Comparing Human Learning and Machine Learning

Proseminar in Cognitive Science: 185:500:01

12/8/20

By: Priya Jain

Introduction

In the early days of artificial intelligence - dating back to the 1950's - the architecture was heavily influenced by human cognition and intelligence (Langley, 2006). The goal was to model human thinking and intelligence by understanding the nature of the human mind. Although this is still the goal, computer scientists have distanced themselves from studying and implementing algorithms inspired from actual cognitive processes. Machine learning, an application of artificial intelligence where the machine automatically learns without being explicitly programmed, used to include human-like learning, but in the 1990's it was reduced to statistical learning methods and classification techniques (Langley, 2006). Although these methods do have some resemblance to human learning, engineers and scientists started to diverge from cognitively inspired architecture to efficient, statistical algorithmic solutions. In recent years, the focus has started to re-adjust back to infusing cognitive science perspectives into creating artificial intelligent technologies. This paper will discuss the similarities and differences between human learning and machine learning, particularly deep learning which is a subfield of machine learning.

Human Learning and Deep Learning

Before diving into comparing and contrasting human learning and deep learning, this portion of the paper will briefly establish some ground on how humans and machines learn. This topic in itself is very dense and heavily researched, so this section will provide some insight into ongoing perspectives. To begin, in humans, learning is characterized by understanding the world, explaining phenomena, imagining alternate events, and acting on decisions to achieve desired outcomes or goals. In humans, learning is not just using data to predict outcomes but being able to explain things (Lake et al., 2017). Learning in infants and adults can be explained as model-building. Humans develop models and are constantly updating, generalizing, and creating new models based on experiences and innate inclinations. Humans learn through creating mental representations, using perceptual cues and having mechanisms to "learn to learn". The "learning to learn" ideal is seen in infants who pick up language, numeracy, basics of physics, and more, rather early in their life.

On the other hand, machines mostly learn through prediction through statistical methods rather than interpretation (Lake et al. 2017). Deep learning algorithms are made up of artificial neural networks with hidden layers that transform data into output. Artificial neural networks (ANNs) are inspired by human neural networks that pass information along each other. Artificial neural networks or just neural networks (NNs) basically takes in an input and creates an output based on an activation function. Between the input and output units are one or more hidden layers of units that are connected to the input and output units. Each of these connections are represented by weights which corresponds to brain triggers at synapses between real neurons (Schmidhuber, 2014). Deep neural networks refer to a neural network with multiple hidden layers, making the neural network complex due to its depth or deep nature. Deep learning networks "learn" by changing their weights to achieve a desired output value, which corresponds to how humans adapt to their environment based on prior information (Schmidhuber, 2014). During training, the difference between the output and the desired output is used to modify the weights of the connections between neurons. Backpropagation is characterized by working backwards from the output units to the input units to improve the objective function. This is how deep learning networks "learn" (Lake et al. 2017).

The behavior of deep neural networks (DNNs) being able to process and learn information without being explicitly programmed is similar to how the brain functions in some ways because neurons in the brain transmit electrical signals although it is not explicitly known what those signals encode.

Model Building and Pattern Recognition

One big difference between humans and machines is that humans learn through model building and interpretation, whereas machines learn through pattern recognition and prediction. One example that illustrates learning through model building in humans is learning language. Although language acquisition does trace back to having innate tendencies, there is evidence for learning through developing rules, theories, and models about the world (Stromswold, 2000). Infants are predisposed with a Universal Grammar and innate mechanisms to learn, however they also heavily rely on correlations and new knowledge to update what they already know about language. For example, infants are not born knowing the semantic properties of part-of-speech categories like nouns and verbs. It is hypothesized that they know the universal grammatical rules of language but do not come equipped with the specific rules of their native language. Instead, they use inference models using a combination of techniques to learn syntactic properties of their language. For instance, prosodic cues such as changes in frequency can inform the beginning and end of clauses or phrases. Children make correlations between auditory linguistic inputs in the form of a correlation matrix - in addition to innate mechanisms - to learn all possible semantic and syntactic sentences. Semantic bootstrapping is another theory about how kids use conceptual knowledge to make categories, in other words developing a sort of model to reference (Stromswold, 2000). Numeracy is another important type of cognitive learning. Although children have an innate number sense, higher level, complex mathematical concepts are learned. The hypothesis goes that students, or mathematical learners, are thinkers with “emerging theories” about the world and therefore this kind of knowledge is created by learning experiences and reasoning as opposed to acquiring it directly from some known amount of data (Marmasse et al., 2000).

Machines, unlike humans, are not able to build as robust models to spring interpretations about the world. In fact, most machine learning algorithms employ pattern recognition, classification, or feature extraction to reproduce human learning. In one study, a network called DQN is trained to learn to play Atari video games. One of the games learned was Frostbite; in this game, the player controls an agent that is supposed to build an igloo within a time limit. The agent must build the igloo piece by piece by accessing ice floes and there are also opportunities to gain extra points by catching fish. In order to win this game, the DQN must have a short-term extended plan to accomplish all the sub-goals (Lake et al., 2017). The DQN is made of deep convolutional neural networks for pattern recognition and a model-free reinforcement learning algorithm known as Q-learning. It maps pixel frames onto a policy over a set of actions and is trained to optimize reward (in this case, game points). What is interesting about this study is that this deep learning network had to be trained for 200 frames, which is equivalent to 924 hours of game time. Surprisingly - or perhaps unsurprisingly - it took a professional gamer only 2 hours of practice to learn the game. This means the machine had 500 times more experience than humans in order to learn the game and even then it performed sub-par to humans (Lake et al., 2017).

This contrast in performance can be explained by the fact humans and machines fundamentally learn differently. As discussed above, humans build models to interpret the world through experience. These models are generalized to apply to many situations. Some hypothesize that people infer general schemas to explain the game's goals, interactions, and objects using flexible model-based planning (Lake et al., 2017). Humans may develop theories about the world using bayesian frameworks where events are assigned probabilities based on generalizing past experiences and knowledge about the world. While human learning is limitless due to rich knowledge and mechanisms for generalization, machine learning is quite limited by the quality of data. Machines need plenty of data to make associations and learn patterns and training is very specific to a certain task, like a certain game. If the machine is tasked to play a different game, it must be retrained as it cannot generalize what it has learned in the previous game. This is cumbersome and not as robust as human learning.

Reinforcement learning, Generalizations, and Curriculum learning

Just like the DQN network used a model-free reinforcement learning algorithm to learn Atari Frostbite, it is common for machines to learn with deep reinforcement learning methods. Reinforcement learning allows a computational agent to learn what actions/decisions to take in effort to maximize expected cumulative rewards. Deep reinforcement learning is just the combination of deep learning and reinforcement learning (Schmidhuber, 2014). Reinforcement learning is similar to learning through operant conditioning in humans. Similar to reinforcement learning in machines, operant conditioning is a form of learning where the agent grasps an association through rewards and punishments.

One study by DeepMind is particularly interesting as it uses deep reinforcement learning and aims to create biologically-inspired architectures. This study creates a computational agent that can navigate real world cities without a map, likewise to how humans absorb visual input and learn to navigate the world through experience. The approach is formalized as Markov Decision Process learning problem where the reward function is contingent on the current goal and state. The objective is for the agent to discover the policy that maximizes expected return. This process spawns two types of learning: general learning and locale-specific learning (Mirowski et al., 2018). An example of general learning is scene understanding. Constrastingly, locale-specific learning is learning features or structures unique to a specific location.

This approach is rather exciting because it reproduces how humans learn about their surroundings. Humans create generalizations from past experiences so that each new environment they are exposed to is not an entirely new learning experience. The results of the study show that the agent was able to exploit some lineararities of goal representation because in a new, unseen environment the agent navigated halfway to the masked goal based on past learning. Although it needed to be retrained in order to reach goals in a new city, the fact that it showed evidence of some held out goal representation is beginning to resemble human learning. The study implements a relatively simple architecture to demonstrate that simpler algorithms can memorize large environments. These elements of generalization and simplicity are two salient features of human learning. The idea is that evolutionarily speaking, humans cannot have too complex mechanisms, and therefore are biased towards simpler models, theories, and outcomes. The tendency to generalize follows a similar logic; it is simpler to have few rules and generalize those rules to a vast number of situations rather than have a large number of situation-specific rules. These principles are followed in most aspects of human learning, such as language learning and learning through bayesian modelling. In fact, Chomsky has conducted a lot of research showing that syntactic rules of language inherit generality and simplicity properties

(Baker). Implementing these properties in machine learning models, as done in Deepmind's study, offers another similarity to human cognitive functioning.

Another aspect of deep reinforcement learning used in both navigating cities without maps and playing Atari Frostbite is guiding learning by giving rewards. The DQN algorithm used incremental rewards for completing each sub-goal. If it weren't for incremental rewards the DQN would have to learn how to build an igloo solely through trial and error. People, however, do not need to use incremental rewards to understand the game (Lake et al. 2017). Similarly, in the Deepmind study, reward shaping was tested. In effort to make the navigation task simpler, the algorithm provides early rewards along the way to encourage exploration. However, again, this is not how humans learn how to complete a task. DeepMind found that curriculum learning more closely matches the type of learning that humans do in similar tasks. Curriculum learning gradually increases the complexity of tasks by presenting increasingly difficult examples to the algorithm. This type of learning approach helps the agent learn to discover progressively farther landmarks (Mirowski et al., 2018). There is no surprise that it is an effective method of learning because humans learn in a similar fashion. Despite the difference in possible reward systems implemented in machines and humans, both exhibit goal-oriented learning. Human cognition possesses metacognitive processes like planning, goal-building, and sub-goal building. (Shuell, 1986).

Model-based vs Model-free Learning

This leads to the question of what type of models of reinforcement learning appear in humans versus in machine learning. There are two main classes in deep reinforcement learning - model-based and model-free learning. Model-based RL builds a statistical model that is then used to predict outcomes and actions. On the other hand, model-free RL learns the same predictions as the model-based RL method, but without the help of a statistical model. Instead, predictions about the world are developed based on a consistency constraint that actions that lead to better predicted outcomes are preferred (Montague et al., 2012).

Human learning is a combination of model-based and model-free methods. Model-free learning comes into play in associative learning. One example that offers evidence for this mechanism in human learning is the firing of dopaminergic neurons that update model-free outputs for reward prediction error. Nonetheless, when overly complex cognitive tasks are at hand, such as planning tasks, there is a need for building model-based cognitive maps. These models are richer and more structured than model-free learning models and therefore require more complex inference algorithms (Lake et al. 2017). Just like complex machine learning models require an abundance of data, richer cognitive models require greater complexity, which in turn produces computationally slow models. Due to this reason, in artificial intelligence and machine learning, engineers more often than not employ model-free models as they consume less memory and computational resources than planning based models. Model-free processes depend on past learning rather than present inferences which come with disadvantages like inflexibility to real time environment changes (Montague et al. 2012).

As mentioned above, model based planning is algorithmically incorporated into technology. One study, built a system to improve human-building interactions (HBI). They took a model-based planning approach to modelling human walking navigation. This model predicted where people plant their next steps by representing movement as an inverted pendulum of human bipedal walking. Through this approach, they could plan any number of upcoming steps instead of making impromptu decisions about next step movement. This featured planning and organization to simulate human navigation. Not only does this example demonstrate planning models but it also employs curriculum training. Interestingly enough this study finds that behavior cloning is actually better than reinforcement learning techniques (Kapadia et al., 2020). Behavioral cloning is a training technique where human subjects are recorded performing an action or task. The action and the situation is logged so that the program can learn rules to simulate the same output. The motivation behind this technique is to spur learning by imitation and represent human behavior. According to the study, behavioral cloning was a simpler and more effective approach than reinforcement learning.

“Learning-to-learn” and Transfer Learning

The main idea is that humans use both algorithms of learning in a more sophisticated way than most machines. Nevertheless, the fact that machines are able to replicate these cognitive reinforcement learning processes in artificially intelligent agents is remarkable and represents a gross affinity to human-like cognitive models. Reverse-engineering human intelligence is so crucial to informing human-centered machine learning algorithms. By integrating cognitive processes into artificial intelligence, technologies can become more sophisticated, efficient, and safer for practical use. AlphaGo is an example of a computer program that successfully integrated human cognitive processes into an artificial agent.

DeepMind’s AlphaGo is the first program to win the game Go against a professional player using advanced search trees and deep neural networks. AlphaGo is equipped with a model-based search algorithm that is similar to model-based cognitive maps in human learning. The software also utilizes model-free reinforcement learning for high-level pattern recognition. Akin to humans, this software uses a combination of model-free and model-based planning algorithms to attack the complex nature of the board game. It also exhibits a major ingredient in human learning - the ability of “learning-to-learn” (Lake et al., 2017). People have this mechanism where prior knowledge is acquired through the process of “learning-to-learn” (Harlow, 1949). In machine learning, this is comparable to transfer learning. Transfer learning is about learning constraints, using prior learned information, and transferable inductive skills and applying these things to other tasks. Unfortunately, this mechanism or technique is not quite as strong in machines as it is in humans. With that being said, AlphaGo did an incredible job with learning-to-learn and incorporating intuitive psychology skills that humans possess to win the game. This feat sets an example for how other machine learning technologies can infuse conventional cognitive learning processes to better their algorithms.

Conclusion

Machine learning is generally inspired by human learning. The main task at hand for scientists and engineers is to build an artificial network that is as robust, efficient, and as accurate as the human brain. This is an ambitious goal as the brain is a tremendously complex puzzle that has yet to be solved. If machines and the brain (which is nothing more than a biological machine) are studied in tandem, there is hope for very fruitful technologies and discoveries. In this paper, the similarities and differences between human learning and machine learning were discussed. Machines and humans both use models to decide actions and behaviors. Deep learning is inspired by the biological architecture of neural networks. Machines and humans even possess similar learning tactics like how transfer learning and reinforcement learning are akin to “learning-to-learn” and operant conditioning, respectively. The list of differences is longer with empirical contrasts in interpretation versus prediction, ability for humans to generalize, and the difference is computational efficiency. By closing the gap between machine and human learning, artificial intelligence can reach a whole new level. Looking at these kinds of comparisons can inspire future algorithmic solutions, creative thinking, and human-centered design. In fact, human-centered approaches have infiltrated various companies like Google, Amazon, and IBM as well as in academia in recent years, bringing artificial intelligence back to its roots of a primary application of cognitive science.

References

Lake, B. M., Ullman, T. D., Tenenbaum, J. B., & Gershman, S. J. (2017). Building machines that learn and think like people. Behavioral and brain sciences, 40.

Langley, P. (2006). Intelligent behavior in humans and machines. In *American Association for Artificial Intelligence*.

Marmasse, N., Bletsas, A., & Marti, S. (2000). Numerical mechanisms and children's concept of Numbers. *MIT media laboratory online publication. Accessed: November*, *22*, 2009.

Mirowski, P., Grimes, M., Malinowski, M., Hermann, K. M., Anderson, K., Teplyashin, D., ... & Hadsell, R. (2018). Learning to navigate in cities without a map. Advances in Neural Information Processing Systems, 31, 2419-2430.

Montague, P. R., Dolan, R. J., Friston, K. J., & Dayan, P. (2012). Computational psychiatry. Trends in cognitive sciences, 16(1), 72-80.

Schmidhuber, J. (2015). Deep learning in neural networks: An overview. *Neural networks*, *61*, 85-117.

Shuell, T. J. (1986). Cognitive conceptions of learning. *Review of educational research*, *56*(4), 411-436.

Stromswold, K. (2000). The cognitive neuroscience of language acquisition. *The new cognitive neurosciences*, *2*, 909-932.

Usman, M. (2020). Analyzing Human-Building Interactions in Virtual Environments Using Crowd Simulations.