AI, Machine Learning and Deep Learning Development and Applications

Yew Kee Wong

School of Information Engineering, HuangHuai University, Henan, China.

Abstract

In the information era, enormous amounts of data have become available on hand to decision makers. Big data refers to datasets that are not only big, but also high in variety and velocity, which makes them difficult to handle using traditional tools and techniques. Due to the rapid growth of such data, solutions need to be studied and provided in order to handle and extract value and knowledge from these datasets. Machine learning is a method of data analysis that automates analytical model building. It is a branch of artificial intelligence based on the idea that systems can learn from data, identify patterns and make decisions with minimal human intervention. Such minimal human intervention can be provided using machine learning, which is the application of advanced deep learning techniques on big data. This paper aims to analyse some of the different machine learning and deep learning algorithms and methods, as well as the opportunities provided by the AI applications in various decision making domains.

1. Introduction

Resurging interest in machine learning is due to the same factors that have made data mining and Bayesian analysis more popular than ever. Things like growing volumes and varieties of available data, computational processing that is cheaper and more powerful, and affordable data storage. All of these things mean it's possible to quickly and automatically produce models that can analyse bigger, more complex data and deliver faster, more accurate results – even on a very large scale. And by building precise models, an organization has a better chance of identifying profitable opportunities – or avoiding unknown risks[1].

Because of new computing technologies, machine learning today is not like machine learning of the past. It was born from pattern recognition and the theory that computers can learn without being programmed to perform specific tasks; researchers interested in artificial intelligence wanted to see if computers could learn from data. The iterative aspect of machine learning is important because as models are exposed to new data, they are able to independently adapt. They learn from previous computations to produce reliable, repeatable decisions and results. It's a science that's not new – but one that has gained fresh momentum. While many machine learning and deep learning algorithms have been around for a long time, the ability to automatically apply complex mathematical calculations to big data - over and over, faster and faster – is a recent development [2]. This paper will look at some of the different machine learning and deep learning algorithms and methods which can be applied to big data analysis, as well as the opportunities provided by the AI applications in various decision making domains.

2. BENEFITS OF ARTIFICIAL INTELLIGENCE

Artificial intelligence (AI) is a branch of computer science. AI technologies aim to reproduce or surpass abilities in computational systems that are generally deemed intelligent if performed by a human [1]. These abilities include:

- learning
- reasoning
- pattern-recognition
- problem-solving
- visual perception
- language-understanding

2.1. Different Types of AI

There are two main types of AI [2]:

- **Applied AI:** is more common and includes systems designed to intelligently carry out a single task, eg move a driverless vehicle, or trade stocks and shares. This category is also known as 'weak' or 'narrow' AI.
- **Generalised AI:** is less common and includes systems or devices that can theoretically handle any task, as they carry enough intelligence to find solutions to unfamiliar problems. Generalised AI is also known as 'strong' AI. Examples of true strong AI don't currently exist, as these technologies are still in very early stages of development.

Many modern AI applications are enabled through a sub-field of AI known as 'machine learning'. What is machine learning? The roots of machine learning (ML) are in statistics. ML uses algorithms and statistical models to perform a specific task without using explicit instructions, instead relying on patterns and inference [3]. For example, ML applications can:

- read a text and decide if the author is making a complaint or a purchase order
- listen to a piece of music and find other tunes to match the mood
- recognise images and classify them according to the elements they contain
- translate large volumes of text in real time
- accurately recognise faces, speech and objects

2.2. How are AI and Machine Learning used in Business

Over the years, AI research has enabled many technological advances [4], including:

- virtual agents and chatbots
- suggestive web searches
- targeted advertising
- pattern recognition
- predictive analytics
- voice and speech recognition
- face recognition
- machine translation
- autonomous driving
- automatic scheduling

Many of these are now commonplace and provide solutions to a great number of business challenges and complex, real-world problems.

2.3. How are Businesses using AI

AI is steadily passing into everyday business use. From workflow management to trend predictions, AI has many different uses in business [5]. It also provides new business opportunities.

Application of AI in business:

- **Improve customer services** – e.g. use virtual assistant programs to provide real-time support to users (for example, with billing and other tasks).
- **Automate workloads** – e.g. collect and analyse data from smart sensors, or use machine learning (ML) algorithms to categorise work, automatically route service requests, etc.
- **Optimise logistics** – e.g. use AI-powered image recognition tools to monitor and optimise your infrastructure, plan transport routes, etc.
- **Increase manufacturing output and efficiency** – e.g. automate production line by integrating industrial robots into your workflow and teaching them to perform labour-intensive or mundane tasks [6].
- **Prevent outages** – e.g. use anomaly detection techniques to identify patterns that are likely to disrupt your business, such as an IT outage. Specific AI software may also help you to detect and deter security intrusions.
- **Predict performance** – e.g. use AI applications to determine when you might reach performance goals, such as response time to help desk calls.
- **Predict behaviour** – e.g. use ML algorithms to analyse patterns of online behaviour to, for example, serve tailored product offers, detect credit card fraud or target appropriate adverts.
- **Manage and analyse your data** – e.g. AI can help you interpret and mine your data more efficiently than ever before and provide meaningful insight into your assets, your brand, staff or customers [7].
- **Improve your marketing and advertising** – e.g. effectively track user behaviour and automate many routine marketing tasks.

3. HOW MACHINE LEARNING WORKS

Machine learning is a branch of artificial intelligence (AI) and computer science which focuses on the use of data and algorithms to imitate the way that humans learn, gradually improving its accuracy [3].

Machine learning is an important component of the growing field of data science. Through the use of statistical methods, algorithms are trained to make classifications or predictions, uncovering key insights within data mining projects. These insights subsequently drive decision making within applications and businesses, ideally impacting key growth metrics [4]. As big data continues to expand and grow, the market demand for data scientists will increase, requiring them to assist in the identification of the most relevant business questions and subsequently the data to answer them.

3.1. Machine Learning Algorithms

Machine learning algorithms can be categorize into three main parts:

- **A Decision Process**: In general, machine learning algorithms are used to make a prediction or classification. Based on some input data, which can be labelled or unlabelled, your algorithm will produce an estimate about a pattern in the data.

- **An Error Function**: An error function serves to evaluate the prediction of the model. If there are known examples, an error function can make a comparison to assess the accuracy of the model.
- **A Model Optimization Process**: If the model can fit better to the data points in the training set, then weights are adjusted to reduce the discrepancy between the known example and the model estimate. The algorithm will repeat this evaluate and optimize process, updating weights autonomously until a threshold of accuracy has been met.

3.2. Types of Machine Learning Methods

Machine learning classifiers fall into three primary categories [5]:

Supervised machine learning
Supervised learning also known as supervised machine learning, is defined by its use of labelled datasets to train algorithms that to classify data or predict outcomes accurately. As input data is fed into the model, it adjusts its weights until the model has been fitted appropriately. This occurs as part of the cross validation process to ensure that the model avoids over fitting or under fitting. Supervised learning helps organizations solve for a variety of real-world problems at scale, such as classifying spam in a separate folder from your inbox. Some methods used in supervised learning include neural networks, naïve bayes, linear regression, logistic regression, random forest, support vector machine (SVM), and more.

Unsupervised machine learning
Unsupervised learning, also known as unsupervised machine learning, uses machine learning algorithms to analyse and cluster unlabelled datasets. These algorithms discover hidden patterns or data groupings without the need for human intervention. Its ability to discover similarities and differences in information make it the ideal solution for exploratory data analysis, cross-selling strategies, customer segmentation, image and pattern recognition. It's also used to reduce the number of features in a model through the process of dimensionality reduction; principal component analysis (PCA) and singular value decomposition (SVD) are two common approaches for this. Other algorithms used in unsupervised learning include neural networks, k-means clustering, probabilistic clustering methods, and more [6].

Semi-supervised learning
Semi-supervised learning offers a happy medium between supervised and unsupervised learning. During training, it uses a smaller labelled data set to guide classification and feature extraction from a larger, unlabelled data set. Semi-supervised learning can solve the problem of having not enough labelled data (or not being able to afford to label enough data) to train a supervised learning algorithm.

3.3. Practical Use of Machine Learning

Here are just a few examples of machine learning you might encounter every day [7]:

Speech Recognition: It is also known as automatic speech recognition (ASR), computer speech recognition, or speech-to-text, and it is a capability which uses natural language processing (NLP) to process human speech into a written format. Many mobile devices incorporate speech recognition into their systems to conduct voice search—e.g. Siri—or provide more accessibility around texting.

Customer Service: Online chatbots are replacing human agents along the customer journey. They answer frequently asked questions (FAQs) around topics, like shipping, or provide personalized advice, cross-selling products or suggesting sizes for users, changing the way we think about customer engagement across websites and social media platforms. Examples include messaging bots on e-commerce sites with virtual agents, messaging apps, such as Slack and Facebook Messenger, and tasks usually done by virtual assistants and voice assistants.

Computer Vision: This AI technology enables computers and systems to derive meaningful information from digital images, videos and other visual inputs, and based on those inputs, it can take action. This ability to provide recommendations distinguishes it from image recognition tasks. Powered by convolutional neural networks, computer vision has applications within photo tagging in social media, radiology imaging in healthcare, and self-driving cars within the automotive industry.

Recommendation Engines: Using past consumption behaviour data, AI algorithms can help to discover data trends that can be used to develop more effective cross-selling strategies. This is used to make relevant add-on recommendations to customers during the checkout process for online retailers.

Automated stock trading: Designed to optimize stock portfolios, AI-driven high-frequency trading platforms make thousands or even millions of trades per day without human intervention.

4. HOW DEEP LEARNING WORKS

Deep learning changes how you think about representing the problems that you're solving with analytics. It moves from telling the computer how to solve a problem to training the computer to solve the problem itself.

A traditional approach to analytics is to use the data at hand to engineer features to derive new variables, then select an analytic model and finally estimate the parameters (or the unknowns) of that model. These techniques can yield predictive systems that do not generalize well because completeness and correctness depend on the quality of the model and its features [3]. For example, if you develop a fraud model with feature engineering, you start with a set of variables, and you most likely derive a model from those variables using data transformations. You may end up with 30,000 variables that your model depends on, then you have to shape the model, figure out which variables are meaningful, which ones are not, and so on. Adding more data requires you to do it all over again.

The new approach with deep learning is to replace the formulation and specification of the model with hierarchical characterizations (or layers) that learn to recognize latent features of the data from the regularities in the layers. The paradigm shift with deep learning is a move from feature engineering to feature representation. The promise of deep learning is that it can lead to predictive systems that generalize well, adapt well, continuously improve as new data arrives, and are more dynamic than predictive systems built on hard business rules. You no longer fit a model. Instead, you train the task.

Deep learning is making a big impact across industries. In life sciences, deep learning can be used for advanced image analysis, research, drug discovery, prediction of health problems and disease symptoms, and the acceleration of insights from genomic sequencing. In transportation, it can help autonomous vehicles adapt to changing conditions [4]. It is also used to protect critical infrastructure and speed response.

Most deep learning methods use neural networks architectures, which is why deep learning models are often referred to as deep neural networks. The term "deep" usually refers to the number of hidden layers in the neural network. Traditional neural networks only contain 2-3 hidden layers, while deep networks can have as many as 150. Deep learning models are trained by using large sets of labelled data and neural network architectures that learn features directly from the data without the need for manual feature extraction.
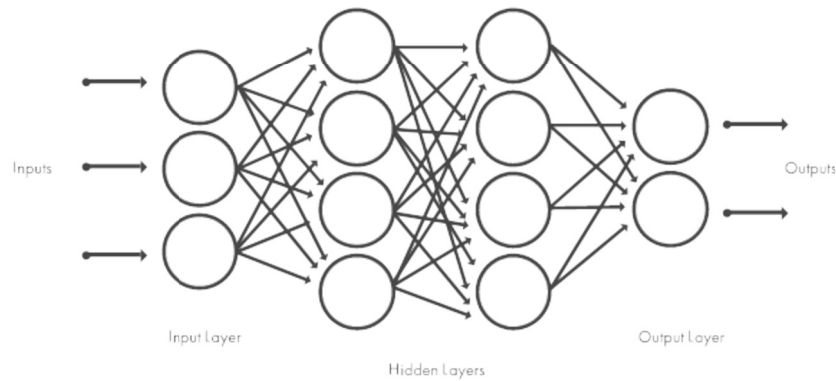

Figure 1. Neural networks, which are organized in layers consisting of a set of interconnected nodes. Networks can have tens or hundreds of hidden layers.

4.2. How Deep Learning Being Used

To the outside eye, deep learning may appear to be in a research phase as computer science researchers and data scientists continue to test its capabilities. However, deep learning has many practical applications that businesses are using today, and many more that will be used as research continues [19]. Popular uses today include:

Speech Recognition
Both the business and academic worlds have embraced deep learning for speech recognition. Xbox, Skype, Google Now and Apple's Siri, to name a few, are already employing deep learning technologies in their systems to recognize human speech and voice patterns.

Natural Language Processing
Neural networks, a central component of deep learning, have been used to process and analyse written text for many years. A specialization of text mining, this technique can be used to discover patterns in customer complaints, physician notes or news reports, to name a few.

Image Recognition
One practical application of image recognition is automatic image captioning and scene description. This could be crucial in law enforcement investigations for identifying criminal activity in thousands of photos submitted by bystanders in a crowded area where a crime has occurred. Self-driving cars will also benefit from image recognition through the use of 360-degree camera technology.

Recommendation Systems
Amazon and Netflix have popularized the notion of a recommendation system with a good chance of knowing what you might be interested in next, based on past behaviour. Deep learning can be used to enhance recommendations in complex environments such as music interests or clothing preferences across multiple platforms.

Recent advances in deep learning have improved to the point where deep learning outperforms humans in some tasks like classifying objects in images [20]. While deep learning was first theorized in the 1980s, there are two main reasons it has only recently become useful:

1. Deep learning requires large amounts of labelled data. For example, driverless car development requires millions of images and thousands of hours of video.
2. Deep learning requires substantial computing power. High-performance GPUs have a parallel architecture that is efficient for deep learning. When combined with clusters or cloud computing, this enables development teams to reduce training time for a deep learning network from weeks to hours or less.

When choosing between machine learning and deep learning, consider whether you have a high-performance GPU and lots of labelled data. If you don't have either of those things, it may make more sense to use machine learning instead of deep learning. Deep learning is generally more complex, so you'll need at least a few thousand images to get reliable results. Having a high-performance GPU means the model will take less time to analyse all those images [21].

4.3. Deep Learning Opportunities and Applications

A lot of computational power is needed to solve deep learning problems because of the iterative nature of deep learning algorithms, their complexity as the number of layers increase, and the large volumes of data needed to train the networks.

The dynamic nature of deep learning methods – their ability to continuously improve and adapt to changes in the underlying information pattern – presents a great opportunity to introduce more dynamic behaviour into analytics [22]. Greater personalization of customer analytics is one possibility. Another great opportunity is to improve accuracy and performance in applications where neural networks have been used for a long time. Through better algorithms and more computing power, we can add greater depth.

While the current market focus of deep learning techniques is in applications of cognitive computing, there is also great potential in more traditional analytics applications, for example, time series analysis. Another opportunity is to simply be more efficient and streamlined in existing analytical operations. Recently, some study showed that with deep neural networks in speech-to-text transcription problems [23]. Compared to the standard techniques, the word-error-rate decreased by more than 10 percent when deep neural networks were applied. They also eliminated about 10 steps of data preprocessing, feature engineering and modelling. The impressive performance gains and the time savings when compared to feature engineering signify a paradigm shift.

Here are some examples of deep learning applications are used in different industries:

Automated Driving: Automotive researchers are using deep learning to automatically detect objects such as stop signs and traffic lights. In addition, deep learning is used to detect pedestrians, which helps decrease accidents.

Aerospace and Defence: Deep learning is used to identify objects from satellites that locate areas of interest, and identify safe or unsafe zones for troops.

Medical Research: Cancer researchers are using deep learning to automatically detect cancer cells. Teams at UCLA built an advanced microscope that yields a high-dimensional data set used to train a deep learning application to accurately identify cancer cells [24].

Industrial Automation: Deep learning is helping to improve worker safety around heavy machinery by automatically detecting when people or objects are within an unsafe distance of machines.

Electronics: Deep learning is being used in automated hearing and speech translation. For example, home assistance devices that respond to your voice and know your preferences are powered by deep learning applications.

4.4. How to Create and Train Deep Learning Models

The three most common ways people use deep learning to perform object classification are:

Training from Scratch
To train a deep network from scratch, you gather a very large labelled data set and design a network architecture that will learn the features and model. This is good for new applications, or applications that will have a large number of output categories. This is a less common approach because with the large amount of data and rate of learning, these networks typically take days or weeks to train [25].

Transfer Learning
Most deep learning applications use the transfer learning approach, a process that involves fine-tuning a pre-trained model. User can start with an existing network, such as AlexNet or GoogLeNet, and feed in new data containing previously unknown classes [26]. After making some tweaks to the network, user can now perform a new task, such as categorizing only dogs or cats instead of 10,000 different objects. This also has the advantage of needing much less data (processing thousands of images, rather than millions), so computation time drops to minutes or hours.

Feature Extraction
A slightly less common, more specialized approach to deep learning is to use the network as a feature extractor. Since all the layers are tasked with learning certain features from images, user can pull these features out of the network at any time during the training process [27]. These features can then be used as input to a machine learning model such as support vector machines (SVM).

5. CONCLUSIONS

So this study was concerned by understanding the inter-relationships between machine learning and deep learning, what frameworks and systems that worked, and how machine learning can impact the AI applications whether by introducing new innovations that foster advanced machine learning process and escalating power consumption, security issues and replacing human in workplaces. The advanced machine learning and deep learning algorithms with various applications show promising results in artificial intelligence development and further evaluation and research using machine learning are in progress.

REFERENCES

[1]   Jonathan Michael Spector, Du Jing, (2017). Artificial Intelligence and the Future of Education: Big Promises – Bigger Challenges, ACADEMICS, No. 7.

[2]   Oscar Sanjuan, B. Cristina Pelayo Garcia-Bustelo, Ruben Gonzalez Crespo, Enrique Daniel France, (2009). Using Recommendation System for E-Learning Environment at degree level, International Journal of Interactive Multimedia and Artificial Intelligence, Vol. 1, No. 2.

[3]   S. M. Patil, T. D. Shaikh, (2014). Implementing Adaptability in E-Learning Management System Using Moodle for Campus Environment, International Journal of Emerging Technology and Advanced Engineering, Vol. 4, No. 8.

[4]   Shraddha Kande, Pooja Goswami, Gurpreet Naul, Mrs. Nirmala Shinde, (2016). Adaptive and Advanced E-learning Using Artificial Intelligence, Journal of Engineering Trends and Applications, Vol. 3, No. 2.

[5]   Ofra Walter, Vered Shenaar-Golan and Zeevik Greenberg, (2015). Effect of Short-Term Intervention Program on Academic Self-Efficacy in Higher Education, Psychology, Vol. 6, No. 10.

[6]   Calum Chace, (2019). Artificial Intelligence and the Two Singularities, Chapman & Hall/CRC.

[7]   Piero Mella, (2017). Intelligence and Stupidity – The Educational Power of Cipolla's Test and of the "Social Wheel", Creative Education, Vol. 8, No. 15.

[8]   Zhongzhi Shi, (2019). Cognitive Machine Learning, International Journal of Intelligence Science, Vol. 9, No. 4.

[9]   Crescenzio Gallo and Vito Capozzi, (2019). Feature Selection with Non Linear PCA: A Neural Network Approach, Journal of Applied Mathematics and Physics, Vol. 7, No. 10.

[10]   Shi, Z., (2019). Cognitive Machine Learning. *International Journal of Intelligence Science*, 9, pp. 111-121.

[11]   Lake, B.M., Salakhutdinov, R. and Tenenbaum, J.B., (2015). Human-Level Concept Learning through Probabilistic Program Induction. Science, 350, pp. 1332-1338.

[12]   Silver, D., Huang, A., Maddison, C.J., et al., (2016). Mastering the Game of Go with Deep Neural Networks and Tree Search. Nature, 529, pp. 484-489.

[13]   Fukushima, K., Neocognitron: (1980). A Self-Organizing Neural Network Model for a Mechanism of Pattern Recognition Unaffected by Shift in Position. Biological Cybernetics, 36, pp. 193-202.

[14]   Lecun, Y., Bottou, L., Orr, G.B., et al., (1998). Efficient Backprop. Neural Networks Tricks of the Trade, 1524, 1998, pp. 9-50.

[15]   Goodfellow, I., Bengio, Y. and Courville, A. (2016). Deep Learning. The MIT Press, Cambridge.

[16]   Fujii, K. (2018). Mathematical Reinforcement to the Minibatch of Deep Learning. Advances in Pure Mathematics, 8, 307-320.

[17]   Xuan, X., Peng, B., Wang, W. and Dong, J. (2019). On the Generalization of GAN Image Forensics. In: Chinese Conference on Biometric Recognition, Springer, Berlin, 134-141.

[18]   Vaccari, C. and Chadwick, A. (2020). Deepfakes and Disinformation: Exploring the Impact of Synthetic Political Video on Deception, Uncertainty, and Trust in News. Social Media + Society, 6, 1-13.

[19]   Wang, F., Xing, L., Bagshaw, H., Buyyounouski, M. and Han, B. (2020). Deep Learning Applications in Automatic Needle Segmentation in Ultrasound-Guided Prostate Brachytherapy. Medical Physics.

[20]   McClelland, J.L., et al., (1995). Why There Are Complementary Learning Systems in the Hippocampus and Neocortex: Insights from the Successes and Failures of Connectionist Models of Learning and Memory. Psychological Review, 102, 1995, pp. 419-457.

[21]   Kumaran, D., Hassabis, D. and McClelland, J.L., (2016). What Learning Systems Do Intelligent Agents Need? Complementary Learning Systems Theory Updated. Trends in Cognitive Sciences, 20, pp. 512-534.

[22]   Wang, R., (2019). Research on Image Generation and Style Transfer Algorithm Based on Deep Learning. *Open Journal of Applied Sciences*, 9, pp. 661-672.

[23]   Krizhevsky, A., Sutskever, I., Hinton, G.E., et al., (2012). ImageNet Classification with Deep Convolutional Neural Networks. Neural Information Processing Systems, 141, pp. 1097-1105.

[24]   Long, J., Shelhamer, E., Darrell, T., et al., (2015). Fully Convolutional Networks for Semantic Segmentation. Computer Vision and Pattern Recognition, Boston, pp. 3431-3440.

[25]   Noh, H., Hong, S., Han, B., et al., (2015). Learning Deconvolution Network for Semantic Segmentation. International Conference on Computer Vision, Santiago, pp. 1520-1528.

[26]   Cheng, Z., Yang, Q., Sheng, B., et al., (2015). Deep Colorization. International Conference on Computer Vision, Santiago, 415-423.

[27]   Mahendran, A. and Vedaldi, A., (2015). Understanding Deep Image Representations by Inverting Them. Computer Vision and Pattern Recognition, Boston, pp. 5188-5196.

Author

Prof. Yew Kee Wong (Eric) is a Professor of Artificial Intelligence (AI) & Advanced Learning Technology at the HuangHuai University in Henan, China. He obtained his BSc (Hons) undergraduate degree in Computing Systems and a Ph.D. in AI from The Nottingham Trent University in Nottingham, U.K. He was the Senior Programme Director at The University of Hong Kong (HKU) from 2001 to 2016. Prior to joining the education sector, he has worked in international technology companies, Hewlett-Packard (HP) and Unisys as an AI consultant. His research interests include AI, online learning, big data analytics, machine learning, Internet of Things (IOT) and blockchain technology.