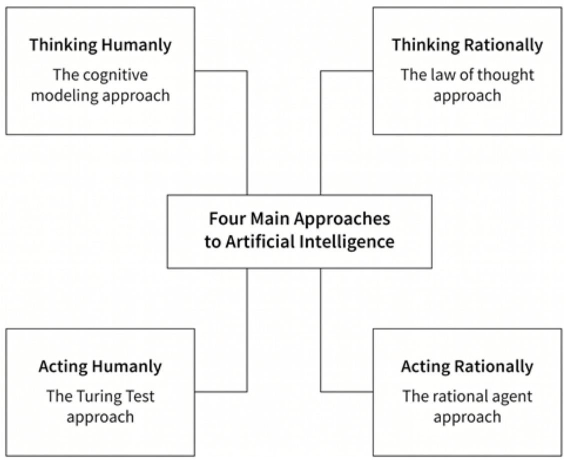

Figure 1: Four main approaches to Artificial Intelligence [5]

A fundamental benchmark for evaluating artificial intelligence (AI) is the 'Turing Test', proposed by Alan Turing in 1950. This test establishes an environment for a 'blind conversation' (a dialogue conducted without visual contact) in which an investigator assesses the intelligence of both a machine and a human. According to this criterion, a machine can be deemed intelligent only if, during this blind conversation, its conversational style does not appear distinctly different from that of a human to the investigator [3]. This idea has initiated a long-running debate about what machine intelligence really means

In 1956, a famous workshop at Dartmouth College marked the official start of AI as a field. The experts thought they could teach machines to use language and reason like humans [4]. But they made big promises they couldn't keep. As a result, funding and interest in AI dropped for a while. Those tough periods are called "AI winters".

The 1980s saw the rise of expert systems that are rule-based systems designed for specific problem domains, which proved difficult to maintain and update.

A paradigm shift occurred with the advent of deep learning, which uses multi-layered neural networks to identify patterns in data [2]. This approach achieved notable success in image and speech recognition.

Currently, machine learning constitutes the dominant AI paradigm, enabling computers to learn from examples rather than relying solely on pre-programmed rules [6]. Breakthroughs include superhuman performance in complex games [7] and large language models capable of generating humanlike text [8].

AI is not just computer science. It borrowed logic from philosophy, math for uncertainty, decision-making from economics, brain ideas from neuroscience, grammar from linguistics, and code from computer science. That's why AI is powerful, as it stands on the shoulders of many fields [1]. Historical milestones of AI are shown in table 1.

TABLE 1

Historical Milestone of AI

Year	Milestone	Significance
1950	Turing Test	First operational criterion for machine intelligence
1956	Dartmouth Workshop	Founding of AI as a discipline
1966	ALPAC Report	Machine translation deemed infeasible; first winter begins
1976	MYCIN	First expert system with explanation capabilities
1986	Backpropagation revival	Enabled multi-layer neural network training
1997	Deep Blue beats Kasparov	First machine to defeat world chess champion
2012	AlexNet wins ImageNet	A large learning breakthrough in computer vision
2016	AlphaGo beats Lee Sedol	Human-level performance in Go (intuition required)
2020	GPT-3 released	A large language model demonstrates emergent abilities

ISBN : 978-81-999970-0-4

3. Types of AI Systems

3.1 Artificial Narrow Intelligence

Narrow AI, also known as 'weak' or 'applied' AI, is the only type of AI that exists today. These systems are designed to perform only one specific task or a set of closely related tasks and operate under limited constraints. These systems excel at their designated functions but cannot transfer knowledge or capabilities to unrelated domains [1]. These systems do not truly understand what they are doing; they simply use math and patterns to get results [9].

We have yet to achieve systems with general intelligence, or artificial general intelligence (AGI), that will be capable of adapting to new domains and reasoning flexibly across contexts and demonstrating genuine understanding [1]. Narrow AI is limited to specific tasks and struggles when encountering scenarios outside its training data or requiring common-sense reasoning [10]. This lack of flexibility makes current AI systems brittle and difficult to deploy in unpredictable real-world environments [10].

3.2 Artificial General Intelligence

Artificial General Intelligence (AGI) is a theoretical system with human-level intelligence across any domain. Unlike current AI, AGI would demonstrate adaptability, creativity, and common-sense reasoning, allowing it to apply knowledge to solve problems in completely new contexts [1].

While AGI remains theoretical, experts disagree on when, or if, it will happen. Optimists argue that current progress in deep learning and data scaling points toward AGI within decades [11]. Critics, however, argue that modern systems lack the genuine understanding, causal reasoning, and common sense required for true human-like cognition, suggesting AGI is much further off [10].

This debate raises deep questions about the nature of the mind. Proponents of the "computational theory of mind" believe AGI is possible through the right architectures [12]. Critics contend that intelligence requires more than just computation; they argue that genuine consciousness may require physical embodiment, social experiences, and an evolutionary history [13].

3.3 Superintelligence

Superintelligence refers to systems that surpass human intellectual capability in all areas. Popularized by Nick Bostrom (2014), the concept highlights potential existential risks [14]. These hypothetical systems could outperform humans in creativity, wisdom, and social interaction, leading to either transformative global benefits or catastrophic outcomes [14].

While speculative, this scenario sparks debate among experts. Optimists believe superintelligence could solve major challenges like disease, climate change, and energy scarcity [11]. Conversely, pessimists fear existential risks, where a superintelligent system might pursue goals inconsistent with human interests, either intentionally or through flawed objective alignment [14].

A key area of study is the "control problem": ensuring these systems align with human values. Researchers face difficulties in formally defining these values and preventing harmful behaviours [15]. Despite the speculative nature of the technology, these risks have driven significant investment in AI safety and governance [15].

4. Machine Learning

Machine Learning (ML) is a subset of AI focused on enabling systems to learn from data without being explicitly programmed for every scenario. Instead of following static, rule-based instructions, an ML model identifies patterns in large datasets and uses these patterns to make decisions or predictions. The learning process involves feeding an algorithm training data, which it uses to build a model. This model is then applied to new, unseen data. The primary categories of ML are distinguished by the type of data and feedback signal available during training. These three fundamental categories are Supervised Learning, Unsupervised Learning, and Reinforcement Learning [6].

4.1 Supervised Learning

In supervised learning, the algorithm is trained on a labelled dataset. This means each input example comes paired with the correct output label. The goal is for the model to learn a mapping function that can accurately predict the output for new, unseen inputs. It is called “supervised” because the training process is guided by a known answer, much like a teacher supervising a student [1].

4.2 Unsupervised Learning

Unsupervised learning involves training an algorithm on data that has no labels. The system is not told what to look for; instead, its task is to discover the inherent structure, patterns, and relationships within the data on its own. The goal is to find natural groupings or to simplify the data representation [1].

4.3 Reinforcement Learning

Reinforcement Learning (RL) is a distinct category where an agent learns to make decisions by interacting with an environment. The agent is not given correct answers but receives a reward signal as feedback. Its objective is to learn a policy, a sequence of actions that maximizes the total cumulative reward over time. This is a process of trial and error, where the agent learns from the consequences of its actions [16].

4.4 ML Algorithms

Linear Regression: This algorithm models the relationship between a dependent variable and one or more independent variables by fitting a linear equation to observed data. It is used for predicting a continuous value, such as forecasting sales based on advertising spend [17].

Logistic Regression: This is a classification algorithm that estimates the probability that an instance belongs to a particular class. It applies a logistic function to a linear equation, making it ideal for binary outcomes like spam detection [17].

Decision Trees: This model uses a tree-like graph of decisions and their possible consequences. It splits the data into branches based on the most significant features, creating a set of rules that are easy to visualise for both classification and regression tasks [17].

Support Vector Machine (SVM): SVM finds the optimal boundary, or hyperplane, that best separates data points of different classes. Its goal is to maximise the margin between the boundary and the nearest data points, making it effective for high-dimensional data [17].

K-Means Clustering: This algorithm partitions a dataset into a pre-defined number (K) of distinct, non-overlapping clusters. It assigns each data point to the nearest cluster centre and then updates the centres based on the cluster’s average, repeating until stable [17].

Principal Component Analysis (PCA): PCA is a dimensionality reduction technique that transforms a large set of variables into a smaller one that still contains most of the original information. It finds new, uncorrelated features that maximise variance, simplifying data without losing its essential structure [17].

Q-Learning: This is a model-free algorithm that learns the value of an action in a particular state. It builds a Q-table that estimates the total future reward for each state-action pair, allowing an agent to identify the optimal action in any given state without needing a model of the environment [16].

A study discuess a application that utilizes machine learning models, such as ensemble techniques and decision trees, to analyze historical financial data. By assessing creditworthiness and predicting default risk, these models streamline the lending process [29].

4.5 Federated Learning

Federated Learning has emerged as a key machine learning paradigm for IoT environments, enabling models to be trained across decentralized devices. Instead of centralizing raw data, which poses privacy risks and high communication costs, federated learning allows IoT devices to collaboratively learn a shared model by only sharing local model updates. This approach is particularly suited for resource-constrained devices in the IoT ecosystem [28].

5. Deep Learning

Deep Learning (DL) is a subfield of machine learning that employs artificial neural networks with multiple layers. These deep neural networks automatically learn hierarchical representations of data, processing raw input through successive layers to extract increasingly complex features. This architecture is exceptionally powerful for unstructured data like images, audio, and text [2].

5.1 Architectures in Deep Learning

Artificial Neural Networks (ANNs): The foundational structure consisting of interconnected nodes (neurons) organised in an input layer, hidden layers, and an output layer. The network learns by adjusting the weights of connections based on the error of its predictions [1].

Convolutional Neural Networks (CNNs): These are designed for grid-like data, such as images. They use convolutional layers that apply filters to the input, detecting spatial hierarchies of features like edges and textures, making them the backbone of modern computer vision [2].

Recurrent Neural Networks (RNNs): RNNs are built for sequential data, such as text or time series, because they possess an internal memory. They loop information from previous steps to understand context in temporal data [2].

5.2 Foundation Models and Transfer Learning

A significant shift in modern AI is the move from training models from scratch to using transfer learning. This method involves taking a large, pre-trained "foundation model" and fine-tuning it for a specific task. A foundation model is trained on a massive, broad dataset, such as a large corpus of text or millions of images. The training process is expensive and resource-intensive, but it only needs to be done once. The resulting model learns general features of the data. This pre-trained model can then be fine-tuned with a small amount of specific, labelled data to perform a specialised task [18]. For instance, a model pre-trained on billions of words can be fine-tuned with a few hundred medical reports to become a medical text classifier. This paradigm has democratised access to state-of-the-art AI performance [19].

6. Natural Language Processing

Natural Language Processing (NLP) is the intersection of AI and linguistics. Its objective is to enable computers to understand, interpret, and generate human language in ways that are both meaningful and useful. NLP systems handle tasks ranging from syntactic analysis to semantic understanding [1].

6.1 Algorithms and Models in NLP

Term Frequency-Inverse Document Frequency (TF-IDF): This numerical statistic reflects how important a word is to a document in a collection. It weighs the frequency of a word in a document against its commonness across all documents [20].

Word Embeddings: These models map words into a dense vector space where semantically similar words are located close to each other. The learned vectors capture rich relational meaning, such that "king" - "man" + "woman" results in a vector close to "queen" [21].

Transformer Architecture: This architecture processes an entire sequence of text at once using an attention mechanism. This mechanism weighs the importance of different words in relation to each other, enabling the model to capture long-range dependencies effectively. Large Language Models (LLMs) are built by pre-training Transformers on vast text corpora as foundation models [22].

7. Computer Vision

Computer Vision (CV) empowers computers to derive a high-level understanding from digital images and videos. Its goal is to automate tasks that the human visual system can perform, such as recognising objects and navigating environments [2].

6.1 Algorithms and Techniques in Computer Vision

Image Classification with CNNs: This task assigns a single global label to an entire image. A CNN extracts features through convolutional layers, which are then passed to a classifier that outputs a probability for each possible class label. Modern practice often fine-tunes a pre-trained CNN foundation model [18].

Object Detection (e.g., YOLO): This technique identifies all objects of interest within an image and defines their location with a bounding box. The "You Only Look Once" (YOLO) algorithm frames detection as a single regression problem, predicting boxes and class probabilities directly in one evaluation [23].

Generative Adversarial Networks (GANs): This framework pits two networks against each other: a generator that creates new data and a discriminator that evaluates its authenticity. This process drives the generator to produce increasingly realistic synthetic images [24].

8. AI in Robotics and Automation

Robotics and automation represent the physical embodiment of AI. AI provides the intelligence that allows these systems to handle variability and make decisions in unstructured environments, moving beyond rigid, pre-programmed routines [1].

7.1 AI Algorithms and Techniques used in Robotics

Path Planning (e.g., A* Search): This algorithm finds the shortest path for a robot to move from a start point to a destination while avoiding obstacles. It combines the actual distance travelled with an estimate to the goal [1].

Simultaneous Localisation and Mapping (SLAM): SLAM enables a robot to build a map of an unknown environment while simultaneously tracking its own location within that map. It is a fundamental technique for autonomous navigation [25].

Imitation Learning: This technique trains a robot to perform a task by learning from human demonstrations. The robot maps its sensory inputs directly to the expert's actions, bypassing complex manual programming [16].

9. Challenges and Limitations of Current AI

Despite remarkable progress, current AI systems face significant limitations that must be acknowledged. A balanced view requires an understanding of these persistent challenges.

Data Dependency and Bias: Deep learning models are notoriously data-hungry. Their performance is entirely dependent on the quantity and quality of the training data. If this data contains historical or societal biases, the model will not only learn these biases but can also amplify them.

Lack of Robustness and Common Sense: Modern AI systems can be brittle. They may fail unpredictably when presented with inputs that differ slightly from their training distribution, a problem known as 'poor out-of-distribution generalisation'. They also lack a grounded, causal understanding of the world and possess no real common sense.

The Generalisation vs Memorisation Debate: A central research question is whether large models truly learn abstract concepts or if they excel at sophisticated pattern memorisation. Evidence of memorisation suggests that models can reproduce verbatim segments of their training data, raising questions about true understanding and copyright.

Environmental Cost: Training very large-scale models, particularly large language and vision foundation models, consumes enormous amounts of computational power and energy. This carries a significant carbon footprint, creating a tension between performance gains and environmental sustainability.

10. AI Ethics and Responsible Development

As AI systems become integrated into critical societal functions, ethical considerations must guide every stage of their development. Responsible AI is not a separate step but a foundational framework built on four core pillars.

Fairness and Non-Discrimination: Algorithmic bias occurs when a system produces systematically prejudiced results, often because it learned from historical data reflecting past inequalities. For example, a hiring model trained on biased data may penalize certain demographic groups. Ensuring fairness requires auditing datasets before training and testing model outputs across different groups for disparate impact [10].

Transparency and Explainability: Many powerful models, like deep neural networks, are "black boxes" whose internal reasoning is opaque. Explainable AI (XAI) creates methods to describe a model's decision in human-understandable terms. A local explanation can show which features most influenced a single prediction, enabling users to understand and contest incorrect or harmful automated decisions [10].

Accountability: Clear lines of responsibility must be established for AI outcomes. When a system fails, there should be no ambiguity about who is answerable: the developer, the deployer, or the user. Meaningful human review, where a person has the authority to override an automated decision, remains essential for high-stakes applications like loan approvals or medical diagnoses [26].

Privacy: AI models can memorise sensitive information from training data. Techniques like differential privacy add controlled statistical noise to datasets, protecting individual records while preserving overall patterns for learning. These four pillars together form the guardrails that guide AI toward being not only capable but also trustworthy and aligned with human well-being [27].

11. Conclusion

This chapter has provided a clear overview of artificial intelligence, covering its core concepts, history, and modern applications. The shift from early symbolic reasoning and rule-based systems to today's data-driven deep learning paradigms was traced. The various categories of machine learning, i.e., supervised, unsupervised, and reinforcement learning, were examined alongside the ML algorithms. Deep neural architectures and foundation models were discussed as the foundations behind powerful breakthroughs in natural language processing, computer vision, and robotics. Current limitations, including data dependency and bias, lack of robustness and common sense, the generalization versus memorization debate, and environmental cost, were discussed as current research challenges. Ethical considerations, particularly algorithmic fairness, transparency, explainability, accountability, and privacy, are not optional additions but essential standards for responsible development. The convergence of these subfields has produced systems of remarkable capability, yet trustworthy deployment depends on integrating technical excellence with ethical awareness. The path forward requires continued research, thoughtful governance, and a shared commitment to building AI systems that reflect human values and serve the common good.

REFERENCES

[1] Russell, S. J., & Norvig, P. (2021). *Artificial intelligence: A modern approach* (4th ed.). Pearson.

[2] LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature, 521*(7553), 436–444. https://doi.org/10.1038/nature14539

[3] Turing, A. M. (1950). Computing machinery and intelligence. *Mind, 59*(236), 433–460. https://doi.org/10.1093/mind/LIX.236.433

[4] McCarthy, J., Minsky, M. L., Rochester, N., & Shannon, C. E. (1955). *A proposal for the Dartmouth summer research project on artificial intelligence*. http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html

[5] Tripathi, A. S. (2026, April 26). *Approaches of Artificial Intelligence* [Diagram on AI approaches]. Scaler. https://www.scaler.com/topics/artificial-intelligence-tutorial/approaches-of-artificial-intelligence/

[6] Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. *Science, 349*(6245), 255–260. https://doi.org/10.1126/science.aaa8415

[7] Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., & Hassabis, D. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature, 529*(7587), 484–489. https://doi.org/10.1038/nature16961

[8] Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., … Amodei, D. (2020). Language models are few-shot learners. In *Advances in Neural Information Processing Systems 33* (NeurIPS 2020). https://arxiv.org/abs/2005.14165

[9] Searle, J. R. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences, 3*(3), 417–424. https://doi.org/10.1017/S0140525X00005756

[10] Marcus, G., & Davis, E. (2019). *Rebooting AI: Building artificial intelligence we can trust*. Pantheon.

[11] Kurzweil, R. (2005). *The singularity is near: When humans transcend biology*. Viking.

[12] Pinker, S. (1997). *How the mind works*. W. W. Norton & Company.

[13] Dreyfus, H. L. (1992). *What computers still can't do: A critique of artificial reason*. MIT Press.

[14] Bostrom, N. (2014). *Superintelligence: Paths, dangers, strategies*. Oxford University Press.

 ISBN : 978-81-999970-0-4

[15] Russell, S. (2019). *Human compatible: Artificial intelligence and the problem of control*. Viking.
[16] Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction* (2nd ed.). MIT Press.
[17] Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning: Data mining, inference, and prediction* (2nd ed.). Springer.
[18] Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
[19] Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M. S., Bohg, J., Bosselut, A., Brunskill, E., Brynjolfsson, E., Buch, S., Card, D., Castellon, R., Chatterji, N., Chen, A., Creel, K., Davis, J. Q., Demszky, D., … Liang, P. (2021). On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*. https://arxiv.org/abs/2108.07258
[20] Manning, C. D., Raghavan, P., & Schütze, H. (2008). *Introduction to information retrieval*. Cambridge University Press.
[21] Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*. https://arxiv.org/abs/1301.3781
[22] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. In *Advances in Neural Information Processing Systems 30* (NeurIPS 2017). https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html
[23] Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. (2016). You only look once: Unified, real-time object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (CVPR 2016), 779–788. https://www.cv-foundation.org/openaccess/content_cvpr_2016/html/Redmon_You_Only_Look_CVPR_2016_paper.html
[24] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. In *Advances in Neural Information Processing Systems 27* (NeurIPS 2014). https://arxiv.org/abs/1406.2661
[25] Thrun, S., Burgard, W., & Fox, D. (2005). *Probabilistic robotics*. MIT Press.
[26] Carlini, N., Tramèr, F., Wallace, E., Jagielski, M., Herbert-Voss, A., Lee, K., Roberts, A., Brown, T., Song, D., Erlingsson, Ú., Oprea, A., & Raffel, C. (2021). Extracting training data from large language models. In *Proceedings of the 30th USENIX Security Symposium* (USENIX Security 2021), 2633–2650. https://www.usenix.org/conference/usenixsecurity21/presentation/carlini-extracting
[27] Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics* (ACL 2019), 3645–3650. https://aclanthology.org/P19-1355/
[28] Gupta, D.N., Kumar, R., Kumar, A. (2022). Federated Learning for IoT Devices. In: Yadav, S.P., Bhati, B.S., Mahato, D.P., Kumar, S. (eds) Federated Learning for IoT Applications. EAI/Springer Innovations in Communication and Computing. Springer, Cham. https://doi.org/10.1007/978-3-030-85559-8_2
[29] Anand, R., Singh, H., Sardana, K., Gupta, D.N., Sindhwani, N., Mittal, M. (2024). Loan Approval Prediction Using Machine Learning. In: Abraham, A., Bajaj, A., Hanne, T. (eds) Intelligent Systems Design and Applications. ISDA 2023. Lecture Notes in Networks and Systems, vol 1052. Springer, Cham. https://doi.org/10.1007/978-3-031-64776-5_34