Cognitive Computing: Architecture, Technologies and Intelligent Applications

Table of Contents

Abstract

This report is designed to inform the reader about the history, implementation, application, and consequences of cognitive computing technologies. Such cognitive computing technologies include machine learning, expert computer systems, and artificial intelligence algorithms. The paper is organized into three sections and an appendix.

The first section broadly introduces cognitive computing to the reader. It is organized into four subsections. The first of which explains the origins of cognitive computing in a field of biology called "cognitive neuroscience", and frames cognitive computing as a means by which to emulate human and animal thought processes digitally. The following sections introduce the core technologies behind cognitive computing in chronological order of invention starting with artificial intelligence and ending with machine learning.

The second section presents examples of industrial uses for the technologies mentioned in section one. It is organized into three subsections. The first explains how cognitive computing can increase the efficiency of quantitative analysis and numerical determine patterns more effectively than humans. GPS algorithms and the uses of machine learning in pandemic epidemiology are discussed in order to illustrate how numerical data is used by cognitive computing programs and the impact of such programs on the human condition.

The second subsection delves into how computers can be made to understand qualitative or opinion-driven data by expressing it numerically or graphically. Medical diagnosis systems and the mining and analysis of qualitative data by marketing professionals are used as case studies to illustrate the process of qualitative data analysis.

The third subsection explains how advanced cognitive computing technologies are beginning to make their way into consumer products and, in many cases, improving people's lives. Intelligent wireless technologies and smart assistant programs are discussed within this subsection.

The third section deals with the consequences incurred by various stakeholders through the use of cognitive computing technologies. Various hypotheticals are also consulted to theorize about possible consequences for future adoption of such technologies. This section is divided into three subsections by stakeholder.

The first subsection deals with costs and benefits to corporations and corporate customers in business to consumer (B2C) and business to business (B2B) transactions respectively. Issues related to profit, productivity, and technological integration are covered in this subsection.

The second subsection mentions costs and benefits to employees of large companies, small businesses, and corporations involved in the making of cognitive computing technologies. Issues related to employment, automation, and wages are covered in this subsection.

The third subsection deals with costs and benefits to consumers.

The Appendices section contains a chart with external resources per the project requirements, a table of figures, and a table of references.

Part 1: An Introduction to Cognitive Computing Disciplines

What is Cognitive Computing?

In this subsection, the foremost tenets of cognitive computing will be discussed. The section starts with an overview of cognitive neuroscience, the field from which cognitive computing draws inspiration. This will be followed by a brief overview of what cognitive computing means to accomplish and an introduction to three main branches of cognitive computing.

An Introduction to Cognitive Neuroscience

In order to grasp what is meant by "cognitive computing", one must first develop familiarity with a branch of science called "cognitive neuroscience." The field of cognitive neuroscience lies somewhere between traditional neuroscience, the study of the tissues found in the human brain and their structure, and psychology, the study of the human psyche or mind. According to an entry in the International Encyclopedia of Social and Behavioral Sciences by J. McClelland, cognitive neuroscientists seek to discern the origin of human thought and its processes from a neuroscience perspective (McClelland, 2001).

Ergo, cognitive neuroscience attempts to determine how large brain structures, such as the hippocampus or the prefrontal cortex, and small brain structures, such as neurons, lead to human thought by means of chemical and electrical signals (McClelland, 2001). One might ask "How does the study of the human brain, something we do not fully understand, contribute to the field of computing?" or "Computers are far simpler structures when compared to the human brain. How could they benefit from this knowledge?", and the answer lies in simplification of known processes in the brain.

Mimicking Human Thought

A computer performs its task in much the same way a human brain does. Coded signals are sent through and processed by the computer as electricity according to an article published by E. H. Chudler of the University of Washington (Chudler, 2008). Chudler states that, allthough a human brain also produces and processes chemical signals, and the electrical signals it produces are used differently than those inside of a computer, one can plainly observe the similarities (Chudler, 2008). A comparison between the human brain and a modern computer can be found at this link containing Chudler's insights (Chudler, 2008).

According to a study by Crapo et al, cognitive neuroscience enables simplification and modeling of the processes in a human brain (Crapo et al, 2000), so early cognitive computer scientists proposed that it would be possible to enable a computer to employ reason, solve problems, and learn as humans do. They were correct, and computers can now perform these tasks through means of various cognitive computing techniques.

The first human-like skill computers were able emulate was the ability to solve simple problems without being explicitly told what to do. This gave rise to artificial intelligence agents. According to Darrell West of the Brookings Institution, Artificial intelligence agents are programs designed to perform a specific task using a general algorithm simulating human thought (West, 2019). An example of a possible task that could be inferred from an article by R.

E. Korf would be finding the shortest path through a maze (Korf, 1988). The second skill computers were given was the ability to employ knowledge-based reason as humans do. A type of program known as an expert system was developed around this concept. According to a web article on the popular website geeksforgeeks, expert systems are composed of heuristic rules developed by field experts and are designed to help make decisions using a context and field-specific knowledge base (Aneja, 2018). The third skill given to computers was the ability of to learn. Based on an article in the Journal of Pharmaceutical and Biomedical Analysis, this feat was enabled through the development of simulated neurons and the use of complex probability algorithms (Agatonovic-Kustrin and Beresford, 2000).

Artificial Intelligence

In this subsection, the basic principles of simple artificial intelligence agents will be explained. An introduction to the concept of artificial intelligence will be presented followed by discussions of intelligent agents and path finding algorithms.

What is Artificial Intelligence?

Artificial intelligence programs are programs that utilize intelligent agents and thought-based algorithms to solve specific instances of a general problem. Often these problems are framed mathematically in order to allow the computer to interpret them. For example, a road map can be expressed as an adjacency matrix of intersections or points of interest. At this link, a YouTube video can be found that explains AI and its sub-disciplines (Ramesh, 2017). It is worth noting that Artificial intelligence is loosely defined and that the views and discipline organizations shown in the video are not a general consensus by the computer science community at large.

Intelligent Agents

According to Charles Dyer of the University of Wisconsin Madison, intelligent agents are pieces of software that observe their surrounding environment and user input using sensors and act based on their observations in an attempt to accomplish a given task (Dyer, 2003). In purely software-based agents, this means observing the problem at hand rather than a physical environment. Actuators and sensors in the case of such agents are simply functions designed to conduct observation of the problem and cause the agent to perform an action. Dyer mentions various types of intelligent agents such as reflex agents, goal-based agents, and utility based agents and their best-case uses (Dyer, 2003).

Simple reflex agents simply observe a condition and perform a predefined corresponding action (Dyer, 2003). For example, a robotic vacuum may observe that there is dirt in its path and start cleaning, whereas if there is not dirt, it may change paths. Internal states in reflex agents make them more efficient by remembering past events (Dyer, 2003). For example, the automated vacuum may decide which path to change to based on a stored database of when various paths were last cleaned.

Goal-based agents, appropriately, work to reach a given goal state says Dyer (Dyer, 2003). This is a more proactive solution, but is also more complicated than a reflex agent. Utility-based agents are similar to goal-based agents, but utilize the principle of utilitarian good to make decisions regarding the goal state (Dyer, 2003).

Path Finding Algorithms

Many artificial intelligence problems can be expressed using graphs. According to a paper by R. E. Korf, many intelligent agents rely on special algorithms related to finding efficient pathways through these graphs (Korf 1988). Such algorithms include A* search and iterative deepening search. Dijkstra's, Kruskal's, and Prim's minimum spanning tree (MST) algorithms are also used often. For example, path finding algorithms are used in the production of the automated directions seen in modern GPS systems. A GPS treats a roadway as a graph connecting various nodes (intersections). Graph algorithms are utilized to find the shortest and least impeded route to one's destination to ensure minimal time is spent driving.

Expert Systems

In this subsection, the fundamentals of expert systems will be discussed. This includes the design and purpose of expert systems as well as an in-depth discussion regarding their construction.

What is an Expert System?

According to an article in Encyclopaedia Britannica by Vladimir Zwass, expert systems are programs designed to emulate the knowledge and decision-making skills of an expert in a particular field (Zwass, 2016). Zwass points out that they do this by consulting a series of rules designed by experts and cross referencing their database of pre-programmed knowledge when necessary (Zwass, 2016).

Parts of an Expert System

An expert system can be broken down into three pieces according to geeksforgeeks (Aneja, 2018). The first is the knowledge base. This piece, as explained by Vladimir Zwass, contains all the knowledge that is deemed relevant to the system (Zwass, 2016). The second is the inference engine. The inference engine is the system's "thought process". The aforementioned geeksforgeeks article states that the inference engine is a series of logical rules or "heuristics" that the system consults in order to solve problems (Aneja, 2018). The third piece is a user interface. It is the place in the program that defines the method by which users query the system. Based on information present on geeksforgeeks, user interfaces for expert systems often make use of natural language processing (NLP) in order to answer questions posed by the user in simple English or another language (Aneja, 2018).

The pieces of an expert system can be compared to the elements of a human expert. The knowledge base is akin to a human expert's subject-specific knowledge. The inference engine can be said to emulate a human expert's intuition, experience, and problem-solving skills in relation to their field. Lastly, a user interface is representative of the senses of a human expert. In order to answer a question, a human must hear or see what has been asked and may use other senses to perceive certain subject matter pertinent to the question. So how exactly are these elements effectively simulated in an expert system?

The Knowledge Base

As we already know, the knowledge base of a given expert system defines all things known by that system. They are most often specific to a given field and are designed to supply the system with enough information to make decisions and solve problems. According to the "Knowledge Base Construction" page on Stanford University's DeepDive project website, facts present in knowledge bases are known as "assertions" and are expressed as sentences ("Knowledge Base Construction," 2017). For example, the phrase, "A dog is an animal", is an assertion. Furthermore, each assertion contains references to entities ("Knowledge Base Construction," 2017).

Entities, as defined by "Knowledge Base Construction," are code objects representing real-world people, places, things, etc. ("Knowledge Base Construction," 2017). These entities contain data of their own and relational data between themselves and other entities ("Knowledge Base Construction," 2017). For example, a dog may have an attribute representing their breed, and a relation to a person who is their owner. The webpage also states there are determinations of how a given entity may be "mentioned" ("Knowledge Base Construction," 2017). This enables automated tools to gleam data from sentences using natural languages processing and form new assertions based on gathered data. The diagram below, courtesy of Stanford University's DeepDive project, expresses how entities are related using the case of Barack and Michelle Obama being married ("Knowledge Base Construction," 2017).

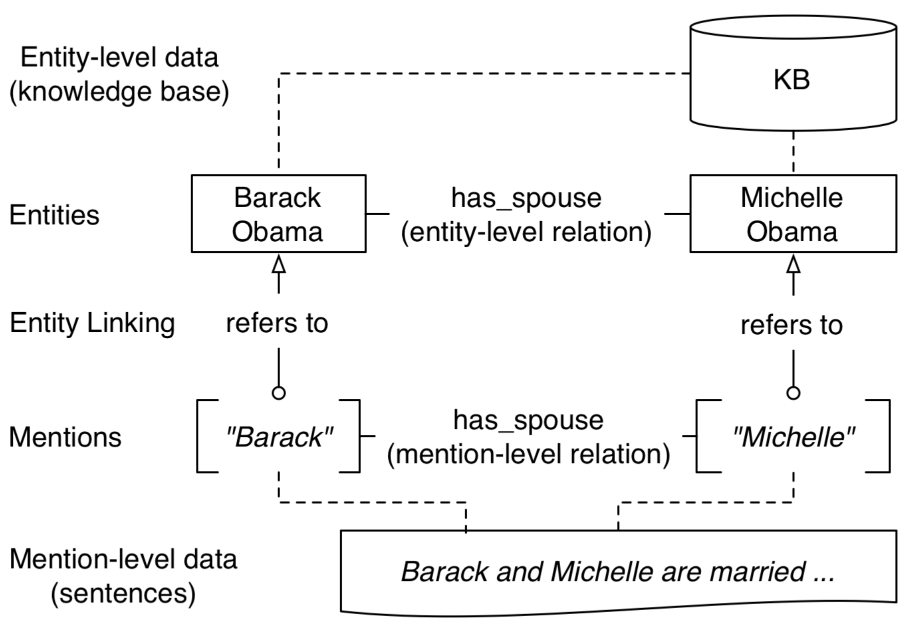

Figure 1

The Inference Engine

The inference engine is utilized by the expert system to make decisions based on its knowledge base. Pursuant to "Inference Engine" (2015), It is composed of many rules pursuant to the knowledge base ("Inference Engine," 2015). The engine can also infer new rules based on other rules according to Encyclopedia Britannica. For example, the transitive property can be used to infer that A leads to C if A leads to B and B leads to C ("Inference Engine," 2020). There are many ways an inference engine can be implemented, but the two main methods, as explained by former Oxford University lecturer of artificial intelligence Jocelyn Ireson-Paine, are a forward chaining inference engine and a backward chaining inference engine (Ireson-Paine, 2011).

According to Ireson-Paine, a forward chaining inference engine takes any new assertion and immediately finds all relationships related to entities present in the assertion using it and any assertions loaded into memory prior to that moment (Ireson-Paine, 2011). An explanation of forward chaining by Ireson-Paine can be found at this link (Ireson-Paine, 2011). Alternatively, Ireson-Paine states that backward chaining works backwords through known assertions when prompted with a question by the user (Ireson-Paine 2011). Ireson-Paine details this process in this link (Ireson-Paine, 2011). These questions can be posed by a user's interaction with the user interface.

The User Interface

An expert system requires a method by which to take input from the user. This input can be in the form of hypotheses for the system to answer or, in some cases, user-provided assertions for use by the system. Examples of interfaces that could be implemented in an expert system are text input utilities with NLP, graphical interfaces, and voice-input utilities with NLP. For example, a forward chaining expert system designed to identify birds for an ecological survey could allow the user to click on pictures of visual features present in a given bird and output a picture and information of the bird the system thinks has been sighted.

Machine Learning

In this subsection, machine learning facilitated by artificial neural networks will be discussed. It contains a summary of machine learning as a whole, a brief explanation of the various types of machine learning, and some background knowledge regarding artificial neurons and neural networks.

What is Machine Learning?

According to an article in MIT technology review, machine learning is the cognitive computing discipline concerned with the ability of computers to learn when presented with large amounts of data (Hao, 2020). There are three main types of machine learning. These, according to the flowchart contained in the aforementioned article, are supervised learning, reinforcement learning, and unsupervised learning (Hao, 2020).

The first type, supervised learning, involves exposing the computer to pre-processed data in hopes that the computer will learn to process similar data by observing patterns in this “training” data (Hao, 2020). A common example is training a machine learning algorithm to recognize dogs and cats in pictures. Images are fed to the algorithm that already indicate which animal is pictured. Following training, the computer should be able to take a picture of a dog or cat, determine which animal is pictured, and sort the image into the corresponding category.

The second type, reinforcement learning, is essentially trial and error (Hao, 2020). Data is fed to the algorithm, and the algorithm attempts to process it. If processed correctly, the algorithm is “rewarded,” or told that it is correct. If it is not correct, the algorithm is “punished” in a similar manner. Eventually, the algorithm will learn how to process the data correctly.

The third, and arguably most complicated type of machine learning, unsupervised learning, has many different approaches. All of them attempt to have the machine learn without any form of “training” (Hao, 2020).

An Introduction to Artificial Neural Networks

All types of machine learning are facilitated using artificial neural networks (Agatonovic-Kustrin and Beresford, 2000). Artificial neural networks are comprised of code objects called neurons that attempt to simulate their biological namesake found in the human brain.

According to a webpage published by the University of Queensland, neurons in humans take in signals from multiple branch-like structures called “dendrites”, process that data in the main body of the cell or “soma”, and send an output based on this processing to other neurons by way of the “axon” (“What is a Neuron?” 2019). Artificial neurons perform in a similar way. According to a previously mentioned study by Agatonovic-Kustrin and Beresford, artificial neurons take in data, evaluate it, and output it to other neurons like biological neurons (Agatonovic-Kustrin and Beresford, 2000). Below is a photo of a biological neuron from medicalxpress news (Why are neuron axons long and spindly? Study shows they're optimizing signaling efficiency, 2018):

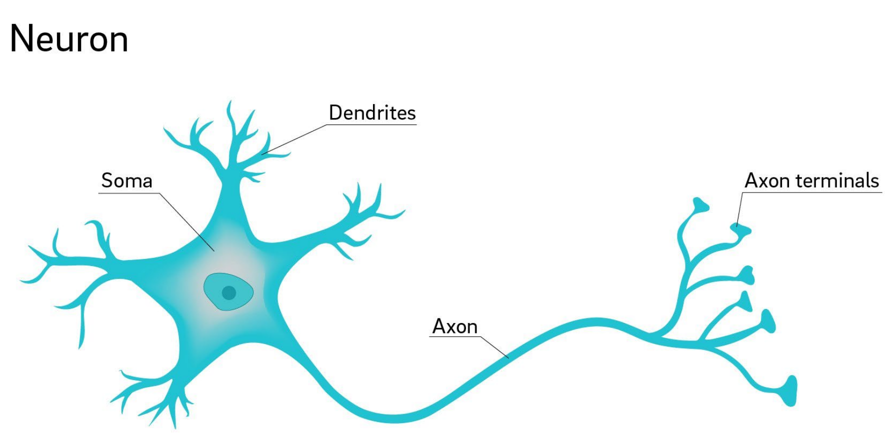

Figure 2

The same study explains that artificial neurons are arranged into layers. Input variables or "factors" are fed into the input layer, the data is processed once by this layer, passed to the hidden layer, and then to the output layer where projected independent variables are output (Agatonovic-Kustrin and Beresford, 2000). Usually, there is only one hidden layer. When there are more, the neural network is deemed a "deep" neural network.

Part 2: Industrial Uses for Cognitive Computing

Analytics and Processing of Numerical Data

Numerical data is ever so prevalent in our lives now. Databases full of such data make possible such things as bank transactions, inventory tracking, and even predictive analysis of stock prices. Humans alone are unable to analyze such vast amounts of information with the level of detail cognitive computing programs can while working alongside their human creators. In this section, two uses of cognitive computing for the analysis of quantitative, numerical data will be summarized as case studies. The case studies will include the uses of blank in blank and the uses found for machine learning technologies within the field of pandemic epidemiology.

AI for GPS Guidance Systems

Many of us could not navigate our daily lives if not for automated GPS navigation. A technology made possible by satellites; GPS navigation systems use path finding algorithms to generate the shortest route from point A to point B with any number of waypoints in between. Although the data points used to get you from place to place such as street names and freeway exits may not seem mathematical, they can be expressed as such. Roadways are expressed as weighted graphs. The weight of a given edge is computed as an aggregate of traffic, distance, and other factors that lengthen a commute. Following the formation of such a graph, established graph traversal or path finding algorithms can be used to compute the shortest route. Common algorithms that could be used by GPS systems include A* (A star) search, heuristic search, and iterative deepening search. According to Dr. Mike Pound of the YouTube channel "Computerphile", some GPS systems may also make use of proprietary algorithms. These are often treated as trade secrets and are not known to the public in many cases (Pound, 2017). The video from Computerphile can be found at this link and explains A* search and places it within the context of GPS navigational systems.

Machine Learning for Pandemic Epidemiology

The field of statistics called "pandemic epidemiology" deals with the numerical data involved in a mid-to-large scale contagious disease outbreak. This includes things such as transmission vector analysis, predictive modeling, and communication of data to the public. Machine learning and artificial intelligence technologies are frequently used to aid epidemiologists due to the massive amount of data these medical experts have to analyze and interpret on a daily basis. A recent example of these technologies being used in this way is that machine learning and artificial intelligence have been used to analyze trends and develop solutions within the context of the worldwide Covid-19 pandemic.

According to a scientific journal entry by Ahmad Alimadadi et al, deep learning has been used to model death and contagion rates to predict Covid-19's impact, develop systems designed to determine a Covid-19 diagnosis, and even to synthesize chemicals that may aid in fighting off the disease (Alimadadi et al, 2020). Alimadadi first shares that many large organizations such as The Allan Institute for AI and the White House worked together with top minds to create data mining techniques in order to study Covid-19 data in as close to real time as possible (Alimadadi et al, 2020). Later, scientists started analyzing Covid-19 on a genetic level and using pharmacogenomics to design models that predict a person's susceptibility to it (Alimadadi et al, 2020). Pharmacogenomics is the use of highly advanced deep learning models to develop personalized medicine or otherwise find relationships between conditions based on a person's genome. The presentation at this link shows how machine learning used by epidemiologists and cognitive computing specialists to fight Covid-19. The figure can also be seen below (Alimadadi et al, 2020):

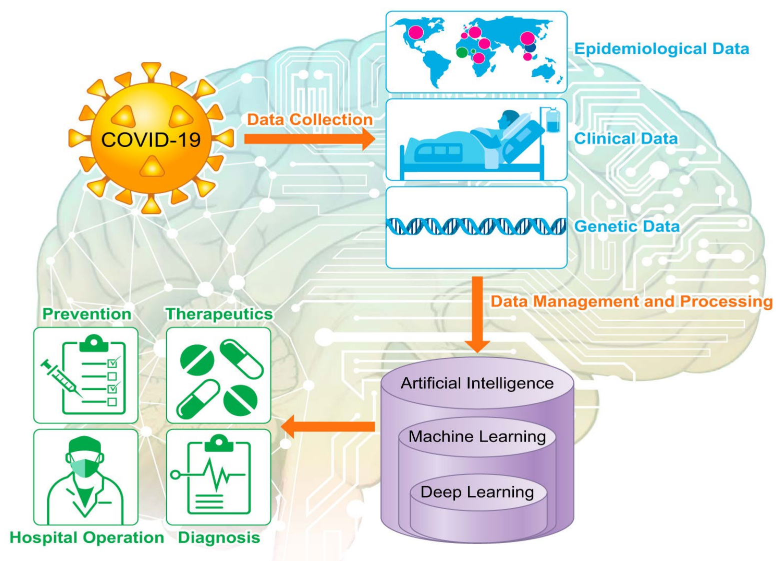

Figure 3

Analytics and Processing of Qualitative Data in a Natural Language

The processes that make up our professional and, certainly, personal lives cannot currently be broken down completely into numerical data and patterns. We humans can analyze data outside the realm of the numerical. Through cognitive computing, so too are computers. Computers are now able to create complex webs of thought regarding data that has been expressed in a natural, or human, language such as English. They can also make decisions based on such information. In this subsection, two case studies will be presented that display the ability of computers to perform analysis on non-traditional sources of data. These case studies include expert systems that can make a preliminary medical diagnosis based on symptoms and machine learning technologies designed to extract and analyze consumer sentiment from social media posts.

Expert Systems for Medical Diagnosis

One of the most common uses of knowledge based expert systems in the modern era is for handling tasks related to the medical industry. This primarily includes expert system-based diagnostic tools. An expert system designed for this purpose utilizes a knowledge base derived from medical knowledge to help health professionals diagnose bacterial, viral, physical, and genetic abnormalities in the human body. This is helpful because it can provide physicians and even more specialized healthcare providers a secondary diagnosis to reinforce their thinking or force them to question it. With expert systems, it is possible to receive this secondary diagnosis without the need for a human expert's second opinion. According to an article by Anitha Kannan, examples of such systems include Mycin, Internest-1, QMR, and DXplain (Kannan, 2019). Another system that is being used for diagnostics is IBM's Watson. Below is a screenshot from the video at this link explaining how Watson is being used in cancer diagnosis (Gupta, 2014):

Figure 4

These systems can also aid in progressing from a differential diagnosis to a more formal diagnosis and diagnostic tests. A differential diagnosis is, according to the article by Kannan, a list of diagnoses that are further refined until one diagnosis remains (Kannan, 2019). Kannan also mentions costs of traditional medical diagnosis that are circumvented or partially circumvented by technological means. These include such things as lack of diagnostic tools and the difficulties experienced by medical professionals when attempting to mine information (Kannan, 2019).

These expert systems are not perfect. They are used as supplementary tools that aid human experts, as the risk is simply too great to allow the systems to make decisions in life or death medical scenarios. However, they are getting better and technologies are even emerging that allow automated knowledge base creation using deep learning. Automated knowledge base creation and expansion are capabilities present in IBM’s Watson for example.

Machine Learning for Sentiment Analysis and Data Mining

One of the most important metrics to keep track of as a marketing professional is public opinion on products, companies, and key people within the context of a marketing campaign. This data is, unfortunately, not easy to express in terms of numerical data. However, it is possible for a computer program to facilitate such expression through a process known as sentiment analysis. As the name suggests and as an article by Shashank Gupta points out, sentiment analysis involves text in a natural language being analyzed and mined for positive and negative connotation or “sentiment” towards particular people, phrases, and things of interest (Gupta, 2018).

According to the article, sentiment analysis is the first step to large-scale subjective analysis of text by computer (Gupta, 2018). The second stage is intent analysis or determining what kind of action the sentence is attempting to perform (Gupta, 2018). The third is contextual semantic searching (Gupta, 2018). That is the act of determining the business context of a sentence (Gupta, 2018). Using these three methods, it is possible to determine such things as “the customers do not like the price” or “the customers love the service we deliver”. Such data can be utilized by marketing and business experts to fix problem areas and further publicize areas of extraordinary performance.

Unfortunately, this type of technology used in data mining also leads to privacy concerns amongst consumers. Social media platforms such as Facebook have come under fire in recent years due to mishandling and unethical use of customer data. With technologies like sentiment analysis and machine learning, it is even possible for companies like Facebook to accurately determine data about customers that has not explicitly been provided.

One example of this is the estimation of base income from other, previously provided data. A customer’s base income can be estimated using geolocation data from shops they may have visited, posted pictures featuring logos, or any other number of other data sources the customer provides. The customer may not have provided a company like Facebook explicit access to their income data, but as the data produced by data mining algorithms is technically a very good estimation rather than official data, it falls under a loophole in privacy law. While this is not illegal or unethical according to most professional organizations, it is cause for concern, as information is increasingly being weaponized.

An example of potential malfeasance is the misuse of a sentiment analysis algorithm that diagnoses manic-depressive disorder based on social media posts and predicts manic episodes. According to an article in Scientific American, such an algorithm does exist (Luerwig, 2019). On the surface, the algorithm may seem helpful to the mental health community, but it could also be used to target ads for casinos, lotteries, and other high-risk activities to those with lower inhibitions due to entering a manic phase. This is highly unethical as it is taking advantage of people’s mental illness to make a profit.

Applications in Consumer Products

Cognitive computing technologies are also used to improve older technologies, create derivative technologies, and add new functionality. This is true as much in consumer products as it is in business-related technologies. This subsection will analyze two examples of consumer products that rely on cognitive computing technologies to improve the user experience and overall product quality. The two examples are the use of AI in smart wireless technologies and voice assistants such as Amazon's Alexa and Google Voice Assistant.

AI for Intelligent Wireless Technologies

A common application of AI technology is optimization. In fact, many advanced machine learning algorithms make use of AI-style optimization. In consumer products, this kind of algorithm can translate into products such as whole-home wireless systems that intelligently manage bandwidth for wireless transmission of internet data to and from client devices. In layman's terms, this means a wireless router or system of wireless routers that adapt to internet usage by devices such as smartphones. The adaptations allow "smart wireless" products to provide more stable and fast access to the wireless network than traditional wireless routers can. This same type of adaptation can also be utilized at other wireless frequencies such as Bluetooth in order to facilitate more intelligent communication between internet of things (IoT) devices.

Smart Voice Assistant Programs

Another example of consumer product that use cognitive computing are voice assistant programs. These include products such as Apple's Siri, Amazon's Alexa, and Samsung's Bixby. These products are marketed as "assistants", but they are voice-based interfaces for existing smart device functions. Voice-based interfaces rely on two major cognitive computing technologies. The first technology is speech recognition. The second is natural language processing.

Speech recognition, according to IBM, is the ability of a computer to convert spoken language as audio waveforms into written language as characters of the relevant alphabet (What is Speech Recognition, 2020). IBM also states that speech recognition solutions can be implemented in multiple ways, but the most impressive of which are using machine learning (What is Speech Recognition, 2020). One potential problem with voice recognition is that some accents are difficult for the system to translate. The Scottish accent, in particular, was in the news rather frequently following the release of Apple's Siri.

Natural Language Processing or NLP is, as previously mentioned, the ability of the computer to interpret meaning from the text in a natural language. Natural language processing is frequently accomplished by means of machine learning; although, in the early days of the technology, it was possible only through the use of expert systems.

A voice assistant utilizes speech recognition to interpret speech as text. That command is then (typically) directed to a server by the device running the assistant. The server interprets the text using NLP and sends a command based on the text back to the device for execution. For example, if one asks Google Voice Assistant "What is two plus two?" The speech is converted into text.

Following conversion, the text is sent to Google's servers. The server then interprets the text as a mathematical problem and sends a command back to the client device. The command tells the device to launch the calculator app and execute 2+2.

Part 3: Possible Costs and Benefits of Cognitive Computing

Corporations

This subsection will delve into examples of potential costs and benefits to large corporations and small businesses incurred using cognitive computing technologies. It will also summarize examples of costs and benefits to corporate consumers that are involved in business to business (B2B) transactions. These costs and benefits will fit within the realms of monetary gain or loss, loss or gain of consumer confidence, and other common metrics for business success.

Costs

Possible monetary costs to corporations using cognitive computing technologies include increased cost of skilled technologist salaries if the technology is made in-house or increased costs to pay an external provider of a technology. Beyond those costs, many companies are not equipped to handle integrating new technologies into their workflow without the expertise of consultants, so consulting or integration fees may be involved as well. In Fact, according to an article by Paolo Del Nibletto, technology integration falls within the top ten technology issues companies face (Del Nibletto, n.d.)

Possible blows to consumer confidence could result if the technology makes heavy use of consumer data, if the technology takes away consumer control, or if the technology causes a major change in the customer support experience. This is especially important if the company caters to a largely technology or change-averse subsection of the population.

An example of a blow to consumer confidence due to data mining occurred, according to a 2020 article by Joe Kukura, when the popular video conferencing app "Zoom" leaked mined data to Facebook without notifying consumers (Kukura, 2020). The data mining technology itself played no part in wrongdoing; however, scandals like this increase the reluctance of consumers to participate in activities which may expose their data to mining whether legitimate or unethical.

Motivating a work force may also prove difficult if said work force is used to interacting heavily with consumers and is now interacting primarily with a computer screen. This could affect a company's bottom line or result in unnecessary layoffs or resignations.

Benefits

Possible monetary gains could result from such things as a more efficient manufacturing process. Increased employee productivity due to automation, lower costs due to layoffs of redundant positions, better analysis of P&L (Profit and Loss) statements, etc. In essence, these things are why businesses pursue automation and cognitive computing technologies in the first place.

According to an article in Forbes, automation technology can be used to quickly perform repetitive tasks (Atchison, 2020). Some of these tasks might be simple, and others might be complex. The article emphasizes the ability of businesses to focus on other, more human tasks such as customer satisfaction when automation technology is used to handle other aspects of the day to day workload (Atchison, 2020).

Consumers could become more confident in businesses that make use of new technology if the technology has a proven track record, increases convenience, or is only a supplementary tool to be used by a human expert. For example, financial institutions experience greater customer satisfaction when cognitive computing technologies are used to increase convenience. An example of this could be seen when OCR technology was employed at major banks to enable the depositing of checks via a smart phone camera. Many smaller banks and credit unions adopted the technology shortly after its inception in order to reap the same benefits.

For companies that create cognitive computing technologies, the obvious benefit will be more business and the ability to employ more people. IBM, for instance, makes a great deal of money and employs many people in both traditional workforce positions and those involved in corporate research.

Employees

This subsection will explore possible costs and benefits to employees of companies that make use of cognitive computing technologies. These examples will cover topics such as layoffs, pay scale changes, work difficulty, and minimum qualification changes.

Costs

One common argument against automation of any kind is the potential for machines to replace humans as low-level workers. In the future, this may indeed be possible. Already there are machines that can perform the duties of such workers. Vending machines that cook fresh pizza in a similar way to chain restaurants are common in Japan for example.

The worry of some is that cognitive computing technologies will make machines even more capable of replacing and even surpassing human workers. This is, of course, not possible with current technology. It is a very real possibility for the distant future though, and the economic implications of mass layoffs of replaced workers may be catastrophic. According to an article published by REBA, some industries, such as healthcare and STEM research, have a minimal probability of being affected, but others, such as manufacturing, could likely be largely automated (Could automation be the key to a better work/life balance?, 2019).

Secondly, workers who have duties that are replaced by machines, may experience a decline in pay even if they are not replaced, as their job becomes less skilled the more it is automated. Although, the likelihood of this outcome was determined by mere speculation.

Benefits

A benefit to employees is that automation often leads to high productivity with a less intensive workload. In fact, an article in dot magazine proposes a maximum 1.4 percent increase by in productivity by 2065 (Comastri, 2020). This productivity increase will enable employees to have a healthier work-life balance. However, employers must be careful not to let their employees fall into bad habits.

Cognitive computing technologies may also allow workers with less qualifications to perform at the same level as other workers who have more qualifications. This would enable greater upward mobility, as a lower level employee could be promoted and perform at the same level as their colleagues with assistance from machines. This would increase the ability of companies to fill available positions and train new employees faster. Perhaps it could even help lower the unemployment rate and increase the employability for historically disadvantaged populations.

Consumers

This subsection will explore possible costs and benefits to consumers who make use of products or services that contain cognitive computing technologies. These examples will cover topics such as product price and quality and convenience.

Costs

Consumers experience very few, if any, costs from the use of cognitive computing technology when it is properly utilized and developed.

Benefits

Customers benefit from lower prices when companies experience lower cost. Because cognitive computing and automation can contribute to lower costs in manufacturing and repetitive business tasks, it also contributes to lower prices for customers. Customers can also experience added convenience. One example of such convenience is the ability of google calendar to automatically add vacation events based on hotel or airline confirmation emails.

Part 4: Appendices

External Resources

Description	Link
An explanation of similarities and differences between the human brain and computers	https://faculty.washington.edu/chudler/bvc.html
A YouTube video explaining the basics of AI	https://www.youtube.com/watch?v=2ePf9rue1Ao
An explanation of forward chaining inference engines	http://www.jpaine.org/students /lectures/lect3/node11.html#SECTION00081000000000000000
An explanation of backward chaining inference engines	http://www.j-paine.org/students/lectures/lect3/node12.html
A YouTube video that explains A* search	https://www.youtube.com/watch?v=ySN5Wnu88nE
A YouTube video that shows how IBM Watson is used in medical diagnosis	https://www.youtube.com/watch?v=hbqDknMc_Bo
A YouTube video that explains machine learning	https://www.youtube.com/watch?v=ukzFI9rgwfU
A YouTube video that explains the differences between various cognitive computing technologies	https://www.youtube.com/watch?v=k2P_pHQDlp0
A YouTube video that explains data mining	https://www.youtube.com/watch?v=R-sGvh6tI04
A YouTube video telling the story of IBM Watson's win against Jeopardy! champions	https://www.youtube.com/watch?v=P18EdAKuC1U

Figures

References

Agatonovic-Kustrin, S., & Beresford, R. (2000). Basic concepts of artificial neural network (ANN) modeling and its application in pharmaceutical research [Abstract]. *Journal of Pharmaceutical and Biomedical Analysis, 22*(5), 717-727. doi:10.1016/s0731-7085(99)00272-1

Alimadadi, A., Aryal, S., Manandhar, I., Munroe, P. B., Joe, B., & Cheng, X. (2020). Artificial intelligence and machine learning to fight COVID-19. *Physiological Genomics, 52*(4), 200-202. doi:10.1152/physiolgenomics.00029.2020

Aneja, A. (2018, August 02). Expert Systems. Retrieved September 7, 2020, from https://www.geeksforgeeks.org/expert-systems/

Atchison, J. (2020, October 14). Council Post: Three Ways Marketing Automation Can Help Grow Your Business. Retrieved October 15, 2020, from https://www.forbes.com/sites/theyec/2020/10/15/three-ways-marketing-automation-can-help-grow-your-business/#6d3819149261

Chudler, E. H. (2008, December 20). Neuroscience For Kids. Retrieved September 7, 2020, from https://faculty.washington.edu/chudler/bvc.html

Comastri, M. (2020, September 23). 2017: The Year of Automation? Retrieved October 15, 2020, from https://www.dotmagazine.online/issues/digital-production/driving-digital-transformation/year-of-automation

Could automation be the key to a better work/life balance? (2019, July 09). Retrieved October 15, 2020, from https://reba.global/content/could-automation-be-the-key-to-a-better-work-life-balance

Crapo, A. W., Waisel, L. B., Wallace, W. A., & Willemain, T. R. (2000). Visualization and the process of modeling. *Proceedings of the Sixth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD 00*. doi:10.1145/347090.347129

Del Nibletto, P. (n.d.). Top 10 technology challenges businesses face. Retrieved from https://channeldailynews.com/news/top-10-technology-challenges-businesses-face/7570

Dyer, C. R. (2003, January 17). Intelligent Agents. Retrieved September 15, 2020, from http://pages.cs.wisc.edu/~dyer/cs540/notes/agents.html

Gupta, A. Joshi, P. [IBMIndiaSS]. (2014, January 6). *IBM Watson Demo Oncology Diagnosis and Treatment in 8 min* [Video]. YouTube. https://www.youtube.com/watch?v=hbqDknMc_Bo

Gupta, S. (2018, January 19). Sentiment Analysis: Concept, Analysis and Applications. Retrieved October 13, 2020, from https://towardsdatascience.com/sentiment-analysis-concept-analysis-and-applications-6c94d6f58c17

Hao, K. (2020, April 02). What is machine learning? Retrieved September 15, 2020, from https://www.technologyreview.com/2018/11/17/103781/what-is-machine-learning-we-drew-you-another-flowchart/

Inference engine. (2015, November 09). Retrieved September 8, 2020, from http://zasoby.open.agh.edu.pl/~15sdrzeznik/en/system-wnioskujacy.html

Inference engine. (2020, September 08). Retrieved September 8, 2020, from https://www.britannica.com/technology/inference-engine

Ireson-Paine, J. (2011, January 14). How the Inference Engine Works. Retrieved September 15, 2020, from http://www.j-paine.org/students/lectures/lect3/node10.html

Kannan, A. (2019, April 15). The science of assisting medical diagnosis: From Expert systems to Machine-learned models. Retrieved October 13, 2020, from https://medium.com/curai-tech/the-science-of-assisting-medical-diagnosis-from-expert-systems-to-machine-learned-models-cc2ef0b03098

Knowledge Base Construction. (2017, August 31). Retrieved September 7, 2020, from http://deepdive.stanford.edu/kbc

Korf, R. E. (1988). Optimal Path-Finding Algorithms. *Search in Artificial Intelligence,* 223-267. doi:10.1007/978-1-4613-8788-6_7

Kukura, J. (2020, April 01). Zoom Video Conferencing Hit With Lawsuit Over Facebook Data Sharing; Other Security Issues Racking Up. Retrieved October 15, 2020, from https://sfist.com/2020/04/01/zoom-video-conferencing-hit-with-privacy-scandal/

Luerweg, F. (2019, March 14). The Internet Knows You Better Than Your Spouse Does. Retrieved October 13, 2020, from https://www.scientificamerican.com/article/the-internet-knows-you-better-than-your-spouse-does/

Mcclelland, J. (2001). Cognitive Neuroscience [Abstract]. *International Encyclopedia of the Social & Behavioral Sciences,* 2133-2140. doi:10.1016/b0-08-043076-7/03406-9

Pound, Mike. [Computerphile]. (2017, February 15). *A* (A Star) Search Algorithm – Computer-phile* [Video]. YouTube. https://www.youtube.com/watch?v=ySN5Wnu88nE

Ramesh, Raj. [Raj Ramesh]. (2017, August 13). *What is Artificial Intelligence? In 5 minutes.* [Video]. YouTube. https://www.youtube.com/watch?v=2ePf9rue1Ao

West, D. M. (2019, October 25). What is Artificial Intelligence? Retrieved September 14, 2020, from https://www.brookings.edu/research/what-is-artificial-intelligence/

What is a neuron? (2019, August 13). Retrieved September 15, 2020, from https://qbi.uq.edu.au/brain/brain-anatomy/what-neuron#:~:text=A neuron has three main,receives input from other cells.

What is Speech Recognition? (2020, September 2). Retrieved October 13, 2020, from https://www.ibm.com/cloud/learn/speech-recognition

Why are neuron axons long and spindly? Study shows they're optimizing signaling efficiency. (2018, July 11). Retrieved October 15, 2020, from https://medicalxpress.com/news/2018-07-neuron-axons-spindly-theyre-optimizing.html

Zwass, V. (2016, February 10). Expert system. Retrieved September 15, 2020, from https://www.britannica.com/technology/expert-system