

Introduction to Artificial Intelligence: Current Developments, Concerns and Possibilities for Education

Rejaul Karim Barbhuiya

Assistant Professor, Central Institute of Educational Technology, NCERT

Email: rejaul.ncert@nic.in

Abstract

Artificial Intelligence is shaping the world at a speed much faster than anticipated. Machine learning models are extensively used in banking, e-commerce, healthcare, weather forecasting, etc. ChatGPT is the latest entry in the bloc but the most engaging one so far. The world is abuzz with concerns about AI/ML rendering people jobless. The who's who in the world of technology is warning of potential danger in the exponential growth in learning capabilities of intelligent systems and the concerns related to privacy, ethics, safety, and security. Educationists should also be equipped with the basic know-how of AI and its related fields to have a considered opinion while adopting teaching-learning-assessment. This article discusses the concept of Artificial Intelligence in Education (AIED) and introduces the branches of AI, such as Machine Learning, Artificial Neural Networks, and Deep Learning, as well as their working with examples that educators can relate with. The article also highlights some of the ethical concerns associated with AI.

Keywords: Artificial Intelligence in Education (AIED), Machine Learning, Artificial Neural Network, Deep Learning, Ethical concerns with AI.

Introduction to Artificial Intelligence

Artificial Intelligence (AI) is the branch of computer science that studies how to use computers to carry out intelligent tasks that are otherwise done by humans (Huang et al., 2019). The term "Artificial Intelligence" was first proposed as early as 1955 (McCarthy et al., 2006), with the premise that if we can precisely describe the features of human intelligence or intellectual activities such as learning, reasoning, judgment, decision-making, etc., then a machine can be made to reproduce it (Wang, 2019). Therefore, early research in AI tried to decode how intelligence processes work in humans so that those could be reproduced or replicated through computers. This was the trend till the twentieth century. But today's AI systems are designed to work in complex domains and scenarios, irrespective of whether they follow a human-like approach or not (Florida,


2014). AI techniques have evolved over the years to carry out tasks that require some intelligence. Today's advanced AI-based systems are similar to the human mind in certain ways but do not necessarily mimic human behaviour, nor are they identical in all aspects. Different approaches to AI research and their application areas include automated game playing, pattern recognition, data mining, expert systems, neural networks, and the latest deep learning. In 1997, Deep Blue (Campbell et al., 2002) - an AI-based chess-playing software, defeated the then-world chess champion Garry Kasparov. Other examples of AI-based systems we often interact with are:

- Google map - predict traffic, determine the best route, and estimate travel time.
- Personalised recommendation - YouTube recommends videos

based on our search history, Netflix recommends movies/shows based on past viewing data, Amazon/

Flipkart recommends products based on shopping history, and many other parameters.

Figure-1: Relationship between AI, ML, ANN and DL



Many AI subdomains exist today, such as knowledge representation, reasoning, machine learning, computer vision, natural language processing, artificial neural network, deep learning, etc. Figure 1 shows the hierarchical relationship between artificial intelligence, machine learning, neural networks, and deep learning.

The recent progress in the development of artificial intelligence techniques is making our life easier in many ways. They are automating various tasks, bringing efficiency, reducing cost, and playing a vital role in human progress. There is a keen belief among educationists that AI can help improve the teaching-learning processes, too (Chen et al., 2022; Luckin, 2017; Ouyang & Jiao, 2021). Therefore, it is essential for the education fraternity to have a basic understanding of artificial intelligence techniques to explore their possible use cases (Lee & Perret, 2022; Chaudhry & Kazim, 2022).

The remaining sections of this article introduce the basic concepts of Machine Learning, Neural Networks, and Deep Learning from the perspective of a common reader. It then discusses the possible use cases for AI in different aspects of education. The article also highlights the concern around the unchecked growth in AI and its potential implications.

Machine Learning

Machine learning (ML) is the science of teaching computers how to learn from experience. It is the branch of AI that aims to train computer algorithms with lots of examples (data or situations) so that they learn during the training and, once trained, can generate results for unknown data or situations. Also, ML models can continuously learn while being used and improve their performance. Such a model takes a large amount of data as input, processes them, and as output, either predicts an upcoming result or classifies the given


data (Valiant, 1984). Today's machine learning techniques exhibit only a limited range of intelligent behaviour, primarily in prediction, classification, clustering, dimensionality reduction, etc.

The statistical techniques used in ML themselves are not new. Instead, they were formulated in the early 20th

century. The Improvement in computing power and other advances in computer science have made it possible to utilise those statistical methods in ML algorithms. The learning techniques used in ML models are prominent in three types (Figures 2 - 7), as detailed below.

Supervised learning

Figure-2: Supervised learning-training process (Raj, 2023)




In the case of supervised machine learning, we have a dataset in which both the input and output are known (called labelled), but the relation (mapping) between input and output is not exclusively defined.

As an example, we have data about student attendance, family details, performance in periodic tests, etc., as input and also the performance in the annual examination as output.


However, we don't know the relation between these input parameters and the output. The machine is trained with a large set of labelled data to find the output for an unknown input value after training. In our case, if we feed the machine with another student's input values, it should be able to predict the student's performance in the upcoming final examination. This kind of learning technique is used for classification and regression problems.

Figure-3: Example of supervised learning (Raj, 2023)



Unsupervised learning

Figure-4: Unsupervised learning-training process (Raj, 2023)

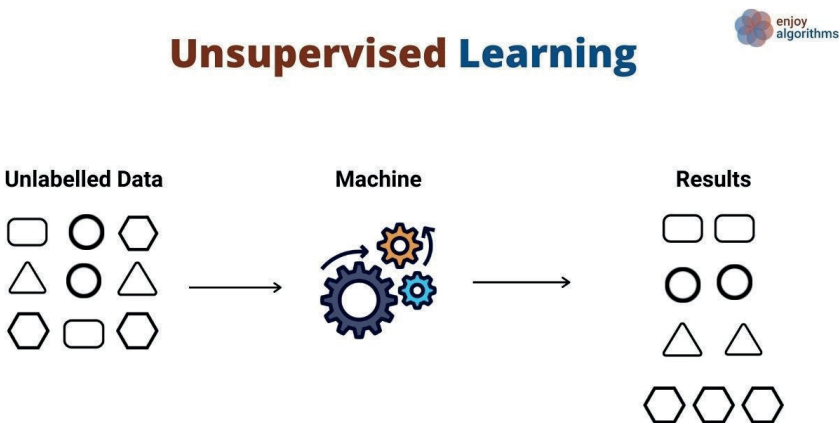


In the case of unsupervised machine learning, we have a dataset in which only the input values are known, but the output and the relation between input-output are not known beforehand (called unlabelled). The machine is trained with a large set of unlabelled data, and during this process, it learns to categorise similar datasets into groups (also called clusters). Once trained, the output is the decision about which cluster the given input best fits into. Human learning is largely unsupervised as we are not exclusively taught everything in life. Instead, we observe the world and discover the structure and phenomena that occur (LeCun &

Hinton, 2015).


For example, we have data about various learners and their learning behaviours using an e-learning platform like DIKSHA (NCERT, 2023). We can input these data into an unsupervised machine learning model, and it can segregate the learners into different categories based on specific parameters in the given data. Such clusters can help us learn many interesting facts about how learners use an e-learning platform, and accordingly, we can either improve the features of the platform or its contents or both. This kind of learning technique is used for clustering and dimensionality reduction problems.

Figure-5: Example of unsupervised learning (Raj, 2023)



Semi-supervised learning


Figure-6: semi-supervised learning-training process (Raj, 2023)



In the real world, most data collected are unlabelled, meaning the output is unknown a priori. It is costly to manually level such data as it is time-consuming and needs domain expertise. As the name suggests, only a tiny portion of

the input data are labelled while the others are not. The model is trained with this small labelled data and then applied to unlabelled data. This process is repeated, and the model starts improving its output.

Figure-7: Example of semi-supervised learning (Raj, 2023)



Supervised learning provides accurate output, but having a vast dataset with labelled output is costly. Unsupervised learning doesn't require any labelled output, but the accuracy is less. Semi-supervised learning takes advantage of both these techniques. Hence, using semi-supervised learning, only a small sample of data is human labelled, and the rest of the unlabelled data are given to the model for labelling.

For example, a student's career counselling should ideally be done based on the performance across the years and the student's interest, aptitude, attitude, etc. To conduct career counselling at the end of schooling, it may be easy to collect such data about learners from the school records and through a questionnaire to be filled out by the student. But it requires the expertise and time of a counsellor to analyse all such data and suggest appropriate choices for the learner.

For example, a student's career


Machine learning can be used in many areas of education, such as predicting student performance (Anozie & Junker, 2006); assessing and grading students (Luckin et al., 2016), improving retention (Đambić et al., 2016); assisting teachers in assessment and other tasks (Celar et al., 2015).

Artificial Neural Network (ANN)

Our brain has many interconnected neurons, which can be considered a neural network. In such neural

networks, when the neurons get excited (or activated), they send signals called neurotransmitters to other neurons connected to them. These neurotransmitters change the electric potentials of the receiver neurons. When the electrical potentials cross a limit called threshold, those neurons get excited, and they in turn, send signals further to the neurons connected to them. In a nutshell, a neuron takes input, processes it, and sends the output to the other connected neurons.


Figure-8: A biological neuron (image source: Wikipedia)



This functioning of the biological neural network has been inspiring computer science researchers since the 1940s (Yadav & Kumar, 2015; Pitts, 1943). Attempts to write algorithms that mimic the interconnected structure of neurons in computer-based data processing have led to the creation of

a major subfield called Artificial Neural Networks (ANN). The basic element of ANN is also called a neuron or unit. An artificial neural network is created by arranging the neurons and their interconnections into layers called a) the input layer, b) the output layer, and c) one or more hidden layers.

Figure-9: A single layer ANN (image source: Wikipedia)



Deep Learning

All neural networks must be trained with large datasets as they learn during training and ultimately become capable of producing near-accurate results. Deep learning adds multiple hidden layers in the artificial neural network (LeCun & Hinton, 2015). These additional layers in deep learning models make it possible for them to generate results with more accuracy than single-layer neural networks. But deep learning models require to be trained with large datasets for hundreds of hours. Earlier, getting access to very large datasets and their storage took a lot of effort. The processing of large datasets with traditional computer processors wasn't easy either. However, with so many people using social media, doing online shopping, and creating and sharing content through the Internet means the easy creation of enormous datasets quickly. This is one of the reasons why deep learning is gaining popularity and attention.

The training phase of any deep learning model is very resource intensive and time-consuming in the case of massive datasets. The multiple layers in deep learning models mean they require more computing power. The availability of cloud computing resources (compute and storage) and the emergence of a new set of processors called Graphics Processing Units (GPU) has made it possible to develop and train deep learning models. Compared to traditional CPUs, GPUs are cheaper, faster, and capable of real parallel processing.

Deep learning-based models are extensively used in speech recognition (Siri, Alexa, Google Voice, etc.), computer vision (face detection, autonomous cars, robots), Natural language processing, fraud detection, image processing, etc.

Natural Language Processing

Humans communicate with each other in many ways - speaking, listening, writing (text, image, graphics), making gestures, and using sign language. Communication between people reading or writing in different languages (natural languages) is difficult as one may not know the other's language. The vast knowledge base available in the electronic or print format is often created in the languages that are most spoken. These are often not accessible to people who do not know those languages. It is a tedious and expensive task to translate between languages manually. Digital computers are binary machines that do not readily process natural languages. There have been efforts to use computers to process natural languages. However, understanding natural language is not easy for computers. This is because human languages are ambiguous and have imprecise characteristics, such as:

- The meaning of a word can vary depending on the context of use.
- Use of different alphabets at different positions of a word to signify plurality.
- Pronunciation of two or more words can be the same, yet their spelling can be different.
- Slang words and other culture-specific words can have different meanings.

The branch of AI that deals with communication between humans and computers in natural languages is called Natural Language Processing (NLP) (Eisenstein, 2018). NLP uses AI techniques such as machine learning and deep learning (called Neural Machine Translation) to read, understand, and make sense of human languages. Popular applications of NLP include

- machine translation between languages
- real-time speech conversion between natural languages,
- conversion of speech to text and vice versa,
- text summarisation,
- grammar and spelling checking,
- email spam filtering,
- Search by search engines based on keywords given,
- chatbots understand human query responding in natural languages,
- personal digital assistants like Siri, Cortana, Alexa, OK Google,
- conversational agents such as ChatGPT.

The data (text or voice) must be preprocessed to apply an AI-based algorithm for NLP. Following are some of the basic pre-processing activities in NLP (Kibble, 2013):

- A. Tokenization - breaking down a sentence and its components into individual words, numbers, punctuation marks, other symbols, and characters. Each of these components is a token.
- B. Part of Speech (POS) tagging - categorises every token as a part of speech or into a grammatical area, such as nouns, pronouns, adjectives, adverbs, etc.
- C. Stemming and Lemmatization - identifying the root word for a set of words. For example, the root word for run, ran, and running is run.
- D. Stop word removal - words such as articles, prepositions, etc. which do not add much value/information to the text.

Language Model

It is a type of machine learning model that predicts the next word(s) in a sequence that is most appropriate to fill in a blank space in a sentence or phrase. It relies upon statistical methods (probability distribution of different available words) and the context to select the next term.

Language models constitute a fundamental part of Natural Language Processing (NLP) as they help machines understand, analyse, and generate human languages. The auto-complete suggestions we see while typing in smartphone work based on language models.

Generative Artificial Intelligence

The application of AI techniques to generate different kinds of texts, audio, video, animation, 2D/3D images, and even programming code snippets is called Generative AI.

A) ChatGPT

A chatbot developed by OpenAI, ChatGPT (2023), has taken the world by storm. Within two months of its launch, it has got 100 million users, which is the fastest technology adoption in history. It can converse with users and generate articles, essays, poems, and even programming codes. The output generated is so impressive that some people find it difficult to distinguish between machine-produced output and those by humans.

Generative Pre-trained Transformer (GPT) that powers the ChatGPT is a language model trained with contents (datasets) collected from billions of websites and web pages.

B) Other Examples of Generative AI

- A. DALL-E 2 (2023) is another tool from Open AI, the same organisation that

has come out with ChatGPT. DALL-E is a tool to generate art using AI. They claim that based on input given in natural languages, the tool can generate images that appear realistic.

- B. Alphacode (2023) from Deepmind can generate professional-grade code in different programming languages.

AI in Education (AIED)

The purpose of bringing AI into education is not to supplant the teacher. Researchers working in Artificial Intelligence in Education (AIED) must focus on areas where AI can alleviate teachers of routine tasks or achieve complex functions which are otherwise difficult. It is a well-known fact that each learner has a unique learning style, preferences, likes/dislikes, varying prior knowledge, and comes from different socio-economic backgrounds. AI can be used to identify these individual traits to build and update a student model for each learner. Such a student model can be used to carry out personalised tutoring by recommending appropriate content, carrying out an adaptive assessment, predicting potential dropouts, and accordingly intervening, etc. (Kučak et al., 2018). The recent advances in NLP in terms of powerful language models are also relevant for education, mainly in the areas of machine translation, speech recognition (for speech-to-text and speech-to-speech conversion), retrieval of information from the web or a content repository, question-answering, text summarisation, etc. This subsection briefly explores some of the areas of education where AI can be of help (Holmes et al., 2002):

- A. Support collaborative learning - AI can refer to individual student models and search for other students having similar learning

preferences and traits and are best suited for a particular collaborative task. Such learners can be suggested for collaborative work.

- B. Monitor and manage student discussion forums - Online courses having more than a few hundred participants can lead to many discussions and questions in forums. AI-enabled chatbots can answer some questions while leaving the critical ones for the teacher to respond to. AI can be used to summarise or club similar questions, thus reducing the number of questions for the course tutor. AI-powered sentiment analysis tools can analyse the discussion posts by students to flag likely dropout cases or inappropriate posts such as racist or sexist remarks or any other emotional traits of a learner that require a teacher's intervention.
- C. Continuous and holistic assessment - the importance of continuous assessment rather than a year/semester end test is well known. A learner should not be assessed merely based on the performance in a high stake examination. Such assessment causes undue stress and anxiety and encourages learners to learn what is to be assessed rather than what is required. Easy accessibility of devices and the Internet has provided learners multiple learning opportunities beyond the classroom. Skills, competencies, and certificates earned by a learner from different sources should be acknowledged and duly recorded in a student portfolio. It can be further authenticated using blockchain technologies. This will help create a robust, verified, and in-depth record of a learner's learning experiences, which will be far more realistic than a collection of certificates/

degrees. AI can help in continuous and holistic assessment by asking suitable questions based on the learner's recent learning history. It can assess the learning gap and guide the learner to the appropriate learning resource, a tutorial or a practice test.

- D. Learning companion - an Intelligent Tutoring System (ITS) can be a great companion for both a learner and a teacher as it can suggest the relevant learning path; recommend appropriate resources; assist in solving a problem or clearing a doubt; provide timely guidance and feedback; continuously record the learner's interest and progress in the student model (Barbhuiya et al., 2011). ChatGPT and voice assistants such as Siri, Alexa, Google Home, etc., can be used for some of these tasks.
- E. Teaching assistant - AI is not going to replace human teachers (Timms, 2016). However, teachers must continuously adapt and evolve to leverage AI in automating some of their work which otherwise takes lots of time. These can include checking the homework, managing attendance, compiling results, etc.
- F. Understanding how people learn - Educationists are still exploring to understand the science of learning. The cognitive science and neuroscience domains are still figuring out how we learn about something or someone, how we recall (retrieve) someone we met years ago, or some incident that happened long back, etc. With the shift to online and digital learning, we are generating lots of data about learners' behaviour and performance during learning. Data mining and other AI-enabled data analytics techniques can be used to analyse such data and advance the

science of learning.

Concerns around the progress of AI

As discussed earlier, the emergence of cloud computing services, advancements in computer processing capabilities, affordable storage options, generation of huge online datasets, etc., have led to exponential growth in the field of AI, particularly in deep learning techniques. The recent success of deep learning has made AI a hot topic attracting much public attention. We are seeing an unprecedented level of automation in various areas. It is becoming difficult to predict the impact of AI in the near future in the areas of work, healthcare, digital surveillance, cyber warfare, etc. Educators struggle to determine the essential future skills that today's school-going learners should be equipped with.

AI has already started showing impressive results in certain domains. Researchers are working hard to leverage AI in Healthcare (early detection of disease, targeted drug delivery), weather forecasting, defense, legal matters, etc. However, there are concerns about the implications of unregulated research and innovation in AI. Many big names in the world of science and technology, including Stephen Hawking, have warned about possible existential threats to humanity due to advancements in AI. Technology leaders like Elon Musk of Tesla, Apple co-founder Steve Wozniak, Geoffrey Hinton from Google, and many others have already urged to make a pause in the development of powerful artificial intelligence (AI) systems until their alleged safety concerns are addressed by deliberating and devising safety protocols.

AI models are data-hungry, and as a result, there is a tremendous interest in collecting data, including our data. Consider big technology corporations

collecting data about our physical activities, medical conditions, and personal and social life, including what we eat, whom we meet, and what we talk about. Such personal data can be exploited and misused if they fall into the wrong hands. In the context of education, people envisage having AI for teaching-learning-assessment, but constant monitoring of student behaviours and achievements can result in severe and far-reaching ethical

questions. Governments worldwide are considering regulations and policies to deal with the dangers and pitfalls of AI systems related to ethics, equity, privacy, and humanity at large. The European Union (Hickman & Petrin, 2021) is taking the lead so far in initiating serious attempts to regulate AI and big technology companies. The aim is to break any potential monopoly and ensure the ethical and safe use of AI and related technologies.

References

- AlphaCode. (2023, June 03). *DeepMind* by Google. <https://alphacode.deepmind.com/>.
- Barbhuiya, R. K., Mustafa, K., & Jabin, S. (2011). Design Perspectives of Intelligent Tutoring System. *In IICAI* (pp. 1484-1495).
- Campbell, M., Hoane Jr, A. J., & Hsu, F. H. (2002). Deep blue. *Artificial intelligence*, 134(1-2), 57-83.
- Celar, S., Stojkic, Z., Seremet, Z., Marusic, Z., & Zelenika, D. (2015). Classification of Test Documents Based on Handwritten Student ID's Characteristics. *Annals of DAAAM and Proceedings of DAAAM Symposium*. 2014. ISBN 978-3-901509-99-5., 782-790.
- ChatGPT. (2023, June 03). *ChatGPT* by Open AI. <https://openai.com/chatgpt>.
- Chaudhry, M. A., & Kazim, E. (2022). Artificial Intelligence in Education (AIEd): A high-level academic and industry note 2021. *AI and Ethics*, 1-9.
- Chen, X., Zou, D., Xie, H., Cheng, G., & Liu, C. (2022). Two decades of artificial intelligence in education. *Educational Technology & Society*, 25(1), 28-47.
- DALL-E 2. (2023, June 03). *DALL-E 2* by Open AI. <https://openai.com/dall-e-2>.
- Đambić, G., Krajcar, M. & Bele, D. (2016). Machine learning model for early detection of higher education students that need additional attention in introductory programming courses. *International Journal of Digital Technology & Economy*, 1 (1), 1-11
- Eisenstein, J. (2018). Natural language processing. *Jacob Eisenstein*.
- Floridi, L. (2014). *The fourth revolution: How the infosphere is reshaping human reality*. OUP Oxford.
- Hickman, E., & Petrin, M. (2021). Trustworthy AI and corporate governance: the EU's ethics guidelines for trustworthy artificial intelligence from a company law perspective. *European Business Organization Law Review*, 22, 593-625.
- Holmes, W., Bialik, M., & Fadel, C. (2023). *Artificial intelligence in education*. Globethics Publications.
- Huang, B., Huan, Y., Xu, L. D., Zheng, L., & Zou, Z. (2019). Automated trading systems statistical and machine learning methods and hardware implementation: a survey. *Enterprise Information Systems*, 13(1), 132-144.
- Kibble, R. (2013). *Introduction to natural language processing*. London: University of London.

Kučak, D., Juričić, V., & Đambić, G. (2018). MACHINE LEARNING IN EDUCATION-A SURVEY OF CURRENT RESEARCH TRENDS. *Annals of DAAAM & Proceedings*, 29.

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *nature*, 521(7553), 436-444.

Lee, I., & Perret, B. (2022). Preparing High School Teachers to Integrate AI Methods into STEM Classrooms. In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 36, No. 11, pp. 12783-12791).

Luckin, R. (2017). Towards artificial intelligence-based assessment systems. *Nature Human Behaviour*, 1(3), 0028.

Luckin, R., Holmes, W., Griffiths, M., & Forcier, L. B. (2016). *Intelligence unleashed: An argument for AI in education*.

McCarthy, J., Minsky, M. L., Rochester, N., & Shannon, C. E. (2006). A proposal for the dartmouth summer research project on artificial intelligence, August 31, 1955. *AI magazine*, 27(4), 12-12.

NCERT (2023, June 03). *DIKSHA*. <https://diksha.gov.in/>. Accessed on 03 June 2023.

Ouyang, F., & Jiao, P. (2021). Artificial intelligence in education: The three paradigms. *Computers and Education: Artificial Intelligence*, 2, 100020.

Pitts, W. (1943). The linear theory of neuron networks: The dynamic problem. *The bulletin of mathematical biophysics*, 5, 23-31.

Raj, R. (2023, June 03). *Supervised, Unsupervised and Semi-supervised Learning with Real-life Use case*. <https://www.enjoyalgorithms.com/blogs/supervised-unsupervised-and-semisupervised-learning>.

Timms, M. J. (2016). Letting artificial intelligence in education out of the box: educational cobots and smart classrooms. *International Journal of Artificial Intelligence in Education*, 26, 701-712.

Valiant, L. G. (1984). A theory of the learnable. *Communications of the ACM*, 27(11), 1134-1142.

Wang, P. (2019). On defining artificial intelligence. *Journal of Artificial General Intelligence*, 10(2), 1-37.

Yadav, N., Yadav, A., & Kumar, M. (2015). *An introduction to neural network methods for differential equations* (Vol. 1, p. 114). Berlin: Springer.