

Introduction

Artificial intelligence (AI) refers to the art and science of creating computer systems that simulate human thought and behavior.

Research in artificial intelligence, or AI, is of interest to many professionals, including computer scientists, psychologists, philosophers, neuroscientists, engineers, data scientists and others who often collaborate on AI research. While AI research is the domain of scientists, AI applications are impacting all industries and touching all of our lives. Most of us use AI technologies every day without even knowing it. Web search engines and personal digital assistants like Siri use AI to provide useful information and search results. Ecommerce websites display products that you are likely to buy based on your previous behavior and purchases. From face recognition for photos in Facebook, to driverless vehicles, and a million other applications in between, AI is providing new, revolutionary services every day, often with complex, controversial ethical considerations.

After years of relative neglect, new levels of computing power has brought AI back as one of the hottest areas of investment and research in the tech industry. Google's \$600 million acquisition of DeepMind was just one of over one hundred acquisitions of AI companies made by companies like Facebook, Google, Apple, and Twitter since 2011. The U.S. government has invested billions of dollars in AI research over the past three years including \$100 million to the "BRAIN Initiative" intended to reverse engineer the brain to find algorithms that allow computers to think more like humans. The recent AI fever isn't just in the U.S.—it's global. Chinese search giant Baidu has invented a powerful supercomputer devoted to an AI technique to provide software with more power to understand speech, images, and written language. AI has brought computers to the cusp of human intelligence and beyond, and every government and company wants to be the first to capitalize on that power.

AI is the power behind most of the biggest areas of technological innovation today. It is essential in mining useful information from big data. For example, one AI algorithm analyzes tweets to determine which restaurants might make you sick. AI is key to making sense of the Internet of Things. For example, Pittsburgh has implemented a smart traffic signal technology that analyzes traffic patterns to control when traffic lights change, reducing travel time for residents by 25 percent and idling time at stop lights by over 40 percent. Google is developing an AI system that will get rid of the need for passwords on Android phones. A phone will recognize its human owner by analyzing usage and location patterns.

Some have estimated that the human brain operates at roughly 30 times faster than today's fastest supercomputers. Today's fastest supercomputers operate at more than 100 petaflops (100 thousand trillion floating point operations per second). That's 100 followed by 15 zeros. While the human brain is not increasing in its processing speed, computers are getting faster every year. Moore's Law states that computers double in speed approximately every two years. By 2020, it is expected that supercomputers will achieve exascale performance—that is 1,000 petaflops. At petaflop speeds, supercomputers are nearing human thought speed. Bill Gates believes that machines will outsmart humans in some areas within a decade. Facebook CEO Mark Zuckerberg thinks that within five to ten years, artificial intelligence could advance to the point where computers can see, hear, and understand language better than people.

What do you think it will take for computers to "think" like humans? What will happen when scientists can simulate the functions of the neurons and synapses in the human brain? Will they be able to create a thinking machine? What is thinking? What is the difference between human intelligence and computer intelligence?

At a Dartmouth College conference in 1956, John McCarthy proposed the term artificial intelligence to describe computers with the ability to simulate human thought and behavior. Many AI pioneers attended that first AI conference; a few predicted that computers would be as smart as people by the 1960s. That prediction was proven wrong by a long shot. The failure of early AI research to produce thinking machines caused many to give up on the field. But now, after years of relative neglect, new levels of computing power and insight has brought AI back as one of the hottest areas of investment and research in the tech

industry. Just as “mobile” has been the hot technology growth area of the past decade, “AI” looks to be the big area of growth for the next decade.

Google’s acquisition of DeepMind in 2014 was just one of over one hundred acquisitions of AI companies made by tech giants like Facebook, Google, Apple, and Twitter since 2011. The U.S. government has invested billions of dollars in AI research over the past three years including a \$100 million “BRAIN Initiative” intended to reverse engineer the brain to find algorithms that allow computers to think more like humans. BRAIN in this case, stands for Brain Research through Advancing Innovative Neurotechnologies.

The recent AI fever isn’t just in the U.S.—it’s global. Chinese search giant Baidu has invented a powerful supercomputer devoted to an AI technique to provide software with more power to understand speech, images, and written language. AI has brought computers to the cusp of human intelligence and beyond, and every government and company wants to be the first to capitalize on that power.

To imbue computers with intelligence, we must be able to recognize the traits of intelligent behavior so that we can recognize them in computers. Certainly our ability to use our senses to observe the world around us and store knowledge (i.e., remember) is an important element of intelligence. Another important element is our ability to reach conclusions and establish new understandings by combining information in new ways. Yet another important element in exhibiting intelligence is our ability to communicate and use language. All of these characteristics of intelligence are being duplicated in machines.

This unit begins looking at the various methodologies used in AI research, and then turns to the most useful AI applications being developed today.

The area of artificial intelligence has been controversial since its origin. The thought of computers gaining intelligence that rivals human intelligence could threaten our very existence. The Singularity, a future time when computers exceed humans in intelligence, promises computers with “superintelligence.” Would these computers continue to serve humankind or enslave us?

Some safeguards against rogue AIs are being developed. The White House has called on AI researchers to share information across the industry so that all involved can make informed decisions. Amazon, Google, Facebook, IBM, and Microsoft have formed a not-for-profit organization, called Partnership on AI to Benefit People and Society, to educate the public about AI technologies and reduce anxieties around its application. Google’s AI research lab in London, DeepMind, teamed up with Oxford University’s Future of Humanity Institute to develop a software “panic button” to interrupt a potentially rogue AI agent.

Still, AI is bound to create ethical dilemmas. Expert systems and robotics can carry out repetitive procedures with pinpoint accuracy, never tiring, getting bored, or complaining. Such systems have displaced many human workers. Some would argue that these systems free up the human workforce for more interesting and fulfilling careers. That argument fails, however, when the unemployment rate rises and government and industry are unable to educate displaced workers with new marketable skills.

As driverless cars enter the market, how will they be programmed to make choices in life or death situations? In an accident scenario in which sacrificing the driver’s life will save a bus full of children, how should the AI navigate the car? Should an AI decide who should live or die?

AI and robotics serve many national defense applications. Drone aircraft are remotely controlled to carry out attacks in cities hundreds of miles away from their operators. Robots hunt down bombs and improvised explosive devices (IEDs), saving many lives. Other military and police robots are being armed with automatic weapons and explosives and being sent into combat situations. Critics of military and surveillance robots are concerned about what happens when these robots malfunction or fall into the wrong hands. As with all technologies, the good that is accomplished with AI is tempered by the associated threats.

Lesson 5.1: The Development of Artificial Intelligence

Lesson 5.1 Introduction

Artificial intelligence describes a wide range of software that is modeled on or imitates human intelligence to various degrees. Human intelligence is notoriously difficult to define, however, so any software development project tasked with perfectly replicating the human capacity for creativity, intuition, and inspiration is immediately faced with daunting technological and philosophical challenges. A much more productive approach to AI has been to choose a smaller subset of human intellectual skills, such as the ability to recognize faces or voices, and leverage the exponentially increasing processing power of computers to meet, and surpass, human abilities in those areas.

Reading: AI Methodologies

AI Methodologies



Artificial-intelligence-2167835 1280. Authored by: 6eotech, License: CC BY 2.0

Artificial intelligence methodologies consist of the various approaches to AI research that generally fall under one of two categories: conventional AI or computational intelligence.

Why This Matters

AI researchers have developed many approaches to AI in order to meet a wide variety of goals. From simple systems that offer suggestions as you type in a key term in Google, to systems that offer professional medical advice, to systems that mimic the biological functioning of the human brain, the goals for AI systems cover a broad spectrum of applications. Understanding AI methodologies helps us to understand the complexities of human thought.

AI systems typically include three components:

A set of logical rules to apply to input in order to produce useful output

Instructions for how to handle unexpected input and still produce useful output—i.e., the ability to “wing it”

The ability to learn from experience to continuously improve performance, referred to as machine learning

Several formal methods are used to create intelligent software that has these components and capabilities. They fall under two categories:

Conventional AI methodologies—such as expert systems, case-based reasoning, Bayesian networks, and behavior-based AI—that rely on the programmer to instill the software with logical functionality to solve problems

Computational intelligence methodologies—such as neural networks, fuzzy systems, and evolutionary computation—that set up a system whereby the software can develop intelligence through an iterative (cyclical) process

It is often the case that AI developers combine methods from different methodologies to achieve their goals.

Reading: Conventional AI

Conventional AI

Conventional AI—also called symbolic AI, logical AI, or neat AI—uses programming that emphasizes statistical analysis to calculate the probability of various outcomes in order to find the best solution.

Why This Matters

Conventional AI has led to the development of software that is more intuitive to use and systems that can automate the decision-making process in certain professional activities. Conventional AI techniques such as expert systems, case-based reasoning, Bayesian networks, and behavior-based AI systems are embedded in many of today's popular software applications, making them easier to use and to control robots and autonomous vehicles.

Essential Information

In conventional AI, advanced programming techniques are leveraged in order to imbue software with human qualities. Heuristics (rules of thumb) are used to assist software in reaching conclusions and making recommendations.

Case-based reasoning is an area of conventional AI in which the AI software maintains a library of problem cases and solutions. When confronted with a new problem, it adjusts and applies a relevant previous solution to the new problem. Case-based reasoning relies on the four Rs for solving new problems:

- Retrieve: Find cases in the case library relevant to the new problem.
- Reuse: Map the solutions from a previous case(s) to the variables of the new problem, adjusting where necessary.
- Revise: Test the solution and revise if necessary.
- Retain: Store the new problem and its newly acquired successful solution as a case in the library.

AI software built on case-based reasoning is useful for help-desk support, medical diagnosis, and other situations where there are many similar cases to diagnose.


A Bayesian network, sometimes called a belief network, is a form of conventional AI that uses a graphical model to represent a set of variables and their relationships and dependencies. A Bayesian network provides a model of a real-life scenario that can be incorporated into software to create intelligence. Bayesian networks are used in medical software, engineering, document classification, image processing, military applications, and other activities. For example, the Combat Air Identification Fusion Algorithm (CAIFA) Bayesian network is used in missile defense software and draws on multiple, diverse sources of ID evidence to determine the allegiance, nationality, platform type, and intention of targeted aircraft.

Behavior-based AI is a form of conventional AI that is popular in programming robots. It is a methodology that simulates intelligence by combining many semi-autonomous modules. Each module has a specific activity for which it is responsible. By combining these simple modules, the resulting system exhibits intelligent behavior—it becomes smarter than the sum of its parts.

Expert systems are very successful forms of conventional AI and are covered next.

Reading: Expert System

Expert System



Artificial Intelligence & AI & Machine Learning. Authored by: mikemacmarketing. License: CC BY 2.0

An **expert system (ES)** is a form of conventional AI that is programmed to function like a human expert in a particular field or area.

Why This Matters

Many professional activities are tedious, redundant, and dangerous. In such cases, expert systems can assist human professionals. In other cases, decisions are made that can be programmed using AI techniques. In these cases, expert systems can provide solutions as a human expert would, freeing the human expert to focus on more creative and productive activities.

Essential Information

Expert systems are created with the assistance of a human expert who provides subject-specific knowledge. The rules applied by the expert to perform some activity are programmed into the expert system software.

Like human experts, computerized expert systems use heuristics, or rules of thumb, to arrive at conclusions or make suggestions. A heuristic provides a solution to a problem that can't necessarily be proven as correct but usually produces a good result.

Many expert systems are used in various professions every day. Computerized expert systems have been developed to diagnose diseases given a patient's symptoms, suggest the cause of a mechanical failure of an engine, predict weather events, assist in designing new products and systems, test and

evaluate chemical reactions, or detect possible terrorist threats. Some expert systems have become commonplace; users think of them simply as helpful software. For example, tax preparation software is yet another implementation of expert systems.

Reading: Computational Intelligence

Computational Intelligence



[aman_geld 03.01.04 \(neural network\)](#). Authored by: cea +, License: CC BY 2.0

Computational intelligence is an offshoot of AI that employs methodologies such as neural networks, fuzzy systems, and evolutionary computation to set up a system whereby the software can develop intelligence through an iterative learning process.

Why This Matters

Rather than programming intelligence into software through the use of statistics and probabilities, as in conventional AI, computational intelligence creates software that can learn for itself. Computational intelligence is where the big tech companies are investing their research dollars in order to develop systems that can go beyond human capabilities in analyzing large amounts of data.

Essential Information

There are many mathematical and logical techniques that contribute to successful computational intelligence algorithms. Two of the more important techniques are fuzzy logic and evolutionary computation.

Computers typically work with numerical certainty: certain input values always result in the same output. However, in the real world, certainty is not always the case. To handle this dilemma, a specialty research area in computer science called fuzzy systems, based on fuzzy logic, has been developed. Fuzzy logic is derived from fuzzy set theory, which deals with reasoning that is approximate rather than precise. A simple example of fuzzy logic might be one in which cumulative probabilities do not add up to 100%, a state that occurs frequently in medical diagnosis.

Evolutionary computation includes areas of AI that derive intelligence by attempting many solutions and throwing away the ones that don't work—a "survival of the fittest" approach. A genetic algorithm is a form of evolutionary computation that is used to solve large, complex problems where a number of algorithms or models change and evolve until the best one emerges. If you think of program segments as building blocks similar to genetic material, this process is similar to the evolution of species, where the genetic makeup of a plant or animal mutates or changes over time. Some investment firms use genetic algorithms to help select the best stocks or bonds.

Neural networks, covered next, are a very popular form of computational AI that attempts to simulate the functioning of a human brain.

Reading: Neural Network

Neural Network

A **neural network**, or neural net, uses software to simulate the functioning of the neurons in a human brain.

Why This Matters

Simulating a human brain in a neural network provides many useful applications. However, an additional, perhaps even more important, result is a deeper understanding of how the human mind functions. Advances in neural networks are providing insight into treatment for mental illness and other brain-related diseases. Ultimately, the study of neural networks may reveal the essence of the human soul.

Essential Information

Just as human brains gain knowledge by generating paths between neurons through repetition, neural networks are programmed to accomplish a task through repetition. A neural network typically begins by running a random program. The output is measured against a desired output, which positively or negatively affects the pathways between simulated neurons. When the neural net begins getting closer to the desired output, the positive effects on the system begin to create circuits that become trained to produce the correct output. Similar to a child's learning, a neural net discovers its own rules. For example, a neural net can be trained to recognize the characteristics of a male face, and once trained, may be able to accurately identify faces in photos as male or female.

Once a neural net is trained, it can process many pieces of data at once with impressive results. Some specific features of neural networks include:

- The ability to retrieve information even if some of the neural nodes fail
- The ability to quickly modify stored data as a result of new information
- The ability to discover relationships and trends in large databases
- The ability to solve complex problems for which all of the information is not present

Neural networks excel at pattern recognition, and this ability can be used in a wide array of applications, including voice recognition, visual pattern recognition, robotics, symbol manipulation, and decision making.

Recently, researchers have developed new neuron modeling techniques in a technology known as deep learning. Deep learning is a class of computer algorithms that use many layers of processing units to mimic the brain's parallel-processing patterns. Each successive layer uses the output from the previous layer as input to tackle more complex problem scenarios. Deep learning was the technique Google's DeepMind unit used in developing its winning Go algorithm (see video). Google and the other big tech companies are investing heavily in developing expertise around these computational methods. They may be the key to unlocking the truly intelligent machine.

Reading: Singularity

Singularity

The **Singularity**, or more specifically, the technological Singularity, is the point in time at which computers exceed humans in intelligence, launching a new era of innovation.

Why This Matters


Should it come to pass, a technological Singularity could completely change life as we know it. Those who subscribe to the theory believe that the Singularity will usher in an era of deeper understanding and rapid advancement. Others believe that the Singularity will never occur or that if machines become more intelligent than humans, they are liable to take over the world.

Essential Information

The theory of a technological Singularity was originally proposed by mathematician and sci-fi author Vernor Vinge. Vinge believed that advances in AI and human biological enhancement would eventually lead to “superintelligences” with the ability to advance science and technology well beyond what is possible with the human brain alone. Since these superintelligences have knowledge beyond humans, it is impossible to predict what might transpire after the Singularity. In his book *The Singularity Is Near*, Ray Kurzweil takes an optimistic view of the near future and the impending Singularity. He classifies the days ahead as the “most transforming and thrilling period in history.” He goes on to say, “It will be an era in which the very nature of what it means to be human will be both enriched and challenged as our species breaks the shackles of its genetic legacy and achieves inconceivable heights of intelligence, material progress, and longevity.”

The rationale behind the theory of the Singularity is tied to Moore’s Law and other technological trends. For the past four decades, processing power has grown exponentially. Unless something arises to interrupt it, it will continue to grow to become more powerful than the human brain.

The Blue Brain project is a bold attempt to simulate a human brain at a molecular level in software running on one of the world’s fastest supercomputers, IBM’s Blue Gene. The project has succeeded in simulating the functioning brain of a mouse. Some speculate that we will have effective software models of the human brain in the next few decades. Kurzweil placed the date for the Singularity somewhere in the mid-2040s. He expects that shortly thereafter a single computer will become smarter than the entire human race. The concept of machines with intelligence that exceeds human intelligence is referred to as strong AI.



Authored by: Courtesy of Wikipedia, License: Fair Use

Reading: Turing Test

Turing Test



Turing Test. Authored by Gwydion M. Williams, License: CC BY 2.0

The **Turing Test** was devised by Alan Turing as a method of determining if a machine exhibits human intelligence.

Why This Matters

How can you tell if a computer has acquired the capability for human-level intelligence? This is an important question, especially now that supercomputers have become as powerful at processing

information as the human brain. Alan Turing thought it was an important question as well when he wrote his famous paper decades ago. Turing believed that thinking machines were just around the corner. Researchers today, over 65 years later, are thinking the same thing.

Essential Information

Alan Turing was an English mathematician, logician, and cryptographer who is well known for his belief that computers would someday be as intelligent as humans. In his 1950 paper entitled “Computing Machinery and Intelligence,” Turing proposed a test, now known widely as the Turing Test, which he claimed would be able to determine if a computer exhibited human intelligence.

Turing described his test as follows: “A human judge engages in a natural language conversation with two other parties, one a human and the other a machine; if the judge cannot reliably tell which is which, then the machine is said to pass the test.” Passing Turing’s test is no simple feat. In June of 2014, a computer simulating a 13-year-old boy was able to fool 33% of the judges and claim the prize at a competition hosted by the Royal Society in London.

The intelligence of any computer passing the Turing Test is likely to be disputed. Many AI researchers have challenged Turing’s assumption, claiming that clever use of language does not necessarily imply intelligence. However, many of us are getting used to carrying on limited conversations with machine intelligences named Siri, Cortana, Alexa, Google, and others. It is no longer a stretch of the imagination to picture full conversations with such artificial personalities as if they were our human friends.

Lesson 5.2: Applications of Artificial Intelligence

Lesson 5.2 Introduction

AI applications are ways in which AI methodologies are applied to solve problems and provide services.

Research and innovation in artificial intelligence are impacting lives in subtle and not-so-subtle ways. From software that anticipates the user’s needs to robots that assist the elderly, to applications that predict trends in the stock market, to cars that drive themselves, artificial intelligence is being used in many applications.

Artificial intelligence has many and varied applications. AI techniques can be integrated into systems and applications to make them more intuitive. AI software is used to automate various types of decision-making logic. AI software can be connected to sensors and mechanical apparatus to create a variety of robots and other devices. Auto manufacturers are testing self-driving cars that utilize AI to detect and automatically respond to their surroundings. Amazon, UPS, and others are developing autonomous drones to deliver packages. This section explores these and many other common AI applications.

Reading: Robotics

Robotics



Robotics. Authored by: TheDigitel Myrtle Beach. License: CC BY 2.0

Robotics involves developing mechanical or computer devices to perform tasks that require a high degree of precision or are tedious or hazardous for humans.

Why This Matters

In recent years, robots have become a form of entertainment in toys for children and a necessity in many manufacturing processes. Contemporary robotics combines both high-precision machine capabilities and sophisticated AI controlling software. In recent years, robot capabilities have improved greatly, and it is likely that robots will soon be commonplace in many environments.

Essential Information

There are many applications of robots, and research into these unique devices continues. For many businesses, robots are used to do the three Ds: dull, dirty, and dangerous jobs. Manufacturers use industrial robots to assemble and paint products. Medical robots enable doctors to perform surgery via remote control. Sitting at a console, the surgeon can replace a heart valve or remove a tumor. Some surgical robots cost more than \$1 million and have multiple surgical arms and sophisticated vision systems.

Some robots can propel themselves by rolling, walking, hopping, flying, snaking, or swimming. Other robots are designed to interact with humans using speech recognition and gesture and facial expression interpretation. Increasingly, robots are being used to transport items in office buildings. Robots have become commonplace in combat zones, flying overhead and traversing mined roads to deactivate bombs. Some have predicted that robots will become common household appliances within the next decade.

Honda claims that its ASIMO is the world's most advanced humanoid robot. ASIMO can navigate a variety of terrains, walk up and down stairs, play sports, and serve drinks. Boston Dynamics has developed a team of biped and quadruped robots that are able to walk, run, and perform all sorts of physical functions. The iRobot company makes robots for the home that vacuum the carpets, clean the gutters, and more. Recently, iRobot spun off its military robotics division into a company named Endeavor Robotics.

Since technologies have advanced to the point where the construction of useful humanoid robots is not only feasible but also affordable, more researchers and companies are investing in robotics research. Many predict that household robots will become commonplace over the next decade. In other cases, robotics is used to augment human capabilities. A Pittsburgh woman who is paralyzed from the neck down has recently taken control of robotic arms that, through a brain implant, she controls with the fluidity of a natural arm.

Reading: Autonomous Vehicles

Autonomous Vehicles



Autonomous Vehicle. Authored by: wuestenigel. License: CC BY 2.0

Autonomous vehicles include aircraft, ground vehicles, and watercraft that operate autonomously, without the aid of a human pilot, driver, or captain, such as driverless cars and drone aircraft.

Why This Matters

Advances in artificial intelligence, processing speed, and other technologies have made it possible for moving vehicles (cars, trucks, aircraft, boats, and ships) to navigate and control themselves through all kinds of environments according to existing rules and laws. The primary thing holding these technologies back from going mainstream is the lack of provisions in laws that would allow autonomous vehicles to join human-controlled vehicles in traffic. Such provisions are being crafted and soon we all will be taking advantage of the opportunities that autonomous vehicles provide.

Essential Information

Of all the AI-driven technologies, autonomous vehicles probably get the most media attention. Self-driving cars from Google, Tesla, and others are already on public roads being tested, with human drivers at the ready should they need to intervene. Amazon, UPS, and others are testing drone-delivery systems that promise same-day delivery. These fantastic inventions are poised to take off, but are humans ready to give up control?

Self-driving, sometimes called driverless, cars have been in development for years. In 2004, DARPA held the first Grand Challenge with a prize of \$1 million for the autonomous vehicle that could traverse a course across the Mojave Desert. None of the entrants were able to complete the challenge. By 2005, five vehicles completed a more challenging DARPA Grand Challenge, and in 2006, six vehicles autonomously navigated through empty city streets in the Urban Challenge. Since then car and truck manufacturers have been working hard to refine the technology so that it is safe for populated city streets.

There are several approaches to developing self-driving cars and vehicles. Small autonomous golf-cart-sized vehicles are available as a cab service on some private campuses. IBM's Watson is the intelligence behind OLLI, the shuttle that transports pedestrians through the streets of National Harbor, Maryland. Some auto manufacturers are adding automation gradually so that drivers gradually get used to giving up control of the vehicle. Many cars now have the ability to park themselves. Tesla evolved its driver assist features to the point where its current models can now completely drive themselves. Others, like Google, are creating fully autonomous vehicles right from the start, hoping consumers will be ready for the dramatic shift. Think of what autonomous vehicles will mean to a company like Uber! The company is testing its self-driving technology on the streets of Pittsburgh. Soon you might be able to call up a driverless Uber ride in your town. Uber is also investing in self-driving semi-trucks.

Self-driving cars offer benefits by reducing driver error and saving fuel and time. Drivers can give up the controls and become productive during their commutes. Autonomous vehicles are believed to be safer than human-driven vehicles. They eliminate concern about distracted drivers, drunk drivers, or incompetent drivers. As self-driving vehicles gain the ability to communicate with each other and traffic signals over wireless mesh networks, savings and safety will increase even more.

Autonomous vehicles will mean huge savings for the commercial delivery industry. Autonomous trucks from traditional manufacturers like Freightliner and new companies like Otto Motors are poised to reshape the trucking industry. Beyond removing human error from the equation, autonomy also provides new ways for trucks to collaborate. In a technique called platooning, several autonomous trucks link together in a tightly knit convoy, with the lead truck in control. The technique is believed to save 20 percent in fuel costs and additional savings in human resources. On the down side, truck drivers are concerned about the future of their livelihood.

In 2013, Amazon CEO Jeff Bezos caused quite a stir when he unveiled a new delivery system in development that involved autonomous drones. It was soon discovered that Google had been working on the same technology for years, and soon everyone from UPS to the convenience store 7-Eleven were working on drone delivery systems. This was during the same period that hobbyist remote-controlled drones were coming down in price and becoming popular. The Federal Aviation Association became quite alarmed at the potential threat to air traffic and public safety and passed new rules that require licenses for flying remote-controlled drones and disallowed autonomous drones. Since then companies have been working to convince the FAA to allow the use of drones for delivery.

In the meantime, companies like Starship Technologies are developing small, six-wheeled robots to deliver packages to businesses and homes. Mercedes has created a "Robovan" to transport eight robots to a neighborhood and launch them into the streets to deliver their packages. The typical volume of deliveries for a human delivery person is 180 packages per nine-hour shift. The Robovan and its robots will more than double that to 400 packages.

In Amsterdam, autonomous vehicles have taken to the water. "Roboats" are programmed to motor pedestrians to their destination via canal routes. The small, two-person platforms can also be joined together to create a bridge across the canal.

It is clear to see from these examples that there is a new world around the corner—one filled with autonomous vehicles and robots. How will these technologies impact our daily routines? Will they increase human safety and efficiency? Or will this constitute the first step of the AI-controlled robotic takeover?

Reading: Computer Vision

Computer Vision



Computer Vision Syndrome. Authored by downloadsource.fr, License: CC BY 2.0

Computer vision combines hardware (cameras and scanners) and AI software that permit computers to capture, store, and interpret visual images and pictures.

Why This Matters

Computer vision enables software to react to visual input. This ability provides all sorts of opportunities for automation. No longer does the computer need to wait for manual input; it can be programmed to automatically react to movement or objects in its environment. This is useful in all forms of robotics, in security systems, in games, and in a wide variety of applications that benefit from visual input.

Essential Information

Although seeing is second nature for most humans, it is a complicated capability to program into computers. To process input from video cameras or simply to analyze a photograph requires complicated AI techniques. Through those techniques, computers have learned to notice irregularities in video and still photos and to distinguish colors, motion, and depth. And they're getting better all the time! At the 2016 Large Scale Visual Recognition Challenge—hosted by Stanford, Princeton, and Columbia Universities—the accuracy in computers' correctly identifying objects in common photos increased to 66.3%. All contestants utilized deep learning techniques in their algorithms.

Computers are used to interpret medical images like x-rays and ultrasonic images. They are used in pattern recognition, such as face recognition in security systems. Computer vision is used in the military to guide missiles, and in space exploration to guide the Mars Rovers.

Microsoft brought computer vision into the living room with the Kinect Xbox 360 controller. The set-top box allows players to put away their handheld controllers and control game action with body action. Kinect applications have been developed to help stroke victims and children with autism, to assist doctors in the operating room, to create new forms of art, to assist the disabled, and to implement other valuable services.

Computer vision systems are also being developed to assist in human sight. Dr. William Dobbelle designed a system that has brought a form of vision to totally blind individuals. The system combines a small video camera mounted in sunglasses, a powerful but small computer worn in a belt, and a brain implant that stimulates the visual cortex. By connecting an array of electrodes to the areas of the visual cortex, the system can provide vision with a grid of. By enabling a video camera to control the intensity of power to the electrodes, one patient has been able to see well enough to drive a car.

Reading: Natural Language Processing

Natural Language Processing

Natural language processing uses AI techniques to enable computers to generate and understand natural human languages, such as English.

Why This Matters

Natural language processing makes it possible to interact with computer systems using spoken words. Apple's Siri, Microsoft's Cortana, Google's Assistant, and Amazon's Alexa act as personal digital assistants that attempt to answer any question spoken by a user using natural language processing. Xfinity's X1 system allows users to speak into the remote to call up any television show, movie, channel, or actor. Natural language processing is useful for individuals who cannot use a keyboard and mouse because they are disabled or because their hands are needed for other work. Increasingly, people are taking advantage of natural language processing to dictate voice input to a computer rather than typing text or commands.



Authored by Business Wire, License: Fair Use

Essential Information

One area of natural language processing that is evident in today's software and business systems is speech recognition. Speech recognition enables a computer to understand and react to spoken statements and commands. With speech recognition, it is possible to speak into a microphone connected to a computer, or into a mobile phone, and have the spoken words converted to commands or text displayed on the screen. Like Siri and Cortana, Alexa on Amazon's Echo device fields spoken questions and commands, but from a shelf in your home. Google Home performs the same function.

Current versions of many desktop and mobile operating systems have utilities that enable you to speak commands to the computer as well. Microsoft Word has a speech utility that enables you to dictate to the computer instead of typing. Dragon NaturallySpeaking, which allows users to dictate text as they would naturally speak, is popular software for both PCs and cell phones. As with all dictation software, you must speak punctuation characters and commands to place the cursor in a new location. For example, you might dictate this statement, closing with "period – new line."

Google Voice is able to deliver voicemail as text transcriptions. Google recently improved its transcription capability by 50% with new technology developed through deep learning technology.

Speech recognition is not to be confused with voice recognition, which is a different technology used in the security field to identify an individual by the sound of his or her voice.


Speech recognition must utilize AI methodologies to overcome some difficult obstacles:

- **Speech segmentation:** If you listen to a person speak from a computer's perspective, you will notice that there is roughly the same length of pause between syllables in words as there is between words in a sentence (roughly none). How is the computer to determine where one word ends and the next begins?
- **Ambiguity:** There is no difference between the sound of the words "their," "there," and "they're." When spoken, how is the computer to know which one is intended?
- **Voice variety:** Everyone speaks differently. The English language has many regional and personal dialect variations. It is difficult for one program to understand both Arnold Schwarzenegger and Kim Kardashian.

Using AI, speech recognition software "learns" a dictionary of words to determine how to parse a sentence into individual words. It must also learn grammatical rules in order to overcome the problem of ambiguity. In addition, it must have some insight into human thinking to understand a sentence such as "Time flies like an arrow." Increasingly, speech recognition relies on context to help determine what is being spoken. Google, for example, has access to massive amounts of user data that it collects through Gmail, Google Search, YouTube, Android OS, and other Google software. This information provides Google with a context in which to interpret and satisfy a user's needs.

Reading: Pattern Recognition

Pattern Recognition



[pattern.. Authored by Curt Smith, License: CC BY 2.0](#)

Pattern recognition is an area of AI that develops systems that are trained to recognize patterns in data.

Why This Matters

Many useful applications are based on computers' ability to recognize patterns in data. Pattern recognition is used to find exceptions to trends in data for businesses in a process called data mining. Such findings can help in discovering fraud or other unusual business activities. Pattern recognition is also useful for training AI computer vision systems to recognize faces, handwriting, or other visual patterns.

Essential Information

Speech recognition, handwriting recognition, and facial recognition all fall under the AI category of pattern recognition. Handwriting recognition uses AI techniques in software that can translate handwritten characters or words into computer-readable data. The technology is useful in mobile devices, like Microsoft's Surface Pro, where a stylus can be used to write on the touch display.

For governments and other organizations concerned with security, facial recognition is a hot area of AI research. Facial recognition uses cameras and AI software to identify individuals by unique facial characteristics. Facial recognition is used in authentication—verifying a person's identity before allowing access to secure areas or systems—and in surveillance. The Australian SmartGate system uses facial-scan technology to compare the traveler's face against the scan encoded in a microchip contained in the traveler's ePassport. Disney is experimenting with foot recognition to identify visitors in its theme parks without infringing on their privacy.

Facial recognition is catching on as a handy tool for organizing photos as well. Apple's iPhoto, Facebook, and other photo software uses facial recognition technology to automatically recognize and tag friends in digital photos. The tool is trained to recognize individuals with the help of the user who corrects the software when it incorrectly identifies a face. Over time the facial recognition gets more competent and can call up all the photos that include a specified person.

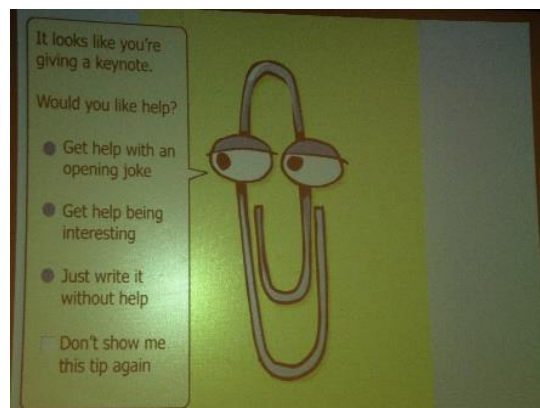
Pattern recognition is used in the iPhone, iPad, and MacBook Touch ID technology that allows users to bypass an access code with the press of a fingertip. Using physical characteristics such as fingerprints to access computers is referred to as biometric authentication.

Perhaps the most valuable use of pattern recognition is in mining big data. The amount of data being generated by online systems has grown beyond our ability to manage it. By 2020, the amount of digital data produced is expected to exceed 40 zettabytes, of which roughly 33 percent of will contain valuable information. Artificial intelligence is needed to explore that data, using pattern recognition techniques, to find trends and outliers, and produce useful, actionable information.

The Internet of Things is expected to add another 400 zettabytes to data generated by 2018. This is data generated by sensors, appliances, cameras, and many other physical objects. Utilizing pattern recognition in AI technologies like deep learning will give computers the ability to make predictions about a variety of environments and human behavior to serve humanity in small ways, like designing homes and products, to critical ways, like saving the planet from greenhouse gasses.

Reading: Personal Digital Assistant

Personal Digital Assistant



Clippy!. Authored by: Matt DiGirolamo, License: CC BY 2.0

A **personal digital assistant** is a program that draws from a knowledge base in order to answer questions and perform tasks for a person.

Why This Matters

One of the earliest personal digital assistants was a little animated paperclip named Clippy who assisted Microsoft Word users. Clippy would say things like “It looks like you’re writing a letter. Would you like some help?” While sometimes helpful, Clippy soon developed a reputation of being annoying. For decades, we predicted that computers would someday have the ability to anticipate a user’s needs and make life easier, but only recently have we seen that technology begin to emerge. Apple’s Siri, Microsoft’s Cortana, Google Assistant, and Amazon’s Alexa have provided us with a glimpse of how personal digital assistants can serve humanity.

Essential Information

Software referred to as intelligent agent software has been used for years to carry out specific tasks such as searching for data on the web. This software has evolved into general-purpose software designed to assist with all types of questions and tasks. Siri, Cortana, Alexa, and Google Assistant have taken intelligent agents to a whole new level.

Apple’s Siri interface ushered in the era of the personal digital assistant for the iPhone. Google followed close behind with its Google Now app (updated to Google Assistant) that runs on Android and iOS devices, and Microsoft followed suit with its Cortana personal digital assistant for Windows devices. Amazon introduced a home device named Echo, which features Alexa as the personal assistant that responds any time her name is called in the house. Google responded with Google Home, a tabletop device that behaves like Echo with the ability to answer questions and carry on more complex conversations. Meanwhile, Cortana expanded into Windows PCs, and Siri into Macs and MacBooks. So now, people have the opportunity to call on their personal digital assistant conveniently from any location to get assistance with Internet-related services, phone services, home-related questions, and personal computer-related tasks. Clearly the battle of the personal assistants is in full swing.

Using natural language processing, these apps allow users to ask questions or issue commands using natural spoken language such as “Remind me to pick up milk on my way home” and “Any good burger joints around here?” The app responds to commands and questions with a human-like voice that answers the question directly or provides a related website. These apps work their magic by sending a query to a server (in the cloud), where AI software dissects the request, runs it through the most appropriate online service, and responds with a solution.

Since the technology is far from perfect, mistakes are inevitable and sometimes humorous. There are several websites dedicated to amusing Siri mistakes. However, the technology continues to improve. Google Assistant remembers the previous questions and is able to reply to follow up questions in a conversational manner. It is likely that these technologies will continue to improve through AI technologies like deep learning, until they can easily pass the Turing Test.

In addition to accepting commands and fielding questions, a personal digital assistant can also take dictation and send it through email or text message. Many people dictate text messages through a digital assistant while driving so that they don’t have to take their hands off the wheel.

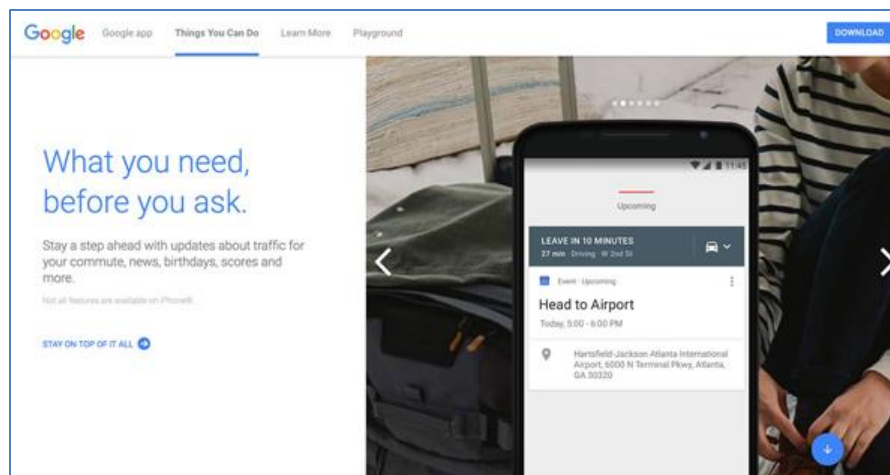
Reading: Context-Aware Computing

Context-Aware Computing

Context-aware computing refers to software that uses artificial intelligence to provide services based on environmental context.

Why This Matters

Traditional software has no awareness of the environment in which it is running. That is all changing. Like Siri, Cortana, and Google Assistant, most software will soon be able to tap into who is using it and where and how it is being used, to provide services specifically designed for that moment and situation. In other words, software is becoming much smarter.



Authored by: Courtesy of Google Inc., License: Fair Use

Essential Information

Context-aware software is able to tap into user preferences, current location and time, social connections, and other attributes to improve services. For example, context-aware security is an early application of this new capability. By understanding the context of a user request, a context-aware security application can adjust its response to better protect the user. Have you ever had a bank or credit card company contact you because your account was accessed in an unusual way—from a foreign location or multiple times within an hour? This is an example of a context-aware security system.

Google is working to do away with the need for passwords on Android phones by using a context-aware technology it calls “trust scores.” Factors including typing speed, vocal inflexions, facial recognition, and proximity to familiar Bluetooth devices and Wi-Fi hotspots are used to calculate the score, and if the score is high enough you can access apps without passwords.

Intelligent software is the product of our increasingly connected lives. Our mobile devices provide detailed information about our location and activities at any given moment. This information informs information systems about our lifestyle and preferences. As the Internet of Things continues to grow, increasing amounts of environmental information will be at the disposal of software and information systems. Although some may consider the collection of all of this information an invasion of privacy, it will give software the ability to provide information and services we need exactly as we need them.