Revista de Instituciones, Ideas y Mercados Nº 69 | Marzo 2020 | pp. 57-86

CAN ENTREPRENEURSHIP BE LEARNED BY INTELLIGENT MACHINES?

Steven E. Phelan[1]

Resumen

Este artículo explora si los empresarios humanos serán suplantados por máquinas inteligentes. Comienza por considerar la capacidad de las máquinas para participar en actividades empresariales utilizando big data y técnicas modernas de inteligencia artificial. Luego se presenta una crítica de la inteligencia artificial (IA) que establece una clara distinción entre la IA estrecha y la IA general. Actualmente, las computadoras son incapaces de realizar una IA general porque carecen de una teoría de la mente y la autoconciencia. Ambos atributos son fundamentales para un emprendimiento exitoso, por lo que es poco probable que las computadoras desplacen a los emprendedores humanos
Palabras clave: emprendimiento, inteligencia artificial, cálculo económico, aprendizaje automático.

Abstract

This paper explores whether human entrepreneurs will be supplanted by intelligent machines. It starts by considering the capacity for machines to engage in entrepreneurial activity using big data and modern artificial intelligence techniques. A critique of artificial intelligence (AI) is then presented that draws a sharp distinction between narrow AI and general AI. Computers are currently incapable of general AI because they lack a theory of mind and self-awareness. Both of these attributes are critical for successful entrepreneurship making it unlikely that computers will displace human entrepreneurs any time soon.

Keywords: entrepreneurship, artificial intelligence, machine learning, economic calculation

---

[1] Steven E. Phelan is Distinguished Professor of Entrepreneurship at Fayetteville State University, a constituent institution of the University of North Carolina. He has published over forty articles at the intersection of strategy, entrepreneurship, and economics. His first book "Startup Stories: Lessons for Everyday Entrepreneurs" was published in 2020.

Introduction

The entrepreneur plays a pivotal role in Austrian theories of the market process (Phelan, 2016). This paper explores whether the role of the entrepreneur in the market process, so critical to economic calculation, might be supplanted by intelligent machines. This is a timely question given the development of quantum computing and new techniques in artificial intelligence that have greatly expanded the capabilities of computing in recent years (Childers, 2019; Cockshott, 2017).

The traditional Austrian argument in the economic calculation debate has been the inability of central planners to gather, collate, calculate, and deploy an optimal allocation of resources in real time to meet constantly changing consumer preferences. Hayek argued that a local knowledge of people, time and space held by a multitude of economic agents motivated by profit (i.e. entrepreneurs) would easily defeat an army of central planners (Hayek, 1945).

The development of distributed computing, combined with massive processing power and high speed data connections, has greatly increased our ability to gather, collate, and analyze local data sources (Cockshott, 2017; Cockshott & Cottrell, 1997). Given information on actual productivity and a list of available techniques, Cockshott (2019) claims that the centralized optimization of production is now possible. The increasing sophistication of artificial intelligence suggests that computers might also be able to predict demand as well, or even better, than humans (Townsend & Hunt, 2019; Yeung, 2019).

Fortunately (for human entrepreneurs), these predictions about artificial intelligence (AI) displacing entrepreneurs are highly premature. As we shall see, the leap from artificial narrow intelligence (ANI), focused on just one domain, to artificial general intelligence (AGI) capable of creative discovery outside its original domain, will require a paradigmatic shift in the design of these systems (Mitchell, 2019). This paper argues that entrepreneurship, as a form of creative discovery, will require a system capable of AGI and thus is in no danger of being replaced in the near term by ANI.

Part of the reason for this skepticism is that ANIs are designed to focus on a narrow set of features in the environment. For instance, a chess playing computer will have representations of the board, the pieces, valid moves, and winning conditions. It is no coincidence that AI has seen the most progress in areas where the domain can be formalized, such as games. While chess playing computers are perfectly capable of considering combinations of valid moves within the rules, they are not capable of 'frame breaking' change that allows them to transfer their knowledge to new domains, such as checkers or Go (Toivonen & Gross, 2015).

However, it is precisely this sort of frame breaking change that is a defining characteristic of entrepreneurship. Early theorists of entrepreneurship within the Austrian tradition saw technological change arriving exogenously, requiring only alertness or willpower to exploit an opportunity (Kirzner, 1973; Schumpeter, 1934). More recent theorists, however, have emphasized that initial opportunity identification is a function of the idiosyncratic experiences of particular individuals (Shane, 2000, 2003) and evolves over time in response to feedback from customers and other stakeholders (Sarasvathy, 2008).

Thus, in order to anticipate customer needs, an AGI must possess a 'theory of mind' about the preferences of others in order to successfully predict their behavior (Hintze, 2016). Indeed, McMullen (McMullen, 2010, 2015) has placed the need for such 'empathic accuracy' at the heart of the entrepreneurial function. However, empathy requires an ability to put oneself in the shoes of others, and imagine the physical, biological, psychological, and social reactions of others to new offerings. This may be a tall order for a computer that has never experienced what it is like to be human.

Entrepreneurs also need knowledge of how to profitably satisfy these customer preferences given the resources at hand (Foss & Klein, 2012). Again, this is not a trivial requirement. Koppl and his colleagues have persuasively argued that it is not possible to pre-state all of the uses for a given resource (Koppl, Kauffman, Felin, & Longo, 2015) nor compute the payoff for a given application (Koppl, 2008). Current computational methods are thus thwarted without a complete list of valid moves and the payoffs from such moves.

Thus, no amount of growth in processing power, data communication, or data storage, can solve this problem. Computer scientists will need to develop a new paradigm to create systems that function like an entrepreneur. It is likely that this new paradigm will need to be phenomenologically based (Clark, 2015; Dreyfus, 1965), echoing Hayek's (1952) earlier work on the sensory order.

The Economic Calculation Debate

The genesis of the socialist or economic calculation debate can be traced to Barone's (1908) conjecture that a centralized planner could solve a Walrasian series of equations through a trial and error process (Bradley & Mosca, 2014)[2]. Neurath (1919) later argued that the managed wartime economy during the Great War was evidence that a centralized government could successfully organize production.

The opposition to this view was sparked by a 1920 paper by Ludwig von Mises entitled "Economic calculation in the socialist economy" (Mises, 2016 [1920]). In the paper, Mises argued that the efficient production and distribution of scarce resources in an economy is theoretically and practically impossible in a socialist system. This is primarily due to the lack of money prices for evaluating production possibilities and the loss of incentives to innovate and economize when the profit motive is eliminated.

It is important to note that the paper was not written in a vacuum. Rather, it was a reaction to the 'war communism' system introduced by the Bolsheviks during the Russian Civil War of 1918-21 (Lavoie, 1981). This system eliminated private property, nationalized all industries, imposed centralized management, and confiscated grain surpluses from peasant farmers. The result was an almost complete collapse of the Russian economy, with production falling to just 20% of 1913 levels along with a famine that killed upward of 5 million people (Richman, 1981). Although some saw the system as a temporary wartime measure, there is evidence that the elimination of the price

[2] Bradley and Mosca (2014) argue that Barone largely favored personal initiative over centralized control except in time of war or disaster where speed and focus were paramount.

system was the intended goal of the new regime. In any case, the policy was reversed in 1921 with the New Economic Policy, which allowed peasants to retain and trade surplus production. The economy recovered quickly and some semblance of a price system was retained even under later reforms by Stalin (Richman, 1981).

However, socialist thinkers continued to attack Mises' argument that a socialist calculation system was impossible. Market socialists such as Lange, Lerner, Taylor, and Dickinson posited schemes to elicit Pareto optimal outcomes in the absence of private ownership of the means of production (Lavoie, 1985). Invariably, these schemes started with a static input-output matrix. A central planning body was then meant to set shadow prices that could be adjusted up or down in response to shortages or surpluses in a process of trial and error (Lange, Taylor, & Lippincott, 1938).

Lavoie (1985) argues that the standard view of the economic calculation debate then has Hayek and Robbins conceding in the 1930s and 1940s that a socialist system could be efficient in theory if not in practice. The need to track and process millions of pieces of decentralized data in real time meant that the ability to conduct the trial and error process required to reach equilibrium was practically impossible (Hayek, 1945; Robbins, 1934)[3].

Big Data and Economic Calculation

Advances in big data (see Appendix) have led to calls to reopen the socialist calculation debate (Cockshott, 2017; Cottrell & Cockshott, 1993; Limas, 2018; Phillips & Rozworski, 2019; Saros, 2014). Cockshott and Cottrell (1993), echoing earlier work by Lange (Lange et al., 1938), conceptualize the production decision as an optimization problem requiring the solution of a massive Leontieff-style input-output model. Instead of money, demand would be measured by 'labor tokens' earned by workers and available to spend on commodities. On the production side, they posit an economy with around 10 million products. The resulting input-output table could be

[3] Lavoie (1985) makes a compelling argument that Mises, Hayek, and Robbins were consistent in their arguments and that no concession was made.

populated with enterprise data collected by remote sensing. They estimate that it would require 10^{21} calculations to invert such a matrix. Currently, the world's most powerful supercomputer can make 10^{18} calculations per second so the production decision could be completed in 1,000 seconds (or approximately 20 minutes). Iterative or linked-list approximations could calculate a solution even faster. Although recent estimates place the number of products sold by Amazon at 120 million (ScrapeHero, 2019), including 44 million books, even errors of several orders of magnitude in the number of products will seem like a rounding error in the era of quantum computing.

Limas (2018) supplements the approach of Cockshott and Cottrell (1993) by suggesting that web searches could be used as proxies for demand in the absence of market prices. Phillips and Rozworki (2019) argue that the extensive supply chain management undertaken by giant firms such as Amazon and Walmart is evidence that economy-wide (centralized) supply chains are now feasible.

Saros (2014), on the other hand, has advocated for a more decentralized form of economic organization, where consumers are allocated points based on several factors including a universal basic income, work differentials, length of time in a job, and prudence. Consumers then allocate their available points to an Amazon-like catalog of goods offered by decentralized worker's councils (firms), who use internal input-output tables to produce a mixture of goods and set prices (in points). A process of "groping" towards equilibrium ensures inputs are allocated efficiently between worker's councils.

In a thoughtful piece, Morozov (2019: 48) acknowledges that "[socialists] initially took their case to be about the difficulty of computing the appropriate price levels, based on the given data—and not about the challenge of gathering and updating the data, which is never automatically 'given'". Big data has the potential to solve this problem by providing real time access to local data erasing the concerns of time and place expressed by Hayek (1945). However, modern computers are capable of much more than data acquisition, transmission, and storage. They can also make data-based decisions that are, in many cases, as good as human experts. This aspect of technology has been overlooked in the calculation debate, but a consideration of the issues involved sheds a great deal of light on the importance of human entrepreneurs in the economy.

Beyond Big Data: Artificial Intelligence and Economic Calculation

Artificial intelligence is the branch of computer science that seeks to design computational agents that can perceive and act on their environment in pursuit of a goal (Poole, Mackworth, and Goebel 1998). The most basic form of intelligent behavior is a stimulus-response model, where a given input (or set of inputs) triggers a given response (Hintze 2016). Thus, when an intelligent thermostat detects that the temperature has risen above 75F, it responds by switching on the cooling function. In an economic context, a shortage of wheat at a given location might trigger an increase in the supply of wheat in the next period[4].

Agents can also specialize in sensing or acting. The internet service IFTTT (If This Then That) enables users to create recipes that tie together sensors to actuators allowing, for instance, lights connected to intelligent switches to be turned on at dusk based on information on sunset times from the National Weather Service (NWS). One could imagine weather data being used by an intelligent agent to increase the production of heating oil when a cold spell was forecast by the NWS. Such computers are clearly capable of decentralized action as well as remote sensing. In fact, in some instances, seeking local optima can result in better system wide outcomes than attempting to calculate a global optimum (Post and Johnson 1997). Intelligent economic agents are thus capable of acting *in loco consilium* (in place of the plan).

Developing, selecting, and implementing the right techniques to convert stimuli into responses has been the focus of a great deal of research in the field of AI (Poole, Mackworth, & Goebel, 1998). In recent years, there has been a surge of interest in machine learning, which uses massive

[4] Exactly how a system converts stimuli to responses depends on the technique being used by the data scientist. For instance, a logistic regression equation might be used to consider multiple inputs that collectively trigger a real-world response if a certain threshold is met.

quantities of historical and simulated data to train intelligent agents on the correct responses to novel combinations of stimuli (see Appendix). The increase in speed and accuracy of these systems has led to speculation that a fourth industrial revolution may be on the horizon (Schwab 2017). While much of the public debate has been focused on job losses due to intelligent machines, the forecasts for the impact on entrepreneurs have been mixed, ranging from AI as a useful complement to entrepreneurial work, to the complete annihilation of entrepreneurs by robots and other intelligent machines.

The remainder of the paper builds the argument for why entrepreneurs will *not* be replaced by learning machines. In the next section, we consider some general arguments against this form of AI and then extend the arguments to the specific case of entrepreneurship.

The Case Against Artificial Intelligence

"It was like claiming that the first monkey that climbed a tree was making progress towards landing on the moon" (Dreyfus, 2012: 92)

Hubert Dreyfus, a philosophy professor in the phenomenological tradition, has been the *bete noire* of artificial intelligence researchers since shortly after the birth of the field in the late 1950s (Dreyfus, 1965, 1972, 2012). Writing initially for the RAND corporation and dismayed at Herbert Simon's 1965 prediction that AI would be able to do any work that a human could do within 20 years, he quickly identified a glaring weakness of all AI systems. "The basic problem facing workers attempting to use computers in the simulation of human intelligent behavior should now be clear: all alternatives must be made explicit" (Dreyfus, 1972: 41). All current AI systems rely on symbolic logic in conditional form: if <these input symbols> then <those output symbols>. This is true whether we are using expert systems or neural networks. For instance, in the case of facial recognition, the inputs are pixels, and the output is the identity of the subject. In the case of a diagnostic expert system, the inputs are symptoms, and the output is the name of the disease. Dreyfus argues that while some problems are amenable to symbolic logic, others are not. In general, the more unstructured (or non-formal) a problem the more difficult it will be for symbolic logic.

One aspect of the Dreyfus critique, known later as the "frame problem", refers to the difficulty in pre-specifying all relevant aspects of the environment (Dennett, 2006; Koppl et al., 2015). Dennett (2006) humorously describes the difficulties of a robot trying to remove a wagon from a room with a bomb and not realizing the bomb was on the wagon. A human would have quickly deduced that the location of the bomb was critical to solving the problem. Mitchell (2019) argues that the computer in this situation lacked "common sense", but it turns out that common sense is a complicated matter.

Mitchell (2019) describes four types of common sense that humans rely on: physical, biological, psychological, and sociological. Physical common sense refers to our background knowledge about the physical world. We know that an object teetering on a cliff is at risk of falling due to gravity. Our biological common sense tells us that a human will start to get hungry, or tired, after some time has elapsed, particularly after exertion. Our psychological common sense tells us that a human who is hungry might also become fearful if food is not readily available. Finally, our sociological knowledge refers to our background knowledge of social interactions between people, such as our ability to discern whether a statement like "I'm going to kill you" is hostile or playful. From these examples, it can readily be seen that humans possess a huge array of background knowledge that they can bring to bear on any given situation to make sense of what is going on.

Dreyfus (1972) divides the sense making challenge for a machine into three steps: where to focus, what to focus on, and when to focus. In structured problems, like Chess or Go, the computer has been highly tuned to a specific domain. Even in the robot example above, we know there is a room with a bomb and a wagon. In unstructured situations, determining where to focus is more problematic. Consider standing on a mountain with a 360-degree view of the horizon. Determining which direction to face is a non-trivial challenge. One solution is to focus on everything, but this is not how humans operate given their limited attention spans. We quickly narrow things down to a relevant area of operation. Computers have a much tougher time with this skill (that most humans seem to accomplish effortlessly).

Once we have focused on where to look, we must determine what is *essential* about the situation. We often conveniently overlook the fact that computers are told what is essential by their programmers but, in an unstructured situation, there is literally an infinite list of relevant factors that could be considered. For instance, consider the room with the wagon. The computer must search its inventory of objects to determine that a wagon exists in the room, that is has functional wheels, that there is an object on the wagon, that the object might be a bomb, and that bombs are not good for the health of the wagon or the wagon operators. At the same time, it must ignore the color of the wagon, the temperature of the room, the sounds emerging from the room (unless the bomb is ticking), and even the number and position of the oxygen molecules floating about the room. While it seems absurd to think that the number of molecules is in any way relevant to solving the problem, it is a very real barrier to a machine that lacks common sense and can simultaneously perceive millions of inputs.

Context is also important. As we saw, the meaning of the phrase "I'm going to kill you" is dependent on the context in which it occurs. The tone of voice, the proximity of the actors, and the history of interactions all provide clues to interpreting the meaning of the statement. There is mounting evidence that perceptual processing is top down as much as bottom up, meaning the brain forms expectations about what it expects to see from limited sensory data and then confirms its intuition with more sensory data (Clark, 2015). This enables the brain to economize on cognitive processing by not processing all sensory input.

Normally, there is high fidelity between expectations and reality but sometimes errors (known as *surprisals*) occur between what is expected and received, triggering a search for a more accurate model. So, a statement of hostile intent would normally be accompanied by other hostile actions such as physical aggression or facial grimaces. If, instead, the other party laughs at the statement (a surprisal) then the predictions would be updated to reflect this new information and create a new hypothesis that the statement was said in jest. Additional cues might be used to confirm or refute the new prediction (such as whether the laugh was a nervous laugh or a chuckle).

Imagine, however, if the recipient of the phrase, "I'm going to kill you", immediately drops to the ground and starts doing pushups. The brain must activate a wider search of relevant factors to derive an explanation for this unexpected behavior. The next thing that comes to my mind is that the subject might be hypnotized, and this is a trigger phrase. But, to the degree that people have different experiences, particularly at the psychological and sociological level, then we would expect them to form different (even competing) hypotheses about what is happening. In fact, the embodiment hypothesis, argues that human-like intelligence is impossible without having bodily experiences to ground (or situate) perceptions (Anderson, 2003).

The upshot is perception is an iterative process, and AI models that rely on only bottom-up inputs (like supervised machine learning) will have a difficult time with ambiguous inputs that require an understanding of context. Using context to decode meaning requires a deep well of common-sense knowledge (and experience) that can be accessed to generate hypotheses or predictions about the world, often by drawing on what Dreyfus (1972) calls fringe consciousness, that is, knowledge not immediately applicable to the situation. The human brain is also capable of creating or combining categories to explain phenomena when existing explanations are insufficient (Mitchell, 2019). For instance, the term "iron horse" was used to describe the first steam locomotives in the early 1800s, combining prior experiences of iron and horses into a neologism. As more examples of locomotives proliferated, it became a new category of transport.

Of course, determining what is essential becomes even more problematic when the issue of time is introduced, as the range of possible actions increases exponentially. Games, like Dota, are stretching the limits of computational power by considering 1,000 possible actions per move, 20,000 moves per game, and causal chains of up to 200 moves. Note, however, that this is a formal game where all the possible moves have been pre-specified by the programmers. In the unstructured (non-formal) world, the set of all possible moves is, once again, effectively infinite, making a brute search of all possible combinations of actions impossible (even if all the possible moves could be listed).

The Dota AI team also trains its algorithm in a simulated environment that enables it to process millions of possible action sequences every second. Similar simulated environments are being used to train robots and autonomous vehicles. Simulation, however, relies on a detailed understanding of the behavior of the world. In video games, the payoff from given actions are pre-specified. Damage from attacks with given weapons and the amount of damage a given character can tolerate are easily tracked, along with other critical in-game metrics. In robotic and driving simulations, the programmers often rely on detailed physical models of the environment. Cars can be taught to avoid objects and robots can learn to manipulate objects. Lessons can then be transferred to the real world. In these environments, so called 'edge cases' loom as a constant threat (Mitchell, 2019). These are situations that were not encountered in the simulation, either due to their rarity or non-representation in the simulated world. A robot that has never seen a bomb, has no conception of a bomb, and lacks the set of actions to dispose, defuse, or avoid a bomb will be helpless in the face of such a threat. Similarly, an autonomous vehicle that is unable to recognize a bicycle crossing its path will not react optimally to the situation.

The formal/non-formal distinction thus acts as a bright dividing line between artificial narrow intelligence (ANI) and artificial general intelligence (AGI). All existing AI is narrow, in the sense that it solves formal problems where the inputs, outputs (actions), and evaluation (payoff) functions are pre-defined by programmers to solve a problem in a specific domain. General AI, on the other hand, must be able to solve problems across multiple domains by: accessing fringe consciousness, making remote associations, solving edge cases, creating new categories, and possessing common sense (Mitchell, 2019). In fact, Hintze (2016) argues that general AI will only be achievable when a computer develops a 'theory of mind', that is, the ability to form predictions about the needs and motivations of others. He also recognizes a second level of general AI, which he calls 'consciousness', where the machine has developed self-awareness of its own needs and motivations.

Such a machine would be able to walk into a room containing a wagon with a bomb on it and surmise that the bomb represented a threat to its own existence, and to the well-being of any humans in the vicinity, perhaps by correctly interpreting the expressions of fear on the faces of people staring at the bomb. It might also determine that the bomb will destroy the wagon, which may be of value to another party and that it might be reasonable to save the wagon as long as it doesn't damage itself or any humans in the process. It might also be aware that it will receive a reward (perhaps in the form of energy) for helping humans.

Clearly, this form of general AI is far removed from the forms of machine learning that we have discussed earlier. Arguably, no amount of machine learning on large sets of historical or simulated data is going to lead to the ability to handle unstructured problems. This will require a theory of mind or consciousness. This is the logic behind the statement by Dreyfus at the start of the section that tree climbing will not get you to the moon. Simply put, the algorithms being used for narrow AI are good for tree climbing but not for moon shots (or general AI). The AI community is assuming that progress is one paradigm signals progress on a different type of problem.

Entrepreneurship as a case of general intelligence[5]

Having established that machine learning is form of ANI that will not lead to AGI, we will now assert and defend the claim that entrepreneurship is a type of general intelligence and thus not amenable to machine learning techniques. In doing so, we will also refute the claims of AI optimists, who see machine learning as a path to entrepreneur-free economic calculation.

In doing so, I will define entrepreneurship as the on-going deployment of resources in the anticipation of valued outcomes (Phelan, 2020). This definition implicitly incorporates many of the Austrian insights about entrepreneurship. First, the entrepreneur perceives an opportunity to move resources from a less desired current state to a more desired future state (Mises, 1949). These beliefs may be widely shared and require rapid action to exploit (Kirzner, 1973) or idiosyncratic

[5] Much of the argument in the next three sections was first presented in Part VII of Startup Stories (Phelan, 2020) and is reproduced with permission.

(Lachmann, 1976). They may involve the redeployment of existing goods or services (Kirzner, 1999) or the creation of novel combinations (Schumpeter, 1934).

As a redeployment typically requires the passage of time there is uncertainty whether the future state will be realized. This exposes the entrepreneur to the prospect of loss if the predictions are not accurate. This risk, in turn, motivates the entrepreneur to make sound judgments (Foss & Klein, 2012) and change plans when unanticipated developments occur (Sarasvathy, 2008). Often it is beneficial for the entrepreneur to own the resources being redeployed as this gives rights to residual income and control (Hart & Moore, 1990).

Demand side considerations

The late Steve Jobs is often held up as the epitome of a successful entrepreneur. His founding of Apple, ousting by his own board, and subsequent return to rescue the company, and then make it the most valuable publicly traded company in the world is the stuff of legend. One of the apparent secrets of his success is a disdain for market research. Isaacson (2011: 349), his biographer, famously quotes him as saying:

> Some people say, "Give the customers what they want." But that's not my approach. Our job is to figure out what they're going to want before they do. I think Henry Ford once said, "If I'd asked customers what they wanted, they would have told me, a faster horse!"' People don't know what they want until you show it to them. That's why I never rely on market research. Our task is to read things that are not yet on the page.

This ability to "read things that are not yet on the page" lies at the heart of the concept of *empathic accuracy* (McMullen, 2015). Empathic accuracy is "the ability to accurately infer the specific content of other people's thoughts and feelings" (Ickes, 1993: 588). McMullen (2015: 666) argues that "the variation represented by the introduction of a new product is…speculative in nature and informed to a greater or lesser extent by the entrepreneur's ability to take the perspective of various stakeholders (...) what is being judged are others' preferences".[6] In the case of Apple, Jobs was able to accurately envisage that products such as the iMac, iPod, iPhone, and iPad would be desired by a sizable number of customers. Importantly, he was able to convince other stakeholders (investors, employees, suppliers, distributors, contractors) of his vision so that he was able to direct a significant chunk of physical, financial, and human capital to realize his beliefs about customer needs.

Of course, the vision or belief does not need to be 100% accurate. There are numerous cases of companies revising their initial beliefs in the face of feedback from customers or other stakeholders. In fact, obtaining rapid feedback on one's plans is a feature of modern theories of entrepreneurship, such as lean startup (Ries, 2011) and effectuation (Sarasvathy, 2008). As Foss and Klein (2012) point out, one of the features of firm ownership is the authority to redirect capital to different activities, which in turn provides an impetus for entrepreneurs to own rather than rent their assets. Companies (and their founders) thus compete on the relative accuracy of their plans and the ability to adjust their plans based on market feedback and the actions of competitors. Empathic accuracy is thus an ongoing process and not a one-shot deal.

Empathic accuracy clearly requires a theory of mind. No amount of information about consumer preferences and buying trends would have enabled an AI to predict the latent demand for the iMac or iPhone. "To infer what product features will resonate with members of a particular group entrepreneurs must simulate (deliberately or automatically) others' decision making by imagining what that target believes and desires" (McMullen, 2015: 670). It is therefore an act of "creative discovery", which Dreyfus (1972, 189) describes as the recognition that a specific solution gratifies a general need. For instance, the iPhone satisfied a need for a technology that was easy to use but that only occurred after the consumer had experienced the device. The genius of Steve Jobs was to

[6] Empathic accuracy is similar to Mises' concept of thymology, which he described as "the cognition of human emotions, motivations, ideas, judgements of value and volitions" (Mises, 1957: 264). However, Mises (1957, 271) believed thymology had "no special relation to praxeology and economics." This conclusion has been questioned by Lavoie and Storr (2011: 222) who argue that, "Any expectational judgment needs to be based on the best available theoretical and historical knowledge."

recognize that need and present an embodiment of it to the consumer. He was not trying to serve a clearly stated preference. A computer without a theory of mind is thus not capable of evaluating the ability of a specific product to meet a general need.

The ability to combine knowledge across multiple domains is also impossible for a machine that is focused on one narrow domain. Psychologists have described “conceptual combination” as a fundamental human cognitive skill that combines two disparate concepts together into a higher level concept (Wu & Barsalou, 2009). For Dreyfus (1965), creative discovery through conceptual combination often involves the engagement of fringe consciousness. Entrepreneurs frequently engage in this process. For instance, Maurya (2012) recommends developing a high-level concept to describe a new product using a mash-up of previous products. The movie Aliens could be understood as “Jaws in Space” or YouTube as “Flickr for videos”. Similarly, Shane (2000) also describes how a single patent was exploited in several different ways by entrepreneurs with different life experiences. He uses the term ‘knowledge corridors’ to describe how different entrepreneurs will see different opportunities to use a particular technology based on their experiences. Accessing these experiences to combine disparate chunks of knowledge is a very human process.

Koppl et al (2015) argue that new opportunities emerge because they are *adjacent* to new discoveries in other fields. The introduction of fundamental technologies like electricity, steam, or the internet opened an array of opportunities for new products that were not foreseeable before key enabling technologies became available. The success of Amazon.com, for example, relied not only on the internet but also a capable browser, an electronic payment system, and a delivery service. Opportunities thus unfold over time in an idiosyncratic way that is often rooted in the experiences of individual entrepreneurs. They do not exist in some pre-formed state awaiting discovery via some clever search algorithm.

Supply side considerations

While the successful navigation of consumer demand requires a theory of mind, the investment of one's time, talent, and treasure into a risky new venture also requires considerable self-awareness, which is Hintze's (2016) second attribute of general AI. According to Hintze (2016, para 23), "Conscious beings are aware of themselves, know about their internal states, and are able to predict feelings of others". It is this requirement for a knowledge of internal states that we turn to next.

As we have seen, machine learning requires large amounts of actual or simulated data to drive its recommendations. However, entrepreneurs do not have the luxury of undertaking a brute search of every possible combination to find the best solution. Actions in the real world are consequential (Foss & Klein, 2012). Entrepreneurs have one shot, or at best, a handful of shots to get it right. They live with the possibility of great returns or great losses, including changes in personal fortune, reputation, or even survival. The upside potential, and the downside risk, jointly serve to discipline the entrepreneur to make the best resource allocations possible. Entrepreneurs are making judgments that their actions will lead to success over failure in an uncertain situation and betting their financial and human capital on the outcome.

But what exactly is a machine risking? In a game of Chess or Go, the machine has an evaluation function that it is trying to optimize. A successful sequence of moves is rewarded because it moves closer to a pre-determined goal. Similarly, a machine is rewarded in a facial recognition task for correctly matching a set of stimuli to a face. Narrow AI, however, has no concept of existential risk, that is, the awareness that a bad decision can diminish (or even extinguish) one's life. This is what makes an entrepreneur's decisions so consequential. A machine, on the other hand, does not care if a decision leads to its demise or destruction.

For Oliver (2017: 18), care is an essential human trait that distinguishes us from machines. He argues that, "[machines] *cannot* die, even if they can be destroyed" (emphasis in original). As humans, we are aware of our own mortality. We understand the abstract concept of death, and the death of others, but we always live in the shadow of our own demise. This awareness informs our actions. Mortality is part of our "being in the world" or *Dasein* to use Heidegger's (1927/1962) term. We are *embedded* in the world because of our temporal nature and *embodied* in the world because of our physical form. We are born, we live, we die. Our experiences from the past, the stories of our ancestors, our present needs for food, shelter, and companionships, and the knowledge that we will die at some unspecified time in the future all inform who we are and what we care about. Machines, at least as currently constructed, have no *Dasein*.

As such, a machine makes a poor entrepreneur because it does not care about the significance of one economic judgment over another. In every AI task to date, the importance of one outcome over another is pre-specified by a human. Humans tell the AI what to care about. We *care* about winning a chess game, we *care* about making a profit, we *care* about not hitting a pedestrian with a vehicle. The AI, on the other hand, places no inherent value on one sequence of moves (or one combination of resources) over another.

The implication is that machines will not naturally demonstrate prudence in assembling resources, having no anxiety (or dread) about negative outcomes. An economy run completely by machines, where machines determine what is valued, would be a disaster. Humans must tell machines what to care about. This prompted Issac Asimov (1950) to formulate the three laws of robotics: a robot may not injure a human being or, through inaction, allow a human being to come to harm; a robot must obey the orders given it by human beings except where such orders would conflict with the First Law; and a robot must protect its own existence as long as such protection does not conflict with the First or Second Laws. While other laws have been suggested over the years the fact remains that there is a void in ethical intelligence that needs to be developed in, or supplied to, machines.

Of course, critics of the free market would argue that entrepreneurs harm humans every day, with Marx famously claiming that profit arises from the exploitation of the proletariat. While not seeking to get into a full-blown debate on the ethics of capitalism, free market economists argue that market capitalism has been more effective in improving quality of life than other forms of economic organization. As Thomas Sowell (2010: 123) puts it

> While capitalism has a visible cost—profit—that does not exist under socialism, socialism has an invisible cost—inefficiency—that gets weeded out by losses and bankruptcy under capitalism (..). profit is a price paid for efficiency. Clearly the greater efficiency must outweigh the profit or else socialism would in fact have had the more affordable prices and greater prosperity that its theorists expected, but which failed to materialize in the real world.

It should be noted that entrepreneurs do not just care about profits. They are embedded within a society that sees profit as a means to an end rather than an end in itself. Simply telling a machine to maximize profits is not going to lead to human flourishing. The profit motive in modern economics has been constrained by laws and conventions that have emerged over time to temper the worst aspects of human greed, including prohibitions on slavery, child labor, price gouging, food safety, public health, and monopolies among others. Knowledge of these institutions, in turn, forms part of the fringe consciousness of every entrepreneur and would have to be incorporated into any AI entrepreneur. Nevertheless, the point is not whether machines can be given an expanded set of rules but rather whether the absence of human emotions (like shame and guilt) and motives (like hunger and pain) make a machine less likely to be a disciplined entrepreneur. I believe this would be the case.

Implications for Computability and Economic Calculation

It should be readily apparent from the previous comments that machine learning is not artificial general intelligence, and that a machine would require artificial general intelligence to effectively function as an entrepreneur in the economy. It is now possible to critically evaluate the proposals of AI optimists like Cockshott (2019) by utilizing this understanding.

First, there is no doubt that machine learning will become a valuable tool for entrepreneurs and managers. The ability to access and process vast amounts of information and find patterns that can be exploited for gain will become an increasingly important part of the economy. Much like previous technologies, we expect that these developments will displace entrepreneurs and other workers who specialize in processing this information in less efficient ways. In many cases, decision making will be delegated to machines, with medical diagnosis, stock trading, and autonomous vehicles being just some of the many examples. However, as we have seen, in every case an entrepreneur is choosing to delegate these decision rights. The liability for error (or original judgment) still rests with the entrepreneur.

Thus, the most important entrepreneurial functions of empathy and judgment will continue to be undertaken by humans. As we saw, economics is not just a matter of producing and distributing existing products more efficiently, but also involves providing incentives to discover ways to create new means and ends to improve our collective prosperity. This involves optimizing the economy in a dynamic rather than static way (Stalebrink, 2004). In fact, the productivity gains from new discoveries over the last two centuries have dwarfed the gains from trade. By some measures, we are almost 400 times better off than our ancestors from two centuries ago (Kling & Schulz, 2011). Moreover, most of these gains have been realized by society at large, with only a small percentage being captured by innovators (Baumol, 2010).

Critically, these gains have relied on an intimate understanding of consumers and the willingness to take calculated risks predicated on expected profits. Even so, current estimates suggest that 75% of equity funded startups will fail (Shane, 2008). While this is a seemingly high failure rate, it pales into insignificance when compared to the millions of (failed) trials needed to train a sophisticated machine learning system. As a result, in unstructured environments like the market for future products, humans will continue to outperform machines by a large margin.

These impediments to an entrepreneur-free, centralized, or machine-led economy have been known for some time. Mises, writing in 1920, stated, “It is now universally agreed that the exclusion of free initiative and individual responsibility, on which the successes of private enterprise depend, constitutes the most serious menace to socialist economic organization”(Mises, 2016 [1920]: 23). Cockshott and Cottrell (1993) while very aware of this venerable Misesian critique are dismissive of its implications. Innovation is just a budgetary matter they argue. Socialist managers will be told to innovate, and innovation shall occur. Their incentive to innovate, will be prestige in the eyes of fellow workers, or perhaps as in Saros’ (2014) scheme, a compensation bonus. The large payoffs for success (or the loss of one's capital in the case of failure) found in a market economy will be absent.

Sowell (2010: 210) argues that this experiment has already been run, quoting Leonid Brezhnev saying that his country's enterprise managers shied away from innovation "as the devil shies away from incense." As machines have no intrinsic motivation to succeed or fail, firms would have to rely on managers to create innovative evaluation mechanisms. However, the prospects that this would be done diligently seem low given historical experience. "Since we are in a position to survey decades of State and socialist endeavor, it is now generally recognized that there is no internal pressure to reform and improvement of production in socialist undertakings, that they cannot be adjusted to the changing conditions of demand" (Mises 1920/2016: 24).

In theory, machines make the optimization of production more tractable. Imagine several worker's councils that produce cars. Remote sensing could enable planners to collect data on the quantity of steel, rubber, skilled labor, and other inputs used in the production process. You could easily average quantities across all facilities and project the required level of inputs for a given demand. You could also insist that every facility match the productivity of the top producer and set aggressive targets. Once again, we have seen this story play out, with production targets being missed, or quantity being favored at the expense quality or variety. The old Soviet joke "So long as the bosses pretend to pay us, we will pretend to work" seems relevant here too. There might be a very good reason why one facility is not as productive as another, or it might be a simple case of shirking. An owner has a strong incentive to determine the true situation, a bureaucrat less so, and a machine none at all. In fact, a machine will only analyze the factors it is told to analyze and will have difficulty considering factors outside its original frame of reference.

Conclusion

This paper raises the simple question, "can entrepreneurship be learned by intelligent machines". Our conclusion is that machine learning, in its current form, is *not* capable of acting in an entrepreneurial capacity. It lacks a theory of mind and the common sense to anticipate consumer needs. A computer also lacks internal motivations, making it indifferent to good (and bad) choices. While such motivations can be imposed from outside, it will require a sophisticated self-awareness based on the experiential concept of *Dasein* (or being-in-the-world) to navigate unstructured problems with the same assuredness as humans.

Machine learning is still in its infancy and will continue to grow and evolve. However, the machine learning paradigm will not lead to the type of artificial general intelligence popularized in story and on screen. There will still be a role for entrepreneurs and the entrepreneurial functions of empathy and judgment well into the foreseeable future.

References

Anderson, M. L. (2003). Embodied Cognition: A field guide. *Artificial Intelligence, 149*(1), 91-130. doi:10.1016/s0004-3702(03)00054-7

Asimov, I. (1950). Runaround. In *I, Robot*. New York: Doubleday.

Baumol, W. J. (2010). *The microtheory of innovative entrepreneurship*: Princeton University Press.

Bradley, M. E., & Mosca, M. (2014). Enrico Barone's 'Ministry of Production': Content and Context. *The European Journal of the History of Economic Thought, 21*(4), 664-698.

Campbell, M., Hoane, A. J., & Hsu, F.-h. (2002). Deep blue. *Artificial Intelligence, 134*(1-2), 57-83.

Chen, J. X. (2016). The evolution of computing: AlphaGo. *Computing in Science & Engineering, 18*(4), 4-7.

Childers, T. (2019). Google's Quantum Computer Just Aced an 'Impossible' Test. Retrieved from https://www.livescience.com/google-hits-quantum-supremacy.html

Clark, A. (2015). *Surfing uncertainty: Prediction, action, and the embodied mind*: Oxford University Press.

Cockshott, P. (2017). *Big Data and Supercomputers: Foundations of Cyber Communism*. Paper presented at the The Ninth International Vanguard Scientific Conference on "100 Years of Real Socialism and the Theory of Post-Capitalist Civilization", Hanoi, Vietnam. http://eprints.gla.ac.uk/148529/

Cockshott, P. (2019). Response to Horwitz. Retrieved from https://paulcockshott.wordpress.com/2019/10/30/response-to-horowitz/

Cockshott, P., & Cottrell, A. (1993). *Towards a new socialism*: Spokesman Pr.

Cockshott, P., & Cottrell, A. (1997). Information and economics: A critique of Hayek. *Research in Political Economy, 16*, 177-202.

Cottrell, A., & Cockshott, P. (1993). Calculation, complexity and planning: The socialist calculation debate once again. *Review of Political Economy, 5*(1), 73-112. doi:10.1080/09538259300000005

Dennett, D. C. (2006). Cognitive wheels: The frame problem of AI. In J. L. Bermudez (Ed.), *Philosophy of psychology: Contemporary readings* (pp. 433-454). New York: Routledge.

Dreyfus, H. L. (1965). Alchemy and Artificial Intelligence.

Dreyfus, H. L. (1972). *What computers can't do: The limits of artificial intelligence* (Vol. 1972).

Dreyfus, H. L. (2012). A History of First Step Fallacies. *Minds and Machines, 22*(2), 87-99.

Foss, N. J., & Klein, P. G. (2012). *Organizing entrepreneurial judgment: A new approach to the firm*: Cambridge University Press.

Hart, O., & Moore, J. (1990). Property Rights and the Nature of the Firm. *Journal of political economy, 98*(6), 1119-1158.

Hayek, F. A. (1945). The use of knowledge in society. *American Economic Review, 35*(4), 519-530.

Hayek, F. A. (1952). *The sensory order: An inquiry into the foundations of theoretical psychology*: University of Chicago Press.

Heidegger, M. (1927/1962). *Being and time*. Oxford, UK: Blackwell.

Hintze, A. (2016). Understanding the four types of AI, from reactive robots to self-aware beings. *The Conversation*. Retrieved from https://theconversation.com/understanding-the-four-types-of-ai-from-reactive-robots-to-self-aware-beings-67616

Ickes, W. (1993). Empathic Accuracy. *Journal of Personality Assessment, 61*(4), 587-610.

Isaacson, W. (2011). *The Best Steve Jobs Quotes From His Biography*. New York: Simon & Schuster.

Kirzner, I. M. (1973). *Competition and entrepreneurship*. Chicago, IL: University of Chicago press.

Kirzner, I. M. (1999). Creativity and/or alertness: A reconsideration of the Schumpeterian entrepreneur. *Review of Austrian Economics, 11*(1), 5-17.

Kling, A., & Schulz, N. (2011). *Invisible Wealth: The Hidden Story of How Markets Work*: Encounter Books.

Koppl, R. (2008). Computable entrepreneurship. *Entrepreneurship Theory and Practice, 32*(5), 919-926.

Koppl, R., Kauffman, S., Felin, T., & Longo, G. (2015). Economics for a creative world. *Journal of Institutional Economics, 11*(1), 1-31.

Kurzweil, R. (2000). *The age of spiritual machines: When computers exceed human intelligence*: Penguin.

Lachmann, L. M. (1976). From Mises to Shackle: an essay on Austrian economics and the kaleidic society. *Journal of Economic Literature, 14*(1), 54-62.

Lange, O., Taylor, F. M., & Lippincott, B. (1938). *On the economic theory of socialism*: University of Minnesota Press, New York and London.

Lavoie, D. (1981). A critique of the standard account of the socialist calculation debate. *Journal of Libertarian Studies, 5*(1), 41-87.

Lavoie, D. (1985). *Rivalry and central planning: The socialist calculation debate reconsidered*: Cambridge University Press Cambridge.

Lavoie, D., & Storr, V. H. (2011). Distinction or dichotomy: Rethinking the line between thymology and praxeology. *The Review of Austrian Economics, 24*(2), 213-233.

Limas, E. (2018). Cybersocialism: A Reassessment of the Socialist Calculation Debate. *Available at SSRN 3117890*.

Maurya, A. (2012). *Running lean: iterate from plan A to a plan that works*: O'Reilly Media, Inc.

McMullen, J. S. (2010). Perspective taking and the heterogeneity of the entrepreneurial imagination. In *What is so Austrian about Austrian economics?* (pp. 113-143): Emerald Group Publishing Limited.

McMullen, J. S. (2015). Entrepreneurial judgment as empathic accuracy: a sequential decision-making approach to entrepreneurial action. *Journal of Institutional Economics, 11*(3), 651-681. doi:10.1017/s1744137413000386

Mises, L. (1949). *Human action*. New Haven: Yale University Press.

Mises, L. (1957). *Theory and History: An Interpretation of Social and Economic Evolution*: Yale University Press.

Mises, L. (2016 [1920]). *Economic calculation in the socialist commonwealth*: Lulu Press, Inc.

Mitchell, M. (2019). *Artificial Intelligence: A Guide for Thinking Humans* (Kindle Edition ed.). New York: Farrar, Straus and Giroux.

Morozov, E. (2019). Digital Socialism? The Calculation Debate in the Age of Big Data. *New Left Review, 116*, 19-46.

Neurath, O. (1919). *Durch die Kriegswirtschaft zur Naturalwirtschaft*: Callwey.

O'Driscoll Jr, G. P., & Rizzo, M. (2014). *Austrian economics re-examined: the economics of time and ignorance*: Routledge.

Olivier, B. (2017). Artificial Intelligence (AI) and being human: What is the difference? *Acta Academica, 49*(1), 2-21.

OpenAI. (2018). OpenAI Five. Retrieved from https://openai.com/blog/openai-five/

Phelan, S. E. (2016). Austrian theories of entrepreneurship: Insights from complexity theory. *The Review of Austrian Economics, 29*(3), 277-297.

Phelan, S. E. (2020). *Startup Stories: Lessons for Everyday Entrepreneurs*: Lupino Press.

Phillips, L., & Rozworski, M. (2019). *The People's Republic of Wal-Mart: How the World's Biggest Corporations Are Laying the Foundation for Socialism*: Verso.

Poole, D., Mackworth, A., & Goebel, R. (1998). *Computational Intelligence: A Logical Approach*: Oxford University Press.

Richman, S. L. (1981). War Communism to NEP: The road from serfdom. *The Journal of Libertarian Studies, 5*(1), 89-97.

Ries, E. (2011). The Lean Startup: How Constant Innovation Creates Radically Successful Businesses. Kindle Edition. In. New York, NY: Penguin Books

Robbins, L. (1934). *The Great Depression. New York: Macmillan.* New York: Macmillan.

Sarasvathy, S. D. (2008). *Effectuation: Elements of entrepreneurial expertise*: Edward Elgar Publishing.

Saros, D. E. (2014). *Information Technology and Socialist Construction: the end of Capital and the transition to socialism*: Routledge.

Schumpeter, J. (1934). *The Theory of Economic Development*. Cambridge, MA: Harvard University Press.

ScrapeHero. (2019). How Many Products Does Amazon Sell? Retrieved from https://www.scrapehero.com/number-of-products-on-amazon-april-2019/

Shane, S. (2000). Prior knowledge and the discovery of entrepreneurial opportunities. *Organization science, 11*(4), 448-469.

Shane, S. (2003). *A general theory of entrepreneurship: The individual-opportunity nexus*: Edward Elgar.

Shane, S. (2008). *The illusions of entrepreneurship: The costly myths that entrepreneurs, investors, and policy makers live by*: Yale University Press.

Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., . . . Bolton, A. (2017). Mastering the game of Go without human knowledge. *Nature, 550*(7676), 354-359.

Sowell, T. (2010). *Basic Economics* (4th Ed.). New York: Basic Books.

Stalebrink, O. J. (2004). The Hayek and Mises controversy: bridging differences. *The Quarterly Journal of Austrian Economics, 7*(1), 27-38.

Toivonen, H., & Gross, O. (2015). Data mining and machine learning in computational creativity. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery, 5*(6), 265-275.

Townsend, D. M., & Hunt, R. A. (2019). Entrepreneurial action, creativity, & judgment in the age of artificial intelligence. *Journal of Business Venturing Insights, 11*, e00126.

Wiggins, G. A. (2006). A preliminary framework for description, analysis and comparison of creative systems. *Knowledge-Based Systems, 19*(7), 449-458.

Wu, L.-l., & Barsalou, L. W. (2009). Perceptual simulation in conceptual combination: Evidence from property generation. *Acta psychologica, 132*(2), 173-189.

Yeung, T. Y.-C. (2019). Is Friedrich Hayek Relevant Anymore in the Age of Artificial Intelligence? *Available at SSRN 3484259*.

Appendix. A Primer on Big Data and Artificial Intelligence

The growth in information technology since Hay Jonikas Stage 1ek wrote about the knowledge problem in 1945 has been nothing short of remarkable. The first digital computer, ENIAC, debuted in 1946, and was capable of 5,000 addition operations per second (about 1,000 times faster than the fastest human can perform additions). Today, the world's fastest supercomputer, Summit, can undertake a quintillion (10^{18}) operations per second. This greatly exceeds Kurzweil's (2000) estimate of the total processing power of the human brain at 10^{16} operations per second[7]. Moreover, scientists working on quantum computing have recently claimed processing times a trillion times faster than Summit for certain problems (Childers, 2019). Clearly, machines are already superhuman when it comes to raw processing power.

Along with the ability to process vast amounts of data, comes the ability to store and transmit equally large quantities of data. Summit has 250 million gigabytes of disk storage, over 3 million gigabytes of random-access memory, and transmits 200 gigabits of data per second between nodes (Wikipedia, 2020). And Summit is just one device. Cloud computing, mobile devices, and high-speed networks have made data ubiquitous around the globe. We are now seeing an Internet of Things (IoT) starting to emerge, where household items such as refrigerators, washing machines, and thermostats can collect data and automatically act based on their owners' preferences. Hayek's 'man [sic] on the spot' having local knowledge of time and place is seemingly a thing of the past. Machines can be omnipresent, all seeing and all knowing. We have entered the era of 'big data'.

But data (or information) is not knowledge. Early AI (known today as Good Old Fashioned AI or GOFAL) tried to capture domain specific knowledge by asking experts (such as doctors) to articulate the rules they used to determine an outcome (such as diagnose a disease). Similarly, on

[7] Note, this includes all operations including cognitive functions (such as addition, language, and perception) and autonomic functions (such as breathing and temperature regulation).

the commercial front, if a trader knew to buy Madeira oranges when the spot price fell below $2 then that association could be programmed into a machine (Mitchell, 2019).

However, expert systems, while competent in narrow domains, lack the nuance and exceptional understanding of seasoned professionals (Dreyfus, 2012). A seasoned trader might have known to buy Madeira oranges when the price fell below $2 *except* when it was raining on Tuesday. This has often been described as the articulation problem (O'Driscoll Jr & Rizzo, 2014). We simply know more than we can say.

Machine Learning

Machine learning is a way of circumventing the articulation problem and underlies much of the excitement in modern artificial intelligence. There are two main approaches to machine learning: supervised learning and reinforcement learning (Mitchell, 2019). In supervised learning, the machine is trained on a set of valid responses to complex stimuli. In the case of facial recognition, a classic use case, thousands (or millions) of pixels from a picture are used as inputs, and the identity of the subject is given as the response. In the simplest terms, the algorithm weights the contribution of each pixel to a given response. By giving thousands or even millions of examples of these associations, the machine "learns" the appropriate response without having to elicit every nuance from a human expert.

The algorithm can also generate responses to novel stimuli. This is because it averages the responses from individual pixels and can thus respond to novel pixel combinations. The technique is often referred to as *deep learning* because a number of 'hidden' layers can exist between the stimulus and response layers. These 'artificial neural networks' can aggregate responses from individual inputs in complex ways, and even provide feedback loops, thus mimicking the neuronal structure of the human brain (while also making the underlying logic of a particular response difficult to dissect).

Reinforcement learning, on the other hand, is a technique that allows a machine to learn a sequence of moves over time to realize a goal. Each task can be represented as a set of elements representing a universe of known states of the world, *U*, the rules of the game, *R*, and an evaluation function, *E* (Toivonen & Gross, 2015; Wiggins, 2006). For instance, the game Tic Tac Toe has nine squares (3 x 3) each capable of holding an X, O, or blank thus yielding 3^9 (or almost 20,000) possible states in its universe. Rules include: there can be only two players, play alternates between those two players, each player has one symbol type (X or O), a player can only play one symbol per turn, symbols can only be placed in blank squares, and a winner must get three symbols in a row (horizontal, vertical, or diagonal). When a winning state is achieved, the evaluation function provides a large payoff or reward to the machine (and, in more complex games, promising intermediate stages might also be rewarded).

The algorithm then observes sequences of moves, either through examining a set of historical games, or more commonly, by simulating many games against itself. Moves that result in large (winning) payoffs receive discounted rewards proportional to their proximity to the final move. So, in Tic Tac Toe, there are nine possible opening moves for the first player. If playing an X is the upper right square is more likely to lead to victory, then that move would be weighted higher by the reinforcement algorithm. The probability of choosing a given move at a given stage would then be proportional to its accumulated rewards.

The same principles can be extended to more complex games like Chess, Go, or online games like Dota2. Table 1 gives an indication of the progressive complexity of these games. Open AI is now defeating human players on Dota 2 using a combination of supervised learning and reinforcement learning that enables the machine to simulate 180 years of human game play per day (OpenAI, 2018). Similarly, the AlphaGo system was trained to the highest professional level in the game of Go and defeated the world champion in 2016 (Chen, 2016; Silver et al., 2017), and it has now been almost a quarter century since DeepBlue defeated the world's best human chess player back in 1997 (Campbell, Hoane, & Hsu, 2002).

Table 1. Comparison of Games

Game	Tic Tac Toe	Chess	Go	Dota2
Breadth (legal moves per position)	9	35	250	1,000
Depth (average moves per game)	3	40	150	20,000
Board Size	3 x 3 = 9	8 x 8 = 64	19 x 19 = 361	Open ended
Position States	3 (X, O, blank)	6 pieces, blank	Black, white, blank	20,000
States of the World	10^{5}	10^{50}	10^{170}	Effectively infinite
Search Depth (moves ahead)	5	6-20	150	180