Artificial Intelligence in the Context of Human Consciousness

Hannah DeFries

A Senior Thesis submitted in partial fulfillment
of the requirements for graduation
in the Honors Program
Liberty University
Spring 2019

Acceptance of Senior Honors Thesis

This Senior Honors Thesis is accepted in partial fulfillment of the requirements for graduation from the Honors Program of Liberty University.

____
Kyung K. Bae, Ph.D.
Thesis Chair

____
Jung-Uk Lim, Ph.D.
Committee Member

____
Mark Harris, Ph.D.
Committee Member

____
James H. Nutter, D.A.
Honors Director

____
Date

Abstract

Artificial intelligence (AI) can be defined as the ability of a machine to learn and make decisions based on acquired information. AI's development has incited rampant public speculation regarding the singularity theory: a futuristic phase in which intelligent machines are capable of creating increasingly intelligent systems. Its implications, combined with the close relationship between humanity and their machines, make achieving understanding both natural and artificial intelligence imperative. Researchers are continuing to discover natural processes responsible for essential human skills like decision-making, understanding language, and performing multiple processes simultaneously. Artificial intelligence attempts to simulate these functions through techniques like artificial neural networks, Markov Decision Processes, Human Language Technology, and Multi-Agent Systems, which rely upon a combination of mathematical models and hardware.

Table of Contents

Artificial Intelligence in the Context of Human Consciousness

The idea of artificial intelligence has existed for millennia, as evidenced by ancient inventions and myths of automatons in Egypt, the bronze Talos of Greece, and robot servants of China. However, the term "artificial intelligence" was coined as late as 1956, during a workshop held at Dartmouth College (Press, 2016).This term has since become a household phrase due to its integration into entertainment, social media, transportation, and medical research. Every year, fresh applications of intelligent devices hit the market, further imprinting AI onto the social consciousness.

Despite the integration of artificially intelligent technology into daily life, few individuals possess a working knowledge of the methods that inspire and provide highly capable technology like smart phones, digital assistants, and outcome prediction software. This dearth of understanding applies to the inspiration of AI as well: the human brain. Yet gaining a deep knowledge of the human brain is essential for the continued development of artificial intelligence. Dreyfus stated in 1965 that any attempt to analyze human intelligence using a computer neglects key aspects of a human's mind, including fringe consciousness (Dreyfus, 1965). Although introspection, social sciences, and biological sciences allow researchers to glimpse the workings of the brain and develop models to replicate certain functions, "significant developments in artificial intelligence…must await computers of an entirely different sort, of which the only existing prototype is the little-understood human brain" (Dreyfus, 1965, p. 111).

In spite of these challenges, artificial intelligence shows great potential for future growth and development. Current projections estimate that computers will exceed the capacity of the human brain by 2045. Ray Kurzweil, renowned computer scientist, has stated that AI will surpass human intellect at a point known as the Singularity, a cycle in which intelligent machines create machines with even higher intelligence. Some notable proponents of the Singularity, including Stephen Hawking and Elon Musk, believe that this exponential spread of artificially intelligent technology poses a danger to the human race (Paasschen, 2017). The key to evaluating this prediction lies in a deep understanding of both the human brain and its progress in creating artificial intelligent technology, namely the mechanisms which attempt to emulate the human being's ability to make rational decisions, assimilate information, and use creativity to solve problems.

Human Intelligence

Learning from the Brain

To understand and develop AI, researchers have labored to study the human brain. Scientists know little about the limitations of human intelligence. While intelligence tests are used to evaluate a subject's ability to solve problems, they cannot accurately predict human intellectual development, as no comprehensive survey of human intelligence exists. Rather, scientists must rely upon other sources to gain insight, such as introspection, social sciences, and biological sciences (Jackson, 2015).

Introspection. Introspection is a technique that allows scientists to gain knowledge of human intelligence by analyzing their own minds and is likely the most common source of AI research. Scientists and developers use their own experiences of solving problems to understand and design intelligent programs. Before testing a solution with a machine, they may naturally choose to test the process themselves. As quoted by Jackson (2015), Descartes' *Discourse on Method* provides one of the most famous notions of introspection: "I think therefore I exist" (p. 6). According to Descartes, people are fundamentally different from machines because of how they think (Jackson, 2015). The human capacity for high-level language, making decisions, and other skills is unparalleled, as will be discussed in future sections.

Introspection supports a concept known as “rooms of consciousness,” one of many characteristics that separates humans from machines. This concept is based upon the perception that the human mind inhabits and traverses multiple rooms. Each room is distinguished by its lighting, which can be interpreted to mean its level of intellect. Within his mind, a human can move into distinct levels of intellectual ability, which may be separated into scientific and spiritual areas. Scientific learning relies on a foundation of rules, whereas spiritual learning does not necessarily follow intellectual reasoning or require language. The rooms can be further divided into areas of thought such as emotional learning, in which knowledge is perceived without reasoning; inspired learning, where knowledge is obtained as if provided by a deity; and paradoxical learning, in which knowledge contradicts itself without regard to logic (Jackson, 2015). Such non-intellectual learning may never be able to be adapted to machines. Thus, although introspection helps scientists understand human minds, it does not provide a full solution for mapping human intelligence to artificial. The goal of AI is to merely simulate the abilities of human intelligence by reproducing the outward behaviors, not necessarily copy the methods themselves.

Social sciences. Social sciences include psychology, anthropology, economics, and sociology. Several conclusions regarding human intelligence have been drawn from these fields. First, human intelligence is indistinguishable across racial and linguistic boundaries. A healthy child raised appropriately in his environment can learn any language, not just the language spoken by his parents. Second, human intelligence takes time to develop and can be greatly impacted by an individual's environment. It is likewise affected by heredity, meaning that identical twins raised in similar environments will show little difference in IQ. Intelligence can also vary with respect to problem domain; individuals can have different aptitudes, such as musical or artistic talents, ease in learning languages, or proficiency in mathematics (Jackson, 2015).

Piaget and other scientists have demonstrated that children's brains develop in stages. In the sensory-motor stage during the first 1.5-2 years of life, a child develops a rudimentary ability to use signals and reason causally. In the preoperational or symbolic-operational stage, from years 2-7, a child learns basic vocabulary in his cultural language and can use sentences to describe his environment. In the concrete operational stage from ages 7-11, the child can express ideas of causality, including invariance, reversibility, and conservation. In the next stage, starting at age 11, formal operations, he develops the ability of abstract reasoning and can develop a logical argument. Of course, these stages are generalizations; exceptions do occur, as seen in the lives of Mozart and other precocious individuals (Jackson, 2015).

Scientists also gather information from behavioristic psychology. Behavioristic psychologists attempt to understand intelligence by viewing subjects as "black boxes"; subjects are presented with standardized situations, and their responses are recorded. These learning experiments frequently involve conditioning, a famous example of which can be found in Pavlov's canine experimentation, in which the scientist's dog was conditioned to salivate when a bell was rung. In human conditioning trials, a stimulus – which the subject recognizes as punishment or reward – is paired with a neutral stimulus, as in a tone or a flash of light. Conditioning occurs when the subject grows to expect the former when presented with the latter. Psychologists have demonstrated that subject conditioning becomes more rapid as the stimulus becomes more intense until the process reaches its limit. Psychologists postulate that conditioning is responsible for all forms of adaptation (Jackson, 2015). By strengthening the connections within the human brain, conditioning improves humans' ability to respond to stimuli and thus adjust quickly to a changing environment.

Biological sciences. Biological sciences include biology, biochemistry, and neurobiology. The human brain holds over 100 billion neurons, millions more of which are found throughout the body in an extensive, complex network. A neuron consists of a main cell body and nucleus. Neurons communicate across synapses using special structures that extend from each neuron. Axons, branches extending from a cell body, carry messages away from the neuron to other neurons, gland cells, and muscle cells. Dendrites, branchy structures, receive messages and relay them toward the cell body. The dendrites and axons are distributed to form powerful neural networks throughout the body (Parker, 2007). Each neuron contains an average of 5600-60000 dendrites, each with a connected axon. Based on these numbers, the human brain may have a processing ability and storage capacity that are several orders of magnitude higher than the typical computer (Jackson, 2015). As a result, humans are capable of performing highly complex tasks, from making informed decisions to learning and applying language, as discussed below.

Key Skills and Techniques of Human Intelligence

Decision-making. Human decision-making is a result of the pathways formed when learning about one's environment, selecting actions, and receiving feedback regarding the action taken. This process is heavily dependent on the formation of memory. To study the creation and function of human memories, scientists undertake comprehensive research through biological study and clinical observations.

As found from biological study, when forming a memory, neurons develop new axons and connections with other cells. Parts of the brain closely monitor their associated data, including facts, feelings, or sensory information. The data selected for input into memory is assigned to its respective area of the brain. When a neuron receives an input in the form of a series of impulses, it transmits corresponding impulses to another neuron. The first and second neurons connect with a third, creating a circuit. As the traffic of data increases, the synapses between neurons are gradually strengthened, and the more the brain recalls the memory, the longer the established connection is retained. Eventually, with repeated activation, the circuit is assimilated into the surrounding network, the collection of neurons preserving the memory (Parker, 2007).

According to psychologists, three forms of memory exist: sensory information storage, short-term memory, and long-term memory. Sensory information is retained for mere tenths of a second until it can be processed by the central nervous system (CNS). For example, when humans shut their eyes, they may view a fleeting afterimage, which is further processed by the CNS to glean any pertinent data (Jackson, 2015). Short-term memory, which occurs in the prefrontal cortex, lasts for approximately one minute and has a capacity of around seven items (Mayfield Brain and Spine, 2019). As a result, a subject may memorize a set of words but quickly forget the sequence if asked to switch to another task, which interrupts the subject's memorization process (Jackson, 2015). Long-term memory is processed in the hippocampus of the temporal lobe (Mayfield Brain and Spine, 2019) and occurs when a memory is retained semi-permanently due to consistent reactivation of the neural synapses. When any two stimuluses assigned to associated cells or networks of cells are temporarily paired through conditioning, the connections between the cells are grown and strengthened (Jackson, 2015), allowing personal memories and important facts to be stored for longer periods of time. This process of memory retention accomplished by neural networks provides an efficient means of storing the information essential for daily decision making.

Neural networks. At the lowest level, every decision is made by the transmission of nerve impulses. Each impulse is based on the movement of sodium and potassium ions through the cell membrane and flows through a neuron as a depolarizing and repolarizing wave. When the neuron is at rest, the outside of the cell membrane holds a greater positive charge than the interior, creating a resting potential of -70 mV. When an area of the cell depolarizes, positive sodium ions flow through channels into the membrane, making the inside slightly positive and creating an action potential of 30 mV. The area then rapidly repolarizes as potassium ions flow out of the cell to restore equilibrium. As each area of the cell experiences these changes in electrical charge, the impulse stimulates subsequent sections until it reaches a synapse and triggers the release of neurotransmitters, which cross the synaptic cleft and either stimulate a new impulse or inhibit one from firing. This decision is based upon the weighted sum of the impulses received by the neuron. If a fresh impulse is fired, neurotransmitters are released into the synaptic cleft once the impulse reaches the synapse. These neurotransmitters include acetylcholine, norepinephrine, and dopamine, which are packaged in vesicles in the pre-synaptic membrane. The neurotransmitters then combine with specific receptors on the post-synaptic membrane to open the ionic gages and stimulate a fresh impulse across the receiving neuron (Parker, 2007).

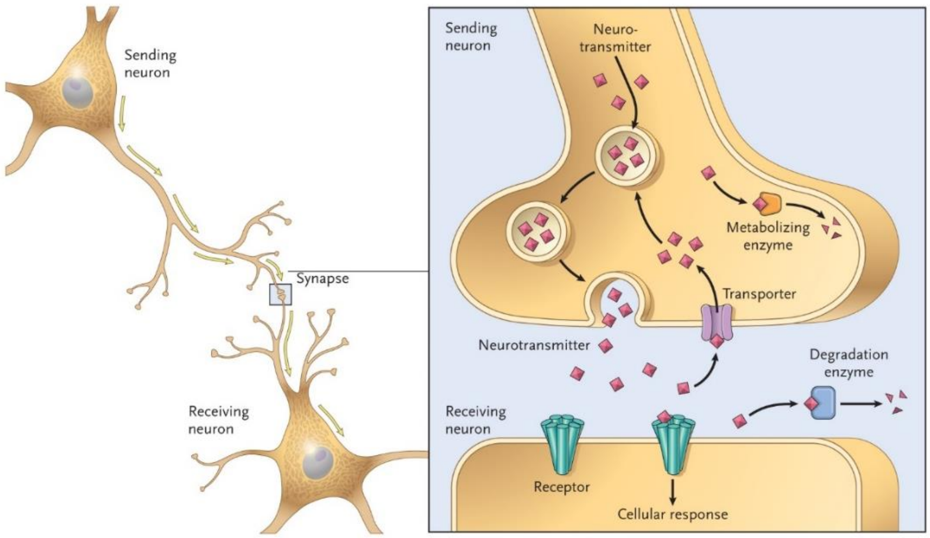

Figure 1: A generic neurotransmitter system. Reprinted from Wikipedia Commons, 2011. Retrieved from https://commons.wikimedia.org/wiki/File:Generic_Neurotransmitter_System.jpg.

The generated impulses carry messages throughout the body's nervous system and are responsible for enacting every decision made by the brain. Peripheral nerves capture sensory information and transmit bioelectric impulses through axons to the command center, which send commands to actuators. These commands, many of which can occur nearly simultaneously, fall under two categories: autonomic and voluntary. Autonomic impulses direct involuntary movements, such as the beating of the heart or the release of adrenaline. Voluntary impulses relate directly to conscious decision-making. They are responsible for activating muscles necessary for movement and speech.

Conditioning processes allow the human nervous system to develop appropriate pathways for impulses to ensure they reach the correct destination, as discussed below.

Basal ganglia. Human conditioning, and thus decision-making, is made possible through basal ganglia. The striatum is the main intake portion of the ganglia; it receives information related to senses, movements, associations, and motivations. Within the striatum, the dorsolateral striatum (DLS), dorsomedial striatum (DMS), and ventral striatum (VS) each play a distinct part in making decisions, as seen in a study by Ito and Doya. These researchers investigated the functions of the DLS, DMS, and VS in decision-making by having rats perform specific tasks at the sound of a tone. When the rats performed the given task correctly, they received a food pellet. The study showed that when the task first began, neuron activity in the VS intensified as the rats explored the options and learned from the results of their actions. When the rats were given the cue to act, the DLS showed increased activity, especially if the command was repeated. When the rats chose an action, the increase of activity in the DMS suggested that the DMS stores information regarding an action and its value, or reward. These results support the authors' hypothesis that the DLS, DMS, and VS hold roles in habitual action, goal-directed action, and motivation, respectively (Ito & Doya, 2015). These areas of the human brain allow individuals to learn about their environment, determine the value of a particular action, and remember the results of that decision for future situations.

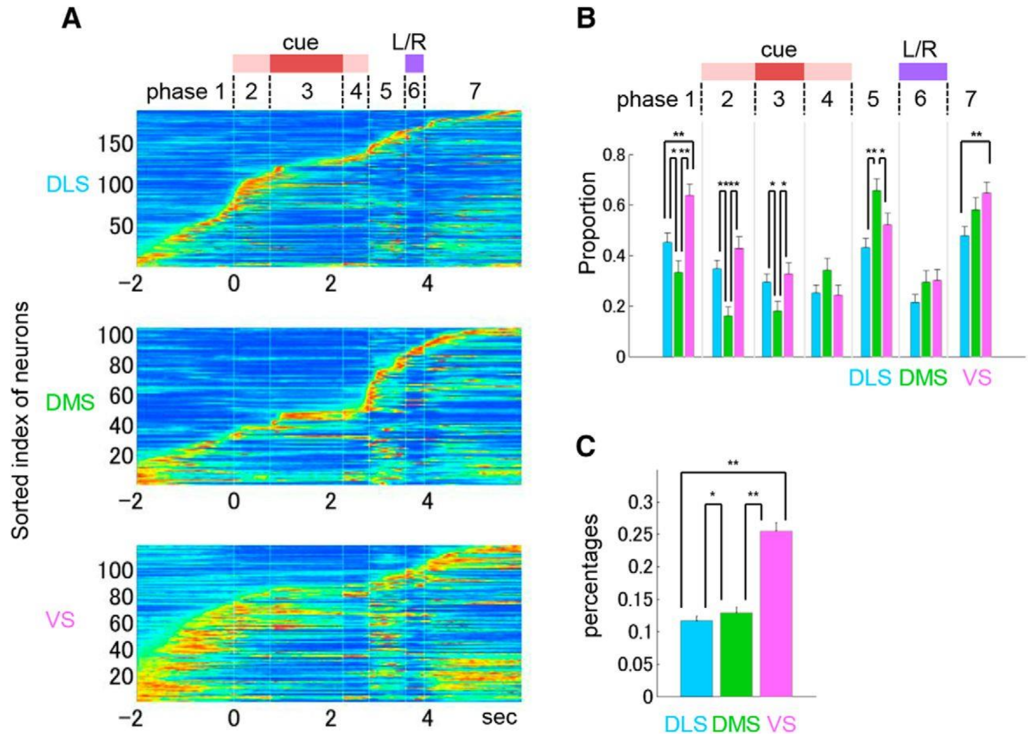

Figure 2: Normalized activity patterns (A), proportion of neurons during trial phases (B), and averaged activity ratio of striatal neurons for each subarea showing the involvement of the DLS, DMS, and VS in the decision-making process. "Figure 4" by Ito and Doya is licensed under CC BY 4.0.

Language.

Areas of the brain. Conditioning, and the neural pathways that result, also enable individuals to learn and apply language according to the context and culture. Multiple structures within the brain are responsible for gathering, storing, and forming language. In the majority of individuals, the left hemisphere handles the majority of language functions. Expressive language is shaped in the posterior inferior frontal gyrus, known as "Broca's area." Within this area lies the Pars triangularis, a region that assumes a role in language semantics. Receptive language is controlled by Wernicke's area, which is connected to Broca's area by a bundle of nerves called the arcuate fasciculus. Damage to this nerve bundle leads to the inability to repeat what another person says. Language information passes through Broca's area and Wernicke's area through the inferior parietal lobule, which is in charge of acquisition and the abstract uses of language. This region allows people to collect and understand words and grammar as well as classify objects using sensory information. Likewise, the fusiform gyrus in the frontal lobe enables word recognition and additional classification. Damage to this area may prevent an individual from recognizing text. Finally, while the right brain is generally the non-dominant hemisphere, it is believed to handle language functions as well. These functions may include comprehending metaphors, recognizing intonation, and understanding poetic meters (Harte, 2018). Both hemispheres and the neural networks within allow individuals to store language and form it into layers of meaning.

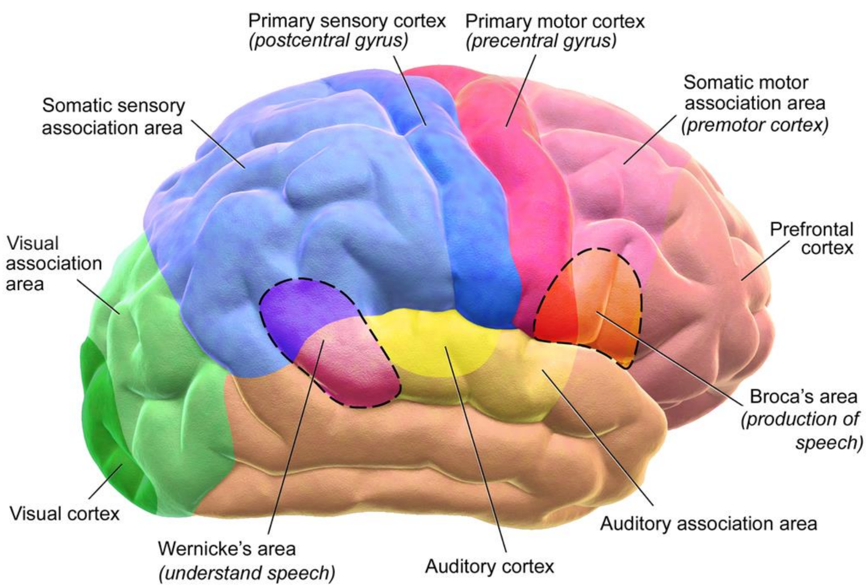

Figure 3: Structural and functional areas of the brain, including those responsible for language. Reprinted from "Medical gallery of Blausen Medical 2014", by Blausen.com staff, 2014. "Motor and Sensory Regions of the Cerebral Cortex" by BruceBlaus is licensed under CC BY 3.0.

Language hierarchy. Human languages contain multiple hierarchical levels.

When learning a language, either from birth or at a later age, humans discover how to form language at each level in the hierarchy until the process of combining morphemes, words, and sentences becomes a fluent process. The basic units of meaning are called morphemes, groups of letters which consist of one or more syllables. At the lexical level, these morphemes combine to form words. The slightest differences in the arrangement of morphemes, whether the addition of a morpheme or a change in morpheme order can significantly impact the meaning of a word. For instance, the morphemes “able” and “ity” combine to form the word “ability”, which assumes a very different meaning with the simple addition of “dis.” Likewise, the syntactical level determines how the order and relationship between words influences a sentence’s meaning. At the semantic level, the meaning of individual words is determined by the context of the words within a sentence. In a properly formed sentence, like “The dog fetched a stick,” a speaker can easily identify the subject and action from the word order, even without knowing the context of the sentence. However, at the discourse level, the meaning of each sentence is determined by its placement among others. For example, the above sentence would not make much sense if followed by a sentence like “Thomas Bayes developed a framework for assessing event probabilities.” When reading the two sentences together, an individual would be unable to ascertain the writer’s intentions. Finally, the pragmatic level includes global knowledge and situational awareness; the meaning of words or sentences is based on likelihood and context within their environment (Varone, Mayer, & Melegari, 2016). Humans are not only able to analyze language at any of these hierarchical levels but can also distinguish between well- and poorly-formed language in order to best ascertain the meaning.

Parallel processing. In addition to decision-making and language, the human brain is extraordinarily capable of parallel processing, the act of performing multiple tasks simultaneously. With billions of neurons, each of which can have 10,000 connections with neighboring neurons, the brain can perform multiple operations, both conscious and subconscious, at once. In contrast to the neuron's thousands of connections, the artificial neurons on current state-of-the-art neuromorphic chips contain just 256 connections. The brain's ability to perform conscious thought while continuously sending and receiving messages throughout the body has been investigated using methods like functional magnetic resonance imaging (fMRI) (fMRI Data, 2014). This imaging allows researchers to monitor activity in parts of the brain through real-time scanning.

FMRI scans track the areas of the brain responsible for specific types of thought by monitoring levels of oxygen in the blood flowing through the brain (MacDonald, 2014). FMRI machines grid the brain into 60x60x30 voxels, or 3D pixels, which allow the activity of the entire brain to be recorded at any instant. The brain's activity can be recorded throughout the duration of a task, resulting in as many as 30 million data points (fMRI Data, 2014).

Harris Georgiou of the National Kapodistrian University sifted through fMRI data to determine the human brain's capacity for parallel processing. Using a signal processing technique called independent component analysis, which he first tested on synthetic data, he ascertained the number of parallel processes in the human brain. Like CPU cores in a computer, these processes are activated as an individual engages in a task. The brain can utilize as many as fifty of these parallel processes at once when a person performs in a complicated visuo-motor activity. In Georgiou’s studies, this activity included watching red or green boxes and holding up a finger in response to the image (MacDonald, 2014). With simple tasks, like indicating when an image appears twice in a series, fewer parallel processes are necessary.

Georgiou’s research showed that parallelism occurs at a level higher than individual neurons; each task is performed by a complex group of neurons working together to achieve a goal. This realization could lead to the development of improved parallel processing in computers; as quoted, “This means that, in theory, an artificial equivalent of a brain-like cognitive structure may not require a massively parallel architecture at the level of single neurons, but rather a properly designed set of limited processes that run in parallel on a much lower scale” (MacDonald, 2014, para. 14). Such brain-inspired processes may soon become prevalent in computer chips, enabling the development of superior parallel processing systems. In this area, as in many others, the human brain has proven to be a source of innovation for AI technology.

Artificial Intelligence

Artificial intelligence was born from a desire to emulate the functions that humans learn to perform from birth. Since its creation, AI has been improved due to knowledge gained from ongoing research of the human brain. However, AI technology does not necessarily use the same techniques as the brain to accomplish a task, as the processes that drive artificial learning must rely on a combination of mathematical models and circuitry. In fact, although introspection and other investigative methods help scientists apply certain human abilities to machines, “machines will often work more efficiently on certain problems when they operate in ways that may seem quite foreign to human reasoning patterns" (Jackson, 2015, p. 8). The goal of AI is to simulate the abilities of human intelligence by reproducing the outward behaviors, not necessarily the methods themselves. As a result, AI has shown impressive progress in a variety of areas; intelligent technology is now capable of solving problems, playing games, recognizing patterns, proving theorems, understanding language, and adapting situational behavior (Jackson, 2015).

A Brief History of AI

The current capabilities of AI are the result of centuries of scientific and technological development. In the mid eighteenth century, Thomas Bayes developed a framework for assessing event probabilities. His technique, known later as Bayesian inference, would become a leading method in machine learning. Less than two centuries later, in 1914, Leonardo Torres y Quevedo demonstrated a chess-playing machine that was capable of applying strategies without human intervention. In 1950, Alan Turing published "Computing Machinery and Intelligence," in which he proposed an idea of assessing a machine's level of humanity through a technique later known as the Turing Test. In 1951, Marvin Minsky and Dean Edmunds designed SNARC (Stochastic Neural Analog Reinforcement Calculator), the first artificial neural network using three thousand vacuum tubes to simulate a network of forty neurons. The next year, Arthur Samuel developed the first self-learning computer program as well as the first checkers-playing program (Press, 2016).

The term "artificial intelligence" itself was coined in a 1955 proposal by Minsky, John McCarthy of Dartmouth College, Nathanial Rochester of IBM, and Claude Shannon of Bell Telephone Laboratories. These men proposed a two-month, ten-man study of AI for an upcoming workshop which was considered the official birthdate of AI. In the December of 1955, Herbert Simon and Allen Newell developed the first artificial intelligence program called the Logic Theorist which would prove thirty-eight theorems of *Principia Mathematica*. In 1957, Frank Rosenblatt developed an artificial neural network called the Perceptron which could recognize patterns (Press, 2016).

The next several decades produced great advances in pattern recognition, natural language processing, heuristics, decision-making, reading music, reasoning, and many other areas. Three years before the turn of the millennium, a computer program called Deep Blue became the first to ever beat a world chess champion. In 2009, Google began developing an autonomous car, which passed its state self-driving test in Nevada in 2014. Likewise, in 2011, a computer named Watson defeated two former champions on the TV show *Jeopardy!* thanks to its ability to process and respond to natural language (Press, 2016). AI continues advancing to this day, thanks to the continued improvement of electrical models and circuitry.

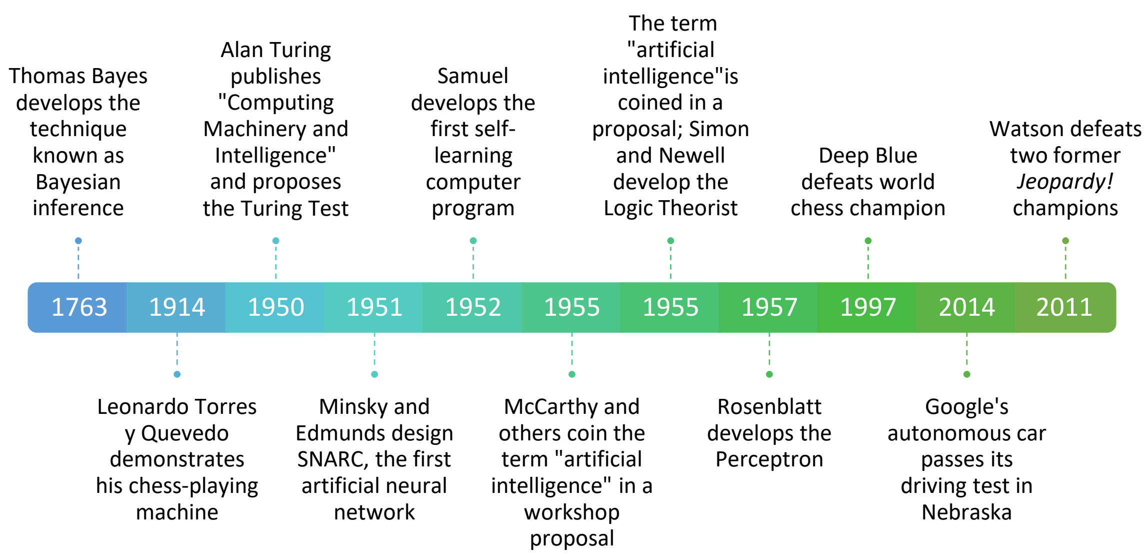

Figure 4: A summary of key advancements in AI.

Key Skills and Techniques of AI

Decision-making. To make decisions, machines rely on a variety of mathematical models that allow them to learn the correct response to a stimulus. The methods for machine learning can be placed into three categories: supervised, unsupervised, and reinforcement learning. In supervised learning, the learning algorithm receives labeled data and the desired outcome. In unsupervised learning, the algorithm is asked to identify patterns in the unlabeled input data. In reinforcement learning, the algorithm receives positive or negative feedback from an interactive, dynamic environment depending upon its level of performance in a task (Internet Society, 2017). Reinforcement learning is considered the optimal method behind decision-making, as it produces information regarding cause and effect. The decision-making body, known as an agent, makes an action and receives this feedback from his environment. From the type or degree of reward received, the agent learns whether the correct decision was made and can improve its future decisions (Ashraf, 2018).

Artificial neural networks. Introduced by McCullock and Pitts in 1943, artificial neural networks are based on the interaction of neurons within the human brain. McCullock and Pitts modeled the neuron as a switch which selects an output according to the sum of weighted inputs (Krogh, 2008). The resulting networks are designed to emulate the brain's natural data-processing methods.

Neural networks provide a decision-making tool that lowers the statistical deviation between the desired values and the result. One of the most popular types of artificial neural networks is the feedforward network, which consists of three layers: the input layer, or data; the hidden layer, which holds process data; and the output layer, or the predicted value. Within the input layer, input nodes represent the data used to fit the model to the target variable. Each independent node is connected to another neuron layer that hold the hidden nodes, which modifies the data using processes obscured from users. The hidden nodes connect to the output layer, which represents the target variable. The virtual pathways between layers are analogous to dendrites and axons within biological neural networks, which carry messages to and from the neuron. Because the connections within biological networks are far more complex, developers of artificial neural networks must view the processes abstractly to incorporate the most relevant and beneficial properties of human neurons (Francis, 2001).

Neural networks can incorporate either supervised or unsupervised learning to lower the statistical deviation between the desired values and the fitted result. In networks trained using supervised learning, a function is used to predict the target variable. Unsupervised learning lacks this target variable; instead, the network finds characteristics within the data and uses these characteristics to group similar data. Neural networks can learn by modifying the input signal strength from connected nodes. The connections between nodes are strengthened as the artificial neuron “memorizes” the correct pathway and learns how to predict the desired value (Francis, 2001).

Within the hidden layer, threshold or activation functions modify the input signal. In the early days of neural networks, these functions resulted in 1s or 0s depending on the sum of the input values. As was believed to occur in human neural networks, the neuron, or node, fired if this sum exceeded the threshold value. However, biological neurons do not necessarily follow this straightforward, yes-or-no process, as scientists have discovered. Using this information, designers of artificial networks base activation functions on sigmoids, so that the results fall on any value between 0 and 1 or -1 and 1, depending on the chosen function. The most common sigmoid function is the logistic function $f(Y) = \frac{1}{1+e^{-Y}}$, which assumes values between 0 and 1. Another common sigmoid function is the hyperbolic tangent $f(Y) = \frac{e^{Y}-e^{-Y}}{e^{Y}+e^{-Y}}$, which results in values between -1 and 1. The selected activation function is performed on the weighted sum represented by the following equation:

$$Y = w_0 + w_1X_1 + w_2X_2 + \cdots + w_nX_n.$$

The modified signal then passes to the output nodes which use activation functions as well. Throughout the process, the pattern being learned is encoded in the signals passing throughout the layers, allowing the relationship between the data and target variables to be mapped (Francis, 2001).

Markov processes. Like artificial neural networks, Markov processes attempt to emulate the human brain's abilities to learn information and make value-based decisions. Markov processes are stochastic learning strategies for automated systems. A Markov process is a sequence of random states [S_1, S_2,…S_t] which each contain the Markov property. This property states that, given the present state or condition, a future state is independent of the past. In other words, once a computer knows the current state, the historical data can be discarded, and the known state characterizes future possibilities just as well as the entire history. Mathematically, state S_t contains this Markov property if and only if the probability of a future state, given the present state, is the same as the probability of that future state given all past states:

$$P[S_t + 1|S_t] = P[S_t + 1|S_1, \cdots, S_t]$$

For each state and its successor, the state transition probability function is described by the following equation:

$$P_{ss'} = \mathbb{P}[S_{t+1} = s'|S_t = s]$$

This distribution represents the probability that an agent will transition to a certain state, given its current state. For instance, if the agent is in state S_2, it determines the likelihood that the agent will pass into S_1, S_3, S_{100}, etc. (Ashraf, 2018).

A Markov Reward Process (MRP) is a Markov process that uses value judgements after determining how much of a reward is gained from a sequence. Mathematically, an MRP is a tuple represented by finite state space S, state transition probability function P, and reward function R:

$$R_s = E[R_{t+1}|S_t = S]$$

which finds how much reward is immediately expected from state S. Likewise, the total discounted rewards are represented by the return G_t in the following equation:

$$G_t = R_{t+1} + \gamma R_{t+2} + \gamma^2 R_{t+3} = \sum_{k=0}^{\infty} \gamma^k R_{t+k+1}$$

The γ term is a discount factor between 0 and 1, inclusive. As time passes, the reward decreases by a factor of γ. This factor tells the agent whether or not to sacrifice an immediate reward for a future reward. The agent’s goal is to maximize the rewards received (G_t) by waiting or accepting rewards as they come (Ashraf, 2018). If γ=0, the agent is short-sighted; it accepts the reward immediately to reap the most profit. When

γ=1, the agent waits as long as possible to receive the maximum reward. This factor is modeled from human cognitive patterns, in which humans prefer an immediate reward over a future one (Ashraf, 2018). However, to receive the most reward, a rational person weighs the benefit of a quick prize over the disadvantage of waiting for a larger profit. For instance, a person is offered \$100 with γ = 0.9, making $G_t = 100 + (0.9)(100) + (0.9)^2(100) + \ldots$. After one unit of time, say an hour, the person would receive \$190; after two hours, \$271, and so on. In this case, the person should decide to wait to receive a greater reward than the initial \$100. Another person is offered \$100 right now with γ = 0 so that $G_t = 100 + (0)(100) + \ldots = \$100$. This person should accept the \$100 up front, since there is no additional reward after waiting any number of hours.

Markov Decision Processes (MDPs) are reward processes that, unlike MRPs, include policies for making decisions. They provide a straightforward technique for framing interactive learning problems, in which an agent selects actions while its environment responds to the actions and provides the agent with fresh situations. Mathematically, MDPs are tuples with a state space (S), a finite set of actions (A), state transition probability function (P), reward function (R), and discount factor (γ). The state transition probability function and reward functions are given by the following equations:

$$P_{ss'}^{a} = \mathbb{P}[S_{t+1} = s' \mid S_t = s, A_t = a]$$

$$R_s^a = E[R_{t+1} \mid S_t = s, A_t = a]$$

True to the Markov property, MDPs do not depend on the history, only the current state. Whenever an agent transitions to a certain state, it performs the action defined by a previously found solution, or policy, represented by π, which is a probability distribution over actions given a state S:

$$\pi(a|s) = \mathbb{P}[A_t = a|S_t = s]$$

Likewise, value function v_π finds the worth of being in state S while following policy π:

$$v_\pi(s) \doteq \mathbb{E}_\pi[G_t \mid S_t = s] = \mathbb{E}_\pi\left[\sum_{k=0}^{\infty} \gamma^k R_{t+k+1} \,\middle|\, S_t = s\right], \text{ for all } s \in \mathcal{S}$$

This function finds the expected reward for following a particular policy in a state. Finally, the action-value function q_π produces the expected reward if an agent begins in state S and takes action a from policy π:

$$q_\pi(s,a) \doteq \mathbb{E}_\pi[G_t \mid S_t = s, A_t = a] = \mathbb{E}_\pi\left[\sum_{k=0}^{\infty} \gamma^k R_{t+k+1} \,\middle|\, S_t = s, A_t = a\right]$$

(Ashraf, 2018). In short, these functions for MDPs attempt to capture the environment in which an agent makes choices in order to find an optimal solution that maximizes the rewards received.

Artificial language. As in artificial decision-making, AI relies upon the development of models and processes in order to approximate human language. Human language technology (HLT) can be defined as the methods, programs, and devices used to analyze human speech. The goal of HLT is to ease the interaction between computers and users by providing language-related services, such as spell-checking, automatic response generation, and assistance through dialogue. HLT can be broken into two separate areas: natural language processing (NLP) and computational linguistics (CL).

Natural language processing. Natural language processing can be defined as the relationship between human language and computers. NLP is an emergent form of AI that analyzes human language to allow devices to form conclusions, predict user commands, and perform other tasks based on contextual patterns (Mills, 2018).

Three types of approaches can be applied to an NLP system: symbolic, statistical, and connectionist. The symbolic approach records generally accepted rules of human speech and has them entered into a computer for analysis. In the statistical approach, large textual samples are run through statistics-based models to recognize recurring themes. After mathematically identifying trends, the system develops speech rules in order to improve future analyses and even generate its own output. The connectionist approach combines the symbolic and statistical approaches by adjusting speech rules for specific applications based on decisions that have been statistically inferred (Varone et al., 2016).

The statistical approach became popular after the development of hidden Markov models (HMMs) in the 1980s and 90s due to the influx of speech data and machine-readable text. The approach was – and is – designed to overcome scalability issues and has been applied in a variety of speech technology such as speech recognition systems, part-of-speech taggers, parsers, and MT and QT summarization systems (Shubert, 2014). As the preferred method of speech recognition and a subset of the Markov decision process discussed in a previous section, HMMs are a flexible, powerful means of modeling word sequences. These models estimate the likelihood that a query is related to a specific topic depending on the contextual environment. HMMs also provide a paradigm for estimating model parameters and offer the possibility of more powerful models. They are proficient in modeling noisy sequences where the intended sequence is obfuscated by random data, like static received during a phone call.

HMMs use the Bayes decision theory to choose among multiple possibilities in NLP applications. Once provided data, these models compute the probability that each possibility is true, then select the most probable possibility. Although the computations rarely result in the correct possibility, parameters can be estimated from the given data to create more accurate models (Chou & Juang, 2003). These parameters allow the models to be "trained" and used to analyze new data using dynamic programming algorithms like the Viterbi algorithm (Shubert, 2014). Due to this inherent flexibility, HMMs can successfully adapt to specific situations, especially in dialogue applications that have a wide variety of input data.

Computational linguistics. Computational linguistics combines the fields of computer science and linguistics using knowledge gained from philosophical logic, cognitive science, theoretical linguistics, and computer science. Though the term "computational linguistics" is frequently interchanged with NLP, the former specifically addresses processes intended to discover linguistic facts, while NLP focuses on developing practical technologies. As such, CL provides a linguistic background for NLP technology (Varone et al., 2016).

Computational linguistics is applied to both theoretical and practical applications. In the realm of theory, CL formulates the grammatical and semantic frameworks that characterize languages in order to enable computational implementation. Likewise, it creates computational models of human language processing and develops techniques and principles that utilize the structural and statistical properties of language. More practically, CL has applications in effective machine translation (MT), question answering (QA), analysis of texts or speech for topics, and human-like dialogue (Shubert, 2014).

Computational linguists help machines process natural language by formulating models for simulations. In the early years of CL, from the mid-1950s to 1970, linguists focused on practical solutions for MT and QA. For MT, they worked to characterize sublanguages for particular topics, such as weather or factories. In QA, the issues involved characterizing question patterns for these topics, and forming relationships between the question and its answer stored in a database. In the 1960s, linguists began gathering information from the human brain itself in attempts to understand genuine dialogue. J. Wizenbaum's ELIZA program, for example, was said to emulate a known psychiatrist; it matched user inputs to stored word sequences and returned the corresponding output just as humans respond to specific social situations with "preprogrammed" phrases. ELIZA is now regarded as an ancestor of the chatbot. In 1968, M. Quillian proposed a semantic memory model that attributed word sense disambiguation to spreading activation within a relational network of concepts. R. Schank's work in the same time frame recognized that understanding language depends heavily on background knowledge. Since the 1970s, CL has trended away from procedural approaches to more modular and reusable ones, enabled by syntactic-semantic frameworks. These frameworks focus on the computational tractability of parsing and mapping syntax to semantics (Boyd, 2018). As a result, linguistic technology is becoming more flexible and, in turn, adaptable to a wider variety of grammars and languages.

Multitasking. Like decision-making and language processing, the skill of multitasking within artificially intelligent systems is being sharpened by ongoing research into the workings of the brain. In the context of AI, multitasking is defined as a machine's ability to perform various tasks at once with limited or no human input (Tweedale, 2013). This skill is implemented through Multi-Agent Systems (MASs) which seek to emulate the human capacity for data collection and problem-solving. Agent-based systems are described as a "new paradigm for conceptualizing, designing, and implementing software systems" (Carnegie Mellon University, 2012, para. 1). Each agent is a sophisticated program that acts independently on the user's behalf. In multi-agent systems, these agents form networks that solve problems more complex than can be solved by a lone agent (Tweedale, 2013).

In a MAS, each agent's perception and decision is represented in every action taken. The actions taken by one agent affect both its own environment and its neighbors', meaning that each agent must be able to predict the action of the others in order to take the best option for the situation. In distributed systems, the agents are located at separate points around the environment; as each agent's sensors have a limited range, a MAS must base its decisions upon the partial observations made by its agents, meaning that its final decision may be less than optimal. Thus, it is essential to quantify how much information each agent must collect in order to produce a satisfactory model of the environment (Balaji & Srinivasan, 2010).

The agents within MASs can be a variety of programs, each of which has its own goals and methods. In such heterogeneous networks, these individual desires can conflict, making a conflict-resolution mechanism essential. An established hierarchy can help ensure the ultimate goal of the network is met by delegating authority to one or multiple agents. Within a hierarchical structure, data from the lower levels flows upward to levels with higher authority while control data is transmitted from high to low levels. In a simple hierarchy, a single agent is granted the authority to distribute instructions to other members, whereas in a uniform hierarchy, authority is distributed among all agents to increase efficiency (Balaji & Srinivasan, 2010).

When a conflict of interest arises between agents, namely in heterogeneous networks, information is exchanged until the system reaches a consensus algorithmically. The Paxos algorithm, proposed in 1988 by L. Lamport, implements a fault-tolerant distributed system using a synod consensus algorithm. Within the Paxos algorithm, each agent is assigned a role as a proposer, acceptor, or learner. If the majority of nodes are available in a system, the algorithm ensures that a single decision will be reached based on the agents' roles. Likewise, the average-consensus algorithm is an adaptive distributed algorithm that requires communication only among neighboring agents. Every agent receives an initial value. When a consensus must be reached, an average is taken of all available agents' initial values and is used to make a final decision (Xie & Liu, 2017). Using this technique, a machine can emulate the human ability to receive information from multiple sources and decide an appropriate outcome.

Man vs. Machine

Challenges for Mankind

Multitasking. To ascertain the extent of progress in artificial intelligence, the abilities of AI are compared to those traditionally ascribed to human intelligence. Due to the brain's amazing capacity for acquiring, storing, and applying knowledge – even so far as to attempt replicating this intelligence artificially – humans are regarded as masters over technology. However, when compared to technology, human intelligence does possess certain limitations, one of which is the inability to multitask. While the brain possesses an innate ability for parallel processing, its efforts to multitask often harm performance. The American Psychological Association defines multitasking as an attempt to perform multiple tasks simultaneously, switch between tasks, or perform two tasks in rapid succession. Evidence suggests that human control processes have two distinct stages which allow people to shift between tasks without conscious effort. These stages include goal shifting – moving to a new project – and rule activation – transitioning from the rules of one task to those of another task. Whenever the brain performs such a shift, a switch cost is paid. These switch costs can be as low as a few tenths of a second per switch, but the total cost can be significant if a person switches repeatedly between tasks. According to Meyer, the mental blocks caused by task shifting can demand as much as 40% of an individual's time spent multitasking. For this reason, researchers recommend that people avoid multitasking, especially in complex tasks (American Psychological Association, 2006). Losing focus due to multitasking does not only lower productivity on a project; it can even play a role in causing accidents while performing activities like driving a vehicle or preparing food. For this reason, research into multitasking and its hidden costs is being performed and studied by everyone from education experts to legislators.

Machine control. The increasing complexity of computers has led to the challenge of ensuring humankind can retain control over their creations. Jackson (2015) writes, "It is not difficult to envision actualities in which an artificial intelligence would exert control over human beings, yet be out of their control" (p. 395). How are humans meant to ensure that the machines benefit humanity? To address this question, Asimov presented his Three Laws of Robotics in 1950. First, a robot may not harm or allow harm to befall a human being. Second, a robot must obey human orders except when the instructions conflict with the former law. Third, a robot must protect its existence as long as it does not violate laws one or two (Jackson, 2015).

With the advent of increasingly intelligent machines, controlling a machine may bring the added challenge of checking a machine's motives. The computer's thought process can quickly become too complex for human understanding, especially in fast-paced environments. Even if a user demanded real-time analysis of a computer's decision-making, the human may not reach an effective understanding before the computation's time limit is exceeded. In light of this, a scientist may opt to program other machines with the sole duty of verifying the decision process of the first. However, the users would need to investigate any side effects experienced by these reasoning checkers (Jackson, 2015). In other words, the more that humans might try to analyze an artificially intelligent machine, the more their methods might skew the result. Thus, people may never be sure if their machines' motives are as pure as desired.

Limitations of AI

Contextual decision-making. One weakness of machines is that they lack the ability to glean peripheral details from their environment which would make a certain decision more relevant than another (Dreyfus, 1965). Humans undergo this gleaning process subconsciously due to a phenomenon known as fringe consciousness. William James coined this term to address humans' background awareness of a change in environment, such as noticing the sudden lack of birdsong in a forest or vaguely recognizing the faces in a crowd when searching for a particular face (Dreyfus, 1965).

This fringe information provides context until an individual subconsciously perceives data of interest.

Fringe consciousness gives humans an advantage over machines by limiting the amount of pertinent data for which the human's mind is consciously responsible. In a game of chess, the player's focus is drawn to an area that used to lie in the fringe of his consciousness, but which has been made interesting by other areas within the fringe (Dreyfus, 1965). In fact, according to Newell, Shaw, and Simon, a player considers far less than 100 positions when analyzing a move, though he still may analyze other moves before making his final decision. Without needing to consciously examine each of the opponents' pieces and analyze his own defense until he finds a potentially effective move, a human player may simply look at the board and "zero in" on an undefended piece. Once he notices a possible weakness, he then begins to analyze future situations, or count out explicit alternatives. Conscious counting is used only when a human must refine a general process in order to analyze the details, once he has zeroed in on a single area of interest from among those contained within his fringe consciousness (Dreyfus, 1965).

Language differences. Wide variations among languages – and even within a single language – significantly challenge the ability of computers to comprehend and reproduce human language. Although humans have developed formal grammars for most languages, these systems do not cover every possibility or variation within the language; regardless, most humans can automatically distinguish between well-formed and poorly-formed grammar. Humans use their intuition and gained experience in a language to process colloquialisms, different dialects and accents, flaws like stuttering or backtracking, and written errors like run-on sentences and misspellings.

As computers gain the ability to understand and reproduce language, they face similar linguistic challenges. Unlike humans, however, computers lack a person’s skills – gained by years of experience – to automatically re-format and extrapolate meaning from language so that it can be better understood. To simulate this ability in computers, linguists work to create grammars that allow flexibility in language while still accurately representing the root language. Typically, linguists accomplish this feat by implementing a high number of language rules – as many as tens of thousands. Each rule is programmed into the computer by having the computer read annotated transcripts written by the linguist. As a result, a flawed phrase or sentence, which humans process every day, would simply require more uncommon rules to be used by the computer in its analysis (Shubert, 2014). The significance of this development is that computers are becoming better at recognizing and compensating for mistakes in language, leading to more human-like comprehension.

The Singularity

As discussed, AI is developing in its ability to emulate traditionally human skills. Although humans still have the advantage in processing language, gleaning fringe information, and experiencing a wide variety of perceptions, emotional and otherwise, AI technology is quickly improving due to advancements in scientific research and modern electronics. Due to AI’s rapid progress, proponents of the phenomenon known as the Singularity suggest that this particular tool will bring about humanity’s final destruction. The Singularity is described as a “self-stoking cycle of machines using their own AI to make even smarter machines" (Van Paasschen, 2017, para. 6). The foremost proponent of the Singularity, Ray Kurzweil, fears that AI will soon reach and surpass human capacity. As Stephen Hawking once said, "The development of full AI could spell the end of the human race" (Van Paasschen, 2017, para. 1).

Mathematically, a singularity is a point at which the exact properties of an event cannot be deciphered. According to Kurzweil, the same situation may apply to a point in time as near as 2045, the year he predicts that machines will surpass human intelligence (Tzezana, 2017). This is the basis of Kurzweil's prediction: since technology is increasing at an exponential rate, its progress will become virtually instantaneous and result in a singularity (Satell, 2016), the effects of which humankind has no hope of predetermining.

Kurzweil's arguments.

Moving beyond Moore's law. Kurzweil believes that humanity is moving beyond Moore's law, which predicted that the number of transistors on a microchip would double every 18 months. The exponential curve is now approaching its theoretical limit, since transistors can only shrink so much before quantum effects at the atomic level inhibit function. In fact, Moore's law may become virtually irrelevant as early as 2020. However, Kurzweil argues that just as microprocessors replaced legacy technology like vacuum tubes and electro-mechanical relays, something else will replace human dependency on chip technology. After all, "the numbers of transistors on a chip is a fairly arbitrary way to measure performance" (Satell, 2016, para. 6). Microprocessors are merely the fifth paradigm of information processing; the sixth, Kurzweil suggests, should be the number of calculations per $1000. Indeed, scientists are finding more ways to improve performance through such techniques as quantum computing and neuromorphic chips (Satell, 2016), meaning that the progress on AI will continue to accelerate for the foreseeable future.

Employing robots. Machines are being substituted for human workers in an increasing variety of occupations, from bank tellers to military personnel. The first industrial robots appeared on the GM assembly line in 1962. Now, servicemen are developing emotional attachments to robots that are becoming popular in the battlefield. Also, while machines are starting to take over low-skill jobs, machines are being incorporated into more creative areas like art, music, and literature – in fact, a book written by AI was recently submitted for the Hoshi Shinichi Literary Award in Japan. This trend of automation will not end any time soon. The Department of Defense has begun experimenting with chips that are embedded in soldiers' brains to improve performance and field monitoring (Satell, 2016). Machines will replace more and more human workers in the decades to come as technology continues to advance and will simultaneously threaten and benefit humanity.

Counterarguments. Though proponents like Kurzweil fear that the Singularity will bring about humanity's doom within the next few decades, this event may prove to be unrealistic due to the intrinsic properties of human consciousness and AI itself. The selected counterarguments of the Singularity use the ideas of existential risks, machines' lack of free will, reciprocal understanding, and Descartes' error to demonstrate the inability of AI to comprehend human intelligence, much less surpass it. As Popoveniuc

(2013) states, “All these contentions…prove one thing: it is unlikely to have a Singularity as the result of a manufactured superior AI” (p. 5).

Existential risks. The argument of existential risk uses the assumption that artificially intelligent systems are designed to perform helpful services for humankind. Since AI machines are a human creation, there should be no reason to assume that they would engage in behaviors contrary to their programming. Unfortunately, this argument does not consider the presence of flaws in any human-made system that could allow a machine to take actions beyond the scope of its creators’ desires: “For example,…we tell it to solve a mathematical problem, and it complies by turning all the matter in the solar system into a giant calculating device, in the process killing the person who asked the question” (Popoveniuc, 2013, p. 4). Though highly extreme, this quote emphasizes the risks associated with any error in the design of an AI system. Ultimately however, any evolvement of AI would be human-caused, either directly or indirectly, meaning that the actuality of self-creating AI –an idea central to the Singularity theory – would be highly implausible.

Lack of agency. The lack of agency counterargument given by Lynne Baker asserts that machines are not true agents since they lack free will. According to this argument, in order to be an agent, an entity must be able to formulate its intentions. But to formulate its intentions, the entity must hold a first-person perspective. Since machines lack this perspective, they are not actually agents. Therefore, machines cannot perform behaviors intentionally; they cannot lie, act malevolently, or even use language to make assertions, since they lack a will of their own (Popoveniuc, 2013).

Reciprocal misunderstanding. The counterargument of reciprocal misunderstanding uses humanity's lack of understanding of intelligence to assert that machines would have even less of an understanding of conscious beings. First, people lack an understanding of their own intellectual background. According to Popoveniuc (2013), "we don't understand the real origin and the requirements of our own culture, which manifests itself through us as its agents" (p. 3). Secondly, humans are unable to comprehend the intricacies of ultra-logical processes within machines. Since the creators of AI cannot fully understand either human or artificial intelligence, a machine's comprehension of consciousness would be even poorer. As quoted by Popoveniuc (2013), "If we fail to understand what a conscious being is and how it functions, even more so an AI, AI+, AI++…will be unaware about other ways of thinking and completely lack of understanding about these" (p. 4). It is unlikely that any form of artificial intelligence could ever comprehend the many levels of consciousness within human beings.

Descartes' error. A crucial counterargument to the Singularity is based upon Descartes' error: the tendency of researchers to focus on logic and reasoning while neglecting the emotions and affectivity of human intelligence. Studies show that "reason alone is insufficient for the proficient operation of the human intellect" (Popoveniuc, 2013, p. 4) and that emotions and will are essential in a full first-person perspective of the world. Thus, an AI system that is fully rational would not work perfectly in real environments, in which humans' emotional biases and even genetic preprogramming are essential for making decisions (Popoveniuc, 2013). Due to the ramifications of Descartes' error, without such subjective components, AI could never match human intelligence.

Conclusion

Over the last several decades, artificial intelligence has developed to the point where machines can emulate certain human skills using a combination of mathematical knowledge and electrical systems. The development of techniques like artificial neural networks and MDPs for making value-based choices, HLT systems for processing language, and multi-agent systems for allocating tasks, has allowed machines to approximate natural processes of the brain. As a result, machines are becoming increasingly incorporated in traditionally human occupations, from assembly line work to the military. Kurzweil's Singularity theory predicts that AI's progress will continue to accelerate to the point of self-development, leaving human intelligence in the dust. However, critics argue that computers lack the programming, sense of will, understanding, and non-rational components of human consciousness necessary to surpass its creators. Though the prospect of a Singularity is equal parts thrilling and terrifying, these counterarguments suggest that it is not as likely as its key proponent suggests.

References

Abbott, D. (2016, September 25). *What brain regions control our language? And how do we know this?* Retrieved from The Conversation: http://theconversation.com/what-brain-regions-control-our-language-and-how-do-we-know-this-63318

American Psychological Association. (2006, March 20). *Multitasking: Switching costs*. Retrieved from American Psychological Association: https://www.apa.org/research/action/multitask.aspx

Ashraf, M. (2018, April 7). *Reinforcement Learning Demystified: A Gentle Introduction*. Retrieved from Towards Data Science: https://towardsdatascience.com/reinforcement-learning-demystified-36c39c11ec14

Ashraf, M. (2018, April 10). *Reinforcement Learning Demystified: Markov Decision Processes (Part 1)*. Retrieved from Towards Data Science: https://towardsdatascience.com/reinforcement-learning-demystified-markov-decision-processes-part-1-bf00dda41690

Balaji, P., & Srinivasan, D. (2010). An Introduction to Multi-Agent Systems. In D. Srinivasan, & L. C. Jain, *Innovations in Multi-Agent Systems and Applications - 1* (pp. 1-27). Berlin: Springer-Verlag.

Boyd, C. (2018, January 10). *The Past, Present, and Future of Speech Recognition Technology*. Retrieved from The Startup: https://medium.com/swlh/the-past-present-and-future-of-speech-recognition-technology-cf13c179aaf

Carnegie Mellon University. (2012). *Multi-Agent Systems*. Retrieved from Intelligent Software Agents : https://www.cs.cmu.edu/~softagents/multi.html

Chou, W., & Juang, B.-H. (2003). *Pattern Recognition in Speech and Language Processing*. CRC Press.

Dreyfus, H. L. (1965). Alchemy and Artificial Intelligence. Santa Monica, CA: RAND Corporation. Retrieved January 5, 2019, from Rand Corporation: https://www.rand.org/pubs/papers/P3244.html

fMRI Data Reveals the Number of Parallel Processes Running in the Brain. (2014, November 5). Retrieved from MIT Technology Review: https://www.technologyreview.com/s/532291/fmri-data-reveals-the-number-of-parallel-processes-running-in-the-brain/

Francis, L. (2001). The Basics of Neural Networks Demystified. *Contingencies*. Retrieved from Semantics Scholar.

Harte, E. (2018, May 17). *How Your Brain Processes Language*. Retrieved from Brain World: https://brainworldmagazine.com/how-your-brain-processes-language/

Internet Society. (2017, April 18). *Artificial Intelligence and Machine Learning: Policy Paper*. Retrieved February 18, 2018, from Internet Society: https://www.internetsociety.org/resources/doc/2017/artificial-intelligence-and-machine-learning-policy-paper

Ito, M., & Doya, K. (2015, February 25). Distinct neural representation in the dorsolateral, dorsomedial, and ventral parts of the striatum during fixed- and free-choice tasks. *Journal of Neuroscience, 35*(8), 3499-3514. Retrieved from http://www.jneurosci.org/content/35/8/3499#sec-18

Jackson, P. C. (2015). *Introduction to artificial intelligence* (2 ed.). New York: Dover Publications, Inc.

Krogh, A. (2008). What are artificial neural networks? *Nature Biotechnology, 26*(2), 195-197. doi:10.1038/nbt1386

Levin, E., Pieraccini, R., & Eckert, W. (1997). Learning dialogue strategies within the Markov decision process framework. *IEEE Workshop on Automatic Speech Recognition and Understanding Proceedings* (pp. 72-79). Santa Barbara, CA: IEEE.org. doi:10.1109/ASRU.1997.658989

Licklider, J. (1960, March). *Man-Computer Symbiosis.* Retrieved from IRE Transactions on Human Factors in Electronics: https://groups.csail.mit.edu/medg/people/psz/Licklider.html

MacDonald, F. (2014, November 14). *This Is How Many Tasks Your Brain Is Processing Right Now*. Retrieved from Science Alert: https://www.sciencealert.com/scientists-have-worked-out-how-many-tasks-our-brains-are-processing-at-once

Mayfield Brain and Spine. (2019). *Anatomy of the Brain*. Retrieved from MayfieldClinic.com: https://mayfieldclinic.com/pe-anatbrain.htm

Mills, T. (2018, July 2). *What Is Natural Language Processing And What Is It Used For?* Retrieved from Forbes: https://www.forbes.com/sites/forbestechcouncil/2018/07/02/what-is-natural-language-processing-and-what-is-it-used-for/#77b7d0405d71

Newell, A., Shaw, J. C., & Simon, H. A. (1959). *The processes of creative thinking.* Santa Monica, CA: RAND Corporation.

Parker, S. (2007). *The human body book.* New York: DK Publishing.

Popoveniuc, B. (2013). Pro and cons singularity: Kurzweil's theory and its critics. *ACM International Conference Proceeding Series*.

Press, G. (2016, December 30). *A Very Short History Of Artificial Intelligence (AI)*. Retrieved from Forbes: https://www.forbes.com/sites/gilpress/2016/12/30/a-very-short-history-of-artificial-intelligence-ai/#5bb2791e6fba

Satell, G. (2016, June 3). *3 Reasons To Believe The Singularity Is Near*. Retrieved from Forbes: https://www.forbes.com/sites/gregsatell/2016/06/03/3-reasons-to-believe-the-singularity-is-near/#2fec5ac67b39

Shubert, L. (2014, February 26). *Computational Linguistics*. Retrieved from The Stanford Encyclopedia of Philosophy: https://plato.stanford.edu/entries/computational-linguistics/#GoaComLin

Tweedale, J. W. (2013). Using Multi-Agent Systems to Pursue Autonomy with Automated Components. *17th International Conference in Knowledge Based and Intelligent Information and Engineering Systems* (pp. 1369-1378). Elsevier B.V. doi:10.1016/j.procs.2013.09.226

Tzezana, R. (2017, March 3). *Singularity: Explain It to Me Like I'm 5-Years-Old*. Retrieved from Futurism: https://futurism.com/singularity-explain-it-to-me-like-im-5-years-old

Van Paasschen, F. (2017, January 16). *The Human Brain vs. Computers*. Retrieved from Thrive Global: https://journal.thriveglobal.com/the-human-brain-vs-computers-5880cb156541

Varone, M., Mayer, D., & Melegari, A. (2016, May 17). *Computational linguistics for today and tomorrow*. Retrieved from Expert System: https://www.expertsystem.com/computational-linguistics-today-tomorrow/

Varone, M., Mayer, D., & Melegari, A. (2016, July 26). *What is Natural Language Processing?* Retrieved from Expert System: https://www.expertsystem.com/natural-language-processing/

Xie, J., & Liu, C.-C. (2017). Multi-agent systems and their applications. *Journal of International Council on Electrical Engineering, 7*(1), 188-197.