Synthetic Minds: Constructing Artificial General Intelligence through Layered Learning

Mrs.P.Sathya[1], Assistant Professor
Department of Computer Science and Engineering, Sri Shanmugha College of Engineering and Technology, Pullipalayam, Morur (Post), Sankari (Tk), Salem, Tamil Nadu, India
sathya.p@shanmugha.edu.in
Ms.R.Jaseenash[2], Assistant Professor
Department of Computer Science and Engineering, Sri Shanmugha College of Engineering and Technology, Pullipalayam, Morur (Post), Sankari (Tk), Salem, Tamil Nadu, India
jaseenash.r@shanmugha.edu.in
Mr.Saravanakumar Pichumani [3], Assistant Professor
Department of Computer Science and Engineering, Sri Shanmugha College of Engineering and Technology, Pullipalayam, Morur (Post), Sankari (Tk), Salem, Tamil Nadu, India
saravanakumar@shanmugha.edu.in
Mr.E.Sivarajan[4], Assistant Professor
Department of Computer Science and Engineering, Sri Shanmugha College of Engineering and Technology, Pullipalayam, Morur (Post), Sankari (Tk), Salem, Tamil Nadu, India
sivarajan.v@shanmugha.edu.in

Abstract- This report describes the current state of the art in artificial intelligence (AI) and its potential impact for learning, teaching, and education. It provides conceptual foundations for well-informed policy-oriented work, research, and forward-looking activities that address the opportunities and challenges created by recent developments in AI. The report is aimed for policy developers, but it also makes contributions that are of interest for AI technology developers and researchers studying the impact of AI on economy, society, and the future of education and learning.

Keywords – artificial intelligence, potential impact for learning, teaching, and education, conceptual foundations.

INTRODUCTION

All human actions are based on anticipated futures. We cannot know the future because it does not exist yet, but we can use our current knowledge to imagine futures and make them happen. The better we understand the present and the history that has created it, the better we can understand the possibilities of the future.[1] To appreciate the opportunities and challenges that artificial intelligence (AI) creates, we need both good understanding of what AI is today and what the future may bring when AI is widely used in the society. AI can enable new ways of learning, teaching and education, and it may also change the society in ways that pose new challenges for educational institutions. It may amplify skill differences and polarize jobs, or it may equalize opportunities for learning[2]. The use of AI in education may generate insights on how learning happens, and it can change the way learning is assessed. It may re-organize classrooms or make them obsolete, it can increase the efficiency of teaching, or it may force students to adapt to the requirements of technology, depriving humans from the powers of agency and possibilities for responsible action[3]. All this is possible. Now is a good time to start thinking about what AI could mean for learning, teaching, and education. There is a lot of hype, and the topic is not an easy one. It is, however, both important, interesting, and worth the effort. Since 2013, when Frey and Osborne estimated that almost half of U.S. jobs were at a high risk of becoming automated, AI has been on top of policymakers' agendas. Many studies have replicated and refined this study, and the general consensus now is that AI will generate major transformations in the labour market [4]. Many skills that were important in the past are becoming automated, and many jobs and occupations will become obsolete or transformed when AI will be increasingly used. At the same time, there has been a tremendous demand for people with skills in AI development, leading to seven figure salaries and sign-up fees. China has announced that it aims to become the world leader in AI and grow a 150 billion AI ecosystem by 2030. The U.S. Department of Defense invested about 2.5 billion USD in AI in 2017, and the total private investment in the U.S [5] is now probably over 20 billion USD per year. In 2017, there were about 1200 AI start-ups in Europe, and the European Commission aims to increase the total public and private investment in AI in the EU to be at least 20 billion euros by the end of 2020. In limited tasks, AI already exceeds human capabilities. Last year, with just about one month of system development, researchers at Stanford were able to use AI to diagnose 14 types of medical conditions using frontal-view X-ray images, exceeding the human diagnostic accuracy for pneumonia [6]. In 2017, given no domain knowledge except the game rules, an artificial neural network system, AlphaZero, achieved within 24 hours a superhuman level of play in the games of chess, shogi, and Go. In May 2018, Google CEO Sundar Pichai caused a firestorm when he demonstrated in his keynote an AI system, Duplex, that can autonomously schedule appointments on the phone, fooling people to think they are discussing with another human. In the midst of self-driving cars, speaking robots, and the flood of AI miracles, it may be easy to think that AI is rapidly becoming superintelligent, and gain all the good and evil powers awarded to it in popular culture [7]. This, of course, is not the case. The current AI systems are severely limited, and there are technical, social, scientific, and conceptual limits to what they can do [8].

A THREE-LEVEL MODEL OF ACTION FOR ANALYZING AI

Cultural-historical theory of activity distinguishes three hierarchically linked levels of human behaviour. First, behaviour can be analysed as socially meaningful activity directed by culturally and socially constructed motives. Activity is realized through goal-oriented acts that essentially are ways of solving problems at hand that need to be solved to accomplish the activity. Operations, in turn, implement the acts in the present situation and concrete context, using the tools available. An important aspect of this three-level hierarchy is that the levels cannot be reduced to each other [9]. We can explain the meaning of an activity only using social, cultural and historical terms that do not make sense at the level of acts or operations. For example, we can explain the object and motive of activity by saying that we are teaching children so that they become citizens, realize their potential as human beings, and get good jobs. Activity how it is translated into concrete acts depends on social institutions, norms, social division of labour and knowing, the ways in which social production is organized, and many other similar things [10]. Most importantly, we rarely are explicitly aware of all those social factors that shape our activities. Cultural norms, values, expectations, social institutions, and other essentially contextual factors shape our activities and provide a tacit normative, emotional, and anticipatory background that allows the ongoing stream of activity to go on. This is also the level that provides the foundation for ethics of action[11]. The relation between acts and activity is, thus, similar to the relation between words and utterances. We need words to express utterances, and acts to express activity. It is, however, impossible to understand the meaning of an utterance by adding up definitions of words. On the contrary, the sense of the word depends on its role in the context of an utterance [19-21]. A written sentence needs words, and words need letters, but the meaning of a sentence cannot be found by studying letters or words. This, in effect, says that it is not possible to build models of human activity from bottom up, simply combining some elementary behavioural components[12]. Activity, properly understood, requires social and inter-generational learning, and the level of human activity cannot be accessed simply by empirical observation of human behaviour. The level of acts, in contrast, consists of externally and internally observable behaviour. Whereas the level of activity answers a socially, culturally, and historically meaningful question "why", the level of acts answers the question "what". This is also the level where we think with concepts, plan, and solve problems. If we call the level of activity a “cultural” level, the level of acts could perhaps be called “cognitive.” A description of teaching at this level could be, for example, that “I am authoring course material for the class.” The third level of operations addresses the question "how." It implements acts in concrete settings [13]. For example, there are many ways to assess student skills, many kinds of homework, and many ways to deliver homework to students. This is the level where technology operates as a tool, and where behaviour can be best understood as routine and habit. A description of teaching activity at this level could be, for example, that “I'm inserting a picture on a slide.”

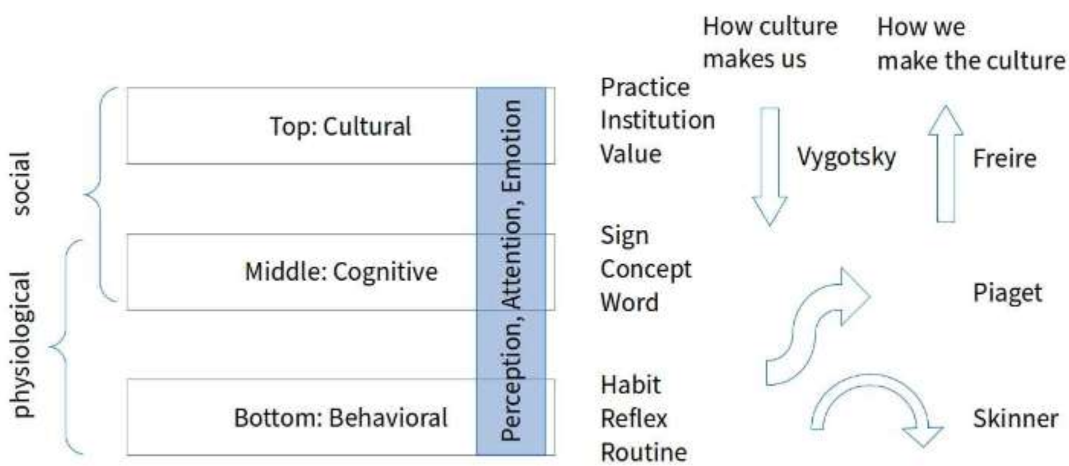

Figure 1- Three levels of human and machine learning[6]

MODELS OF LEARNING IN DATA-BASED AI

Almost all current neural AI systems rely on what is called a supervised model of learning. Such “supervised learning” is based on training data that has been labelled, usually by humans, so that the network weights can be adjusted when the labels for training data are wrongly predicted. After a sufficient number of examples are provided, the error can in most cases be reduced to a level where the predictions of the network become useful for practical purposes. For example, if an image detection program tries to differentiate between cats and dogs, during the training process someone needs to tell the system whether a picture contains a cat or a dog. A practically important variant of supervised learning is called "transfer learning[14]." A complex neural network can be trained with large amounts of data, so that it learns to discern important features of the data. The trained network can then be re-used for different pattern recognition tasks, when the underpinning features are similar enough. For example, a network can be trained to label human faces with millions of images. When the network has learned to recognize the faces that have been used for its training, its deep layers become optimized for face recognition. The top levels of the network can then relatively easily be trained to detect new faces that the system has not seen before [15]. This drastically reduces the computational and data requirements. In effect, AI developers can buy pre-trained networks from specialized vendors, or even get many state-of-the-art pre-trained networks for free and adapt them to the problem at hand. For example, the GloVe vectors, available from Stanford University, are commonly used as a starting point for natural language processing, and Google's pre-trained Inception image processing networks are often used for object recognition and similar image processing tasks [12]. Supervised learning systems can produce statistical guesses of which of possible pre- given class a specific given input data pattern belongs. Supervised learning, thus, assumes that we already know what categories input patterns can represent. This is the most frequently used learning model in AI today because for practical purposes it is often enough to classify patterns into a set of pre-defined classes [13]. For example, a self-driving car needs to know whether an object is a cyclist, truck, a train, or a child. Technically, supervised learning creates machines that map input patterns into a collection of output classes. Their intelligence, thus, is similar to simplest living beings that can associate environmental conditions with learned behaviours. In psychology, these learning models underpin the Pavlovian theory of reflexes and, for example, Skinnerian reinforcement learning. As Vygotsky pointed out in the 1920s, this type of learning represents the developmentally simplest model of learning, and both pigeons and humans are well capable of it [11].

RECENT AND FUTURE DEVELOPMENTS IN AI

The recent interest in AI results from three parallel developments. First, increasingly realistic graphics card manufacturer Nvidia published the CUDA programming interface to its graphics accelerator cards in 2007, fast parallel programming became possible at low cost. This allowed researchers to build neural network models that had many connected layers of artificial neurons and large numbers of parameters that the network could learn. Second, huge amounts of data have become available as computers and computer users have been networked. The digitalization of images, videos, voice and text has created an environment where machine learning can thrive. This has allowed AI researchers to revisit old artificial neural network models, training them with very large datasets. Somewhat surprisingly [12], these huge data sources have proven to be enough for some of the hard problems of AI, including object recognition from digital images and machine translation. Whereas it was earlier believed that computers need to understand language and its structures before they can translate text and speech from one language to another, for many practical uses it is enough to process millions of sentences to find out the contexts where words appear. By mapping words into high-dimensional representational spaces, enough of this contextual information is retained so that translation can be done without linguistic knowledge [13]. A common approach is to use the publicly available GloVe word representations that have been developed using text corpora that contains up to 840 billion word-like tokens found on documents and content on the Internet, subsequently translated to a vocabulary of over 2 million words.32 Using this dataset and machine learning algorithms, the words have been mapped into points in a 300-dimensional vector space [.33 The location and geometric relations between words in this space capture many elements of word use, and can be also used as a basis for translation from one language to another. Although such a purely statistical and data- based approach is not able to comprehend new or creative uses of language, it works surprisingly well in practice [14-17].

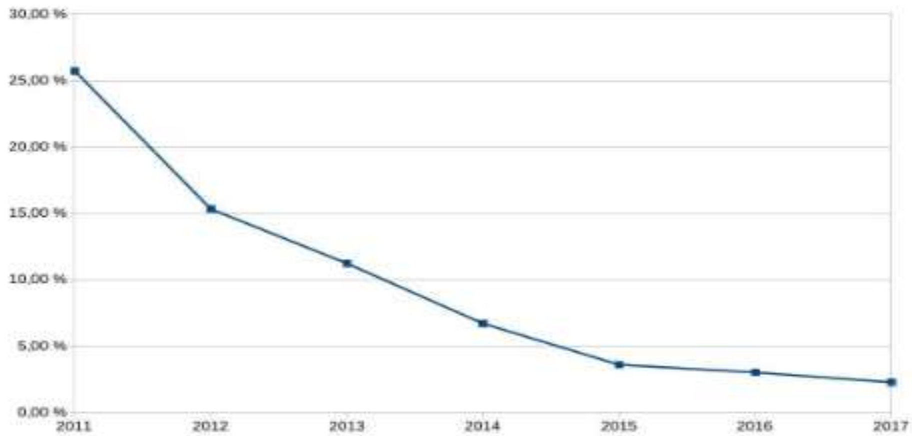

Figure 2- Error rates in the ImageNet ILSRC object recognition competition

Third, specialized open source machine learning programming environments have become available that make the creation and testing of neural networks easy. In most current neural AI models, learning occurs by the gradual adjustment of network weights, based on whether the network makes right predictions with the training data. A central task in such learning is to propagate information about how important each neuron's activity is to right and wrong predictions made by the network. When an active neuron is associated with a wrong prediction, the activity of the neuron is decreased by decreasing the weights of its incoming connections. As there can be very many layers of neurons and many connections between neurons, this is a task that is difficult even for powerful traditional computers [15-18]. The influence of each neuron to the prediction can, however, be computed using the chain rule of calculus, propagating the information from the output layer of the network layer-by-layer towards the input layer. This is known as "backpropagation" of error.34 Although the computation of network weights using this method may involve hundreds of millions of computations in state-of-the-art networks, current neural AI development environments can do this with a couple of lines of program code.

CONCLUSION

The current excitement about AI easily leads to technology push, where AI is viewed as a solution to a wide variety of problems in education and learning. It is probably fair to say that the potential and challenges of AI in education are still not adequately understood. AI can be understood as a general-purpose technology, and it can be applied in many different ways. Although the characteristics of technology itself may push development towards specific directions, it is always possible to use technology in many ways and for many different purposes, also in education. For policy development, it is therefore probably more important to understand why and for what we use technology than how it is used. The future promises of technology, in this view, have to be justified by making explicit the motivation of using the technology, as well as the key assumptions that underpin the stated motivation. This lifts technology to a level of policy, and we have to ask what are the objectives and goals of using it. Only if we have such a birds-eye view on technical development, we can say where we want to go and how technology can help us on the way. When the assumptions and motivations are made explicit, they can also be critically assessed.

REFERENCES

[1]. Acemoglu, Daron, and Pascual Restrepo. 2023. "The Race Between Machine and Man: Implications of Technology for Growth, Factor Shares and Employment." Working Paper 22252. National Bureau of Economic Research. https://doi.org/10.3386/w22252.
[2]. Agrawal, Ajay, Joshua Gans, and Avi Goldfarb, eds. 2023. "Introduction to : 'Economics of Artificial Intelligence.'" In Economics of Artificial Intelligence. Toronto: nber.org. http://www.nber.org/chapters/c14005.pdf.
[3]. Anderson, James, A., and Edward Rosenfeld, eds. 2024. Neurocomputing: Foundations for Research. Cambridge, MA: The MIT Press.
[4]. Autor, David H. 2023. "Why Are There Still So Many Jobs? The History and Future of Workplace Automation." Journal of Economic Perspectives 29 (3): 3–30. https://doi.org/10.1257/jep.29.3.3.
[5]. Autor, David H., Frank Levy, and Richard J. Murnane. 2023. "The Skill Content of Recent Technological Change: An Empirical Exploration." The Quarterly Journal of Economics 118 (4): 1279–1333. https://doi.org/10.1162/003355303322552801.
[6]. Bass, A.S. 2024. "Non-Tech Businesses Are Beginning to Use Artificial Intelligence." Financial Times, March 31, 2018.
[7]. Boccanfuso, Laura, Erin Barney, Claire Foster, Yeojin Amy Ahn, Katarzyna Chawarska, Brian Scassellati, and Frederick Shic. 2024. "Emotional Robot to Examine Differences in Play Patterns and Affective Response of Children with and Without ASD." In The Eleventh ACM/IEEE International Conference on Human Robot Interaction, 19–26. HRI '16. Piscataway, NJ, USA: IEEE Press. http://dl.acm.org/citation.cfm?id=2906831.2906837.
[8]. Boden, Margaret A. 2023. AI: Its Nature and Future. Oxford, New York: Oxford University Press.
[9]. Sanjay Kumar Suman, Dhananjay Kumar and L. Bhagyalakshmi, "Non Cooperative Power Control Game with New Pricing for Wireless Ad hoc Networks", International Review on Computers and Software, vol. 9, no. 1, pp. 18-28, 2014. ISSN: 1828-6003,
[10]. S. Porselvi, Sanjay Kumar Suman and L. Bhagyalakshmi, "Harvesting RF energy for mobile charging", Australian Journal of Basic and Applied Science, vol. 9, no. 20, pp. 454-465, June 2015.
[11]. K. Swapna, P. Rajalakshmi and Sanjay Kumar Suman, "Security Enhancement in MANET using Game Theory", Middle East Journal of Scientific Research, vol. 23, pp. 190-195, 2015.
[12]. VinaySrivatsan, Sanjay Kumar Suman, L. Bhagyalakshmi and S. Porselvi, "Non radiative wireless power transfer", Journal of Advances in Natural and Applied Sciences, vol. 10, no. 16, pp. 147-153, Nov. 2016.
[13]. Sujeetha Devi, Bhagyalakshmi L and Sanjay Kumar Suman, "Cluster based energy efficient joint routing algorithm for delay minimization in wireless sensor networks", International Journal of Pure and Applied Mathematics, vol. 119, no. 15, 307-313, 2018
[14]. Bostrom, Nick. 2024. Superintelligence: Paths, Dangers, Strategies. 1 edition. Oxford: Oxford University Press.
[15]. Brown, J.S., A. Collins, and P. Duguid. 2023. "Situated Cognition and the Culture of Learning." Educational Researcher 18 (1): 32–42.
[16]. Bruner, J. 2023. Actual Minds, Possible Worlds. Cambridge, MA: Harvard University Press.
[17]. Brynjolfsson, Erik, and Andrew McAfee. 2023. Race Against the Machine: How the Digital
[18]. Revolution Is Accelerating Innovation, Driving Productivity, and Irreversibly Transforming Employment and the Economy. Brynjolfsson and McAfee.
[19]. Cole, M. 2024. Culture in Mind. Cambridge, MA: Harvard University Press.

[20]. Cole, M., and J.V. Wertsch. 2024. “Beyond the Individual-Social Antinomy in Discussions of Piaget and Vygotsky.” Human Development 39 (5): 250–56.
[21]. Demiaux, Victor, and Yacine Si Abdallah. “How Can Humans Keep the Upper Hand? The Ethical Matters Raised by Algorithms and Artificial Intelligence.” Report on the public debate led by the French Data Protection Authority (CNIL) as part of the ethical discussion assignment set by the Digital Republic Bill. Paris: CNIL, December 2024.