

# Machine Learning: Advances, Views, and Opportunities

Paras Nath Mishra, Assistant Professor

Department of Computer Science, Arka Jain University, Jamshedpur, Jharkhand, India

Email [Id-paras.m@arkajainuniversity.ac.in](mailto:Id-paras.m@arkajainuniversity.ac.in)

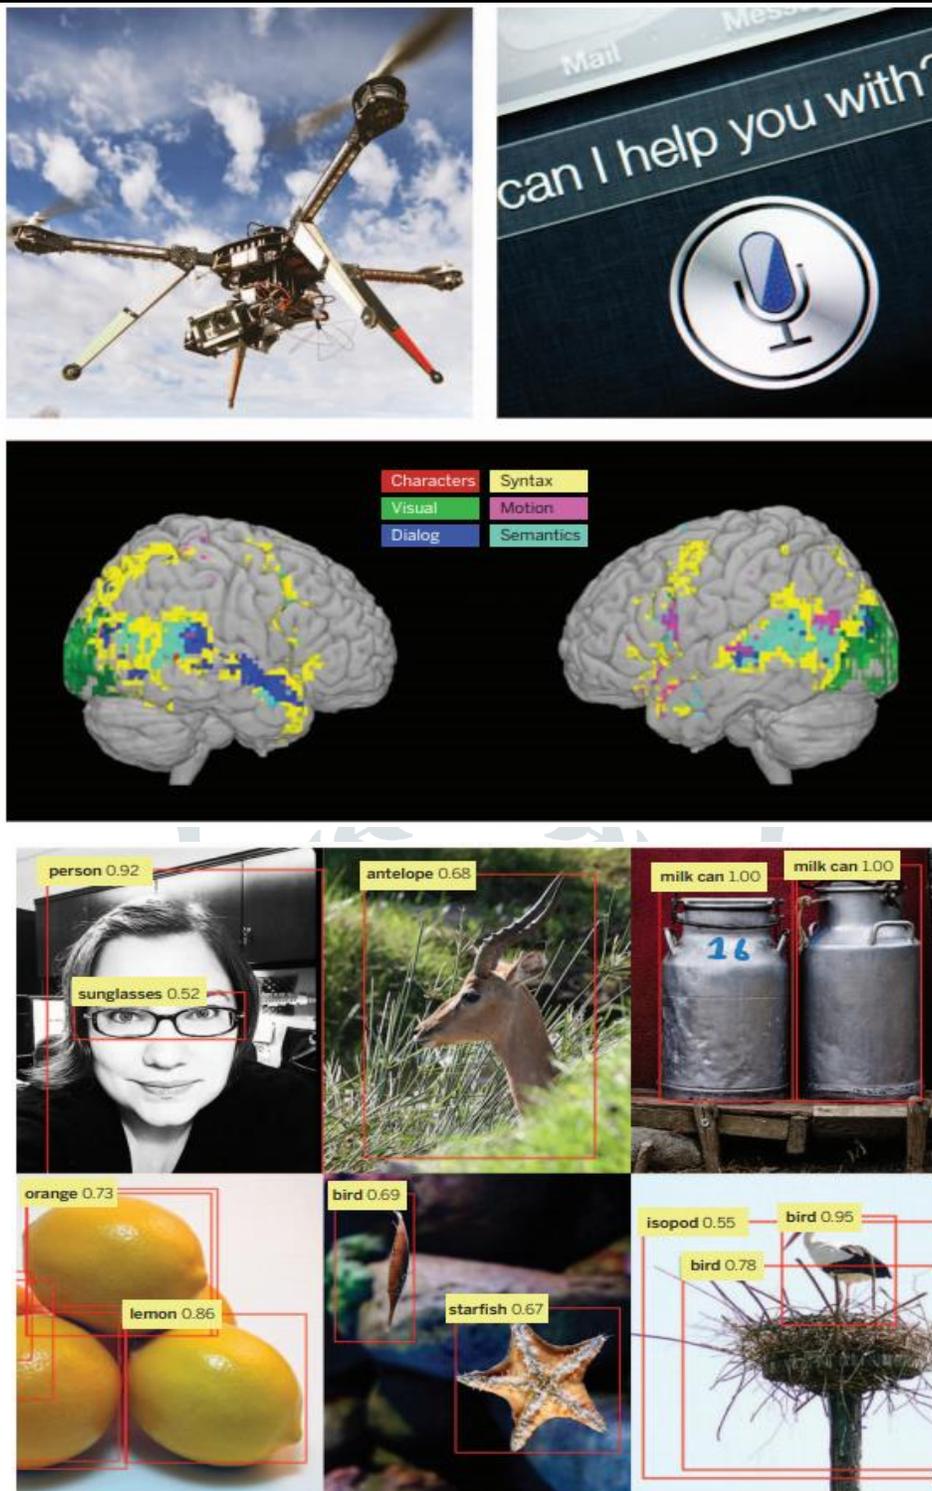
**ABSTRACT:** Machine learning research is essential both for answering these basic scientific and technical issues and for the very practical computer software it has developed and deployed in a variety of applications. Machine Learning (ML) aims to answer the issue of how to create machines that learn on their own. It is at the confluence of computer science and statistics, as well as at the heart of artificial intelligence and data science, and is one of today's fastest developing technological areas. The development of novel learning algorithms and theory, as well as the continuing growth in the availability of online data and low-cost computing, have fuelled recent advances in machine learning. Data-intensive machine-learning techniques are increasingly being used in research, technology, and commerce, resulting in more evidence-based decision-making in a variety of fields such as medical care, industry, education, financial planning, police, and advertisement. In this paper we have discussed the current advances, different views and future opportunities of the ML.

**KEYWORDS:** Deep Learning, Learning Algorithms, Machine Learning, Supervised Learning, Un-Supervised Learning.

## 1. INTRODUCTION

Machine learning (ML) is a branch of computer science that focuses on two interrelated questions: how to build computer systems that improve themselves over time, and what are the fundamental, statistical, and computational information-theoretic laws that govern all learning systems, including computers, humans, and organizations [1], [2]. Machine learning has come a long way in the last two decades, from a laboratory curiosity to a viable technology with broad commercial use. Machine learning has emerged as the preferred technique for creating practical software for computer vision, voice recognition, natural language processing, robot control, and other applications in artificial intelligence (AI) [3]. Many AI system developers now realize that, for many purposes, training a system by giving it instances of desirable input-output behaviour may be much simpler than programming it manually by predicting the correct response for all potential inputs. Machine learning has had a widespread impact on computer science and a variety of sectors that deal with data-intensive problems, such as consumer services, defect diagnostics in complex systems, and logistics chain management. Machine-learning techniques have been created to evaluate high throughput experimental data in new ways, resulting in a similarly wide variety of impacts throughout empirical disciplines, from biology to cosmology to social science. Some current areas of machine learning use are shown in Fig. 1.

A learning issue is described as the difficulty in increasing some measure of performance while doing a task via some kind of training experience. The goal of learning to identify credit-card fraud, for example, is to classify each credit-card transaction as "fraud" or "not fraud." The accuracy of this fraud classifier may be enhanced, and the training experience could consist of a collection of past credit-card transactions, each classified as fraudulent or not in hindsight. Alternatively, a new performance measure may be defined that imposes a greater penalty when "fraud" is mistakenly labelled "not fraud" rather than when "not fraud" is incorrectly labelled "fraud." A new kind of training experience might be defined by adding unlabelled credit-card transactions with labelled instances, for example.



**Fig. 1: Illustrates machine learning applications [4]. Machine learning is having a significant impact on a wide range of technological and scientific fields.**

To address the broad range of data and issue types encountered in machine-learning challenges, a varied set of machine-learning algorithms has been created. Machine-learning algorithms may be thought of as searching through a vast space of candidate programs to find one that maximizes the performance measure, guided by training experience. The way in which machine-learning algorithms represent candidate programs (e.g., mathematical functions, decision trees, and general purpose programming languages) and the way in which they explore throughout this collection of programs differ greatly. We'll concentrate on methods that have shown to be especially effective in the past. Many algorithms focus on function approximation problems, in which the task is embodied in a function (e.g., output a "fraud" or "not fraud" label given an input transaction), and the learning problem is to enhance the precision of that feature, with expertise comprising of a specimen of known input-output combinations of the function [5]. In some instances, the function is explicitly expressed as a

parameterized functional form; in others, the function is implicit and derived via a search process, factorization, optimization, or simulation. Even though the function is implicit, it is usually dependent on parameters or other adjustable degrees of freedom, and training is the process of determining the best values for these parameters to maximize the performance measure.

A major scientific and practical objective, regardless of the learning algorithm, is to conceptually define the capabilities of particular learning algorithms as well as the intrinsic complexity of each given learning problem: How well can an algorithm learn from a certain kind and quantity of training data? What is the algorithm's resiliency to mistakes in its modelling assumptions or in the training data? Is it feasible to develop a successful solution for a learning issue with a certain amount of training data, or is this learning problem inherently intractable? Statistical decision theory and computational complexity theory are often used in such theoretical characterizations of machine-learning algorithms and issues. In fact, attempts to theoretically characterize machine-learning algorithms have resulted in a mix of statistical and computational theory in which the goal is to simultaneously characterize the sample complexity (how much data is required to learn quickly and precisely) and the computational complexity (how much calculation is needed) and to explicitly state how these rely on features of the learning algorithm such as the number of steps in the learning algorithm. Optimization theory, with upper and lower limits on rates of convergence of optimization methods combining nicely with the framing of machine-learning issues as the optimization of a performance measure, has proven especially helpful in recent years.

Machine learning is an area of research that straddles computer science, statistics, and a number of other disciplines concerned with continuous improvement, inference, and decision-making under uncertainty. The psychological study of human learning, evolution, and adaptive control theory, the study of educational methods, neuroscience, organizational behaviour, and economics are all related subjects. Despite increasing crosstalk with these other disciplines over the last decade, we are just now starting to harness the potential synergies and variety of formalisms and experimental techniques employed across these various fields for investigating systems that develop with training.

## 2. DISCUSSION

### 2.1. Drivers of Machine-Learning Progress:

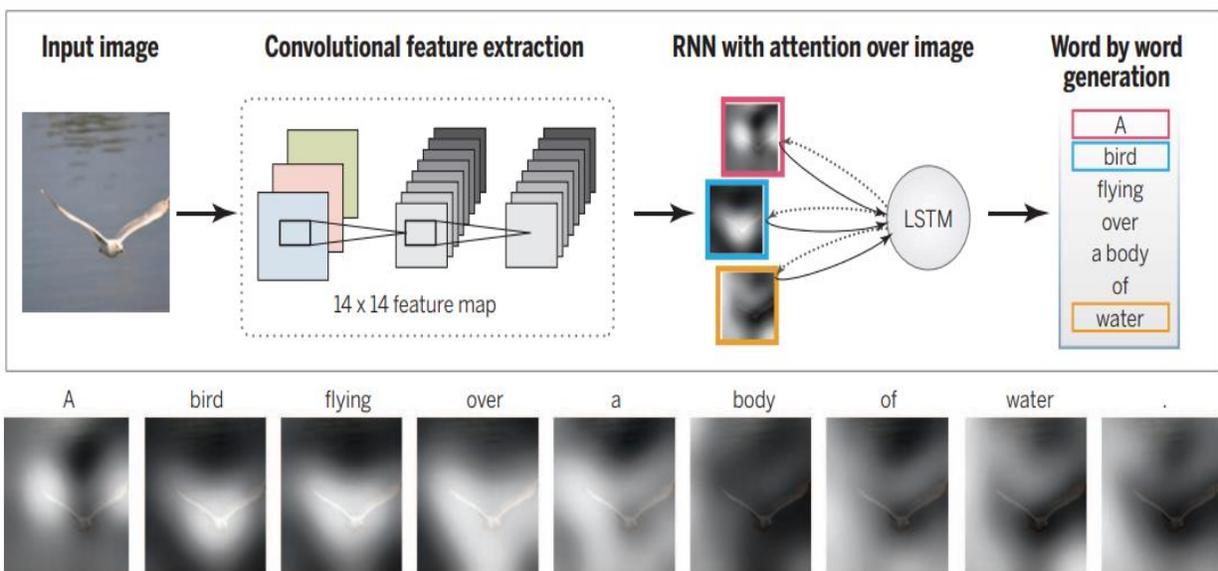
The capacity of networked and mobile computer systems to collect and transmit large quantities of data has rapidly increased over the last decade, a phenomenon known as "Big Data" [6]. Machine learning has frequently been used by scientists and engineers who gather such data to solve the issue of getting valuable insights, predictions, and choices from such data sets. Indeed, the sheer quantity of the data necessitates the development of scalable methods that include both computational and statistical considerations, but the problem is more than just the size of contemporary data sets; it is the granular, individualized character of much of it. Large quantities of data on individual people may be collected via mobile devices and embedded computers, and machine-learning algorithms can learn from this data to tailor services to the requirements and circumstances of each individual. Furthermore, these customized services may be linked to form a larger service that takes use of the richness and variety of data from many people while still tailoring to their specific wants and situations. Many areas of business, research, and government have examples of this tendency toward collecting and processing huge amounts of data to enhance services and productivity. Historical medical records are used to figure out which patients will respond best to which treatments; historical traffic data is used to improve traffic control and reduce congestion; historical crime data is used to help local police officers assign themselves to specific locations at specific times; and large experimental data sets are captured and curated to speed up progress in biology, astronomy, and neuroscience. Many areas of research, business, and government seem to be at the start of a decades-long trend toward more data-intensive, evidence-based decision-making.

Large-scale data's growing importance in all fields of human activity has resulted in a slew of increasing requirements on the underpinning machine learning techniques. Huge data sets, for example, need computationally tractable algorithms, highly personal data necessitates algorithms that minimize privacy implications, and the availability of massive amounts of unlabelled data necessitates the development of learning algorithms to take use of it. The next sections look at how these needs have influenced current progress in machine-learning techniques, theory, and implementation.

## 2.2. Core Methods and Recent Progress:

Supervised learning techniques are the most commonly utilized machine-learning approaches. The function approximation problem is exemplified by supervised learning systems, such as e-mail spam classifiers, face recognizers over images, and medical diagnosis systems for patients, where the training data is in the form of a collection of  $(x, y)$  pairs and the goal is to produce a prediction  $y^*$  in response to a query  $x^*$  [7]. The  $x$  inputs may be simple vectors or more complicated things like texts, pictures, DNA sequences, or graphs. Likewise, several other types of output  $y$  have been investigated. Many advances have been made by focusing on the simple binary classification problem, in which  $y$  has one of two values, for example, “spam” or “not spam”, but there has also been a lot of research on problems like multiclass classification (where  $y$  has one of  $K$  labels), multi-label classification (where  $y$  is labelled by several of the  $K$  labels at the same time), and ranking problems (where  $y$  is a combinatorial object such as a graph, whose components may be required to satisfy some set of constraints). Part-of-speech (POS) tagging is an example of the latter issue, in which the objective is to identify every word in an input phrase  $x$  as a noun, verb, or other part of speech at the same time. Cases in which  $y$  contains real-valued components or a combination of discrete and real-valued components are likewise covered by supervised learning.

In most cases, supervised learning systems make predictions using a learnt mapping  $f(x)$ , which generates an output  $y$  for each input  $x$ . Decision trees, decision forests, logistic regression, support vector machines, neural networks, kernel machines, and Bayesian classifiers are all examples of mapping  $f$ . To estimate these many kinds of mappings, a number of learning algorithms have been suggested, as well as general methods like as boosting and multiple kernel learning, which integrate the results of several learning algorithms. The particular structure of machine learning issues (e.g., that the objective function or function to be integrated is typically the sum over a large number of terms) drives advances in procedures for learning  $f$  from data, with concepts from optimization theory or numerical analysis often used. This diversity of learning architectures and algorithms reflects the varying needs of applications, with different architectures capturing different kinds of mathematical structures, offering varying levels of amenability to post-hoc visualization and explanation, and providing varying trade-offs between computational complexity, data size, and performance.



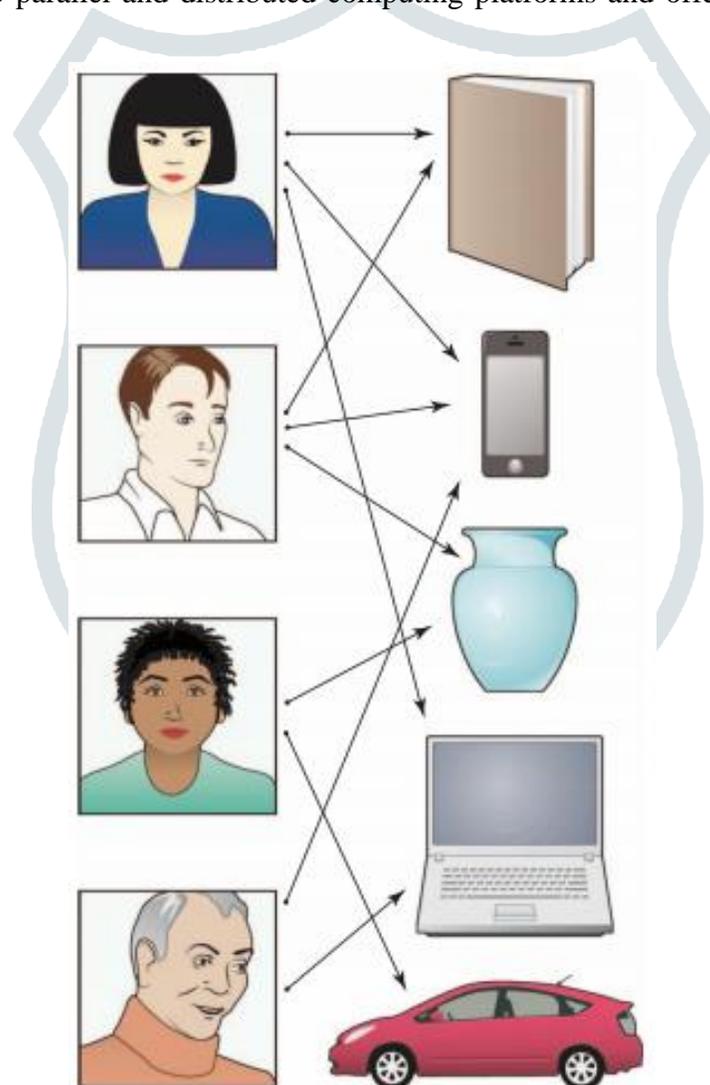
**Fig. 2: Deep networks are used to automatically generate text descriptions for pictures [4]. The output of a convolutional neural network trained to analyse pictures is then utilized by a recurrent neural network trained to produce a text caption. The bottom clip illustrates the network's word-by-word emphasis on various portions of the input picture as it produces the caption word-by-word.**

Deep networks, which are multilayer networks of threshold units, each of which computes some basic parameterized function of its inputs, have been a high-impact area of development in supervised learning in recent years [8]. Gradient-based optimization methods are used by deep learning systems to change parameters across a multi-layered network depending on mistakes at its output. It has been feasible to develop deep learning systems with billions of specifications and that can be provided with training on a wide ranges of pictures, video

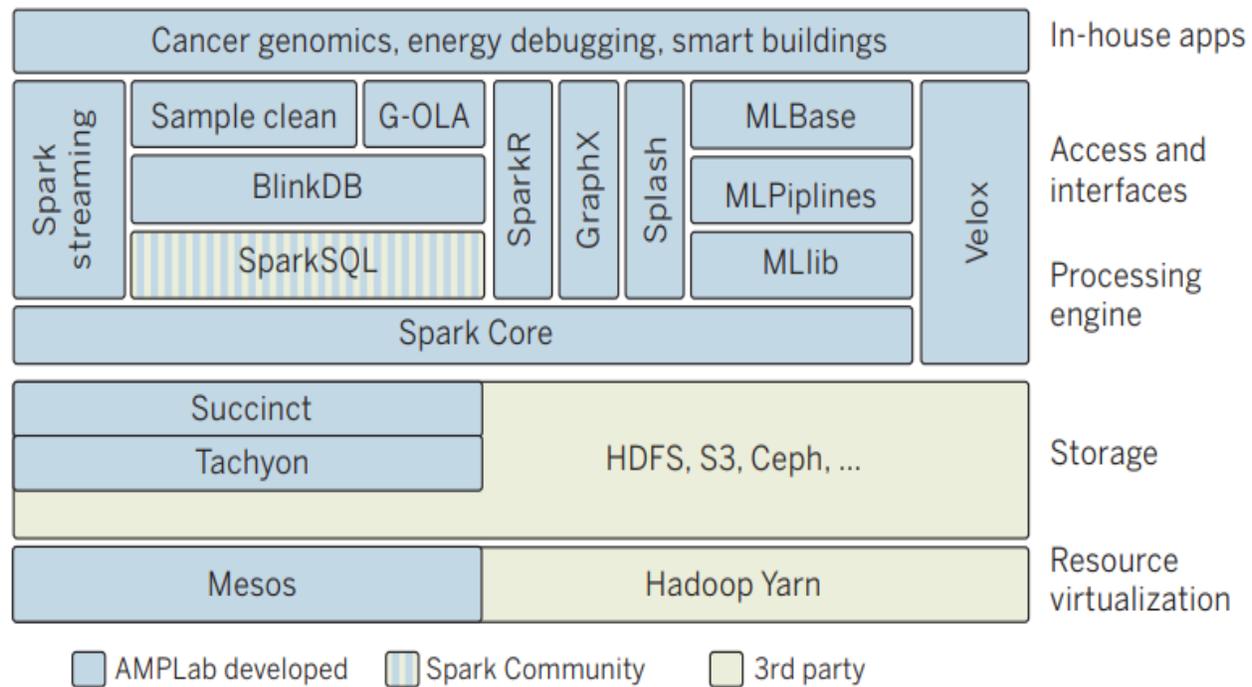
files, and speech samples available on the web by utilizing modern parallel computing frameworks, such as graphics processing units initially created for computer games. As illustrated in Fig. 2, large-scale deep learning systems have had a big impact in recent years in computer vision and voice recognition, where they have produced significant performance gains over prior methods. From natural language translation to collaborative filtering, deep network techniques are being actively explored in a range of other applications.

### 2.3. Emerging Trends:

Machine learning is still a very new subject, and it is quickly growing, typically by creating new formalizations of machine-learning issues motivated by practical applications. The evolution of recommendation systems, as shown in Fig. 3, is an example [9]. A rising worry for the environment in which a machine-learning algorithm works is one key trend driving this growth. The term "environment" refers to the computing architecture in part; whereas a traditional machine-learning system consisted of a single application running on a single machine, it is now prevalent for machine-learning techniques to be implemented in architectures with several thousand or tens of thousands of processing units, posing communication constraints and parallelism and distributed process issues. Machine-learning systems, as shown in Fig. 4, are rapidly becoming sophisticated collections of software that operate on large-scale parallel and distributed computing platforms and offer a variety of methods and services to data analysts.



**Fig. 3: Shows how a recommendation system works [4]. A recommendation system is a machine-learning system based on data that shows connections between a group of users and a group of products.**



**Fig. 4: This diagram depicts the data analytics stack [10]. Scalable machine-learning systems are developed on parallel and distributed computing platforms and have layered architectures.**

The term "environment" also refers to the data's source, which could be a group of people with privacy or ownership concerns, an analyst or decision-maker with specific requirements for a machine-learning system, and the social, legal, or political framework in which the system is deployed. Other machine learning systems or agents may be present in the environment, and the total collection of systems may be cooperative or hostile. Environments, in general, offer different resources to a learning algorithm while also imposing restrictions on those resources. Machine-learning researchers are increasingly formalizing these connections, with the goal of creating algorithms that are provably successful in a variety of settings and explicitly enable users to express and manage resource trade-offs.

### 3. CONCLUSION

Despite its pragmatic and commercial achievements, machine learning is still a relatively new subject with a lot of untapped research potential. By comparing existing machine-learning methods to the kinds of learning we see in naturally occurring systems such as people and other animals, organizations, economics, and biological evolution, we can identify some of these possibilities. Machine learning, like any sophisticated technology, raises issues about which of its potential applications society should promote and which should be discouraged. As previously stated, the current drive to gather new types of personal data, motivated by its economic worth, has resulted in apparent privacy concerns. The growing value of data poses a second ethical question: who will have access to and ownership of internet data, and who will profit from it? Currently, businesses gather a lot of data for particular purposes that lead to increased profits, with little or no incentive for data exchange. However, even from current internet data, the potential benefits to society would be significant if that data were made accessible for public use.

#### REFERENCES

- [1] B. B. Benuwa, Y. Zhan, B. Ghansah, D. K. Wornyo, and F. B. Kataka, "A review of deep machine learning," *International Journal of Engineering Research in Africa*. 2016, doi: 10.4028/www.scientific.net/JERA.24.124.
- [2] V. Kumar and M. L., "Deep Learning as a Frontier of Machine Learning: A Review," *Int. J. Comput. Appl.*, 2018, doi: 10.5120/ijca2018917433.
- [3] A. Holzinger, "From machine learning to explainable AI," 2018, doi: 10.1109/DISA.2018.8490530.
- [4] M. I. Jordan and T. M. Mitchell, "Machine learning: Trends, perspectives, and prospects," *Science*. 2015, doi: 10.1126/science.aaa8415.
- [5] I. Portugal, P. Alencar, and D. Cowan, "The use of machine learning algorithms in recommender systems: A systematic review," *Expert Systems with Applications*. 2018, doi: 10.1016/j.eswa.2017.12.020.
- [6] A. L'Heureux, K. Grolinger, H. F. Elyamany, and M. A. M. Capretz, "Machine Learning with Big Data: Challenges and Approaches,"

*IEEE Access*, 2017, doi: 10.1109/ACCESS.2017.2696365.

- [7] V. Nasteski, "An overview of the supervised machine learning methods," *HORIZONS.B*, 2017, doi: 10.20544/horizons.b.04.1.17.p05.
- [8] S. Zhou, H. Zou, C. Liu, M. Zang, Z. Zhang, and J. Yue, "Deep extractive networks for supervised learning," *Optik (Stuttg.)*, 2016, doi: 10.1016/j.ijleo.2016.07.007.
- [9] C. Lee, O. Kwon, M. Kim, and D. Kwon, "Early identification of emerging technologies: A machine learning approach using multiple patent indicators," *Technol. Forecast. Soc. Change*, 2018, doi: 10.1016/j.techfore.2017.10.002.
- [10] S. Mazumder, "Fast data analytics stack for big data analytics," in *The Human Element of Big Data: Issues, Analytics, and Performance*, 2016.

