International Journal of Aquatic Science
ISSN: 2008-8019
Vol 11, Issue 01, 2020

A Three-Level Model of Action for Analyzing AI and Models of Learning in Data-Based AI

Kola Vasista[1], S. Narasimha Reddy[2]

[1]*Director, Dark Horse Analytics Pvt Ltd, Hyderabad, India*
[2]*Associate professor & Head, Department of Computer Science & Engineering, KITS, India*

Abstract: 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. The classical planning representation talks about what to do, and in what order, but the representation cannot talk about time: how long an action takes and when it occurs. For example, the planners, could produce a schedule for an airline that says which planes are assigned to which flights, but we really need to know departure and arrival times as well. This is the subject matter of scheduling. This paper provides a three-level model of action for analyzing AI and Models of learning in data-based AI.

Index Terms: Machine Learning, AI, models, data-based AI

1. 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. 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. 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. 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. 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. 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. 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 super intelligent, and gain all the good and evil powers awarded to it in popular culture. 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. As one recent author noted, AI may be riding a one-trick pony as almost all AI advances reported in the media are based on ideas that are more than three decades old. A particular challenge of the currently dominant learning models used in AI is that they can only see the world as a repetition of the past. The available categories and success criteria that are used for their training are supplied by humans. Personal and cultural biases, thus, are an inherent element in AI systems. A three-level model of human action presented in the next section suggests that norms and values are often tacit and expressed through unarticulated emotional reactions. Perhaps surprisingly, the recent successes in AI also represent the oldest approach to AI and one where almost all the intelligence comes from humans.

Instead of a beginning of an AI revolution, we could be at the end of one. This, of course, depends on what we mean by revolution. Electricity did not revolutionize the world when Volta found a way to store it in 1800 or when Edison General Electric Company was incorporated in 1889. The transformative impact of general purpose technologies becomes visible only gradually, when societies and economies reinvent themselves as users of new technologies. Technological change requires cultural change that is reflected in lifestyles, norms, policies, social institutions, skills, and education. Because of this, AI—now often called the "new electricity"—may revolutionize many areas of life when it is taken into use even if it keeps on driving its "one-trick" pony for the foreseeable future. Many interesting things will happen when already existing technologies will be adopted, adapted, and applied for learning, teaching, and education. For example, AI may enable both new learning and teaching practices, and it may generate a new social, cultural, and economic context for education.

Below we ask simple questions that illustrate the relevance of AI for educational policies and practices. Which vocations and occupations will become obsolete in the near future? What are the 21st Century skills in a world where AI is widely used? How should AI be incorporated in the K-12 curriculum? How will AI change teaching? Should real-time monitoring of student emotions be allowed in classrooms? Can AI fairly assess students? Do we need fewer classrooms because of AI? Does AI reduce the impact of dyslexia, dyscalculia, or other learning difficulties? These questions are simple to ask, and relevant for understanding the future of learning, teaching, and education. The answers, of course, are more complex.

The main aim of this report is to put these and other similar questions in a context where they can be meaningfully addressed. We do not aim to provide final answers; instead, we hope to provide background that will facilitate discussion on these and other important questions that need to be asked as AI becomes increasingly visible in the society and economy around us. To do this, we have to first open the "black box" of AI and peek inside. There are several things AI can do well, and many things it cannot do. At present there is an avalanche of reports and newspaper articles on AI, and it is not always easy to distinguish important messages from noise. It is, however, important to understand some key characteristics of current AI to be able to imagine realistic futures. In the next sections, we put AI in the context of learning, teaching, and education, and then focus on the specific form of AI, adaptive artificial neural networks, that have generated the recent interest in AI.

A three-level model of action for analyzing AI and its impact

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. 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. The "content" of this 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. 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.

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. 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. 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. 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.”

Psychologists and learning theorists have focused on different levels of this three-level hierarchy during the last century. Behaviouristic and associationist theories of learning have addressed mainly the level of operations. Cognitivist and constructivist theorists have mainly addressed the cognitive level, with constructionists also emphasizing the material, affective, and social context. Socio-cultural theorists, in turn, have often focused on the social, cultural and materially embedded dimensions of knowing and learning. Human learning occurs on all three levels of the activity hierarchy. When habit and routine hits an obstacle, we become aware of it, operation ceases, and action replaces it. We start to interpret the problem, and try to find a solution. At this level, learning consists of problem solving, creative reframing, and formation of new anticipatory models. New ways of doing and thinking emerge, can be internalized, and can become the basis for new habits and routines. Lev Vygotsky, the founder of cultural-historical theory, however, also pointed to the importance of the social and cultural level of activities that shape human thinking and learning. Advanced forms of thought are made possible because they rely on culturally and historically developed stocks of knowing. Cognitive level acts, thus, use resources from both the top level of activity and the bottom level of operations. Whereas Vygotsky emphasized the influence of social and cultural factors in cognitive development, critical pedagogists such as Paulo Freire and newer activity theorists such as Yrjö Engeström have emphasized the role of learning in changing existing social practices. Engeström, in particular, has highlighted the role of learning in the creation of new educational practices.

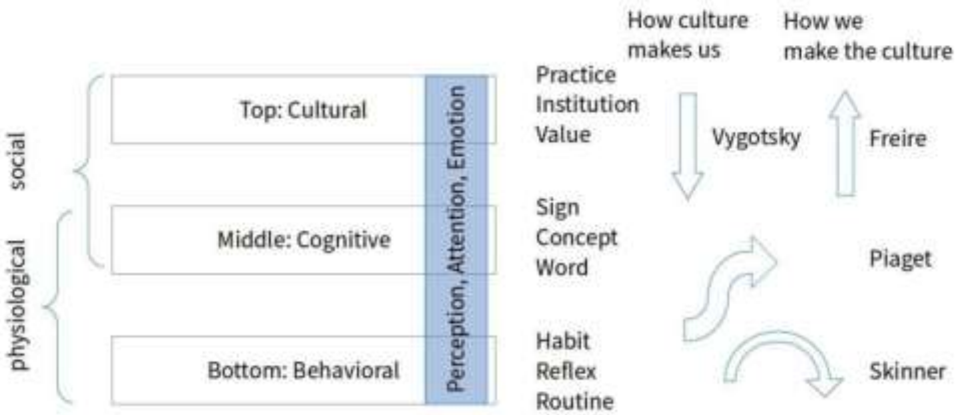

Figure 1. Three levels of human and machine learning

In this conceptual frame, learning at the level of activity can be understood as innovation and realization of imagined futures.[2] Possibilities that have been figured out at the level of cognition can start to change social practices and systems of activities, eventually leading to new motives and reasons that start to organize the society. Much of this activity-level development, however, is also emergent and unintended. Social structures, practices and institutions get their shape as a result of complex ongoing social interaction and highly diversified interests and interpretations, and to a large extent remain unobservable for the members of society.

This three-level model provides a useful entry point for understanding artificial intelligence and its potential impact on human activities. When AI enters social practices at the level of operations, it augments and complements them, increasing the efficiency and effectiveness of current ways of doing things. When it enters at the level of acts, it replaces, substitutes, and automates acts that were previously done by humans. When it enters social practice at the level of activity, it transforms the system of motives, making current activities and specializations redundant and obsolete. For example, technical and routine skills emphasize the level of operations. Vocational education has traditionally focused on this level, teaching students how to use tools and domain-specific knowledge. The recent calls for competence-based education, in turn, emphasize problem solving, critical thinking, decision-making and analytical skills, focusing on the cognitive level. Entrepreneurial and innovation competences, highlighted in frameworks for key competences and 21st century skills, mainly address the opportunities for social and cultural change at the level of activities.

Consequently, learning at the level of operations requires data on the current concrete environment. This data can be generated using perception and physical interaction. Learning at the level of socially motivated activity, in contrast, requires knowledge about social systems of meaning. To gain such knowledge, communication, language, and dialogue become necessary. An important indicator of the current change in the dynamics of development is that whereas technology in the industrial age focused on tools for automating and supporting operations, the focus is now increasingly on technologies for social change. The three levels of activity have complex dependencies. In the course of historical development, what originally was a means may become an end in itself. “Zooming in” to modern social life, therefore, we may see a rather fractal structure or activities and acts. Using this three-level model of activity, it becomes, however, clear that different types of artificial intelligence and machine learning systems operate on different layers of this hierarchy. Most importantly, the level of meaningful activity, which according to socio-cultural theories of learning underpins advanced forms of human intelligence and learning, remains beyond the current state of the AI art. This paradigm is currently being explored in the field of Child-Robot Interaction and social robotics. In the next section, we briefly outline the main characteristics of three different types of AI to locate their capabilities in this hierarchy, and discuss their potential impact.

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." 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. 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.

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. 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.

A particular challenge of supervised learning models is that they can only see the world as a repetition of the past. The available categories and success criteria that are used for their training are supplied by humans. Personal and cultural biases, thus, are an inherent element in AI systems that use supervised learning. The three-level model presented above suggests that norms and values are often tacit and expressed through unarticulated emotional reactions. It is, therefore, to be expected that supervised learning models materialise and hardwire cultural beliefs that often remain otherwise unexplored. In somewhat provocative terms, supervised learning creates machines that are only able to perceive worlds where humans are put in pre-defined boxes. From ethical and pedagogic points of view this is problematic as it implies that in interactions with such machines, humans are deprived of agency powers that allow them to become something new and take responsibility of their choices.

Many unsupervised or partially supervised neural learning models have been developed since the 1960s, some of which are also currently being developed and applied. Increasing computational power has also allowed researchers to use simple pattern- matching networks as components in higher-level architectures. For example, Google's AlphaZero game AI uses “reinforcement learning” where the system generates game simulations and adjusts network weights based on success in these games. Inspired by Skinnerian models of operant conditioning, reinforcement learning amplifies behaviour that leads to outcomes that are defined as positive. A variant of reinforcement learning is known as generative adversarial networks, or GANs, where one network tries to fool another to believe that the data it generates actually comes from the training data set. This approach has been used, for example, to create synthetic images of artworks and human faces that an image recognition system cannot distinguish from real images. It is also commercially used for product design, for example in the fashion industry. A variation of GAN is called "Turing learning," where the system that learns is allowed to actively interact with the world in trying to guess whether the data comes from the real environment or from a machine.

Towards the future

As some economists, philosophers, and scientists have made high-profile statements about the forthcoming emergence of super-intelligent AI systems that eventually may replace humans in many areas of human life, it is perhaps useful to note that most current AI learning models represent cognitive capabilities that most closely resemble biological instincts. Many predictions about the future of AI have been based on extrapolations of historical technical development, and in particular estimates of the continuation of "Moore's Law" in computing, with little

concern about differences between advanced forms of human learning and the more elementary capabilities of association. Human learning requires many meta-level competences. In particular, for humans it is important to know what counts as knowledge, how to go on in acquiring, creating, and learning knowledge, how to regulate cognition, attention and emotion in learning processes, and what the social and practical motivation for learning is. As Luckin has recently well pointed out, at present AI lacks most of these meta-cognitive and regulatory capabilities.
It is important to note that the future of the current AI boom will to an important extent be determined by developments in chip design. For almost fifty years, developments in processor and memory chips were driven by rapid continuous improvements in miniaturization of component features on semiconductor chips. During the last ten years it has become increasingly accepted that this development is about to end, and new approaches are needed to keep the semiconductor industry growing. Neural AI addresses this "post-Moore" era by shifting development towards new computing models, including analog computing. This represents a major discontinuity in the technological foundations of knowledge society.

2. CONCLUSION

In practice, most AI experts work with "narrow AI," in contrast with "general AI" that would have capabilities similar to humans. In setting up the first Dartmouth summer project on artificial intelligence, the leading researchers believed that computers will soon be intelligent. Such expectations seem to be unrealistic also today. Although it might be possible to develop AI systems that have capabilities that more closely resemble human intelligence, current AI systems use rather simplified models of learning and biological intelligence. Most current AI systems rely on essentially reflexological and behaviouristic models of learning, popularized by Pavlov and Thorndike at the beginning of the 20th century. They could perhaps therefore better be described as mechanical instincts, instead of artificial intelligence. Despite these limitations, the potential of AI in education has been widely recognized during the last three decades. Although the impact on classrooms has been relatively minor, the recent developments suggest that the situation may change. In particular, AI-based systems can become widely used as systems that support teachers and learners. AI can also rapidly change the economy and job market, creating new requirements for education and educational systems. This paper provided a three-level model of action for analyzing AI and Models of learning in data-based AI.

3. REFERENCES

[1] KOLTER, J. Z., AND JOHNSON, M. J. REDD: A public data set for energy disaggregation research. In Workshop on Data Mining Applications in Sustainability (SIGKDD), San Diego, CA (2011), vol. 25, Citeseer, pp. 59–62.
[2] KUBIATOWICZ, J., BINDEL, D., CHEN, Y., CZERWINSKI, S., EATON, P., GEELS, D., GUMMADI, R., RHEA, S., WEATHER- SPOON, H., WEIMER, W., ET AL. Oceanstore: An architecture for global-scale persistent storage. ACM Sigplan Notices 35, 11 (2000), 190–201.

[3] KUNIAVSKY, M. Smart Things: Ubiquitous Computing User Ex- perience Design: Ubiquitous Computing User Experience Design. Elsevier, 2010.
[4] LEE, E. A., RABAEY, J., BLAAUW, D., DUTTA, P., FU, K., GUESTRIN, C., HARTMANN, B., JAFARI, R., JONES, D., KUBI- ATOWICZ, J., ET AL. The Swarm at the Edge of the Cloud.
[5] MIDDLETON, P., KJELDSEN, P., AND TULLY, J. Forecast: The internet of things, worldwide, 2013. Gartner Research (2013).