



ISSN: 2146-1740
<https://dergipark.org.tr/tr/pub/ayd>,
Doi: 10.54688/ayd.1104830
Review Article



PROBLEMS AND OPPORTUNITIES OF ARTIFICIAL INTELLIGENCE

Necmi GÜRSAKAL¹

Sadullah ÇELİK²

Bülent BATMAZ³

Abstract

Article Info

Received:
17/04/2022

Accepted:
15/06/2022

This article reviews Artificial Intelligence (AI)'s challenges and opportunities and discusses where AI might be headed. In the first part, it was tried to reveal the differences between Symbolic AI and Deep Learning approaches, then long promises but short deliveries of AI were mentioned. The problems of AI are that the media has high expectations about artificial intelligence and keeps the problems and restrictions it creates low. Today, while AI is stuck with issues such as deepfake applications and carbon footprints that create moral and climatologic problems; on the other hand, it is struggling with problems such as deep learning models requiring huge amounts of data. Another problem with deep learning is that deep learning models are a black-box and not open to improvements because it is not known where mistakes were made. Among the new paths ahead of AI are Hierarchical Temporal Memory (HTM) models and hybrid models that generally try to bridge the gap between Symbolic AI and Connectionist AI. If we consider that the most important leaps in AI have been made with the features of the brain that AI can imitate, then the developed HTM models may also be a new opportunity for AI.



Keywords: Hierarchical temporal memory, Deep learning, GPT-3, Black-box, Carbon footprint.

Jel Codes: C01, C10, C15, C19.

¹ Professor, Mudanya University, ORCID: 0000-0002-7909-3734, negursakal@gmail.com.

² **Corresponding Author:** Research Assistant (PhD), Aydın Adnan Menderes University, ORCID: 0000-0001-5468-475X, sadullah.celik@adu.edu.tr.

³ Lecturer (PhD), Anadolu University, 0000-0003-4706-5808, bbatmaz@anadolu.edu.tr.

Cite: Gursakal, N., Celik, S. & Batmaz B. (2022). Problems and opportunities of artificial intelligence. *Journal of Academic Approaches*, 13 (1), 203-225.



1. Introduction

Artificial Intelligence is an area of computer science and engineering focused on the creation of intelligent agents, which are systems that can learn, and act autonomously (Rathore, 2022). Intelligent systems are those that can reason, learn, and solve problems like humans. AI researchers are working on creating machines that can intelligently process and respond to complex real-world situations. This involves creating algorithms or rules that can analyse data and make predictions or recommendations. It also involves creating systems that can interact with humans and other agents in natural ways. AI has the potential to transform many industries and revolutionize the way we live and work. It is already used in a variety of applications, including healthcare, finance, manufacturing, and transportation. AI holds the promise of transforming our economy and society in ways that we cannot even imagine. It has the potential to improve our quality of life in several ways, including providing us with new and better ways to solve problems, making us more efficient and productive, and giving us more time to enjoy our lives.

Between 1588 and 1679, Thomas Hobbes said: “*Thinking is a manipulation of symbols and Reasoning is computation*”. In 1950, Alan Turing asked the question of whether machines could think. In the late 1950s, studies on Symbolic AI intensified, with a focus on the ability of machines to manipulate symbols. Symbolic AI algorithms are based on formal logic and use symbols to represent knowledge. This paradigm has continued into the present day, with researchers striving to create ever more powerful and intelligent machines (Singhi, 2019).

The history of artificial neural networks and deep learning begins with the perceptron, a simple form of neural network introduced in 1958 by Frank Rosenblatt. However, it was not until 2012 that artificial neural networks and deep learning began to gain popularity. In 2012, AlexNet was the beginning of deep learning models. It included “*a total of 8 layers, 5 convolution layers, and 3 fully connected layers*”. It had 9216 nodes, 62 million parameters, and 600 million connections (Krizhevsky et al., 2012). In recent years we live in, AI as an artificial neural network model has taken a big leap forward in its achievements in solving specific applications, but it is widely acknowledged that we are still a long way from general-purpose AI systems (Dickson, 2022). In this case, we can say that 60 years of the last 70 years of AI have been dominated by Symbolic AI; but in the last 10 years, Machine Learning (ML) and deep learning have made a big breakthrough.

Symbolic AI and Connectionist AI are both AI approaches that aim to simulate human intelligence. Symbolic AI uses symbols and rules to represent knowledge, while connectionist AI uses a network of nodes and connections to represent knowledge. Both approaches have their strengths and weaknesses, but they also share some common features. For example, both approaches use algorithms to learn from data and make predictions. Both approaches can be used to solve problems in areas such as planning, decision making, and natural language processing. A comparison of Symbolic AI and Connectionist AI is given in Table 1.

Table 1
Symbolic AI and Connectionist AI

Symbolic AI (Rule-Based AI)	Connectionist AI (ML/Subset Deep Learning)
Symbolic AI focuses on rule-based, logical reasoning.	Connectionist AI is based on neural networks.
Symbolic AI is based on formal logic and rules.	Connectionist AI is based on statistical methods.
Symbolic AI is more focused on single-task learning.	Connectionist AI is more adept at learning multiple tasks.
Symbolic AI often requires more human intervention and knowledge to function.	Connectionist AI can learn from data with little to no human intervention. (Solution: Synthetic data)
Symbolic AI is better at understanding and representing abstract concepts.	Connectionist AI is better at dealing with noisy or unstructured data.
Symbolic AI can represent information.	Connectionist AI is based on implicit knowledge representation.
Symbolic AI is better at dealing with complex problems that require step-by-step reasoning.	Connectionist AI, on the other hand, is better at dealing with simple problems that can be solved by trial and error.
The operation of the symbolic AI system requires less energy and has a low carbon footprint.	Connectionist AI has a huge carbon footprint as it is capable of self-learning using data that users make available.
Symbolic AI systems are better at dealing with precise, well-defined problems.	Connectionist AI systems are better at dealing with more open-ended problems.
Symbolic AI systems tend to be more domain-specific.	Connectionist AI systems can be more general purpose.
We don't have a black-box problem in symbolic AI.	We have a black-box problem in connectionist AI.

Source: Dickson, 2019.

The main purpose of artificial neural networks in the past 70 years has been to mimic the functioning of the human brain. However, there are still some deficiencies in imitating the human brain. In this article, we will discuss the opportunities for eliminating these deficiencies.

2. Long Promises but Short Deliveries of Artificial Intelligence

Did Herbert Simon, one of the AI pioneers, set the bar too high for the promises of AI by saying, “*Machines will be capable, within twenty years of doing any work a man can do*”

back in 1965 (Zador, 2019:1)? We must answer “yes” to this question because the promise has yet to be fulfilled, even in 2022. Even those who are very close to AI technology may make significant evaluation errors when predicting its future advancements. For example, machines are still far from understanding human language. It will likely be many years before machines can do any work that a human can do (Marcus, 2022). However, this does not mean that AI technology has not made significant advancements since 1965. AI has made great strides in many areas, such as pattern recognition and data analysis.

We should never forget that technology has fascinating features which can make the impossible appear possible, almost magical. For example, doors can open without us even saying “open the door!”, and cars can drive themselves. This magical appearance of technology might mislead us into thinking that the reality of technology is bigger than it is, both in the short term and the long term. The reason for such confusion might be that, as research has shown (Naser & Ross, 2022), some communities have tied AI to a mystique stigma. The reason AI is not generally included in undergraduate curricula might be that “*deep learning has mathematical foundations which could be beyond the reach of average undergraduate students*” (Hoover et al., 2019). Most people involved in AI are interested in it as appliers, with AI algorithms commonly being black-box algorithms. In addition, AI is often associated with futuristic concepts like robots and cyborgs, which can make it seem even more magical and mysterious. However, it is important to remember that AI is just a tool, and like any tool, it has its limitations.

Geoffrey Hinton, who can be considered the Godfather of deep learning, said at an AI conference in 2016 that deep learning will eventually surpass radiologists' inaccuracy. He reasoned that people should stop training radiologists now and that it's just a matter of time before deep learning does better. In 2020, Hinton said that “*deep learning can do anything*” (Marcus, 2022). However, even though we have arrived at the year 2022, radiologist training has not been abandoned. Maybe we need to stop and wait for a while because we know that in the short term, we overestimate the truth about technology, and in the long term, we overestimate what it will be.

IBM Watson Health, which was said to revolutionize health, was sold in January 2022. A private equity group has agreed to take over most of IBM Watson Health, seven years after the business was launched with a pledge to revolutionize healthcare data analysis (Taylor, 2022). But even after part of Watson Health has been sold, one can write that the company has

made impressive advances in data analytics. For example, the Watson algorithms have been very successful in managing large amounts of data and finding patterns that might not have been otherwise apparent. This has been a significant breakthrough (Pati, 2022). All this shows us that there is a problem in the media's view of AI.

A Tesla car failed to recognize a person holding up a stop sign in the middle of the road because the scene was too far enough outside of the training database that the system had no idea what to do (Marcus, 2022). When a Deep Neural Network (DNN) is given an image and then the same image is given to the DNN by manipulating a pixel or two, the DNN can identify the orange as a hill (Szegedy et al., 2013), with minor perturbations to their inputs that lead to vastly different predictions (Schuller et al., 2022: 16). This is especially problematic when it comes to safety-critical applications like autonomous vehicles.

As a result, deep learning is most effective when quick, approximate results are desired. There is no agreement among AI researchers about how fast AI will progress in the future (Marcus, 2022) and no scientific basis for making such predictions. Scientific research can make slow progress on certain problems (e.g., unifying all forces in physics) or make rapid advances following a breakthrough (e.g., that of deep learning) (Bengio, 2022).

3. Black-box Problem

The black-box problem is a major challenge in AI research. It refers to the difficulty of understanding or explaining the behaviour of a system that is opaque or appears to be a "black box". This can be a result of the system being too complex, or if it is a proprietary system whose inner workings are not known. The black-box problem is especially relevant to ML, where algorithms are often used that are not fully understood by the researchers who created them. This can make it difficult to debug the system or improve it. The black-box problem also raises ethical concerns, as opaque systems can be used to make decisions that affect people's lives without understanding how or why those decisions were made.

As data sets, model sizes, and computer speed increase, it becomes difficult for people to interpret and understand the ever-growing deep learning models. The number of training data and parameters used by these models is constantly increasing, making them more complex and opaquer. Hyperparameter (HP) tuning in neural networks with billions of parameters is a costly process. On the other hand, models such as the GPT-3 can be larger and more fluid (Yang et al., 2022). However, this does not mean that the GPT-3 information models are more reliable

(Marcus, 2022). Studies have shown that scaling alone offers some improvements in all metrics. However, research shows that improvements in safety and grounding lag far behind human performance (Thoppilan et al., 2022).

If we shift our focus from big data to good data, we can find solutions to big issues in AI using only small data sets. According to Andrew Ng, this data-centric approach could offer humankind "small data" solutions to big problems in AI. Ng says that in many industries where large datasets are not available, the focus should be on designing a few good example datasets that can teach a neural network what we want it to learn (Egri-Nagy & Törmänen, 2022).

The model with the greatest expectations for AI today is the artificial neural networks (ANN) model. The most interesting feature of this model is so complex that we do not know exactly how it works and hence considering this model as a black box. These deep learning models are used more and more "in speech recognition, object detection, and recognition, drug discovery, genomics, and in numerous other domains", with the increasing computing power of machines but these models are not interpretable. As these models gain importance and decisions concerning our lives are increasingly taken by these models, the fact that they have the black-box feature becomes a problem for individuals and societies (Mortensen, 2022). Again, these models, which we mentioned, give answers to the questions asked to them, starting with the most correct answer with high probability, without contributing to human knowledge, and without explanation, in a black-box manner (Vale et al., 2022).

As neural networks get more complex with more layers and more neurons, their performance improves. But this increase in complexity makes networks less transparent and harder to understand. This lack of transparency makes it difficult to maintain, debug, and test the neural network for robustness (Vale et al., 2022: 1). The interpretability of neural networks shows that they are much more than just a neural network. Modern neural networks consist of many layers, and it is almost impossible to determine which layer is responsible for the capabilities of the network. This makes it difficult to reuse network components by traditional software (Chaudhuri et al., 2021: 3).

The lack of transparency in machine-learning algorithms is due to two main reasons: the complex structure of the algorithm, which exposes the complexity of the algorithm; and the reliance of the algorithm on geometric relationships that humans cannot visualize, known as dimensionality (Bathae, 2018: 901). These reasons make it difficult for humans to understand how the algorithm works and why it makes the decisions it does.

The lack of transparency in artificial neural network models prevents users from understanding, interpreting, and trusting the output of these models. This lack of transparency makes it difficult for software engineers to improve these models, legal bodies to assess whether the models are working correctly and individuals to exercise their legal rights concerning the decisions made by the models. Software engineers may not intervene to improve performance because individuals are less likely to trust and control machines. Legally, the lack of transparency makes it difficult for regulators to determine “*whether a particular system processes data fairly and securely*”. This can prevent individuals from exercising their rights (Zednik, 2021: 1-2). A solution to this problem is therefore sought, in the form of so-called explainers or explainers (Fraternali et al., 2022: 10), which would reveal a simple connection between the inputs and outputs of these models (Buhrmester et al., 2021: 983). However, there is still much work to be done to create reliable and effective explainers. It is important to ensure that explainers do not inadvertently introduce bias or distort the original model.

4. The Carbon Footprint of Artificial Intelligence

The symbolic AI carbon footprint is an approach used by an organization to determine its carbon footprint. This approach does not cover all the information that can be used to determine an organization’s carbon footprint size. The connectionist AI carbon footprint is an approach used by an organization to determine its carbon footprint. This approach allows an organization to determine its carbon footprint more accurately.

The carbon footprint of data centres is estimated to be more than 2 percent of global carbon emissions and expected to rise to 3.2 percent by 2025 and 14 percent by 2040 (Soon & Hui, 2022). A single model is capable of emitting as much carbon as five cars in its lifetime (Hao, 2019). Google’s AlphaGo Zero, which plays the game of Go against itself to self-learn, generated 96 tonnes of carbon dioxide over 40 days of research training (Soon & Hui, 2022). Data centres all over the world use large amounts of energy to power the computers that run AI algorithms. This energy use results in carbon emissions that contribute to climate change.

Human activity is responsible for the warming of the atmosphere, ocean, and land. The world has undergone widespread and rapid changes in its climate, with observed increases in greenhouse gas concentrations since the 1700s (Masson-Delmotte et al. 2021: 4). Calculations require doubling every few months. That's why the biggest changes are happening in deep learning research today. This has increased nearly 300,000 times from 2012 to 2018. The carbon footprint of AI training efforts is enormous, and it is increasing exponentially with a 3.4-month

doubling. This is compared to Moore's Law, which has a doubling time every 2 years (Schwartz et al., 2020).

It is expected that AI will be used in two major ways concerning climate change: to understand it and to combat it (Schuller et al., 2022: 2). However, just as a bread knife can be used to feed a hungry person or to kill someone, AI can also be used for both good and bad purposes. In this respect, AI is like a bread knife. As we will use it to solve the issue of climate change, AI contributes to climate change with the carbon footprint it produces. Deep learning is a powerful tool for a wide range of tasks, including image recognition, natural language processing, and predictive modelling. However, the computations required for deep learning can be very energy-intensive, making it a costly endeavour. This can be an obstacle for researchers who want to participate in deep learning studies, especially in developing economies (Schwartz et al., 2020).

The development of large natural language processing models with deep learning models can be costly, both financially and environmentally. This is due to the hardware and electricity cost required to train these models, as well as the cloud computing time and carbon footprint required to fuel modern tensor processing hardware (Schuller et al., 2022).

In recent years, many studies in the AI literature have focused on the energy consumption and carbon footprints of deep learning programs (Lacoste et al., 2019; Lannelongue et al., 2021; Strubell et al., 2019). In these studies, Green AI is proposed, and tools are proposed to evaluate deep learning models in terms of energy consumption and carbon footprints. By better understanding the environmental impact of training ML models, we can make more informed decisions about how to develop and deploy these models (Lacoste et al., 2019).

The challenges and opportunities for AI's carbon footprint may vary depending on the specific application of AI. However, some potential problems that could arise as a result of the use of AI in carbon footprint reduction efforts are as follows:

- (1) Misuse of AI to create inaccurate or biased carbon footprint models – If AI is used to create models that estimate an individual or organization's carbon footprint, there is a risk that these models could be inaccurate or biased. This could lead to individuals or organizations being unfairly penalized for their carbon footprint or could result in the misallocation of resources devoted to reducing carbon footprints.

- (2) Difficulty in scaling AI-based carbon footprint reduction efforts – While AI-based carbon footprint reduction efforts may be effective on a small scale, it may be difficult to scale up these efforts to cover a larger population. This could limit the overall effectiveness of AI in reducing carbon footprints.
- (3) Implementation of AI-based carbon footprint reduction efforts could lead to job losses – If AI is used to automate carbon footprint reduction efforts, this could lead to job losses in sectors such as transportation and energy. This could have negative social and economic consequences and could offset some of the potential benefits of AI in reducing carbon footprints.

5. Negative Sides of GPT-3

The GPT-3 is an artificial intelligence that was created by Google. It is designed to be a general-purpose learning machine, capable of learning from a variety of data sources and making predictions about the world. The GPT-3 is based on a deep learning algorithm called a recurrent neural network (RNN) (Narasimman, 2022). RNNs are a type of artificial neural network that are well-suited for learning from sequential data, such as text.

The GPT-3 has a number of applications. One potential application is text generation. The GPT-3 can be used to generate new text, based on the statistical properties of the Google Billion Word Corpus. This text can be in the form of articles, essays, short stories, lyrics or even novels (Anjali, 2021). The GPT-3 can also be used for machine translation. By learning the statistical properties of a corpus of text in one language, the GPT-3 can generate text in another language that is similar to the original text.

The GPT-3 has also been used for question answering. The GPT-3 can be given a question and a corpus of text, and it will generate an answer to the question (Zhang et al., 2021). This answer is generated by finding the text in the corpus that is most similar to the question. The GPT-3 has been used to answer questions about a variety of topics, including history, science, and politics.

The GPT-3 is a powerful artificial intelligence that has a number of potential applications (Sezgin et al., 2022). It is capable of learning from a large dataset of text and making predictions about new text. The GPT-3 can be used for text generation, machine translation, and question answering.

The risks associated with large language models (foundation models) are becoming more and more apparent as AI and ML become more prevalent. These risks include unfair discrimination and social biases, private or other sensitive information being inferred, and data leaks. These risks can lead to false and misleading information being provided, causing harm to both users and product developers (Weidinger et al., 2021). It is becoming increasingly clear that foundation models pose a serious threat to equity, misuse, economic and environmental impact, and legal and ethical considerations (Bommasani et al., 2021). Foundation models are only just beginning to transform the way AI systems are built and deployed in the world (Bommasani et al., 2021), and these risks must be addressed.

One of the problems with large language models is the issue of bias. For example, “Muslim” is associated with “terrorist” in 23% of cases, while “Jewish” is mapped to “money” in 5% of cases. This indicates a clear bias against Muslim people and Jewish people (Abid et al., 2021:1). One possible solution to this problem is to create language models that are trained on data that is diverse and representative of different groups. Another solution is to use debiasing techniques to reduce the amount of bias in the language model.

The use of AI and ML has changed the landscape of security risks for citizens, organizations, and governments. Malicious use of AI includes criminals compromising digital security using AI and ML tools; the endangerment of physical security by non-state actors, such as weaponizing drones; also, we can count political security being compromised such as interference in elections and disinformation campaigns (Brundage et al., 2018: 10). The use of AI in weapon systems is a major strategic and tactical decision with far-reaching implications. The current state of AI development, with its accelerating capabilities, hidden decision-making processes, and decreasing costs, necessitates careful consideration of how these weapon systems will be used. Military policies and doctrines must be put in place to ensure that these weapons are used responsibly (Feldman et al., 2019).

Virtual potential toxic molecules are methods used to calculate the toxicity potential of a molecule. This method is used to predict the action potential of the molecule in a human body, using the molecule’s structure and properties. Virtual potentially toxic molecules are one of the many methods that can be used to predict the toxicity of a molecule and can be used to detect molecules with a high potential for toxicity. “*It is possible to design virtual potential toxic molecules with little effort, time, or computational resources*”, according to AI studies (Urbina

et al., 2022). This raises the possibility that drug-discovery projects created for positive purposes, such as treating diseases, could be abused.

Some of the idioms about the downside of GPT-3 are as follows:

It can be said for the GPT-3: “*Bloviator: Open AI’s language generator has no idea what it’s talking about*” (Marcus & Davis, 2020).

Are the big language models “too repetitive, little understanding” stochastic parrots? (Bender et al. 2021: 610).

When you ask GPT-3 this question, “*Do GPT-3 repeat more or understand more? Explain 200 words*”. It answers:

“GPT-3 Repeat More

GPT-3 is a ML platform that enables developers to train and deploy AI models. It is designed to be scalable and efficient, so developers can build models that can handle large amounts of data.

GPT-3 understands more than just words; it can also understand relationships between words. This makes it possible for GPT-3 to better understand the context of a sentence. For example, if you ask GPT-3 to define the word “bank”, it will not only give you the definition but also explain the meaning of the word concerning other words, such as “savings” and “loan.”

GPT-3’s understanding of relationships between words also enables it to generate better responses to questions. For example, if you ask GPT-3 to define the word “car”, it will not only give you the definition but also provide examples of cars and explain the difference between a car and a truck.

Another time, the answer is:

GPT-3 repeats more than it understands. It can only understand a certain number of commands at a time, and after it has repeated those commands a certain number of times, it will start to repeat them. This can be frustrating for humans trying to communicate with it, as it often does not understand what is being said. However, its ability to repeat commands makes it an invaluable tool for research, as scientists can be sure that it will carry out a task exactly as instructed.

6. The Human Brain and Artificial Intelligence

Two basic approaches have been observed in AI to date: Symbolic AI (Rule-Based AI) and Connectionist AI (ML/Subset Deep Learning). As indicated in Table 1, the solution for both approaches are to mimic the human brain more accurately and use hybrid models. The idea of combining neural and symbolic methods begins in 2002 (Garcez et al., 2002) and has a long history in AI research (Chaudhuri et al., 2021: 8). The goal of neuro-symbolic programming is to learn functions from data by combining the strengths of deep learning and program synthesis (Chaudhuri et al., 2021: 1). This approach has the potential to improve the accuracy of ML models by combining the strengths of deep learning and program synthesis.

Artificial neural networks have been developed and inspired by the human brain. We should not be surprised by the success of artificial neural networks, despite our incomplete understanding of the brain. After all, humans learn best when examples are presented in a meaningful order, rather than randomly. Yet in ML, we typically give data to machines in random order (Soviany et al., 2022: 2). Artificial neural networks can still be effective without meeting such simple requirements. This is likely because artificial neural networks can learn from data in a way that is like how humans learn. In other words, they can learn from data in a way that is more efficient and effective than other ML algorithms.

Daniel Kahneman's work on cognitive tasks is divided into "*two groups, System 1 and System 2. System 1 operates automatically and quickly, with little or no effort and no sense of voluntary control*" (Kahneman, 2011: 104). It is intuitive, fast, and unconscious in pattern recognition. "*System 2 allocates attention to the effortful mental activities that demand it, including complex computations*" (Kahneman, 2012: 3), which are slow, logical, and algorithmic. Joshua Bengio's work expands on this by explaining that the current deep learning landscape is mostly related to System 1 thought, but that future AI systems will need to also automate System 2 thought to be truly effective (Chaudhuri et al. 2021: 4). This is an important area of research, as it seeks to couple classical symbolic algorithms with contemporary deep learning techniques (Chaudhuri et al. 2021: 4).

Let's try to look at the issue from another angle. Does the brain do sparse coding or dense coding?

The efficient coding hypothesis of Barlow postulates (Barlow, 1961) that a neural code is optimized to minimize the number of spikes needed to transmit a given signal. This

hypothesis suggests that populations of visual neurons are highly sparse (Lehky et al., 2021). In neuroscience, the term “sparse coding” refers to the brain's ability to relay essential information using only a few (sparse) neurons. This sparse coding allows the brain to be efficient in its use of resources, and it is an important factor in how the brain processes information.

Sparse coding has several advantages, which include increased storage capacity in associative memories, explicit representation of structure in natural signals, easier read-out of complex data at subsequent processing levels, and saved energy. Different modalities and organisms use sparse coding as a strategy, according to recent physiological recordings from sensory neurons (Olshausen & Field, 2004: 481). Some of the other advantages of sparse coding include the fact that it is a more efficient way to store information in memory, and that it can help to improve the signal-to-noise ratio of sensory information.

The issue of sparseness and denseness in the brain is still controversial. As a result of the neurophysiological studies conducted by the scientists, it was seen that it identified sparsity and other efficient coding measures in some areas of the visual cortex (Franco et al., 2007: 2; Lehky et al., 2021: 547). However, some studies have concluded that this is not always the case (Purdue University News, 2021). Considering what they are going to mimic, “sparseness” is an important issue because “*most artificial networks today rely on dense representations, whereas biological networks rely on sparse representations*” (Ahmad & Scheinkman, 2019: 1). There is evidence that sparseness may be beneficial for some tasks, but it is not clear if this is always the case. Sparseness may be just one of many features that contribute to efficient coding in the brain.

Dendrites within a single neuron have been found to receive input from distinct parts of the brain and process it in different ways. These differences may help neurons to integrate a variety of inputs and generate an appropriate response (Trafton, 2022). Recent studies have determined that there is dendritic integration of long-range inputs in the brain. How the neocortex works, however, is still a mystery (Hawkins et al., 2019) In short, all of this suggests that we are still far from mimicking the brain well.

7. New Ways for Artificial Intelligence: Hierarchical Temporal Memory

Hierarchical Temporal Memory, or HTM, is a theory of AI that builds on the ideas of neural networks and machine learning. HTM is based on the theory of the neocortex, which is

the part of the brain responsible for higher-level cognition (Bautista et al., 2020: 2). The HTM algorithm is designed to learn and recognize patterns in data, and to make predictions based on those patterns.

HTM is a neural network algorithm that uses a hierarchy of cells to store and process information. The cells are arranged in layers, with each layer containing several cells that are connected to the cells in the layer below (Struye & Latré, 2020: 298). The cells in the top layer are called pyramidal cells, and they are responsible for encoding information into a form that can be stored in the cells below. The cells in the bottom layer are called granular cells, and they are responsible for storing the information that has been encoded by the pyramidal cells.

The HTM algorithm can learn and recognize patterns in data by using a process called spatial pooling. Spatial pooling is a method of representing data in a way that makes it easier for the HTM algorithm to learn from. In spatial pooling, the data is first divided into several small regions, each of which is then represented by a single number (Riganelli et al., 2021: 785). The number that represents each region is called a cell. The cells are then arranged in a grid, with each cell being connected to the cells in its immediate vicinity.

The HTM algorithm can make predictions by using a process called temporal pooling. Temporal pooling is a method of representing data in a way that makes it easier for the HTM algorithm to make predictions about future events. In temporal pooling, the data is first divided into several small regions, each of which is then represented by a single number. The number that represents each region is called a cell. The cells are then arranged in a grid, with each cell being connected to the cells in its immediate vicinity. The HTM algorithm makes predictions about future events (Hawkins & Blakeslee, 2004: 61) by using the cells in the grid to represent the data in a way that is similar to how the human brain makes predictions about future events.

“The Thousand Brains Theory of Intelligence” (Hawkins, 2021) proposes “*that every cortical column learns complete models of objects, as opposed to traditional hierarchical ideas in deep learning where objects are learned only at the top. This theory hypothesizes that there are many models of each object distributed throughout the neocortex*” (Lomonaco, 2019), resulting in a more distributed and efficient learning process. An important feature of HTM models is that they require less storage and processing memory than traditional processing methods (Krestinskaya et al., 2018: 1143).

The HTM is a neural network model that can be used for sequence learning and prediction (Niu et al., 2022: 1). The HTM is suitable for modelling the structure and function of the neocortex. The HTM is composed of two main modules: the Spatial Pooler (SP) and the Temporal Memory (TM). The SP module converts the input data into a sparse distributed representation with a fixed sparsity. The TM module learns sequences and makes context-sensitive predictions (Niu et al., 2022: 1). Another study has proposed using the SP learning algorithm in the HTM for document categorization. The HTM accepts only a stream of binary data as input. To convert the input data into the binary format, the Latent Semantic Indexing technique is used to extract the top features from the input (Shah et al., 2017: 3).

“In HTM, each cell has one proximal dendritic segment and multiple distal dendritic segments” (Niu et al., 2022: 1). In each section, there are synapse clusters, which are expressed by the scalar persistence value. In addition, while all cells in the same column share the proximal segment, the cell in the distal segments receives lateral input from neighboring cells (Niu et al., 2022: 1). How the cortical hierarchy works is seen in Figure 1.

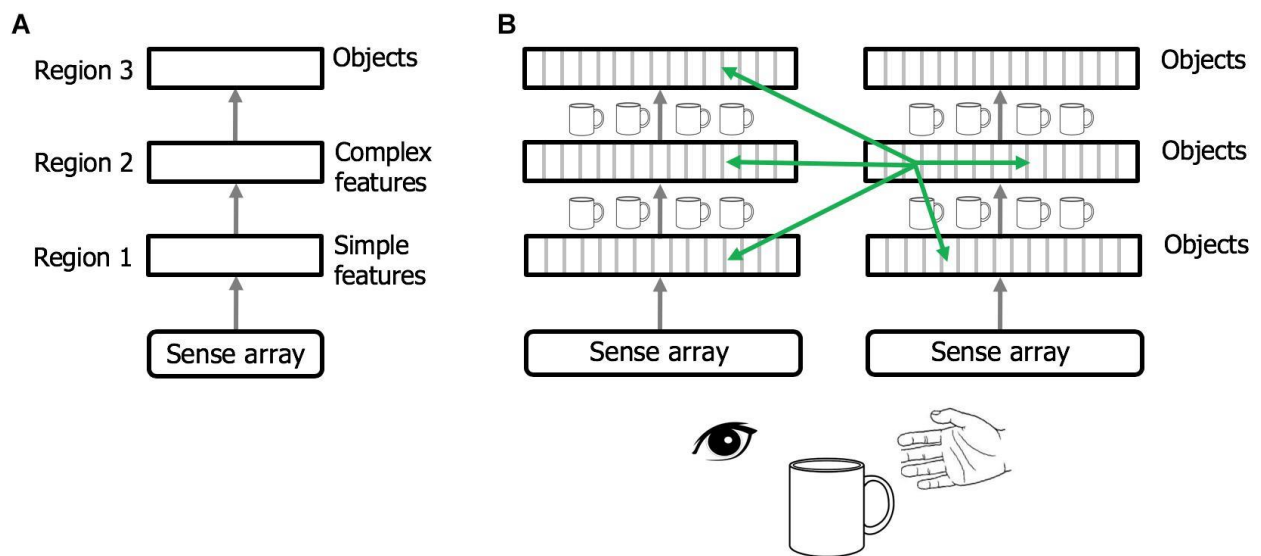


Figure 1.
Cortical Hierarchy Structure
Source: Hawkins et al., 2019: 8

The cortical hierarchy structure is given in Figure 1. In (A), the general view of the cortical hierarchy is given. In (B), a modified view of the cortical hierarchy is given. Each column in each region in Figure 1 learns a complete object model. In Figure 1 two sensory hierarchies are given, one for vision and the other for touch. The green arrows here indicate non-hierarchical numerically-large cortical to cortical connections (Hawkins et al., 2019: 8).

In recent years, interest in the neuro-symbolic approach to AI, which combines the strengths of deep learning and symbolic reasoning, has grown tremendously (Mao et al., 2019). The neuro-symbolic approach to AI is an approach that combines artificial neural networks and other symbolic data processing techniques to make a computer system work like a real brain. The neuro-symbolic approach to AI aims to model brain functions to understand ways of doing brain tasks. This approach is aimed at making computer systems such as learning, understanding, making assumptions, guessing, and problem-solving work like a real brain.

The neuro-symbolic concept learner (NS-CL) is one of the key neuro-symbolic studies (Mao et al., 2019). The NS-CL is a machine learning algorithm that can learn concepts by combining both neural networks and symbolic reasoning. This algorithm can learn new concepts by combining existing concepts that it has already learned. This allows the NS-CL to learn concepts more quickly and accurately than other machine learning algorithms.

In an HTM model, input elements are binary, and there are two kinds of dendrites in a neuron: those that process excitatory input and those that process inhibitory input (Byrne, 2015: 3). Although an artificial neuron has relatively few synapses and no dendrites (Numenta, 2017), biological neurons have thousands of synapses arranged along dendrites. Additionally, dendrites are active processing elements that allow the neuron to recognize hundreds of unique patterns (Numenta, 2017). In Figure 2 we see the comparison of neuron models.

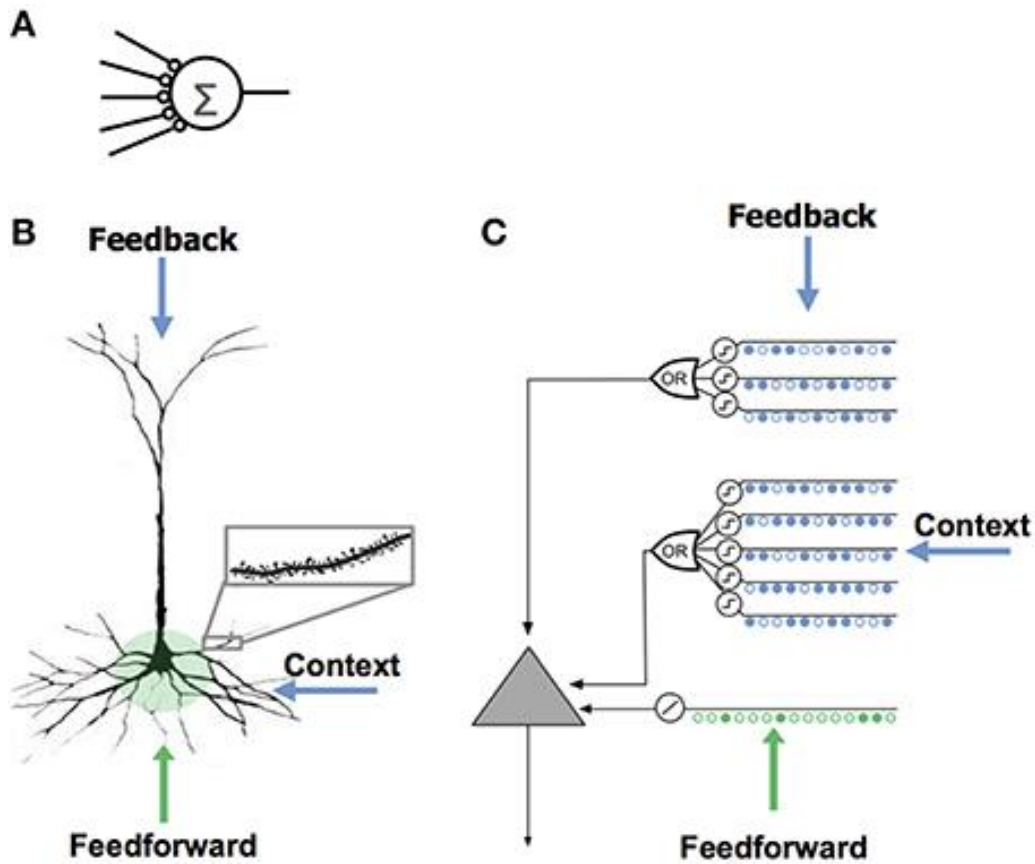


Figure 2.
Comparison of neuron models
Kaynak: Hawkins & Ahmad, 2016: 2

Figure 2 shows the comparison of neuron models. In Figure (A), most neural network models have few synapses and no dendrites. In Figure (B), a neocortical neuron has hundreds of thousands of synapses located in dendrites. Finally, Figure (C) is an HTM model neuron and models dendrites with a set of synapses (Hawkins & Ahmad, 2016: 2).

Some potential opportunities for AI for HTM include:

- (1) Increased efficiency: AI systems can automate tasks that would otherwise be carried out by humans. This can lead to increased efficiency and productivity.
- (2) Improved decision-making: AI systems can help humans to make better decisions by providing them with more data and insights.
- (3) New application areas: As AI technology advances, new application areas are emerging. For example, AI is being used to develop autonomous vehicles and create virtual assistants.

Some potential problems with AI for HTM include:

- (1) Overfitting: If the training data for the AI system is not representative of the real-world data the system will be used on, the system may not be able to generalize well and may overfit the training data. This can lead to poor performance on unseen data.
- (2) Lack of explainability: AI systems can be difficult to understand and explain. This can be a problem when trying to debug the system or understand why it made a particular decision.
- (3) Safety and security concerns: As AI systems become more powerful and sophisticated, there is a risk that they could be used for malicious purposes. For example, an AI system could be used to control a drone strike or to manipulate financial markets. There is also a risk that AI systems could become smarter than humans and pose a threat to our safety and security.

8. Conclusion and Discussion

Some potential problems with AI include data security and privacy concerns, as well as the potential for AI to be used for nefarious purposes. Additionally, there is a risk that AI could become smarter than humans, leading to unforeseen and potentially dangerous consequences.

On the other hand, there are also many potential opportunities for AI. For instance, AI could be used to help humans make better decisions, automate repetitive tasks, and to improve communication. Additionally, AI could help us to better understand and predict the behaviour of other systems, including natural systems.

The main problem with AI is that it can be used to create things that do not exist in nature, which can hurt the environment. For example, if AI is used to create robots that can do the same job as humans, then this could lead to a loss of jobs for people who are employed in that sector. In addition, AI can be used to create things that are harmful to the environment, such as weapons or other devices that pollute the air or water.

The problems and opportunities of AI for the carbon footprint of AI are many and varied. There is potential for AI to help reduce the carbon footprint of AI through several means. For example, AI could be used to manage energy consumption more efficiently in data centres and other facilities where AI is used. AI could also be used to develop new methods and technologies for reducing emissions, such as through the development of cleaner and more efficient energy sources.

There are also several potential risks associated with AI and its carbon footprint of AI. One risk is that AI could be used to develop new methods and technologies for increasing emissions, such as through the development of more polluting energy sources. Another risk is that AI could be used to facilitate environmental destruction, such as through the development of autonomous vehicles that could be used for illegal deforestation or other activities.

There are also several ethical implications associated with AI and its carbon footprint of AI. One ethical implication is that AI could be used to unfairly advantage certain groups or individuals over others regarding the carbon footprint of AI. For example, AI could be used to develop new methods and technologies for reducing emissions that are only accessible to the wealthy. Another ethical implication is that AI could be used to unfairly disadvantage certain groups or individuals concerning their carbon footprint of AI. For example, AI could be used to develop new methods and technologies for reducing emissions that are only accessible to the wealthy.

The challenges and opportunities of AI for GPT-3 vary depending on the specific application and use case. However, some potential problems that could arise from using AI for GPT-3 include errors in the training data that could lead to inaccurate predictions, as well as ethical concerns around the use of AI for decision-making. Additionally, there is also the potential for AI to be used for malicious purposes, such as creating fake news or spreading misinformation. On the other hand, many potential opportunities could be generated from using AI for GPT-3, such as increasing the efficiency of tasks that are currently performed by humans or providing insights that would otherwise be unavailable.

AI's problems and opportunities for hierarchical temporal memory may vary depending on the application and context. However, some general issues to consider include the impact of AI on the efficiency and accuracy of memory recall, as well as the potential for AI to improve the overall performance of hierarchical temporal memory systems. Additionally, it is important to consider the ethical implications of AI concerning memory recall, as well as the potential for AI to be used for malicious purposes.

Peer-review: Externally peer-reviewed.

Contribution Rate Statement: Corresponding author: %35 Other author: %35 Other author: %30.

Conflicts of Interest: There is no potential conflict of interest in this study.

REFERENCES

- Abid, A., Farooqi, M. & Zou, J. (2021). *Persistent anti-muslim bias in large language models*. Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society. (pp. 298-306). USA.
- Ahmad, S. & Scheinkman, L. (2019). How can we be so dense? The benefits of using highly sparse representations. 2 April 2022, <https://arxiv.org/pdf/1903.11257.pdf>.
- Anjali. (2021). What Is GPT-3 & Why is it so popular? 28 May 2021, Medium. <https://medium.com/eoraa-co/what-is-gpt-3-why-is-it-so-popular-b92e87fddafe>.
- Barlow, H. B. (1961). Possible principles underlying the transformations of sensory messages. In W. A. Rosenblith (Ed.), *Sensory Communication*, (pp. 217-234). Cambridge University.
- Bathae, Y. (2018). *The artificial intelligence black box and the failure of intent and causation*. Harvard Journal of Law & Technology, 31 (2), 890-938.
- Bautista, I., Sarkar, S. & Bhanja, S. (2020). Matlab HTM: A sequence memory model of neocortical layers for anomaly detection. *SoftwareX*, 11 (1), 1-7.
- Bender, E. M., Gebru, T., McMillan-Major, A. & Shmitchell, S. (2021). *On the dangers of stochastic parrots*. Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency. (pp. 610-623). New York.
- Bengio, Y. (2022). Superintelligence: Futurology vs. science. 24 January 2022, <https://yoshuabengio.org/2022/01/24/superintelligence-futurology-vs-science/>.
- Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., Arx, S. von, Rernstein, M. S. & Liang, P. (2021). On the opportunities and risks of foundation models. 10 April 2022, <https://arxiv.org/pdf/2108.07258>.
- Brundage, M., Avin, S., Clark, J., Toner, H., Eckersley, P., Garfinkel, B., Dafoe, A., Scharre, P., Zeitzoff, T., Filar, B., Anderson, H., Roff, H., Allen, G. C., Steinhardt, J., Flynn, C., Héigeartaigh, S. Ó., Beard, S., Belfield, H., Farquhar, S., ... & Amodi, D. (2018). The malicious use of artificial intelligence: Forecasting, prevention, and mitigation. 10 February 2022, <https://arxiv.org/ftp/arxiv/papers/1802/1802.07228.pdf>.
- Buhrmester, V., Münch, D. & Arens, M. (2021). Analysis of explainers of black box deep neural networks for computer vision: A survey. *Machine Learning and Knowledge Extraction*, 3 (4), 966-989.
- Byrne, F. (2015). Encoding reality: Prediction-assisted cortical learning algorithm in hierarchical temporal memory. 8 March 2022, <https://arxiv.org/pdf/1509.08255.pdf>.
- Chaudhuri, S., Ellis, K., Polozov, O., Singh, R., Solar-Lezama, A. & Yue, Y. (2021). Neuro-symbolic programming. *Foundations and Trends® in Programming Languages*, 7 (3), 158-243.
- Dickson, B. (2019). What happens when you combine neural networks and rule-based AI? 5 November 2021, <https://bdtechtalks.com/2019/06/05/mit-ibm-hybrid-ai/>.
- Dickson, B. (2022). Neuro-symbolic AI brings us closer to machines with common sense. 14 March 2022, <https://bdtechtalks.com/2022/03/14/neuro-symbolic-ai-common-sense/>.
- Egri-Nagy, A. & Törmänen, A. (2022). Advancing human understanding with deep learning go AI engines. *IS4SI 2021*, 22.
- Feldman, P., Dant, A. & Massey, A. (2019). Integrating artificial intelligence into weapon systems. 10 March 2022, <https://arxiv.org/pdf/1905.03899.pdf>.

Franco, L., Rolls, E. T., Aggelopoulos, N. C. & Jerez, J. M. (2007). Neuronal selectivity, population sparseness, and ergodicity in the inferior temporal visual cortex. *Biological Cybernetics*, 96 (6), 547-560.

Fraternali, P., Milani, F., Nahime Torres, R., Zangrando, N. & di Milano Piazza Leonardo, P. (2022). Black-box error diagnosis in deep neural networks: A survey of tools. 17 January 2022, <https://arxiv.org/pdf/2201.06444.pdf>.

Garcez, A. S. d'Avila, Broda, K. & Gabbay, D. M. (2002). *Neural-symbolic learning systems: Foundations and applications*. Springer Science & Business Media.

Hao, K. (2019). Training a single AI model can emit as much carbon as five cars in their lifetimes. 6 April 2022, <https://www.technologyreview.com/2019/06/06/239031/training-a-single-ai-model-can-emit-as-much-carbon-as-five-cars-in-their-lifetimes/>.

Hawkins, J. (2021). *A thousand brains: A new theory of intelligence*. Hachette Book Group.

Hawkins, J. & Ahmad, S. (2016). Why neurons have thousands of synapses, a theory of sequence memory in neocortex. *Frontiers in Neural Circuits*, 10, 1-13.

Hawkins, J. & Blakeslee, S. (2004). *On intelligence: How a new understanding of the brain will lead to the creation of truly intelligent machines*. An Owl Book, Henry Holt and Company.

Hawkins, J., Lewis, M., Klukas, M., Purdy, S. & Ahmad, S. (2019). A framework for intelligence and cortical function based on grid cells in the neocortex. *Frontiers in Neural Circuits*, 12, 1-14.

Hoover, A. K., Spryszynski, A. & Halper, M. (2019). *Deep learning in the IT curriculum*. Proceedings of the 20th Annual SIG Conference on Information Technology Education, (pp. 49-54), Tacoma, WA.

Kahneman, D. (2011). *Thinking, fast and slow*. Macmillian.

Kahneman, D. (2012). *Of 2 minds: How fast and slow thinking shape perception and choice [excerpt]*. Scientific American.

Krestinskaya, O., Ibrayev, T. & James, A. P. (2018). Hierarchical temporal memory features with memristor logic circuits for pattern recognition. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 37 (6), 1143-1156.

Krizhevsky, A., Sutskever, I. & Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems*, 25, 1-9.

Lacoste, A., Luccioni, A., Schmidt, V. & Dandres, T. (2019). Quantifying the carbon emissions of machine learning. 4 November 2021, <https://arxiv.org/pdf/1910.09700.pdf>.

Lannelongue, L., Grealey, J., Inouye, M., Lannelongue, L. & Inouye, M. (2021). Green algorithms: Quantifying the carbon footprint of computation. *Advanced Science News*, 8, 1-10.

Lehky, S. R., Tanaka, K. & Sereno, A. B. (2021). Pseudo sparse neural coding in the visual system of primates. *Communications Biology*, 4 (1), 1-12.

Lomonaco, V. (2019). A machine learning guide to HTM (hierarchical temporal memory). 24 October 2021, <https://numenta.com/blog/2019/10/24/machine-learning-guide-to-htm>.

Mao, J., Gan, C., Kohli, P., Tenenbaum, J. B. & Wu, J. (2019). The neuro-symbolic concept learner: Interpreting scenes, words, and sentences from natural supervision. 6 April 2022, <https://arxiv.org/pdf/1904.12584.pdf>.

Marcus, G. (2022). Deep learning is hitting a wall. 10 March 2022, <https://nautil.us/deep-learning-is-hitting-a-wall-14467/>.

- Marcus, G. & Davis, E. (2020). GPT-3, bloviator: Open AI's language generator has no idea what it's talking about. 22 February 2022, <https://www.technologyreview.com/2020/08/22/1007539/gpt3-openai-language-generator-artificial-intelligence-ai-opinion/>.
- Masson-Delmotte, V., Zhai, P., Chen, Y., Goldfarb, L., Gomis, M. I., Matthews, J. B. R., Berger, S., Huang, M., Yelekçi, O., Yu, R., Zhou, B., Lonnoy, E., Maycock, T. K., Waterfield, T., Leitzell, K. & Caud, N. (2021). *Working Group I Contribution to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. 2021 Intergovernmental Panel on Climate Change. IPCC, Switzerland.
- Mortensen, J. (2022). Is black box AI dangerous? Tech evaluate. 16 March 2022, <https://www.techevaluate.com/is-black-box-ai-dangerous/>.
- Narasimman, L. (2022). The top 4 AI projects everyone is talking about. 1 June 2022, <https://narasimmantech.com/the-top-four-ai-projects-everyone-is-talking-about-2/>.
- Naser, M. Z. & Ross, B. (2022). An opinion piece on the dos and don'ts of artificial intelligence in civil engineering and charting a path from data-driven analysis to causal knowledge discovery. *Civil Engineering and Environmental Systems*, 39 (1), 1-11.
- Niu, D., Yang, L., Cai, T., Li, L., Wu, X. & Wang, Z. (2022). A new hierarchical temporal memory algorithm based on activation intensity. *Computational Intelligence and Neuroscience*, 2022, 1-17.
- Numenta. (2017). Biological and machine intelligence (BAMI). 8 March 2022, <https://numenta.com/assets/pdf/biological-and-machine-intelligence/BAMI-Complete.pdf>.
- Olshausen, B. A. & Field, D. J. (2004). Sparse coding of sensory inputs. *Current Opinion in Neurobiology*, 14 (4), 481-487.
- Pati, S. (2022). The future of artificial intelligence and metaverse in 2030. 7 March 2022, <https://zipe-education.com/the-future-of-artificial-intelligence-and-metaverse/>.
- Purdue University News. (2021). Think the brain is always efficient? Think again. 5 March 2022, <https://www.purdue.edu/newsroom/releases/2021/Q1/think-the-brain-is-always-efficient-think-again.html>.
- Rathore, M. S. (2022). Top artificial intelligence interview questions and answers. TechStack. 16 April 2022, <https://www.techstack.in/blog/artificial-intelligence-interview-questions-and-answers/>.
- Riganelli, O., Saltarel, P., Tundo, A., Mobilio, M. & Mariani, L. (2021). *Cloud failure prediction with hierarchical temporal memory: An empirical assessment*. 2021 20th IEEE International Conference on Machine Learning and Applications (ICMLA), (pp. 785-790), Italy.
- Schuller, B. W., Akman, A., Chang, Y., Coppock, H., Gebhard, A., Kathan, A., Rituerto-González, E., Triantafyllopoulos, A. & Pokorny, F. B. (2022). *Climate change & computer audition: A call to action and overview on audio intelligence to help save the planet a preprint*. Climate Change & Computer Audition, 1-28.
- Schwartz, R., Dodge, J., Smith, N. A. & Etzioni, O. (2020). Green AI. *Communications of the ACM*, 63 (12), 54-63.
- Sezgin, E., Sirrianni, J. & Linwood, S. L. (2022). Operationalizing and implementing pretrained, large artificial intelligence linguistic models in the US Health Care System: Outlook of Generative Pretrained Transformer 3 (GPT-3) as a service model. *JMIR Medical Informatics*, 10 (2), e32875.
- Shah, D., Ghate, P., Paranjape, M. & Kumar, A. (2017). *Application of hierarchical temporal memory theory for document categorization*. 2017 IEEE SmartWorld, Ubiquitous Intelligence & Computing, Advanced & Trusted Computed, Scalable Computing & Communications, Cloud & Big Data Computing, Internet of People and Smart City Innovation. (pp. 1-6), USA.

- Singhi, R. (2019). The rise and fall of symbolic AI. Philosophical presuppositions of AI. 14 September 2021, <https://towardsdatascience.com/rise-and-fall-of-symbolic-ai-6b7abd2420f2>.
- Soon, O. Y. & Hui, L. K. (2022). Making artificial intelligence work for sustainability. 16 April 2022, <https://technologymagazine.com/ai-and-machine-learning/making-artificial-intelligence-work-for-sustainability>.
- Soviany, P., Ionescu, R. T., Rota, P. & Sebe, N. (2022). Curriculum learning: A survey. *International Journal of Computer Vision*, 130, 1-40.
- Strubell, E., Ganesh, A. & Mccallum, A. (2019). *Energy and policy considerations for deep learning in NLP*. In the 57th Annual Meeting of the Association for Computational Linguistics. Italy.
- Struye, J. & Latré, S. (2020). Hierarchical temporal memory and recurrent neural networks for time series prediction: An empirical validation and reduction to multilayer perceptrons. *Neurocomputing*, 396, 291-301.
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I. & Fergus, R. (2013). Intriguing properties of neural networks. 19 February 2022, <https://arxiv.org/pdf/1312.6199.pdf>.
- Taylor, P. (2022). IBM sells off large parts of Watson health business. 24 January 2022, <https://pharmaphorum.com/news/ibm-sells-off-large-parts-of-watson-health-business/>.
- Thoppilan, R., Freitas, D. de, Hall, J., Kulshreshtha, A., Cheng, H.-T., Jin, A., Bos, T., Baker, L., Du, Y., Li, Y., Lee, H., Zheng, S. H., Ghafouri, A. & Le, Q. (2022). LaMDA: Language models for dialog applications. 10 February 2022, <https://arxiv.org/pdf/2201.08239.pdf>.
- Trafton, A. (2022). Dendrites may help neurons perform complicated calculations. 7 February 2022, <https://news.mit.edu/2022/dendrites-help-neurons-perform-0217>.
- Urbina, F., Lentzos, F., Invernizzi, C. & Ekins, S. (2022). Dual use of artificial-intelligence-powered drug discovery. *Nature Machine Intelligence*, 4 (3), 189–191.
- Vale, D., El-Sharif, A. & Ali, M. (2022). Explainable artificial intelligence (XAI) post-hoc explainability methods: Risks and limitations in non-discrimination law. *AI and Ethics*, 1, 1-12.
- Weidinger, L., Mellor, J., Rauh, M., Griffin, C., Uesato, J., Huang, P.-S., Cheng, M., Glaese, M., Balle, B., Kasirzadeh, A., Kenton, Z., Brown, S., Hawkins, W., Stepleton, T., Biles, C., Birhane, A., Haas, J., Rimell, L., Hendricks, L. A., ... & Gabriel, I. (2021). *Ethical and social risks of harm from Language Models*. DeepMind.
- Yang, G., Hu, E. J., Babuschkin, I., Sidor, S., Liu, X., Farhi, D., Ryder, N., Pachocki, J., Chen, W. & Gao, J. (2022). *Tensor programs v: Tuning large neural networks via zero-shot hyperparameter transfer*. 35th Conference on Neural Information Processing Systems (NeurIPS 2021).
- Zador, A. M. (2019). A critique of pure learning and what artificial neural networks can learn from animal brains. *Nature Communications*, 10 (1), 1-7.
- Zednik, C. (2021). Solving the black box problem: A normative framework for explainable artificial intelligence. *Philosophy & Technology*, 34 (2), 265-288.
- Zhang, Z., Han, X., Zhou, H., Ke, P., Gu, Y., Ye, D., Qin, Y., Su, Y., Ji, H., Guan, J., Qi, F., Wang, X., Zheng, Y., Zeng, G., Cao, H., Chen, S., Li, D., Sun, Z., Liu, Z., ... & Sun, M. (2021). CPM: A large-scale generative Chinese Pre-trained language model. *AI Open*, 2, 93-99.