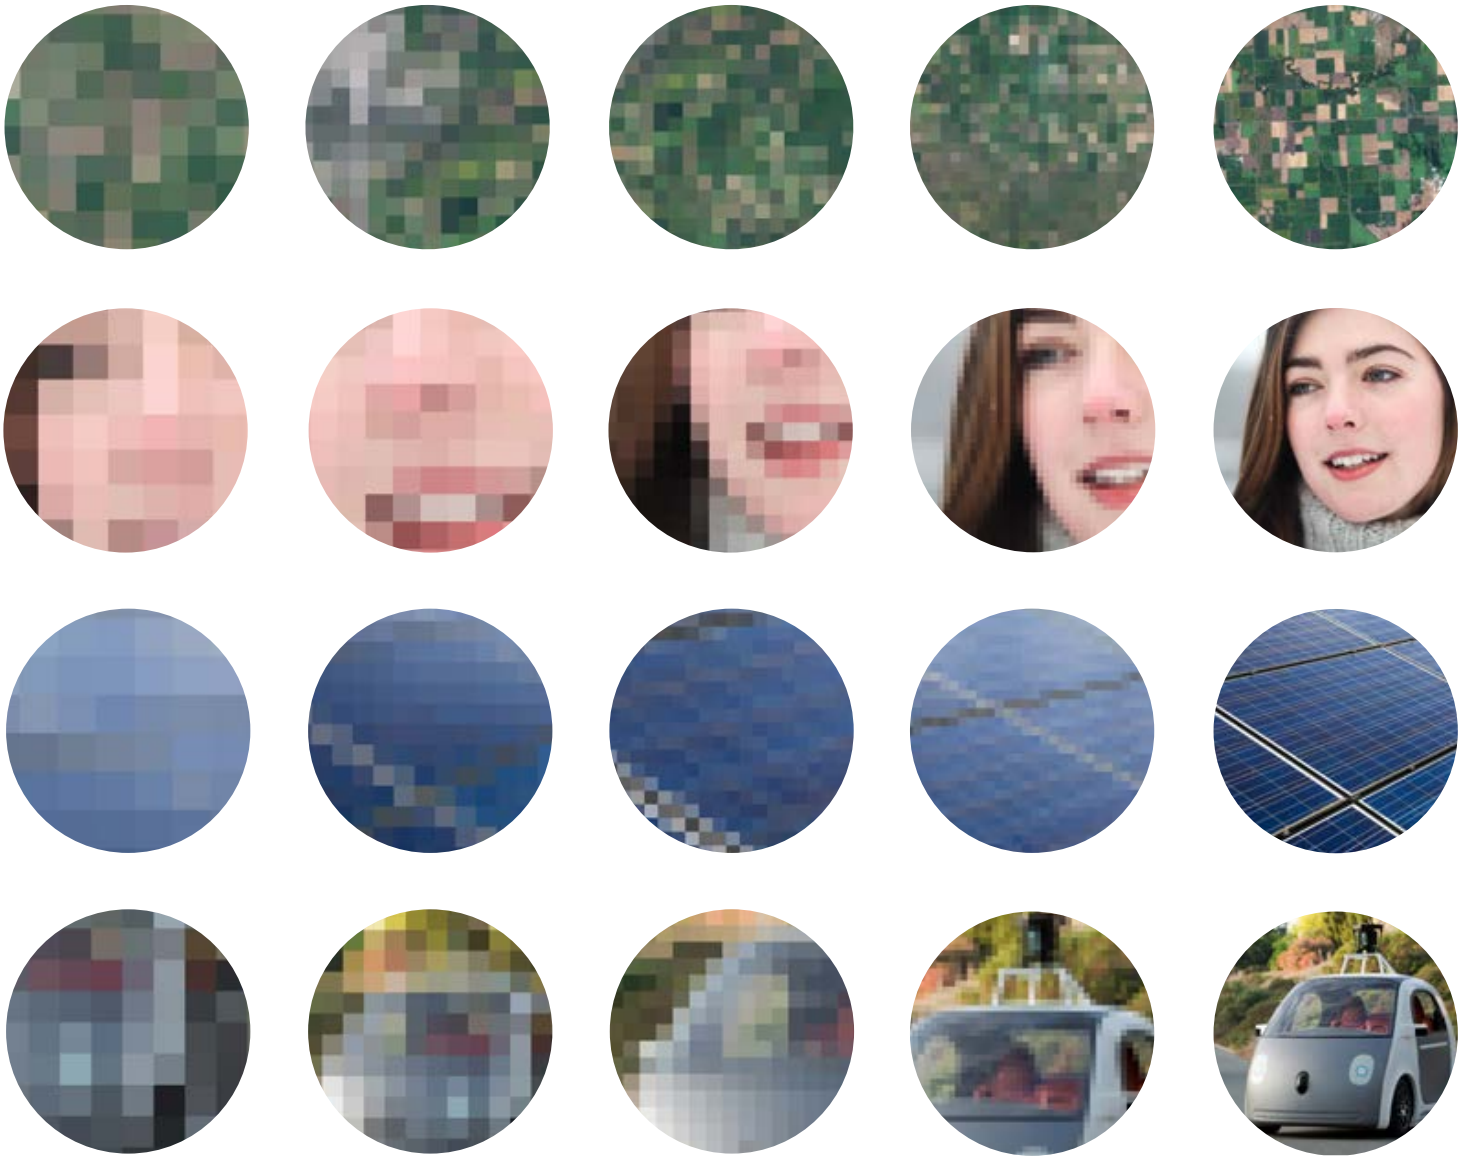

The Promise of Artificial Intelligence

Daniel Castro and Joshua New

CENTER FOR DATA INNOVATION

The Promise of Artificial Intelligence

Daniel Castro and Joshua New

October 2016

Table of Contents

Introduction

Artificial intelligence (AI) is on a winning streak. In 2005, five teams successfully completed the DARPA Grand Challenge, a competition held by the U.S. Defense Advanced Research Projects Agency to spur development of autonomous vehicles.[1] In 2011, IBM's Watson system beat out two longtime human champions to win *Jeopardy!*[2] In 2016, Google DeepMind's AlphaGo system defeated the 18-time world-champion Go player.[3] And thanks to Apple's Siri, Microsoft's Cortana, Google's Google Assistant, and Amazon's Alexa, consumers now have easy access to a variety of AI-powered virtual assistants to help manage their daily lives. The potential uses of AI to identify patterns, learn from experience, and find novel solutions to new challenges continue to grow as the technology advances.

Moreover, AI is already having a major positive impact in many different sectors of the global economy and society.[4] For example, humanitarian organizations are using intelligent chatbots to provide psychological support to Syrian refugees, and doctors are using AI to develop personalized treatments for cancer patients. Unfortunately, the benefits of AI, as well as its likely impact in the years ahead, are vastly underappreciated by policymakers and the public. Moreover, a contrary narrative—that AI raises grave concerns and warrants a precautionary regulatory approach to limit the damages it could cause—has gained prominence, even though it is both wrong and harmful to societal progress.[5]

To showcase the overwhelmingly positive impact of AI, this report provides a description of the major uses of AI as well as details on 70 real-world examples of how AI is already generating social and economic benefits. Policymakers should consider these benefits as they evaluate the steps they can take to support the development and adoption of AI.

What Is Artificial Intelligence?

AI is a field of computer science devoted to creating computing machines and systems that perform operations analogous to human learning and decision-making. As the Association for the Advancement of Artificial Intelligence describes it, AI is "the scientific understanding of the mechanisms underlying thought and intelligent behavior and their embodiment in machines."[6] As described below, the level of intelligence in any particular implementation of AI can vary greatly, and the term does not imply human-level intelligence.

AI involves many functionalities, including but not limited to: a) learning, which includes several approaches such as deep learning (for perceptual tasks), transfer learning, reinforcement learning, and combinations thereof; b) understanding, or deep knowledge representation required for domain-specific tasks, such as cardiology, accounting, and law; c) reasoning, which comes in several varieties, such as deductive, inductive, temporal, probabilistic, and quantitative; and d) interaction, with people or other machines to collaboratively perform tasks, and for learning from the environment.

The promise of artificial intelligence has been around since the era of electromechanical computing began after WWII. But hope has long raced far ahead of reality. At the first AI conference, held at Dartmouth College in 1956, the view was that AI could almost be achieved in one summer's worth of work. In the 1960s and 1970s, some computer scientists predicted that within a decade we would see machines that could think like humans. In 1965, the Nobel laureate Herbert Simon predicted that "machines will be capable, within 20 years, of doing any work a man can do."[7] Two years later, AI pioneer Marvin Minsky predicted, "In from three to eight years we will have a machine with the general intelligence of an average human being."[8] Around the same time, legendary computer scientist John McCarthy told the U.S. Department of Defense that it would be possible to build "a fully intelligent machine" in the space of a decade.[9]

Needless to say, the timeline was vastly underestimated and likely still is today. Even the minimal progress that was achieved came slowly, including two distinct periods those in the field call "AI winters," when progress slowed and funding decreased. However, AI has seen a resurgence in recent years as a result of the development of machine learning—a branch of AI that focuses on designing algorithms that can automatically and iteratively build analytical models from new data without explicitly programming the solution. Before machine learning, computer scientists had to program a wide array of functions into a system for it to mimic intelligence. This capability has emerged for a number of reasons, including better hardware, such as faster processers and more abundant storage, more data, and better algorithms. Machine learning is now commonplace: Pandora learns how to make better music recommendations based on its users' preferences; Google learns how to automatically translate content into different languages based on translated documents found online; and Facebook learns how to identify people in photos based on its database of known users.[10]

An important development in machine learning is deep learning. Deep-learning algorithms use statistical techniques to develop a model to solve problems from large, complex datasets with very little guidance from programmers. As computer scientist Amit Karp writes, "deep learning relies on simulating large, multilayered webs of virtual neurons, which enable a computer to learn to recognize abstract patterns."[11] It is called "deep" because it automatically generates multiple layers of abstractions of the data and uses these abstractions to identify patterns. As deep learning and other machine-learning techniques continue to improve, they will bring significant benefits to individuals and societies along the way.

The cause of many misconceptions about AI, particularly its potential harms, is the difference between weak AI and strong AI. Weak AI, which is the focus of this report, is adept at performing particular types of tasks, but only those types of tasks—somewhat like a technological savant.[12] For example, the AI that makes Apple's Siri virtual assistant so good at interpreting voice commands does not enable Siri to drive a car, predict weather patterns, or analyze medical records. While machine-learning algorithms exist that can accomplish those tasks, they too are narrowly constrained—the AI powering a self-driving Uber will not be able predict a hurricane's trajectory or help doctors diagnose a patient with cancer.

Conversely, strong AI, often referred to as artificial general intelligence (AGI), is a hypothetical type of AI that can meet or exceed human-level intelligence and apply this problem-solving ability to any type of problem, just as the human brain can easily learn how to drive a car, cook food, and write code.[13] Many of the dystopian fears about AI—that it will eliminate most jobs or decide to wipe out humanity, for example—stem from the notion that AGI is feasible, imminent, and uncontrollable.[14] While AGI could in theory exist eventually, not only is there no reason to believe that these concerns will manifest, but it will likely be a very long time before humanity can develop AGI, if it can develop it at all. Computers have grown exponentially more powerful since technologists began predicting AGI was just a few years away in the 1960s, and while AI has progressed dramatically since then, human-level AI may be just as far off as it was decades ago.[15]

The Value of AI

Because AI will continue to evolve and work its way into a wide variety of applications, it is difficult to predict just how much value AI will generate. The International Data Corporation estimates that in the United States the market for AI technologies that analyze unstructured data will reach $40 billion by 2020, and will generate more than $60 billion worth of productivity improvements for businesses in the United States per year.[16] Investors in the United States are increasingly recognizing the potential value of AI, investing $757 million in venture capital in AI start-ups in 2013, $2.18 billion in 2014, and $2.39 billion in 2015.[17] The McKinsey Global Institute estimates that by 2025 automating knowledge work with AI will generate between $5.2 trillion and $6.77 trillion, advanced robotics relying on AI will generate between $1.7 trillion and $4.5 trillion, and autonomous and semi-autonomous vehicles will generate between $0.2 trillion and $1.9 trillion.[18] A report from Accenture examining the potential impact of AI on economic growth found that by 2035, AI could increase the annual growth rate of the U.S. and Finnish economies by 2 percentage points, the Japanese economy by 1.9 points, and the German economy by 1.6 points.[19] The report also found that, for the 12 countries surveyed, AI would boost labor productivity rates by 11 to 37 percent.[20]

The social benefits of AI are similarly substantial, though harder to quantify. As Facebook's chief technology officer Mike Schroepfer puts it, "The power of AI technology is it can solve problems that scale to the whole planet," such as climate change and food insecurity.[21] AI is already delivering valuable social benefits today, such as by helping authorities rapidly analyze the deep web to crack down on human trafficking, fighting bullying and harassment online, helping development organizations better target impoverished areas, reducing the influence of gender bias in hiring decisions, and more.[22] Just as AI can help businesses make smarter decisions, develop innovative new products and services, and boost productivity to drive economic value, it can achieve similar results for organizations generating social value.

Typology of AI Applications

A vast and diverse array of applications use AI, with algorithms powering everything from smartphone apps that help consumers with their holiday shopping, to accelerating the process of discovering new lifesaving drugs.[23] Most uses of AI have at least one of seven functions: monitoring; discovering; predicting; interpreting; interacting with the physical environment; interacting with humans; and interacting with machines.

This typology can serve as a basis for comparing different uses of AI. It is also important to note that many applications rely on more than one of these functions to complete their tasks, and these functions can be interdependent. For example, an application developed by researchers from the Houston Methodist Cancer Center helps doctors predict a patient's risk of developing breast cancer by both analyzing unstructured data from mammograms and patient health records, and producing predictive diagnostic information 30 times faster than human analysis.[24]

Monitoring

AI can rapidly analyze large amounts of data and detect abnormalities and patterns. Because AI can do this far more quickly and accurately than humans—often in real time—AI is very well suited for monitoring applications, such as detecting credit-card fraud, cybersecurity intrusions, early warning signs of illnesses, or important changes in the environment.

Discovering

AI can extract valuable insights from large datasets, often referred to as data mining, and discover new solutions through simulations. In particular, because AI uses dynamic models that learn and adapt from data, it is very effective at uncovering abstract patterns and revealing novel insights that traditional computer programs cannot.[25]

Predicting

AI can forecast or model how trends are likely to develop in the future, thereby enabling systems to predict, recommend, and personalize responses. Many consumers are likely familiar with these types of applications, such as Netflix's recommendation algorithm, which analyzes users' viewing histories, stated preferences, and other factors to suggest new titles that they might like.[26] Data-intensive applications, such as precision medicine and weather forecasting, stand to benefit from this use of AI.

Interpreting

Until recently, most data analytics has focused on structured data—information that is well organized according to a specific framework, such as a spreadsheet of survey responses. Because AI can learn and identify patterns, it can interpret unstructured data—information that is not easily classifiable, such as images, video, audio, and text. As a result, computer systems are now capable of analyzing dramatically more kinds of information about the world. For example, AI helps smartphone apps interpret voice instructions to schedule meetings, diagnostic software to analyze X-rays to identify aneurysms, and legal software to rapidly analyze court decisions relevant to a particular case.[27]

Interacting with the Physical Environment

AI can facilitate a diverse range of machine-to-environment interactions that allow autonomous systems to directly engage with the physical environment. In particular, AI enables robotic systems that can navigate and manipulate the world around them. For example, autonomous vehicles analyze huge amounts of real-time data from an array of sensors, cameras, GPS systems, and maps to determine a safe and efficient route down a street.

Interacting with People

AI can allow humans to interact more easily with computer systems. Humans typically interact with machines by adjusting their behavior to meet the needs of the computer, such as by typing on a keyboard, pressing a button, or adjusting a dial. With AI, humans can interact with computers the way they do with other people, as computer systems can respond to speech, gestures, and even facial expressions. For example, individuals can ask questions of AI-powered chatbots by having a conversation or beckon a robot to come over with a nod or wave.

Interacting with Machines

AI can automatically coordinate complicated machine-to-machine interactions. For example, a control system for a data center can use AI to continuously monitor computing activity, internal temperature, and environmental conditions, and make adjustments to cooling systems to optimize performance while minimizing energy costs.[28] This ability also allows for multiple separate AI systems to coordinate with each other, such as a fleet of autonomous trucks managing themselves in a platoon formation to reduce fuel consumption, or autonomous robots in a warehouse that communicate with each other to sort and retrieve items.[29]




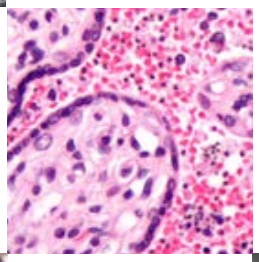
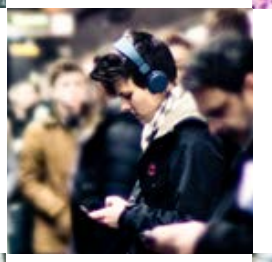
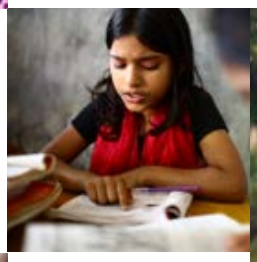




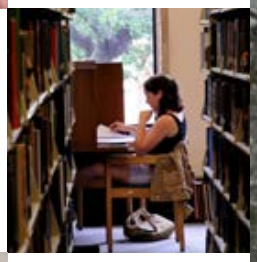

Examples of AI

AI is advancing rapidly, and there is no shortage of examples that show its many benefits. The rest of this report briefly illustrates 70 examples of AI in 14 categories: accessibility; agriculture; business operations; consumer convenience; disaster prevention and response; education; energy; environment; health care: prevention and screening; health care: treatment and monitoring; industrial operations; public safety; social good; and transportation. Each example also lists which AI functions are involved.

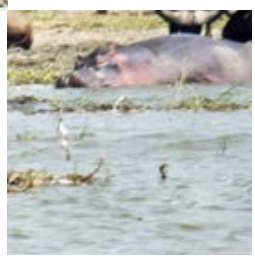

Accessibility

Making the Internet More Accessible for People with Visual Impairments

Facebook has implemented computer-vision algorithms to automatically generate descriptive text captions for images and use text-to-speech software to describe the photos to blind users.[30] Though many people with visual impairments use screen-reading technology to browse the Internet, image-heavy websites such as Facebook are often difficult for screen readers to interpret, limiting accessibility.[31]

Helping People Understand Sign Language

Researchers at the International Islamic University Malaysia have developed a system called the Automatic Sign Language Translator (ASLT) that uses machine learning to interpret sign language and convert it into text.[32] The researchers trained ASLT's algorithms on a database of images and videos of native signers, so it could learn to identify distinct signs as well as the facial and upper body movements that often accompany certain signs to improve its accuracy.[33]

Gamifying Emotional Recognition

Wearable camera company Narrative has developed a game called Autimood, powered by Microsoft's Project Oxford Emotion application program interface (API), which uses machine learning to interpret emotions in facial expressions, which allows parents of children with autism to help their children improve their ability to recognize emotions.[34] Autimood uses a wearable camera that takes pictures every 30 seconds, and at the end of the day the Project Oxford Emotion API tags emotions in each picture so children can attempt to identify emotions and receive feedback.[35]

Making It Easier to Get Around in a Wheelchair

A European Union-funded research initiative called Robotic Adaption to Humans Adapting to Robots (Radhar) has developed a robotic wheelchair that uses computer-vision algorithms to automatically help users navigate their environments.[36] Radhar's wheelchair scans the environment 10 times a second to map obstacles and generate potential routes the wheelchair could take, and as the user steers the wheelchair, it will automatically avoid objects and make the route as smooth as possible.[37]

Identifying Dangerous Sounds

Audio-recognition company One Llama Labs has developed a smartphone app called Audio Aware that uses machine learning to identify sounds associated with dangerous situations, such as sirens or squealing tires, and warn hard-of-hearing users about the noise.[38] Audio Aware can identify a variety of dangerous sounds through a smartphone's microphone, as well as allow users to record and share their own, and when it detects one, it will play an amplified version of the noise through headphones, which can help partially deaf users stay more aware of their surroundings.[39]

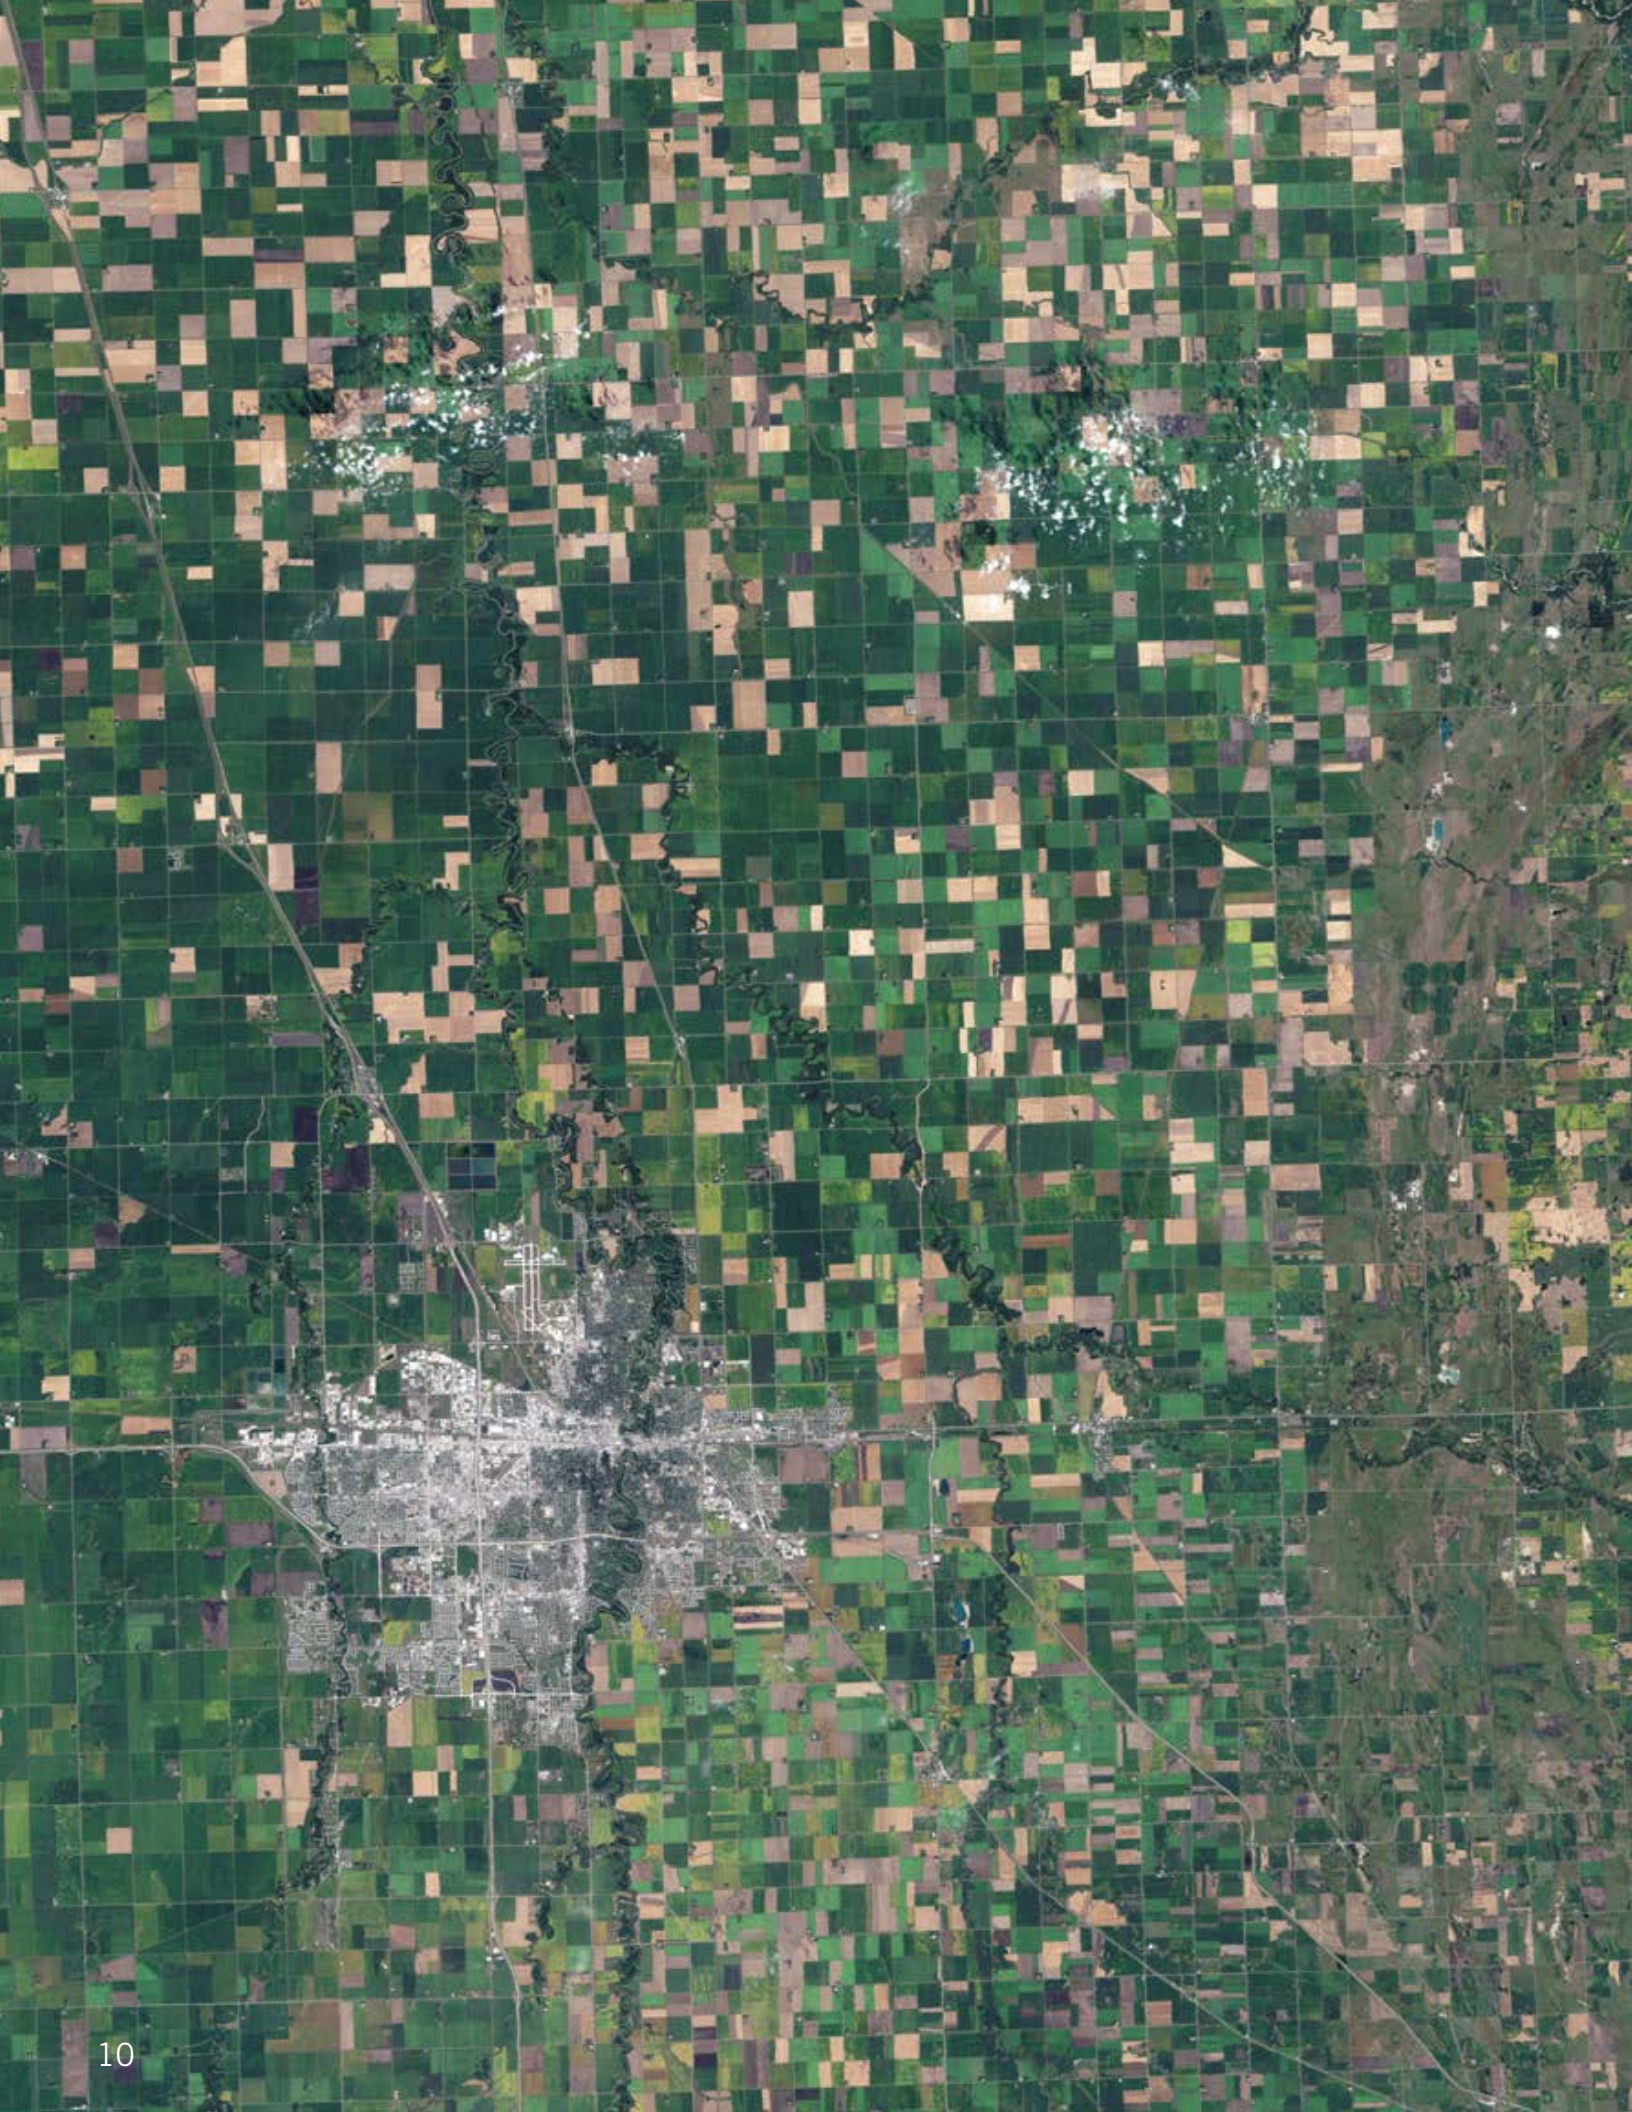

Agriculture

Farming Indoors Autonomously

Montreal startup Motorleaf has developed an autonomous indoor farming system that uses networked sensors and machine learning to constantly monitor an indoor farm's environment and plant growth and adjust lighting, temperature, humidity, water, and soil nutrient levels to maximize a farm's productivity.[40] The system also reports 40,000 data points per user per week to Motorleaf, which uses machine learning algorithms to analyze this data and predict problems that could arise to prompt farmers to take preventative action.[41]

Learning as Soon as Plants Get Sick

Israeli start-up Prospera has developed a system that uses networked cameras, sensors, and a machine-learning algorithm to monitor crops and warn farmers as soon as a plant is sick.[42] The algorithm can identify signs of mold, bacteria, or insect damage in plant images, as well as analyze plant nutrition data, weather forecasts, and other data to predict if a particular plant is at risk of becoming sick.[43] As Prospera's system gathers more data about each farm, it can learn from historical data and refine its predictions.

Forecasting Crop Yields From Space

Descartes Labs has adapted deep-learning image-analysis software to analyze satellite photos of farmland to forecast crop yields faster and more accurately than official government estimates.[44] The software can produce estimates of crop production on a weekly basis by comparing daily photographs of 3 million square kilometers of corn farms with less than a one percent margin of error, allowing farmers, insurers, commodities traders, and governments to make more informed decisions.[45]

Spot-Treating Crops

A company called Blue River Technology has developed a robotic system called LettuceBot that can drive through a field, take 5,000 photos of plants per minute, and use computer-vision algorithms to identify weeds and spray targeted bursts of herbicide directly on them, rather than the whole field.[46] LettuceBot's algorithms can differentiate between photos of lettuce sprouts and weeds in 0.02 seconds.[47]

Making Vegetable Sorting Easy

A Japanese cucumber farmer has built a robotic system powered by TensorFlow, Google's open-source machine-learning software, that can automatically sort cucumbers based on their visual differences such as size and shape.[48] The farmer trained the system's computer-vision algorithms to recognize nine different categories of cucumbers. The system can sort cucumbers on a conveyer belt with 70 percent accuracy and can do so substantially faster than manual sorting, which took a person up to eight hours per day on a small family farm.[49]

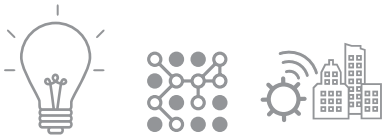

Business Operations

Predicting Area-Specific Weather Implications

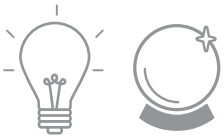

IBM has created a machine-learning tool called Deep Thunder that analyzes historical weather and environmental data to provide targeted weather analysis, with a resolution as high as 0.2 miles, so companies can factor weather into business-planning decisions.[50] For example, a retailer can use Deep Thunder to estimate the impact weather will play on consumer purchasing decisions and preemptively adjust its stock of items accordingly.[51]

Learning How to Keep Customers Happy

Zendesk has developed a tool called Satisfaction Prediction that uses machine learning to evaluate customer-service interactions and predict how satisfied a customer is, prompting an intervention if a customer is at risk of leaving a business.[52] Satisfaction Prediction analyzes data such as customer-support ticket text, wait times, and the number of replies it takes to resolve a ticket to estimate customer happiness and adjusts its predictive models in real time as it learns from more tickets and customer ratings.[53]

Reducing Gender Bias in the Office

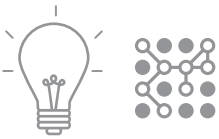

Workforce analytics company Kanjoya uses natural language processing and machine learning to analyze language used in the workplace, such as employee performance evaluations, interview notes, and office communications, to identify signs of implicit gender bias, so companies can treat employees more fairly.[54] Kanjoya's software can identify potentially discriminatory or abusive language and emotions or intent in language, as well as identify biased decision-making, such as by revealing that "assertiveness" is often associated with positive reviews and promotions for men, but negative reviews for women.[55]

Automating Office Assistants

Start-up X.ai offers a virtual-assistant service named Amy that can analyze employee calendars and emails to automatically schedule meetings and adjust calendar appointments.[56] Users can copy Amy in emails when they want to set up meetings, and Amy will analyze email text to determine the topic and time of a meeting, determine if there are any conflicts, and automatically schedule calendar appointments.[57] Amy can also search for and add relevant phone numbers, reschedule meetings by conversing with users, and learn users' preferences over time.[58]

Making Customer Support Multilingual

AI translation start-up Unbabel offers a translation service that uses a combination of crowdsourcing and machine learning to translate businesses' customer-service operations into 14 different languages substantially faster and cheaper than traditional translation services, which can make it easier for businesses to reach international audiences.[59] Unbabel's algorithm automatically translate web pages, customer service emails and chats, and social-media posts for as little as $0.02 per word, and a team of human editors reviews the translations for grammar and consistency.[60]

Consumer Convenience

Deep Learning for Finding the Right Restaurant

Yelp has developed a machine-learning system that can analyze user-submitted photos of restaurants to identify a restaurant's characteristics, such as the style of food and ambiance, with an accuracy of 83 percent.[61] This system can help populate restaurant pages with too few user reviews to classify the restaurant with relevant data so Yelp can improve its search results and provide better recommendations.[62]

Automating Personalized Financial Advice

San Francisco start-up Wallet.AI developed a smartphone app that allows users to specify financial goals, such as "saving enough money for rent," and link their financial information, so its algorithms can learn spending patterns and activities to offer personalized, contextually relevant financial advice.[63] For example, if Wallet.AI detects a user has walked into a coffee shop where he or she tends to spend a lot of money, it will suggest that the user should cut back to have enough money for an upcoming utility payment.[64]

Household Robots Get Emotional

Japanese company Softbank Robotics has used AI to train its household robot Pepper to recognize emotional cues in human voices and facial expressions with computer vision and natural language processing.[65] Pepper, which is four feet tall and can move around on wheels, can use these cues to inform how it interacts with humans, such as when it answers questions or greets guests.[66]

Helping Consumers Buy What They Like

Image-based social media site Pinterest has developed a machine-learning feature that can identify items in pictures that users have favorited and find out where they can be bought online.[67] For example, if a user favorites a picture with a particular pair of shoes, Pinterest's algorithms can automatically send that user a link to buy the shoes online.

Making It Easier to Sort Photos

The Google Photos app uses machine learning to recognize the contents of users' photographs and their metadata to automatically sort images by their contents and create albums.[68] The Google Photos' algorithms can identify pictures taken in the same environment, recognize and tag landmarks, identify people who appear in multiple photographs, and generate maps of a trip based on the timestamps and geotags of a series of photographs.[69]

Disaster Prevention and Response

Predicting Where Earthquakes Do the Most Damage

A start-up called One Concern has developed a system that uses AI to model a town's buildings, based on factors such as building age, density, and construction materials, and predict how they would react to seismic activity.[70] If an earthquake hits, One Concern can plug in the new seismic data to estimate which areas are likely the most damaged.[71] Using these predictions, first responders can prioritize disaster response efforts to target the hardest hit areas after an earthquake.

Keeping Emergency Responders Out of Danger

NASA has developed a safety tool for firefighters called the Assistant for Understanding Data through Reasoning, Extraction, and sYnthesis (AUDREY), which uses AI to monitor data about firefighters' environments to detect signs of danger and help them recover from disorientation to exit a building safely.[72] AUDREY can monitor sensors in firefighting gear and warn firefighters, for example, if it is becoming dangerously hot or if dangerous gases are present, map the surrounding environment, and communicate with firefighters to guide them to safety if they need to evacuate.[73] The more AUDREY is used in the field, the more it can learn about how firefighters operate and make more effective predictions.[74]

Detecting Disease Outbreaks

Researchers at Boston Children's Hospital developed a disease-tracking website called HealthMap that uses AI to automatically analyze news stories in multiple languages.[75] HealthMap's algorithm identified reports of Ebola-linked symptoms on Kenyan and Guinean news sites despite the articles not mentioning Ebola, more than seven days before the outbreak was determined to be Ebola.[76] HealthMap now monitors news stories for signs of a wide variety of diseases, including Zika and West Nile Virus, and the sooner health officials can identify an outbreak of these diseases, the sooner they can take preventative measures to stop its spread.[77]

Understanding a Crisis With Social Media

The Qatar Computing Research Institute has developed an open-source tool called Artificial intelligence for Disaster Response (AIDR) that uses machine learning to monitor and analyze Twitter posts and automatically compile Twitter activity related to a particular crisis to aid humanitarian response.[78] In a test during the 2013 flooding in Pakistan, volunteers trained AIDR on tweets related to the crisis, and it could determine if new tweets were related to the Pakistan floods, based on their text, time stamp, and geotag, with 80 percent accuracy.[79] Getting as much information about a crisis, with minimal effort, as quickly as possible can help emergency responders better prioritize resources and make more informed decisions.

Avoiding Dangerous Solar Flares

Researchers at Stanford University have developed a machine-learning system that can predict M- and X-class solar flares, which produce dangerously high levels of radiation that could harm airline passengers, damage power grids, and disrupt communication satellites.[80] The system analyzes data about the sun, such as the topology of its magnetic field and its atmosphere, and can identify regions of the sun likely to produce solar flares.[81] More advanced warning of when and where a solar flare could occur will allow airlines, power grid operators, and others to take precautionary measures to avoid danger.

Education

Personalizing Math Class

IBM has developed a new tool based on its Watson cognitive computing platform called Teacher Advisor to help third-grade math teachers in the United States develop personalized lesson plans.[82] Teacher Advisor analyzes Common Core education standards, which sets targets for skills development, and student data to help teachers tailor instructional material for students in the same class but with varying skill levels, which can make traditional, static lesson plans ineffective.[83] IBM will make Teacher Advisor freely available to third-grade math teachers by the end of 2016, and plans to expand the number of subject areas and grade levels it can help with.[84]

Predicting Which Students Will Drop Out

The Tacoma, Washington, school district worked with Microsoft to develop a machine-learning model that can analyze student data, such as demographics and academic performance, and historical data to predict which students were at risk of dropping out and prompt early intervention.[85] After a multiyear pilot of the system, the Tacoma school district was able to boost its graduation rates from 55 percent in 2010 to 78 percent by 2014.[86]

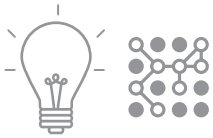

Automating Teacher Assistants

The Georgia Institute of Technology has implemented an automated teaching assistant named Jill, powered by IBM's Watson cognitive computing platform, to help respond to student inquiries for an online course that receives an average of 10,000 messages from students every semester.[87] Jill can analyze and answer student questions, such as where to find course materials, and it was effective enough that the course plans to use Jill to field 40 percent of all student inquiries by the end of 2016.[88] Programs like Jill could help improve retention rates for online courses, which are generally low because students have trouble getting the information they need from professors.[89]

Making It Easier to Learn New Languages

Language-learning software company Duolingo uses machine learning to analyze users' activity and progression to develop personalized lesson plans, as well as regularly test new strategies for instruction to evaluate their effectiveness.[90] Duolingo will structure particular lesson plans differently for various users to learn how the differences correspond to improved outcomes and implement the most effective strategies.[91]

Giving Students Feedback in Real Time

Education software company Turnitin has developed a tool called Revision Assistant that uses machine learning to evaluate students' writing while they draft essays to provide feedback.[92] Revision Assistant evaluates four traits in student writing—language, focus, organization, and evidence—and can detect the use of imprecise language or poor organizational structure to provide both positive feedback and recommendations for improvement.[93] Revision Assistant is designed to evaluate essays from students in grades 6 through 12, and schools can work with Turnitin to define their evaluation criteria for different courses.[94]

Energy

Predicting Renewable Energy Availability

IBM has developed a machine-learning system called Self-Learning Weather Model and Renewable Energy Forecasting Technology (SMT) that analyzes data from 1,600 weather stations, solar plants, wind farms, and weather satellites to generate weather forecasts 30 percent more accurately than the National Weather Service and predict renewable energy availability up to weeks in advance.[95] SMT continuously refines its models as it receives new data, and its predictions can help regional power grids better integrate renewable energy sources, as changes in the weather can cause renewable energy production to vary greatly.[96]

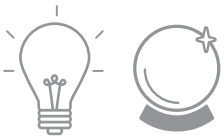

Modeling Energy Consumption for More Efficient Buildings

The U.S. Department of Energy's Oak Ridge National Laboratory has developed a tool called Autotune for its Titan supercomputer that uses machine learning to create highly granular models of a building's energy efficiency with an error rate below one percent and to recommend improvements.[97] Autotune's simulations account for approximately 150 parameters that influence energy efficiency, such as lighting and ventilation, and learn how to optimize these factors to maximize a building's energy efficiency.[98]

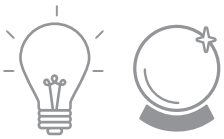

Teaching a Data Center to Make Itself More Efficient

Google has implemented AI software developed by fellow Alphabet subsidiary DeepMind in one of its data centers to automatically optimize energy efficiency while responding to factors such as increased usage and changing weather.[99] The system constantly monitors 120 variables, including server usage and windows, and learns how to adjust equipment performance and cooling systems to run the data center as efficiently as possible.[100] The system has reduced the data center's energy consumption by 15 percent.[101]

Learning How to Manage Home Energy Use

The Nest Learning Thermostat uses AI to learn homeowner's preferences and schedules to optimize home heating and cooling and save consumers an average of $131 to $145 in energy costs per year.[102] The thermostat's sensors can also monitor activity in a house to automatically reduce heating or cooling when nobody is home to avoid wasting energy.[103]

Picking the Best Spot for a Wind Farm

Researchers at the Massachusetts Institute of Technology have developed a machine-learning system that can predict variations in wind speeds over time to help power companies more quickly evaluate potential locations for wind farms.[104] Traditionally, a power company will gather 12 months of wind-speed data to evaluate a potential wind-farm location, but the machine-learning system can produce more accurate models with just three months of data by correlating data from multiple sites and weather stations.[105]


C-Power

Environment

Stopping Deforestation Before It Starts

A start-up called Orbital Insight uses AI to analyze satellite imagery of forests over time to detect early warning signs of illegal logging that can prompt intervention before any trees are cut down.[106] The system can flag changes that might go unnoticed by humans, such as new roads, which could indicate a new logging operation, as well as learn to identify changes that occur before major cutting to improve its warning system.[107]

Predicting Dangerous Air Pollution Levels

IBM researchers have developed a machine learning system that can analyze data about pollution levels in Beijing to forecast changes to air quality up to 72 hours in advance.[108] The system's forecasts are 30 percent more accurate than traditional predictions and have a resolution of one kilometer.[109] The researchers are further developing the system to forecast hypothetical scenarios, such as changes in the number of drivers on the road, and to extend the forecasts up to 10 days in advance.[110]

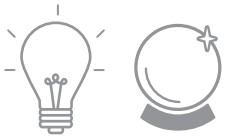

Improving Antipoaching Efforts

The Uganda Wildlife Authority (UWA) is testing a tool called Protection Assistant for Wildlife Security (PAWS) that uses AI to identify poaching routes and optimize resource deployment to prevent poaching in Queen Elizabeth National Park.[111] PAWS analyzes UWA's database of geotagged animal sightings and historical data of poaching activities to predict likely hotspots of poaching activity, which can be difficult to find, as poachers frequently change their habits to elude UWA.[112] PAWS updates its predictions based on new animal sighting and poaching activity data and recommends patrol routes.[113]

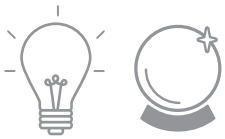

Saving Threatened Birds

Conservation technology start-up Conservation Metrics uses a system of acoustic sensors and machine learning to improve efforts to save the marbled murrelet, a threatened species of bird in California.[114] Networked acoustic sensors throughout California forests collect data about bird calls, and a machine-learning algorithm parses through this audio to identify marbled-murrelet calls and flag when and where they were recorded.[115] With this data, Conservation Metrics can better map the bird's habitat, estimate its population, and guide conservation efforts.

Teaching a Robot to Recycle

Finnish robotics company ZenRobotics has developed an autonomous recycling system called the ZenRobotics Recycler (ZRR) that uses a robotic arm, an array of sensors, and AI to identify recyclable items in waste and separate them for recycling, removing the need for manual sorting.[116] ZRR analyzes trash on a conveyer belt using 3-D scanning, spectrometer analysis, and other methods to determine what a piece of trash is made of, and, if recyclable, a robotic arm will pick it up and move it to a separate container.[117]

Health Care: Prevention and Screening

Preventing Vision Loss in Diabetes Patients

Volunteer data scientists developed an image-analysis algorithm for the California HealthCare Foundation that can analyze retinal scans of diabetes patients and learn to identify subtle signs of diabetes-linked retinal damage with 85 percent accuracy, faster than traditional human analysis and without needing to send scans to a lab.[118] Approximately 80 percent of diabetes patients experience retinal damage, which, if undetected, can cause complete blindness.[119]

Predicting Schizophrenia by Analyzing Speech

Researchers at Columbia University, the New York State Psychiatric Institute, and IBM have developed a machine-learning system capable of predicting if a person at risk of developing psychosis caused by schizophrenia will develop the condition with 100 percent accuracy by analyzing his or her speech, which can exhibit telltale signs of the condition.[120] By learning to identify the speech cues in audio recordings, the system could outperform normal diagnostic models, which are about 79 percent accurate.[121]

Figuring Out How to Prevent Pancreatic Cancer

Biotech start-up Berg has used AI to analyze large amounts of oncological data to create a complete model of how pancreatic cancer functions.[122] Based on this model, Berg identified specific metabolic processes that contribute to pancreatic cancer's rapid growth and developed a drug, currently in phase II trials, that targets these processes to make cancer cells more responsive to chemotherapy.[123]

Automating a Microscope to Diagnose Malaria

An automated microscope called Autoscope uses an artificial neural network to rapidly analyze blood samples in the field and diagnose malaria with 90 percent accuracy.[124] Traditional rapid diagnostic methods can only determine if the malaria parasite is present in a blood sample, which does not necessarily mean a person will contract the disease.[125] Autoscope can differentiate between specific particles in a blood sample to quantify the number of malaria parasites to determine if a diagnosis is warranted.[126]

Diagnosing Voice Disorders

Researchers at the Massachusetts Institute of Technology and Massachusetts General Hospital have developed a system that uses a wearable device to collect data about the movement of a user's vocal cords and uses machine learning to detect subtle signs of abnormal speech that could indicate a person has a voice disorder called muscle tension dysphonia (MTD).[127] People with MTD experience vocal fatigue and deteriorating voice quality, despite no prior physical damage to their vocal cords, which makes the condition hard to diagnose.[128] The system could be particularly useful for collecting vocal performance data to aid diagnostics without having patients spend much time with specialists in person.

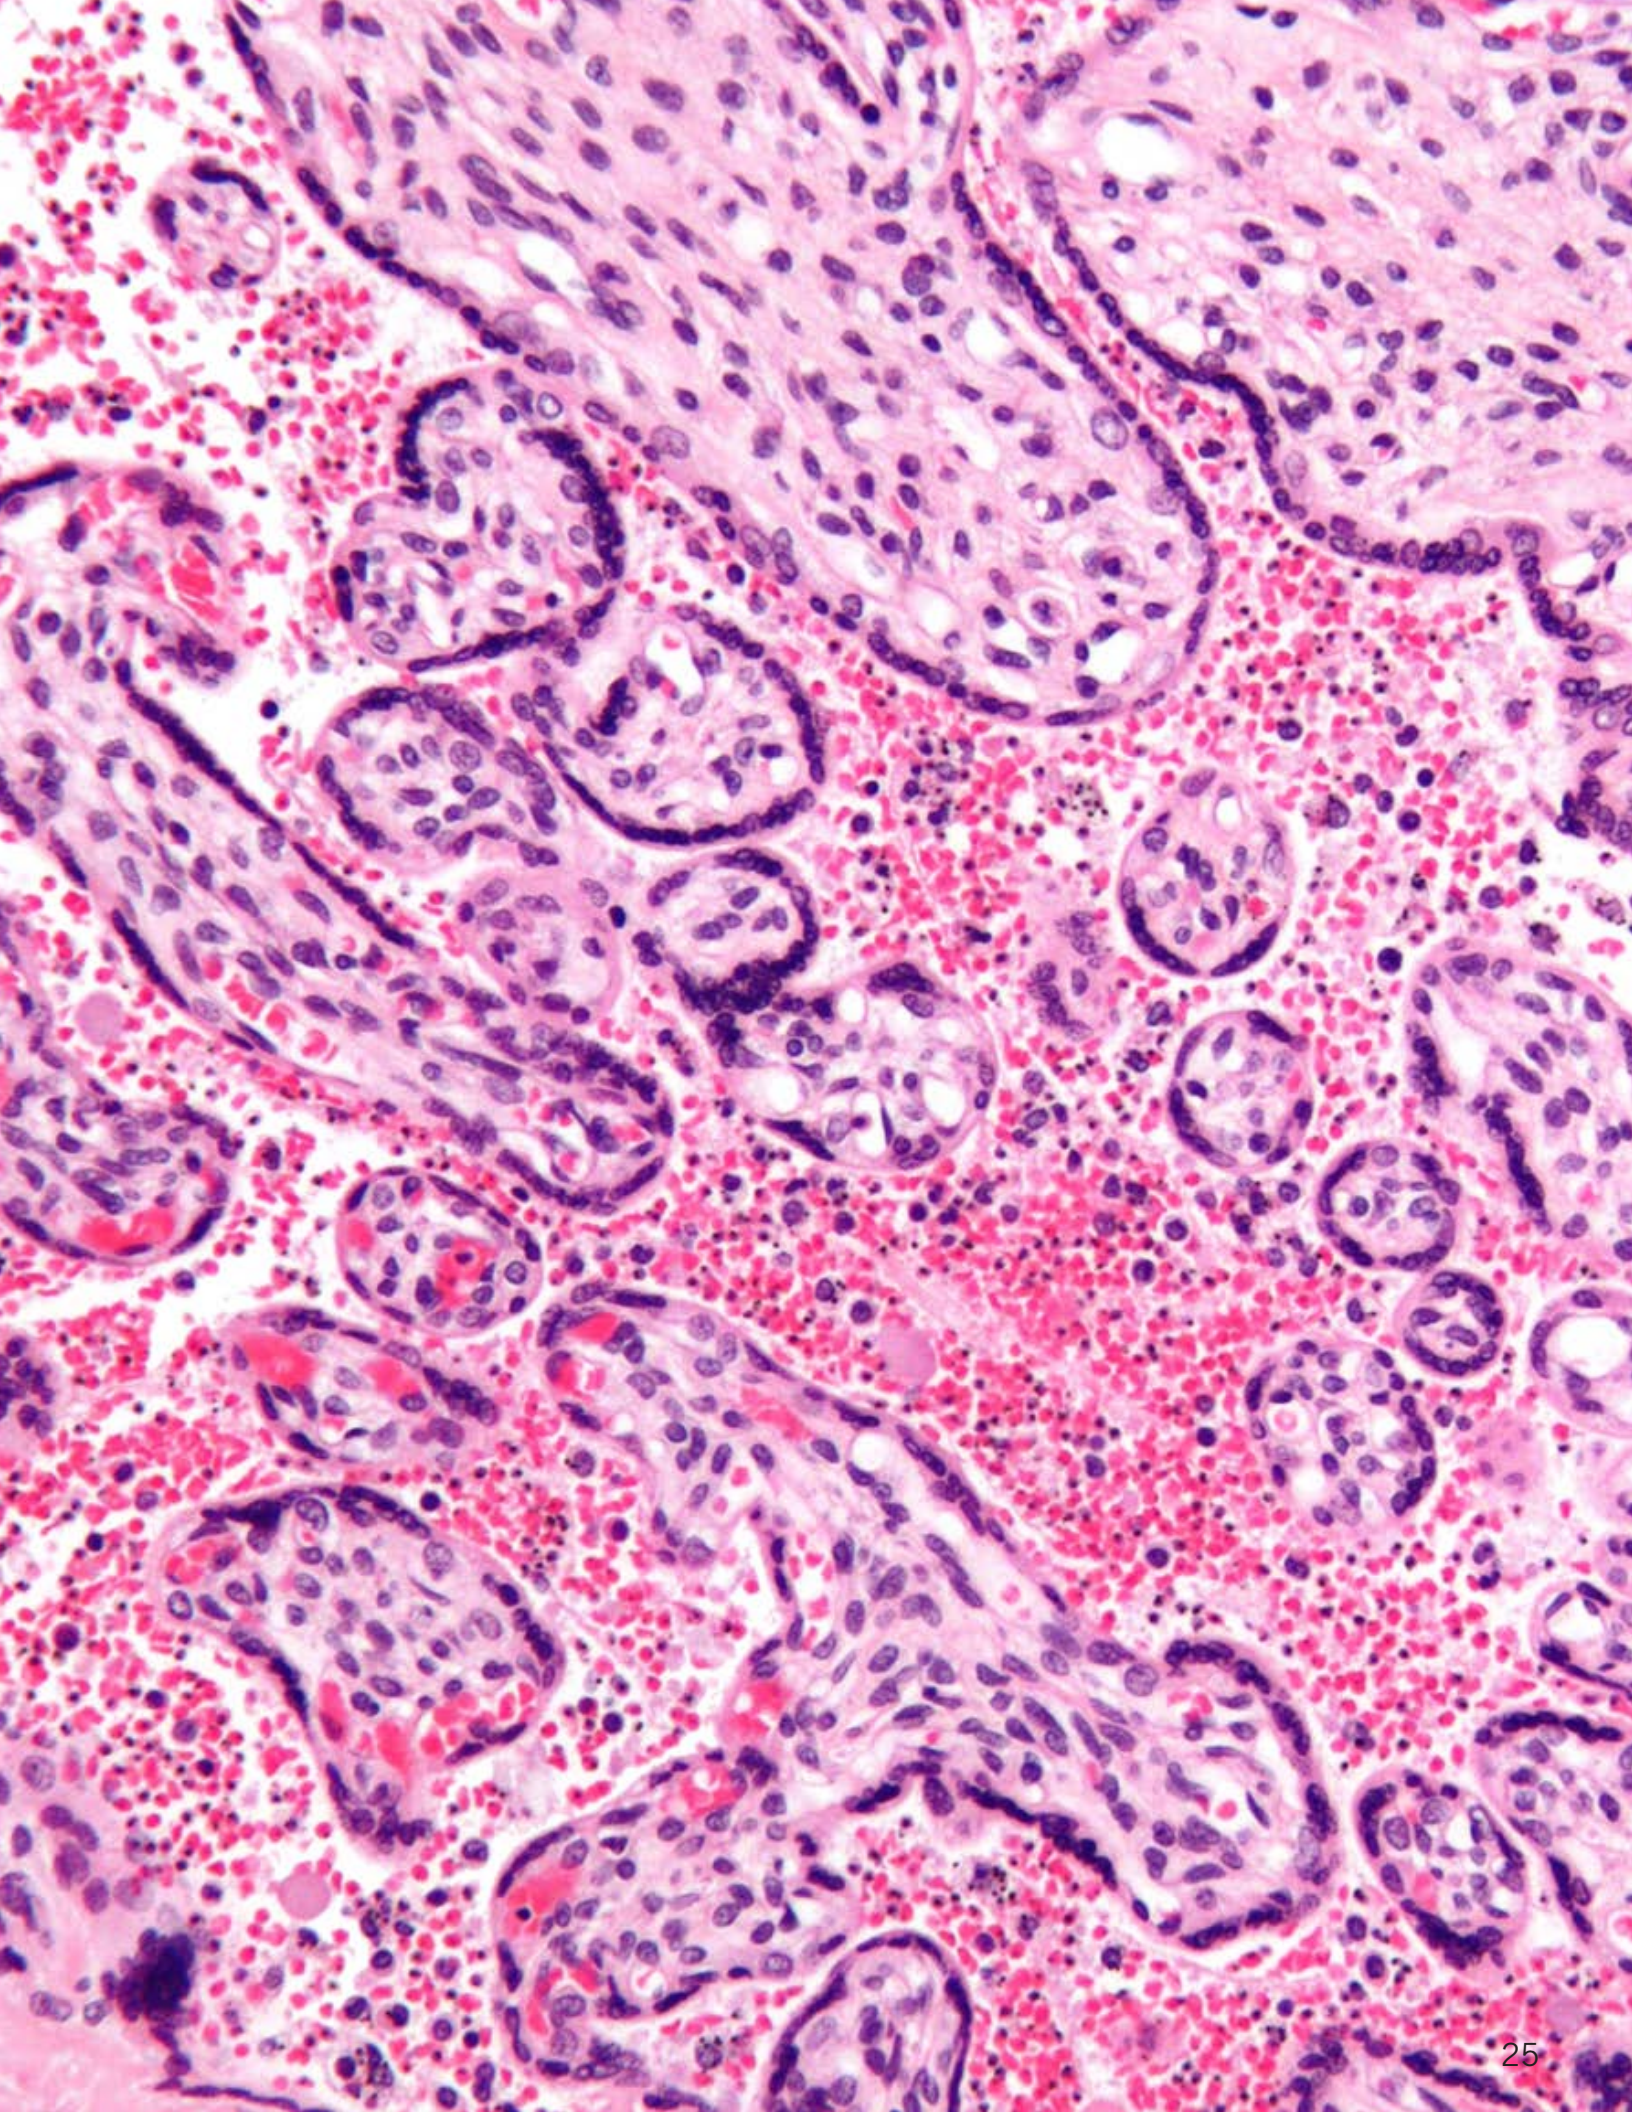


TomoThe
HI·AR

Health Care: Treatment and Monitoring

Helping Diabetes Patients Make Smarter Diet Decisions

Start-up Suggestic has developed a smartphone app for patients with type 2 diabetes that uses AI to analyze medical research and users' behavior to make personalized recommendations about how they can alter their diet to better manage their disease.[129] Suggestic allows users to input their goals and diet plans and then analyzes thousands of medical publications and research papers about diabetes to identify relevant data that can inform personalized nutrition plans and recommend interventions when necessary.[130]

Streamlining Drug Discovery

Researchers at Carnegie Mellon University used a machine-learning system to prioritize which experiments they should conduct to test new drugs, reducing the number of unnecessary tests by up to 70 percent.[131] The system analyzed several baseline experiments of a new drug's effect on a particular protein in a cell and was able to learn which factors influence a drug's interactions, creating a predictive model that could estimate the outcomes of possible variations of the experiment with an accuracy rate of 92 percent.[132]

Making Stitches Safer With Robotic Surgeons

A surgical robot called the Smart Tissue Autonomous Robot (STAR) can administer stitches more precisely than human surgeons.[133] STAR analyzes data from specialized 3D and infrared cameras in real time to generate a plan for an optimal arrangement of stiches and administers them with a robotic arm. In tests, STAR's stiches were more consistent and resistant to leaks than human surgeon-administered ones.[134]

Using AI to Speed Radiotherapy

The U.K. National Health Service has partnered with DeepMind, Google's AI research arm, to train an AI system that can reduce the time needed to provide radiotherapy treatment to patients with head and neck cancers.[135] Doctors normally spend up to four hours in radiological scans, carefully mapping out cancer cells to target with radiation. DeepMind's system will train on 700 anonymized scans to learn how to map cancerous areas automatically, which DeepMind believes can shorten this process to an hour.[136]

Increasing Participation in Clinical Trials

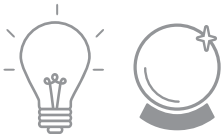

Researchers at Cincinnati Children's Hospital Medical Center have developed a machine-learning system to evaluate whether a patient is likely to participate in a clinical trial.[137] The system analyzed objective and subjective factors about a patient, such as age, race, attitude toward medical research and health conditions.[138] By understanding the factors that make a patient more likely to participate, medical researchers can more effectively target patients for trials that could help develop valuable treatments.

Industrial Operations

Preventing Breakdowns Before They Happen

Analytics company DataRPM has developed a tool called Cognitive Predictive Maintenance that uses machine learning to monitor machine-component performance through networked sensors in real time to detect early warning signs that a machine might be breaking down to prompt preventative maintenance.[139] Servicing machines before they break down is more cost effective than fixing them after they break, and DataRPM can predict when maintenance is needed 300 percent more accurately and 30 times faster than traditional methods, saving its customers 30 percent in maintenance costs.[140]

Intelligent Manufacturing to Automate Factories

Siemens AG has built a manufacturing facility in Amberg, Germany, that relies on 1,000 networked manufacturing units automatically coordinating with one another to retrieve and fabricate components without human oversight.[141] The facility is approximately 75 percent autonomous, allowing its human employees to focus on monitoring the factory floor and operating computer systems.[142]

Maximizing Oil-Well Performance

General Electric and British Petroleum (BP) have partnered to install networked sensors in subsea BP oil wells and use machine-learning software to monitor oil-well performance in real time, allowing BP to maximize production and minimize downtime.[143] The software will analyze data such as pressure, machine vibration, and temperature to provide oil-well workers with relevant insights into how to make wells operate more efficiently, reducing operational costs by up to 25 percent, and how to avoid unexpected failures, which can cost operators over $3 million each week a well is out of commission.[144]

Improving Dairy Supply Chains With Market Forecasting

Italian dairy producer Granarolo implemented a machine-learning tool that analyzes data about sales estimates and planned promotions, which can increase sales 30-fold, to forecast how much its dairy farms should produce and when.[145] Because dairy is perishable, overestimating demand for a particular period can result in large quantities of wasted products, while underestimating demand can cause diaries to miss out on potential sales. With its forecasting tool, which learns to identify the relationship between Granarolo's thousands of promotions per year and demand fluctuations, Granarolo was able to significantly increase its forecast reliability, reduce inventory levels and delivery times by 50 percent, and increase sales.[146]

Making Industrial Design Smarter

Engineering software company Autodesk has developed computer-aided design and drafting software called Dreamcatcher that uses AI to automatically generate industrial designs based on a designer's specific criteria, such as function, cost, and material.[147] Dreamcatcher can generate multiple alternative designs that meet the same criteria and provide designers with performance data for each design, alter designs in real time based on designer feedback, and export finalized designs into formats used for fabrication.[148]

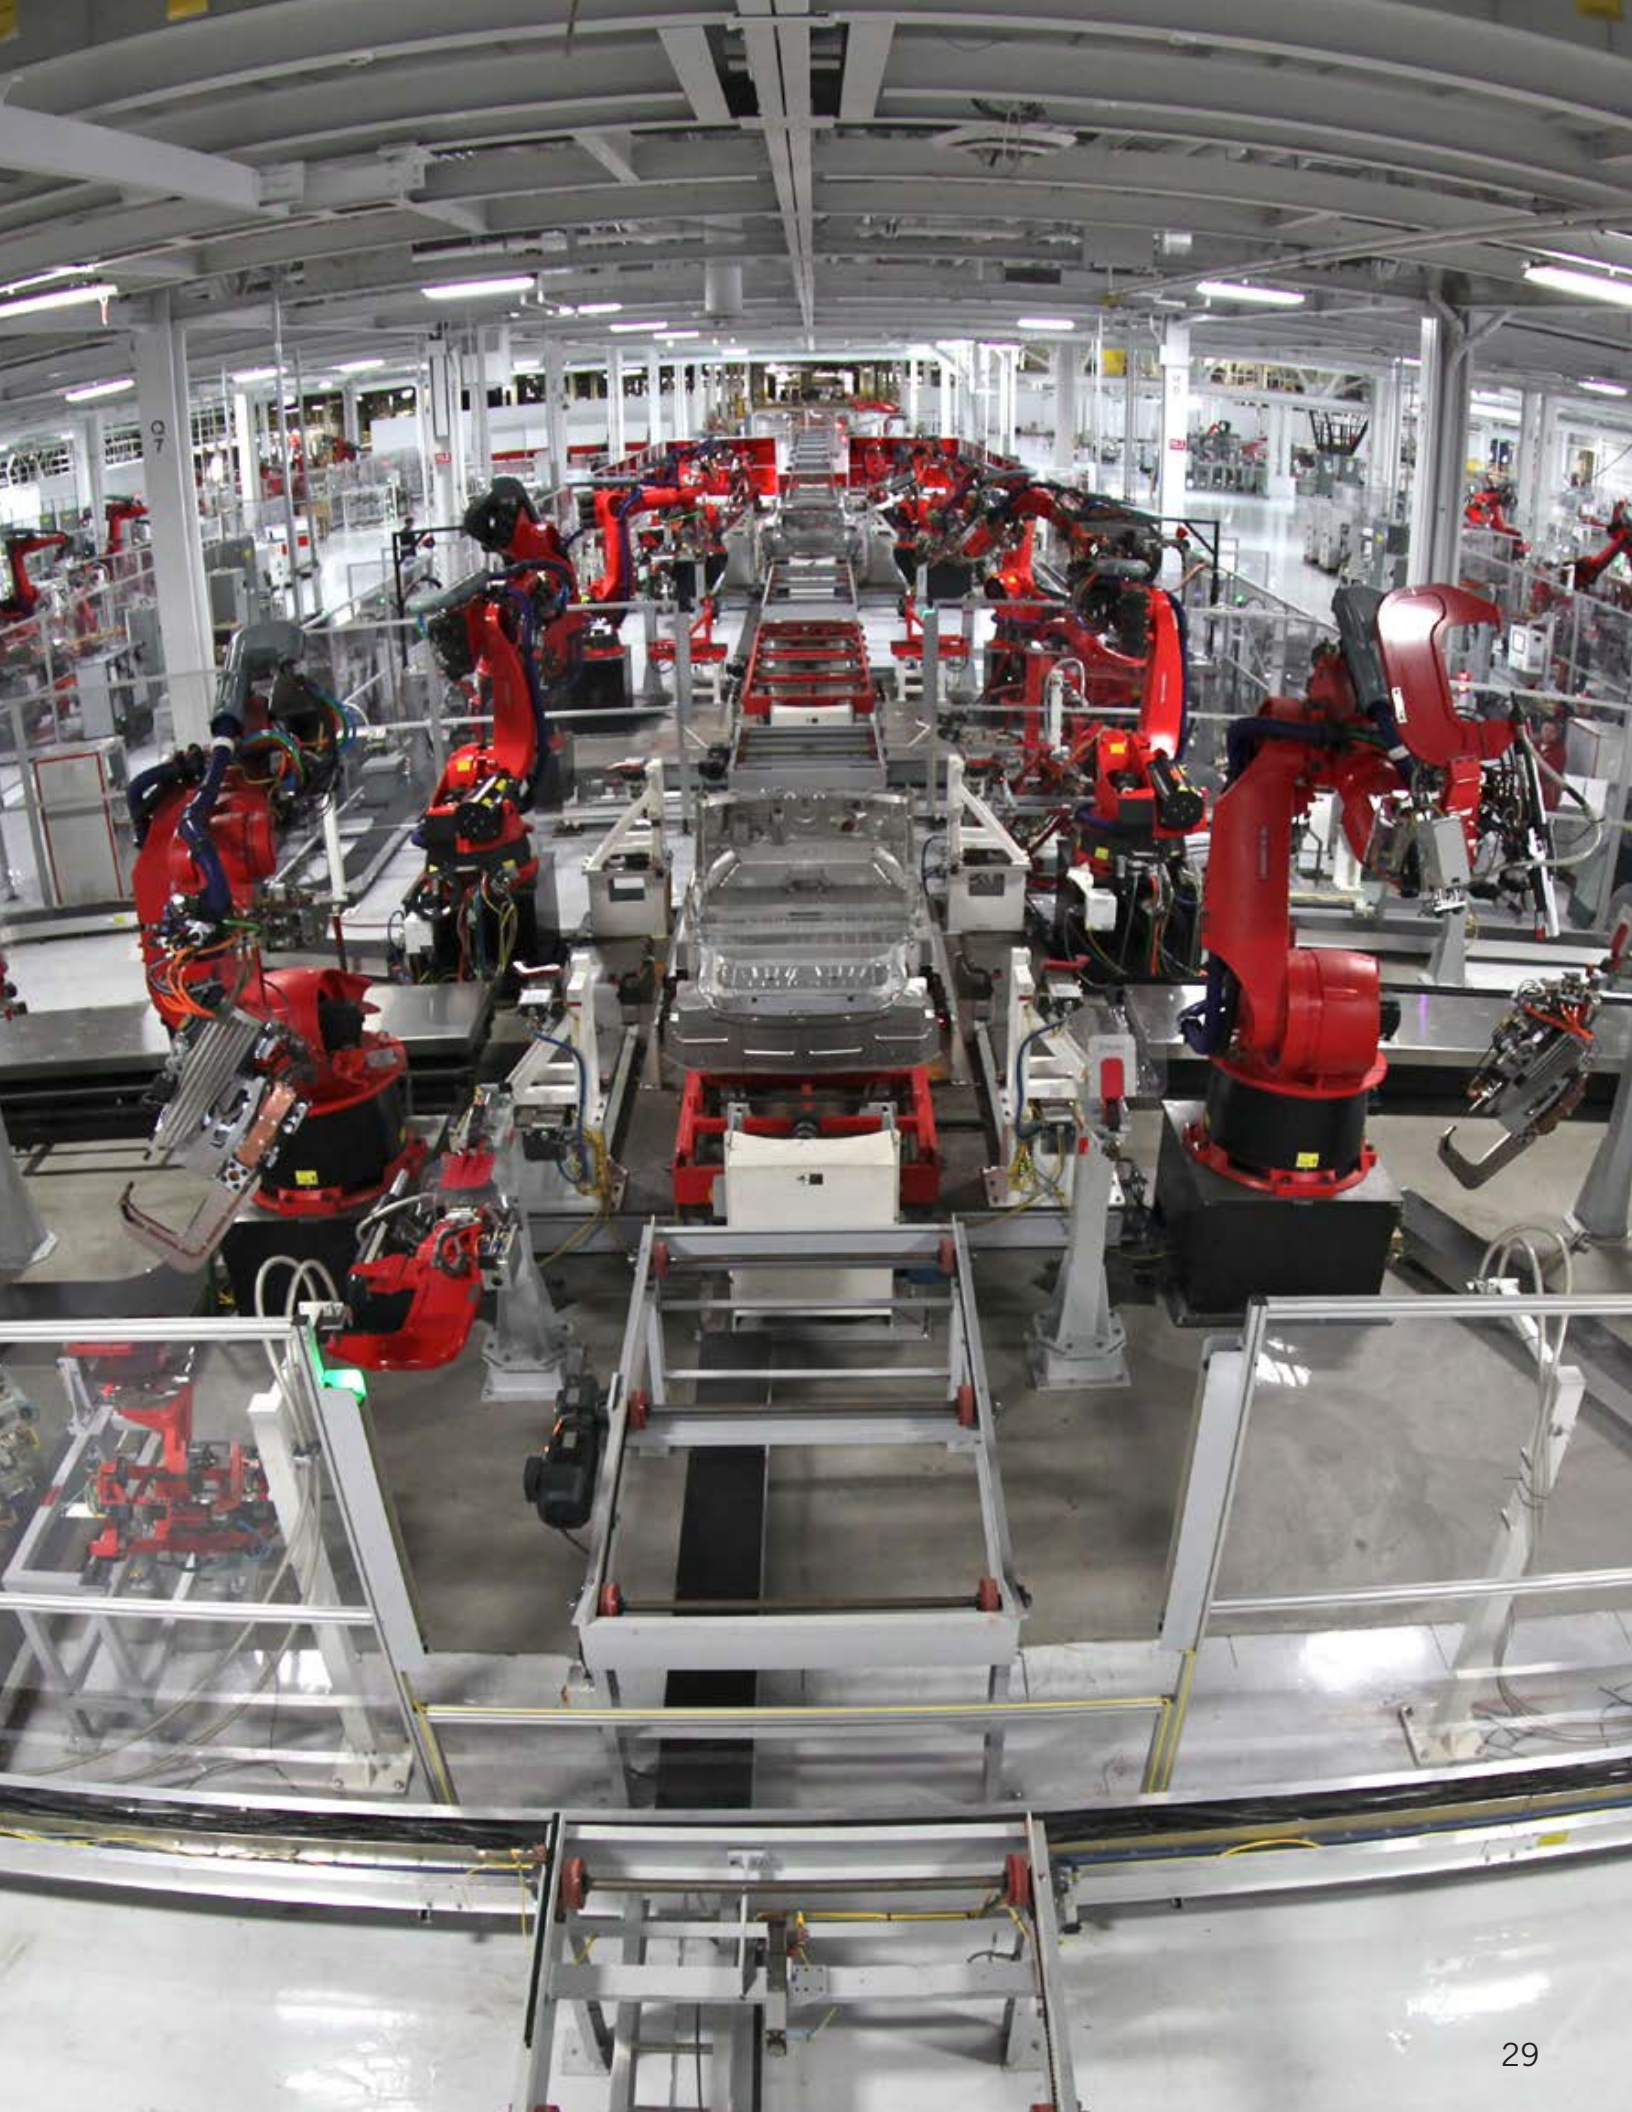

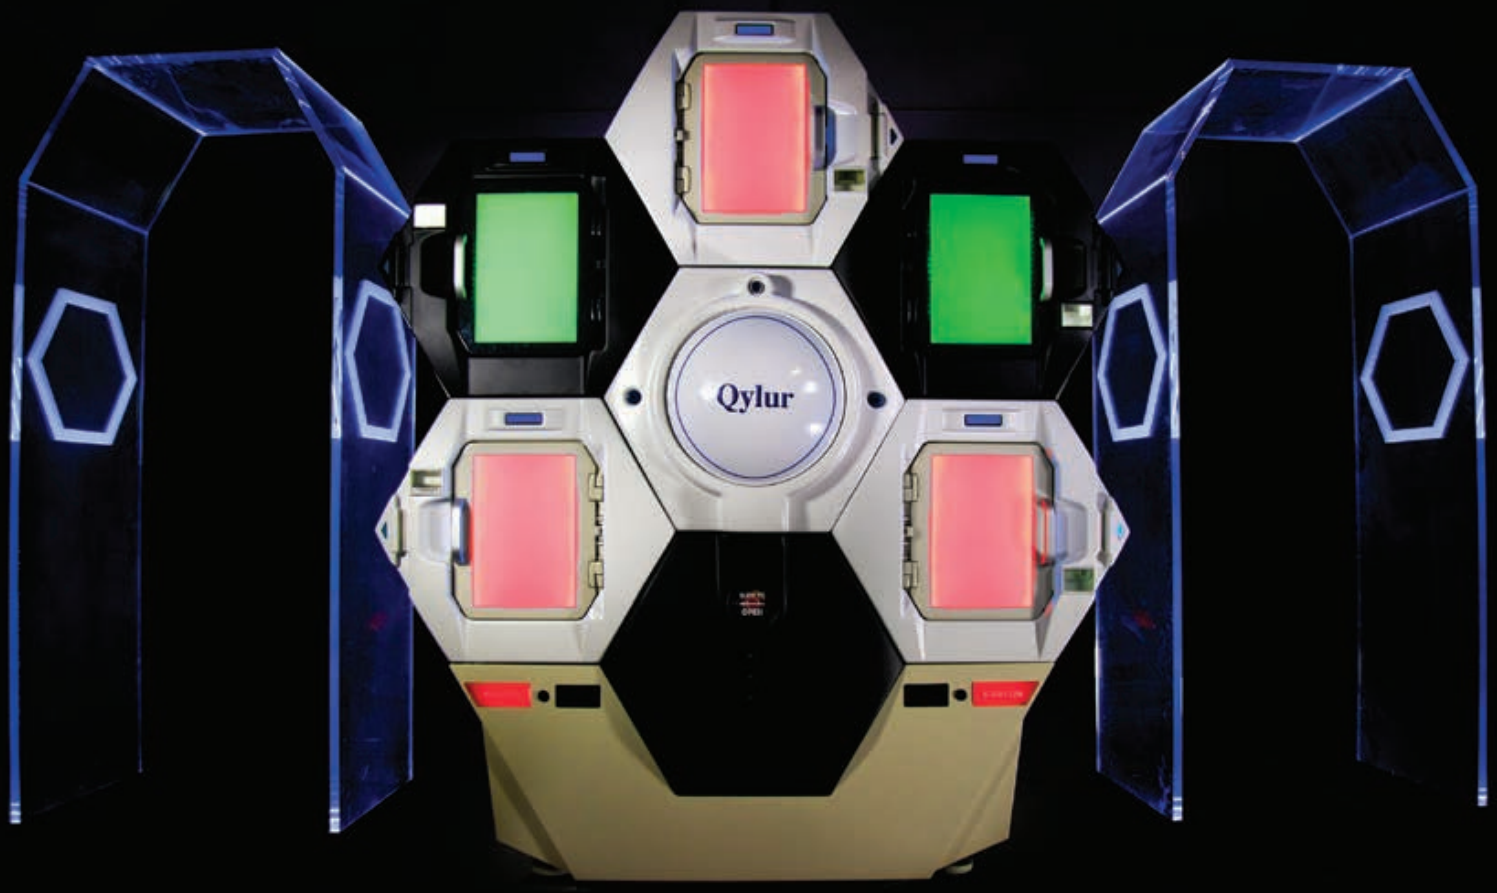
Qylur

Public Safety

Pinpointing Gunshots to Fight Crime

Several cities in the United States have installed a gunshot-detecting system called ShotSpotter that uses networked audio sensors dispersed through city blocks and machine learning to automatically identify the audio signatures of gunshots and report their location to police with high degrees of accuracy.[149] ShotSpotter's algorithms can differentiate between gunshots and other similar noises, such as a car backfiring, triangulate the origin of a gunshot, determine if multiple firearms were involved, and identify the direction a shooter was traveling in.[150]

Predicting Crime Hotspots

Hitachi is piloting software in several U.S. cities that uses natural language processing and machine learning to analyze hundreds of data points, such as social-media activity, 911-call locations, geotagged posts' proximity to schools, and other variables to create heat maps of areas likely to have elevated levels of criminal activity.[151] The software is adept at analyzing public-geotagged social-media activity to identify phrases associated with drug activity and learn anomalous keywords that are likely coded references to drugs.[152]

Autonomously Disposing of Car Bombs

An EU-funded project called the Autonomous Vehicle Emergency Recovery Tool (AVERT) uses a system of four autonomous robotic platforms that coordinate with each other to position themselves under a vehicle suspected of having an explosive device and move it to a safe location.[153] AVERT uses sensor technology called LIDAR to map its environment and automatically develop an extraction route for suspicious vehicles in situations where it is too dangerous or difficult to use normal bomb-disposal tools.[154]

Predicting Buildings' Fire Risk

Data scientists working with the Atlanta Fire Rescuer Department (AFRD) developed a machine-learning model called Firebird that can analyze data to recommend inspection for buildings that pose the highest fire risks, to help AFRD better prioritize inspections.[155] Firebird analyzes past data about fire inspections and fires, commercial property data such as building materials and size, and other information to assign a risk score to each building. Firebird was able to predict which buildings would catch fire with up to 71 percent accuracy in tests on historical data.[156]

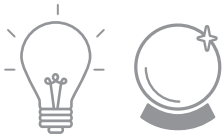

Making Security Screening Less Invasive

Silicon Valley-based Qylur Security Systems has developed a security-screening machine called the Qylatron Entry Experience Solution that uses machine learning to quickly analyze the contents of bags at security checkpoints and identify dangerous objects without humans having to sort through them manually.[157] The machine uses chemical and radiation sensors as well as X-ray to scan a bag, and if the software's algorithms identify objects that look like contraband, such as guns or knives, it will alert security personnel.[158]

Social Good

Mapping Poverty With Satellite Data

Researchers at Stanford University have developed a prototype machine-learning algorithm that can analyze satellite imagery to map impoverished areas, which can help governments and development organizations improve aid efforts.[159] Researchers have previously analyzed nighttime satellite images that use light levels as a measure of economic activity, but this approach is not reliable for low-light or densely populated areas. The Stanford researchers' algorithm learned to identify 4,096 markers of economic activity in daylight images, such as roads and waterways, and by combining this data with light-level analysis, the algorithm was able to predict impoverished areas more reliably than traditional methods.[160]

Measuring Literacy Rates

Norwegian telecom company Telenor's research division has developed a machine-learning algorithm that can determine literacy rates in developing countries by analyzing mobile-phone data.[161] Telenor surveyed thousands of mobile-phone users in a developing country and compared this data to its own phone records, such as number of calls and texts in an area, call length, and number of contacts a user has.[162] Telenor trained its algorithm on this data and was able to identify several factors that can predict literacy rates, such as the difference between incoming and outgoing text messages, and correctly identify if a mobile-phone user is literate with 70 percent accuracy.[163]

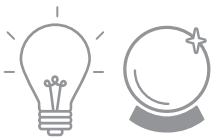

Cracking Down on Human Trafficking

The U.S. Defense Advanced Research Projects Agency's (DARPA's) Memex program has developed a tool that scans pages on the deep web—websites that are not indexed on search engines—and analyzes their contents for signs of illegal solicitations of sex, which are often linked with human trafficking, to aid investigations.[164] Because pages on the deep web can be difficult to access and navigate, they can be attractive covers for criminal activity. With Memex, authorities can gain new insights into specific investigations as well as broad trends about human trafficking, such as by generating heat maps of regions with a high density of illegal solicitation.[165]

Stopping Abusive Internet Trolls

Jigsaw, a subsidiary of Google that develops tools to promote freedom of expression and combat extremism online, has developed an artificial intelligence tool called Conversation AI designed to detect and filter out abusive language online to combat online harassment.[166] Jigsaw trained Conversation AI on 17 million comments on *New York Times* articles and 13,000 discussions on Wikipedia pages to identify abusive language with 92 percent accuracy.[167] Jigsaw will make Conversation AI available as open source to allow web developers to implement it for their sites as a filter that blocks abusive language in comments.[168]

Supporting Refugees' Mental Health

A start-up called X2AI has created an artificial-intelligence program named Karim that uses natural language processing to have conversations with users in Arabic via text message and analyzes their emotional states to provide recommendations that can help improve their mental health.[169] X2AI has partnered with the disaster-relief nonprofit Field Organization Team to raise awareness about X2AI in refugee camps, where many have experienced emotional trauma.[170] Karim is not intended to replace traditional counseling, but rather serve as a supportive "friend" for refugees in camps where mental-health services are rare.[171]


TOPLINE
SCANIA
R620
AXJ 432

Transportation

Making Public Transportation Autonomous

Wageningen University in the Netherlands is piloting a self-driving public shuttle bus called WEpod that uses AI to help navigate public roads and drive among regular traffic.[172] The WEpod can hold up to six people at a time, and passengers use a smartphone app to book their seats and specify their destinations.[173]

Platooning Autonomous Trucks

As part of the Dutch government's European Truck Platooning Challenge, a fleet of autonomous trucks has successfully traveled over 2,000 kilometers by platooning—monitoring each other's speed, proximity, and the road to drive close together to improve efficiency, much like bicyclists drafting off each other.[174] It is difficult for human truck drivers to platoon, given the level of coordination required to do it safely, but platooning autonomous trucks can reduce fuel consumption by up to 15 percent and can also reduce congestion.[175]

Hailing a Self-Driving Taxi

Uber has launched a test fleet of self-driving taxis in Pittsburgh.[176] The taxis, modified Volvo SUVs, have a human in the driver's seat to take over driving if necessary, but can drive around the city to pick up and drop off passengers almost entirely autonomously.[177]

Teaching Trains to Drive Themselves

Australian mining company Rio Tinto has developed an autonomous rail system called AutoHaul to transport iron ore from its mines.[178] AutoHaul trains are equipped with an array of sensors that can detect obstructions on tracks so the software can respond accordingly, and can be safer, more efficient, and operate with less downtime than human-operated trains.[179]

Making Long-Haul Trucking Easier

Researchers at the University of Minnesota have developed a computer-vision system that analyzes truck parking lots along the highway to automatically detect when a spot is available and notify approaching truckers.[180] Roadside parking for truckers is limited and often crowded, making it difficult for long-haul truckers to plan ahead for where they can sleep and raising the risk of overtired drivers staying on the road longer than they should. The system analyzes videos of truck stops and can distinguish between open and occupied spaces with 95 percent accuracy, which is more reliable than systems that rely on sensors embedded in pavement.[181]

Conclusion

The examples in this report only scratch the surface of the many ways that AI is driving innovation in the public and private sectors, generating substantial social and economic value, and transforming everyday life around the globe. But as with any new technology, there will inevitably be detractors who fear change and how it might impact them. While policymakers should respond to legitimate concerns, they should not allow alarmists to delay progress. Instead, they should remain steadfastly focused on accelerating the development and adoption of AI to usher in its many benefits.

Glossary

Artificial intelligence (AI)

AI is the field of computer science devoted to creating computing machines and systems that perform operations analogous to human learning and decision-making.

Machine learning

Machine learning is a subset of AI that focuses on providing algorithms with the ability to learn how to complete tasks without explicit instructions, allowing them to adapt in the presence of new data.[182]

Artificial neural network

An artificial neural network, sometimes called just a neural network, is a type of machine learning that processes data by modeling the neurological structure and function of the brain.[183] Just as a human brain relies on complex webs of neurons to interpret stimuli, an artificial neural network uses layers of processing units acting as simulated neurons, can learn to identify patterns in data, and can adapt and refine themselves as they process new data. AI applications that rely on analyzing unstructured data, such as voice recognition, rely on artificial neural networks.

Deep learning

Deep learning is an advanced type of machine learning adept at identifying abstract or complex patterns in data by using a multilayered artificial neural network.[184] It is called "deep" because it has more layers than simple machine-learning approaches, which allows for more complex processing.

Natural language processing

Natural language processing is a field of AI focused on enabling computers to interpret human language.[185] Though easy for humans to understand, language relies on incredibly complex and abstract relationships between words, syntax, context, and other subtler factors such as emotion, making it difficult for computer systems to interpret.

Computer vision

Computer vision focuses on analyzing unstructured data from digital images and videos to understand their contents, often by simulating the human visual system.[186] Just like natural language processing, computer-vision algorithms must be capable of identifying abstract patterns and learning by example.

Cognitive computing

While AI generally tries to replicate the biological processes of learning and decision-making, cognitive computing focuses specifically on simulating the process of cognition, or human thought through the use of multiple types of analysis and pattern recognition, particularly natural language processing.[187]

References

1. “Grand Challenge ’05,” DARPA, accessed September 29, 2016, http://archive.darpa.mil/grandchallenge05/.

2. “Dave Ferrucci at Computer History Museum: How It All Began and What’s Next,” *IBM Research Blog*, December 1, 2011, https:// www.ibm.com/blogs/research/2011/12/dave-ferrucci-at-computer-history-museum-how-it-all-began-and-whats-next/.

3. Sam Byford, “Google’s AlphaGo AI Beats Lee Se-dol Again to Win Go Series 4-1,” *The Verge*, March 15, 2016, http://www.theverge.com/2016/3/15/11213518/alphago-deepmind-go-match-5-result.

4. Olivia Solon, “Karim the AI Delivers Psychological Support to Syrian Refugees,” *The Guardian*, March 22, 2016, https://www.theguardian.com/technology/2016/mar/22/karim-the-ai-delivers-psychological-support-to-syrian-refugees; Dina Bass, “Microsoft Develops AI to Help Cancer Doctors Find the Right Treatments,” *Bloomberg,* September 20, 2016, http://www.bloomberg.com/news/articles/2016-09-20/microsoft-develops-ai-to-help-cancer-doctors-find-the-right-treatments.

5. Robert D. Atkinson, “‘It’s Going to Kill Us!’ and Other Myths About the Future of Artificial Intelligence,” (Information Technology and Innovation Foundation, June 2016), http://www2.itif.org/2016-myths-machine-learning.pdf.

6. “AI Overview: Broad Discussions of Artificial Intelligence,” AITopics, accessed September 29, 2016, http://aitopics.org/topic/ai-overview.

7. Herbert A. Simon, *The Shape of Automation for Men and Management* (New York: Harper & Row, 1965).

8. Calum Chase, *Surviving AI: The Promise and Peril of Artificial Intelligence* (Three C’s, 2015), 104.

9. John Markoff, *Machines of Loving Grace: The Quest for Common Ground Between Humans and Robots* (New York: Ecco, 2015), 114.

10. Pedro Domingos, *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World* (New York: Basic Books, 2015).

11. Amit Karp, “Deep Learning Will Be Huge—and Here’s Who Will Dominate It,” *Venture Beat*, April 2, 2016, http://venturebeat.com/2016/04/02/deep-learning-will-be-huge-and-heres-who-will-dominate-it/.

12. Irving Wladawksy-Berger, “‘Soft’ Artificial Intelligence Is Suddenly Everywhere,” *The Wall Street Journal*, January 16, 2016, http://blogs.wsj.com/cio/2015/01/16/soft-artificial-intelligence-is-suddenly-everywhere/.

13. Ibid.

14. Atkinson, “‘It’s Going to Kill Us!’ and Other Myths About the Future of Artificial Intelligence.”

15. Ibid.

16. “Cognitive Systems Accelerate Competitive Advantage,” IDC, accessed September 29, 2016, http://www.idc.com/promo/thirdplatform/innovationaccelerators/cognitive.

17. “Artificial Intelligence Explodes: New Deal Activity Record for AI Startups,” *CB Insights*, June 20, 2016, https://www.cbinsights.com/blog/artificial-intelligence-funding-trends/.

18. James Manyika et al., “Disruptive Technologies: Advances That Will Transform Life, Business, and the Global Economy,” (McKinsey Global Institute, May 2013), http://www.mckinsey.com/business-functions/business-technology/our-insights/disruptive-technologies.

19. Mark Purdy and Paul Daugherty, “Why Artificial Intelligence Is the Future of Growth,” (Accenture, September 28, 2016), https://www.accenture.com/us-en/_acnmedia/PDF-33/Accenture-Why-AI-is-the-Future-of-Growth.pdf.

20. Ibid.

21. Will Knight, “Could AI Solve the World’s Biggest Problems,” *MIT Technology Review*, https://www.technologyreview.com/s/545416/could-ai-solve-the-worlds-biggest-problems/.

22. Larry Greenemeier, “Human Traffickers Caught on Hidden Internet,” *Scientific American*, February 8, 2015, http://www.scientificamerican.com/article/human-traffickers-caught-on-hidden-internet/; Davey Alba, “Weeding Out Online Bullying Is Tough, So Let Machines Do It,” *Wired*, July 10, 2015, https://www.wired.com/2015/07/weeding-online-bullying-tough-let-machines/; Michelle Horton, “Stanford Scientists Combine Satellite Data, Machine Learning to Map Poverty,” *Stanford News*, August 18, 2016 http://news.stanford.edu/2016/08/18/combining-satellite-data-machine-learning-to-map-poverty/; Sean Captain, “How Artificial Intelligence is Finding Gender Bias at Work,” *Fast Company*, October 10, 2015, https://www.fastcompany.com/3052053/elasticity/how-artificial-intelligence-is-finding-gender-bias-at-work.

23. Lindsey Kratochwill, “IBM Watson Can Now Help You With Your Christmas Shopping,” *Popular Science*, November 17, 2015, http://www.popsci.com/ibm-watson-can-now-help-you-with-your-christmas-shopping; Glen Martin, “This Company Uses AI to Accelerate Drug Discovery,” *O’Reilly*, July 13, 2016, https://www.oreilly.com/ideas/this-company-uses-ai-to-accelerate-drug-discovery.

24. “Artificial Intelligence Expedites Breast Cancer Risk Prediction,” *ScienceDaily*, August 29, 2016, https://www.sciencedaily.com/releases/2016/08/160829122106.htm.

25. Nelson Sizwe et al., "Integrating Artificial Intelligence Into Data Warehousing and Data Mining," *Proceedings of the World Congress on Engineering and Computer Science 2015 2* (October 2015), pp. 819-823, http://www.iaeng.org/publication/WCECS2015/WCECS2015_pp819-823.pdf.

26. Ben Popper, "How Netflix Completely Revamped Recommendations for Its New Global Audience," *The Verge*, February 17, 2016, http://www.theverge.com/2016/2/17/11030200/netflix-new-recommendation-system-global-regional.

27. Jared Newman, "Everything You Need to Know About AI Assistants, From Siri to Ozlo," *Fast Company*, May 17, 2016, https://www.fastcompany.com/3059719/handicapping-the-ai-assistants-from-siri-to-ozlo; Klint Finley, "Robot Radiologists Will Soon Analyze Your X-Rays," *Wired*, October 27, 2015, https://www.wired.com/2015/10/robot-radiologists-are-going-to-start-analyzing-x-rays/; Chris Weller, "The World's First Artificially Intelligent Lawyer Was Just Hired At A Law Firm," *Tech Insider*, May 16, 2016, http://www.techinsider.io/the-worlds-first-artificially-intelligent-lawyer-gets-hired-2016-5.

28. Joshua Hill, "Google's DeepMind AI reduces Data Center Cooling by 40%," *Clean Technica*, July 22, 2016, https://cleantechnica.com/2016/07/22/googles-deepmind-ai-reduces-data-center-cooling-bill-40/.

29. Eric Limer, "A Fleet of Self-Driving Trucks Just Completed a 1,000-Mile Trip Across Europe," *Popular Mechanics*, April 7, 2016, http://www.popularmechanics.com/cars/trucks/a20310/european-platooning-challenge-self-driving-trucks-1000-miles/; Alexis Madrigal, "Autonomous Robots Invade Retail Warehouses," *Wired*, January 27, 2009, https://www.wired.com/2009/01/retailrobots/.

30. Casey Newton, "Facebook Begins Using Artificial Intelligence to Describe Photos to Blind Users," *The Verge*, April 5, 2016, http://www.theverge.com/2016/4/5/11364914/facebook-automatic-alt-tags-blind-visually-impared.

31. Ibid.

32. Sarah Griffiths, "Automatic Sign Language Tool Can Translate Gestures into 'Readable Language,'" *Daily Mail*, June 23, 2014, http://www.dailymail.co.uk/sciencetech/article-2666166/Automatic-sign-language-tool-translate-gestures-readable-language.html.

33. Ibid.

34. Weronika Jurkiewicz, "Machine Learning Can Help Kids With Autism Recognize Emotions," *PSFK*, March 23, 2016, http://www.psfk.com/2016/03/machine-learning-help-kids-with-autism-microsoft-oxford-project.html.

35. Ibid.

36. Ian Steadman, "Applying Lessons From Self-Driving Cars to Self-Driving Wheelchairs," *NewStatesman*, November 25, 2013, http://www.newstatesman.com/future-proof/2013/11/applying-lessons-self-driving-cars-self-driving-wheelchairs.

37. Ibid

38. Rachel Metz, "App Listens for Danger When You're Not Paying Attention," *MIT Technology Review*, February 26, 2014, https://www.technologyreview.com/s/524971/app-listens-for-danger-when-youre-not-paying-attention/.

39. Ibid.

40. "How Motorleaf is Helping Automate Indoor Farming," *AgFunderNews*, July 20, 2016, https://agfundernews.com/motorleaf-helping-automate-indoor-farming.html.

41. Ibid.

42. Leanna Garfield, "This Robot Has 'Eyes' That Can See Dying Plants Before Farmers Can," *Tech Insider*, July 8, 2016, http://www.techinsider.io/prospera-robot-can-see-dying-plants-before-farmers-2016-7.

43. Ibid.

44. Robinson Meyer, "The Start-Up That Watches Cord Grow, From Orbit," *The Atlantic*, August 10, 2016, http://www.theatlantic.com/technology/archive/2016/08/seeing-corn-with-satellites/495149/.

45. Joshua New, "5 Q's for Mark Johnson, Founder of Descartes Labs," *Center for Data Innovation*, September 2, 2016, https://www.datainnovation.org/2016/09/5-qs-for-mark-johnson-founder-of-descartes-labs/.

46. Matt Simon, "The Future of Humanity's Food Supply Is in the Hands of AI," *Wired*, May 25, 2016, https://www.wired.com/2016/05/future-humanitys-food-supply-hands-ai/.

47. Ibid.

48. Dave Gershgorn, "The Ultimate Promise of Artificial Intelligence Lies in Sorting Cucumbers," *Quartz*, September 1, 2016, http://qz.com/771921/the-ultimate-promise-of-artificial-intelligence-lies-in-sorting-cucumbers/.

49. Ibid.

50. Roland Moore-Colyer, "IBM Launches Deep Thunder to Use Machine Learning and Big Data for Local Weather Forecasts," *V3*, June 16, 2016, http://www.v3.co.uk/v3-uk/news/2461798/ibm-launches-deep-thunder-to-use-machine-learning-and-big-data-for-local-weather-forecasts.

51. Ibid.

52. "Zendesk's Satisfaction Prediction Brings Machine Learning to the Customer Experience," news release, March 15, 2016, https://www.zendesk.com/company/press/zendesks-satisfaction-prediction-brings-machine-learning-customer-experience/.

53. Ibid.

54. Vivian Giang, "The Growing Business of Detecting Unconscious Bias," *Fast Company*, May 5, 2015, https://www.fastcompany.com/3045899/hit-the-ground-running/the-growing-business-of-detecting-unconscious-bias.

55. Ibid.

56. Ben Popper, "The Smart Bots Are Coming and This One is Brilliant," *The Verge*, April 7 2016, http://www.theverge.com/2016/4/7/11380470/amy-personal-digital-assistant-bot-ai-conversational.

57. Ibid.

58. Ibid.

59. Rip Empson, "Unlabel Launches a Human-Edited Machine Translation Service to Help Businesses Go Global, Localize Customer Support," *TechCrunch*, March 23, 2014, https://techcrunch.com/2014/03/23/unbabel-launches-a-human-edited-machine-translation-service-to-help-businesses-go-global-and-localize-customer-support/.

60. Ibid.

61. Ainsley O'Connell, "Your Photo of a Burrito is Now Worth a Thousand Words," *Fast Company*, July 17, 2016, https://www.fastcompany.com/3060884/your-photo-of-a-burrito-is-now-worth-a-thousand-words.

62. Ibid.

63. Preetam Kaushik, "Is Artificial Intelligence the Way Forward for Personal Finance?," *Wired*, February, 2014, https://www.wired.com/insights/2014/02/artificial-intelligence-way-forward-personal-finance/.

64. Ibid.

65. April Glaser, "Pepper, the Emotional Robot, Learns How to Feel like an American," *Wired*, June 7, 2016, https://www.wired.com/2016/06/pepper-emotional-robot-learns-feel-like-american/.

66. Ibid.

67. Tom Brant, "Pinterest Encourages Shopping Frenzy with New AI Features," *PC Magazine*, June 28, 2016, http://www.pcmag.com/news/345688/pinterest-encourages-shopping-frenzy-with-new-ai-features.

68. Max Lewontin, "How Google Photos Uses Machine Learning to Create Customized Albums," *Christian Science Monitor,* March 24, 2016, http://www.csmonitor.com/Technology/2016/0324/How-Google-Photos-uses-machine-learning-to-create-customized-albums.

69. Ibid.

70. Chris Baraniuk, "Earthquake Artificial Intelligence Knows Where Damage is Worst," *New Scientist*, September 30, 2015, https://www.newscientist.com/article/mg22830412-800-earthquake-artificial-intelligence-knows-where-damage-is-worst/.

71. Ibid.

72. Andrew Good, "A.I. Could be a Firefighter's 'Guardian Angel,'" NASA Jet Propulsion Laboratory, August 11, 2016, http://www.jpl.nasa.gov/news/news.php?feature=6590.

73. Ibid.

74. Ibid.

75. Larry Greenemeier, "Smart Machines Join Humans in Tracking Africa Ebola Outbreak," *Scientific American*, September 24, 2014, http://www.scientificamerican.com/article/smart-machines-join-humans-in-tracking-africa-ebola-outbreak/.

76. Ibid.

77. "HealthMap," HealthMap, accessed September 29, 2016, http://www.healthmap.org/en/.

78. Katie Collins, "How AI, Twitter and Digital Volunteers Are Transforming Humanitarian Disaster Response," *Wired*, September 30, 2013, http://www.wired.co.uk/article/digital-humanitarianism.

79. Ibid.

80. Rebecca Merrett, "Machine Learning Used to Predict Hazardous Solar Flares," *CIO*, January 16, 2015, http://www.cio.com.au/article/564056/machine-learning-used-predict-hazardous-solar-flares/.

81. Ibid.

82. Elizabeth Harris, "Next Target for IBM's Watson? Third-Grade Math," *New York Times*, September 27, 2016, http://www.nytimes.com/2016/09/28/nyregion/ibm-watson-common-core.html?_r=0.

83. Ibid.

84. Ibid.

85. “ML Predicts School Dropout Risk & Boosts Graduation Rates,” Microsoft, June 4, 2015, https://blogs.technet.microsoft.com/machinelearning/2015/06/04/ml-predicts-school-dropout-risk-boosts-graduation-rates/.

86. Ibid.

87. Jason Maderer, “Artificial Intelligence Course Creates AI Teaching Assistant,” *Georgia Tech News Center*, May 9, 2016, http://www.news.gatech.edu/2016/05/09/artificial-intelligence-course-creates-ai-teaching-assistant.

88. Ibid.

89. Ibid.

90. Parul Goliani, “Duolingo Looks to Dominate the Mobile Education Market With New Flashcard App TinyCards,” *Forbes*, July 22, 2016, http://www.forbes.com/sites/parulguliani/2016/07/22/duolingo-looks-to-dominate-the-mobile-education-market-with-new-flashcard-app/.

91. Ibid.

92. Carl Straumsheim, “Detecting More Than Plagiarism,” *Inside Higher Ed*, January 21, 2016, https://www.insidehighered.com/news/2016/01/21/turnitin-expanding-beyond-plagiarism-detection-launches-revision-assistant.

93. Ibid.

94. Ibid.

95. Lucas Mearian, “IBM’s Machine-Learning Crystal Ball Can Foresee Renewable Energy Availability,” *Computerworld*, July 16, 2015, http://www.computerworld.com/article/2948987/sustainable-it/ibms-machine-learning-crystal-ball-can-foresee-renewable-energy-availability.html.

96. Ibid.

97. “ORNL Researchers Develop ‘Autotune’ Software to Make It Quicker, Easier, and Cheaper to Model Energy Use of Buildings,” Department of Energy, October 15, 2014, http://energy.gov/eere/buildings/articles/ornl-researchers-develop-autotune-software-make-it-quicker-easier-and.

98. Ibid.

99. Jack Clark, “I’ll Be Back: The Return of Artificial intelligence,” *Bloomberg*, February 3, 2015, http://www.bloomberg.com/news/articles/2015-02-03/i-ll-be-back-the-return-of-artificial-intelligence.

100. Ibid.

101. Ibid.

102. “Programming Isn’t Proven,” Nest, accessed September 29, 2016, https://nest.com/thermostat/real-savings/; “An Introduction to Learning on the Nest Learning Thermostat,” Nest, accessed September 29, 2016, https://nest.com/support/article/An-introduction-to-learning.

103. “An Introduction to Learning on the Nest Learning Thermostat,” Nest.

104. Larry Hardesty, “Siting Wind Farms More Quickly, Cheaply,” *MIT News*, July 17, 2015, http://news.mit.edu/2015/siting-wind-farms-quickly-cheaply-0717.

105. Ibid.

106. Adele Peters, “This AI Watches Satellite Photos and Says ‘It Looks Like You’re About to Cut Down a Forest. Could You Not?,’” *Wired*, May 11, 2015, https://www.fastcoexist.com/3046014/this-ai-watches-satellite-photos-and-says-it-looks-like-youre-about-to-cut-down-a-forest-cou.

107. Ibid.

108. Will Knight, “How Artificial Intelligence Can Fight Air Pollution in China,” *MIT Technology Review*, August 31, 2015, https://www.technologyreview.com/s/540806/how-artificial-intelligence-can-fight-air-pollution-in-china/.

109. Ibid.

110. Ibid.

111. Jackie Snow, “Rangers Use Artificial Intelligence to Fight Poachers,” *National Geographic*, June 12, 2016, http://news.nationalgeographic.com/2016/06/paws-artificial-intelligence-fights-poaching-ranger-patrols-wildlife-conservation/.

112. Ibid.

113. Ibid.

114. Jody Avirgan, “Big Data Is Saving This Little Bird,” *FiveThirtyEight*, October 16, 2015, http://fivethirtyeight.com/datalab/big-data-is-saving-this-little-bird/.

115. Ibid.

116. Arion McNicoll and Stefanie Blendis, “Green Machine: Intelligent Robot System Recycles Waste,” *CNN*, June 10, 2013, http://www.cnn.com/2013/06/07/tech/zenrobotics-recycling-robot/.

117. Ibid.

118. "Now There's an App for That: Computers Can Recognise a Complication of Diabetes That Can Lead to Blindness," *The Economist*, September 19, 2015, http://www.economist.com/news/science-and-technology/21664943-computers-can-recognise-complication-diabetes-can-lead-blindness-now.

119. Ibid.

120. Adrienne Lafrance, "Computers Can Predict Schizophrenia Based on How a Person Talks," *The Atlantic*, August 26, 2015, http://www.theatlantic.com/technology/archive/2015/08/speech-analysis-schizophrenia-algorithm/402265/.

121. Ibid.

122. Laura Lorenzetti, "This Company Wants to Cure Pancreatic Cancer Using AI," *Fortune*, April 22, 2016, http://fortune.com/2016/04/22/berg-pacreatic-cancer-artificial-intelligence/.

123. Ibid.

124. Anna Nowogrodzki, "Artificial Intelligence Offers a Better Way to Diagnose Malaria," *MIT Technology Review*, February 10, 2016, https://www.technologyreview.com/s/600779/artificial-intelligence-offers-a-better-way-to-diagnose-malaria/#/set/id/600798/.

125. Ibid.

126. Ibid.

127. "How Machine Learning Can Help with Voice Disorders," *ScienceDaily*, August 29, 2016, https://www.sciencedaily.com/releases/2016/08/160829140438.htm.

128. Ibid.

129. Lean A. Santos, "Using Artificial Intelligence to Revolutionize Diabetes Treatment," *Devex*, April 8, 2016, https://www.devex.com/news/using-artificial-intelligence-to-revolutionize-diabetes-treatment-87989.

130. Ibid.

131. Chris Wood, "Machine-Learning Robot Could Streamline Drug Development," *New Atlas*, February 10, 2016, http://newatlas.com/machine-learning-drug-development/41759/.

132. Ibid.

133. Eliza Strickland, "Autonomous Robot Surgeon Bests Humans in World First," *IEEE Spectrum*, May 4, 2016, http://spectrum.ieee.org/the-human-os/robotics/medical-robots/autonomous-robot-surgeon-bests-human-surgeons-in-world-first.

134. Ibid.

135. Alex Hern, "Google DeepMind and UCLH Collaborate on AI-Based Radiotherapy Treatment," *The Guardian*, August 30, 2016, https://www.theguardian.com/technology/2016/aug/30/google-deepmind-ucl-ai-radiotherapy-treatment-.

136. Ibid.

137. "Scientists Teaching Machines to Make Clinical Trials More Successful," *Science Daily*, April 27, 2016, https://www.sciencedaily.com/releases/2016/04/160427095057.htm.

138. Ibid.

139. "Offerings," DataRPM, accessed September 29, 2016, http://www.datarpm.com/offerings.php.

140. Ibid.

141. Christopher Alessi and Chase Gummer, "Germany Bets on 'Smart Factories' to Keep Its Manufacturing Edge," *The Wall Street Journal*, October 26, 2014, http://www.wsj.com/articles/germany-bets-on-smart-factories-to-keep-its-manufacturing-edge-1414355745.

142. Ibid.

143. Tomas Kellner, "Deep Machine Learning: GE and BP Will Connect Thousands of Subsea Oil Wells to the Industrial Internet," *GE Reports*, July 8, 2015, http://www.gereports.com/post/123572457345/deep-machine-learning-ge-and-bp-will-connect/.

144. "GE's Industrial Software to Boost Efficiency across BP's Global Operations," news release, July 8 2015, https://www.ge.com/digital/press-releases/ge-industrial-software-to-boost-efficiency-across-bp-global-operations; "GE Oil & Gas Partners with BP, Paradigm," *Offshore*, January 28, 2016, http://www.offshore-mag.com/articles/2016/01/ge-oil-gas-partners-with-bp-paradigm.html.

145. "In the World of Perishables, Forecast Accuracy is Key," *SupplyChainBrain*, August 6, 2014, http://www.supplychainbrain.com/content/industry-verticals/food-beverage/single-article-page/article/in-the-world-of-perishables-forecast-accuracy-is-key/.

146. Ibid.

147. "Project Dreamcatcher," Autodesk, accessed October 3, 2016, https://autodeskresearch.com/projects/Dreamcatcher.

148. Ibid.

149. “Shot in the Dark: New Surveillance Tool Called ShotSpotter Traces and Records Incidents of Gunfire,” *ScienceDaily*, April 16, 2016, https://www.sciencedaily.com/releases/2016/04/160416130850.htm.

150. Maud Rozee, “The NYPD Can Now Pinpoint the Exact Location of a Gunshot,” *Gothamist*, March 17, 2015, http://gothamist.com/2015/03/17/nypd_pinpointing_gunshots.php.

151. Sean Captain, “Hitachi Says It Can Predict Crimes Before They Happen,” *Fast Company*, September 28, 2015, https://www.fastcompany.com/3051578/elasticity/hitachi-says-it-can-predict-crimes-before-they-happen.

152. Ibid.

153. Dan Moren, “Autonomous Robots to Help Remove Car Bomb Threats,” *Popular Science*, April 29, 2015, http://www.popsci.com/autonomous-robots-help-remove-car-bomb-threats.

154. Ibid.

155. Michael Madaio et al., “Firebird: Predicting Fire Risk and Prioritizing Fire Inspections in Atlanta,” (KDD 2016, San Francisco, California, August 13-17, 2016), http://www.kdd.org/kdd2016/subtopic/view/firebird-predicting-fire-risk-and-prioritizing-fire-inspections-in-atlanta.

156. Ibid.

157. Alex Davies, “The Brilliant Machine That Could Finally Fix Airport Security,” *Wired*, July 7, 2014, https://www.wired.com/2014/07/qylur-security-world-cup/.

158. Ibid.

159. Sarah Emerson, “Artificial Intelligence is Predicting Human Poverty from Space,” *Motherboard*, August 18, 2016, http://motherboard.vice.com/en_ca/read/artificial-intelligence-is-predicting-human-poverty-from-space.

160. Ibid.

161. “Mobile Phone Data Reveals Literacy Rates in Developing Countries,” *MIT Technology Review*, July 13, 2016, https://www.technologyreview.com/s/601895/mobile-phone-data-reveals-literacy-rates-in-developing-countries/.

162. Ibid.

163. Ibid.

164. Larry Greenemeier, “Human Traffickers Caught on Hidden Internet,” *Scientific American*, February 8, 2015, http://www.scientificamerican.com/article/human-traffickers-caught-on-hidden-internet/.

165. Ibid.

166. Andy Greenberg, “Inside Google’s Internet Justice League and Its AI-Powered War on Trolls,” *Wired*, September 19, 2016, https://www.wired.com/2016/09/inside-googles-internet-justice-league-ai-powered-war-trolls/.

167. Ibid.

168. Ibid.

169. Olivia Solon, “Karim the AI Delivers Psychological Support to Syrian Refugees,” *The Guardian*, March 22, 2016, https://www.theguardian.com/technology/2016/mar/22/karim-the-ai-delivers-psychological-support-to-syrian-refugees.

170. Ibid

171. Ibid.

172. Matt Hamblen, “Self-Driving Olli Shuttle With IBM Watson Debuts in Washington Area,” *Computerworld*, June 17, 2016, http://www.computerworld.com/article/3085454/internet-of-things/self-driving-olli-shuttle-with-ibm-watson-debuts-in-washington-area.html.

173. Ibid.

174. Eric Limer, “A Fleet of Self-Driving Trucks Just Completed a 1,000-Mile Trip Across Europe,” *Popular Mechanics*, April 7, 2016, http://www.popularmechanics.com/cars/trucks/a20310/european-platooning-challenge-self-driving-trucks-1000-miles/.

175. Ibid.

176. Max Chafkin, “Uber’s First Self-Driving Fleet Arrives in Pittsburgh This Month,” *Bloomberg,* August 18, 2016, *http*://www.bloomberg.com/news/features/2016-08-18/uber-s-first-self-driving-fleet-arrives-in-pittsburgh-this-month-is06r7on.

177. Ibid.

178. Peter Ker, “Rio Tinto Pushes Ahead with Driverless Trains in Pilbara,” *The Sydney Morning Herald*, March 10, 2015, http://www.smh.com.au/business/mining-and-resources/rio-tinto-pushes-ahead-with-driverless-trains-in-pilbara-20150309-13z28v.html.

179. Ibid.

180. Aarian Marshall, "Parking a Truck Is a Pain in the Butt. Tech to the Rescue!" *Wired*, June 6, 2016, https://www.wired.com/2016/06/parking-truck-pain-butt-tech-rescue/.

181. Ibid.

182. "What Is Machine Learning," TechTarget, accessed September 29, 2016, http://whatis.techtarget.com/definition/machine-learning.

183. "A Basic Introduction to Neural Networks," University of Wisconsin Department of Computer Sciences, accessed September 29, 2016, http://pages.cs.wisc.edu/~bolo/shipyard/neural/local.html.

184. Amit Karp, "Deep Learning Will Be Huge—and Here's Who Will Dominate It," *Venture Beat*, April 2, 2016, http://venturebeat.com/2016/04/02/deep-learning-will-be-huge-and-heres-who-will-dominate-it/.

185. Matt Kiser, "Introduction to Natural Language Processing," *Algorithmia*, August 11, 2016, http://blog.algorithmia.com/introduction-natural-language-processing-nlp/.

186. T.S. Huang, "Computer Vision: Evolution and Promise," *CERN School of Computing*, no. 19 (1996): 21-25, http://cds.cern.ch/record/400313/files/p21.pdf.

187. Bernard Marr, "What Everyone Should Know About Cognitive Computing," *Forbes*, March 23, 2016, http://www.forbes.com/sites/bernard-marr/2016/03/23/what-everyone-should-know-about-cognitive-computing/#a64bc145d6e7; John E. Kelly III, "Computing, Cognition, and the Future of Knowing," (IBM, October 2015), http://www.research.ibm.com/software/IBMResearch/multimedia/Computing_Cognition_WhitePaper.pdf.

Image Credits

Cover

Jesse Allen. *Precision Farming in Minnesota*. September 18, 2009. NASA, https://en.wikipedia.org/wiki/Agriculture#/media/File:Precision_Farming_in_Minnesota_-_Natural_Colour.jpg.

Kristen Vogt. *Woman in Grey Knit Sweater*. February 4, 2015. Pexels, https://www.pexels.com/photo/woman-winter-grey-river-61100/.

Chance Angrella. *Solar Panels Seen in the Meadow*. Freerange, accessed October 7, 2016, https://freerangestock.com/photos/36319/solar-panels-seen-in-the-meadow.html.

Flickr user "Smoothgroover22." *Google Self-Driving Car*. May 28, 2014. Flickr, https://www.flickr.com/photos/smoothgroover22/15104006386.

Wikimedia user "Bobjgalindo." *MRI Brain Tumor*. October 28, 2009. Wikimedia, https://commons.wikimedia.org/wiki/File:MRI_brain_tumor.jpg.

Rish MacIntosh. *Juvenile Marbled Murrelet*. February 14, 2004. USFWS – Pacific Region, https://www.flickr.com/photos/usfwspacific/4974768823.

Accessibility

Rob Maccoll. *Kate Nelson, Australian Volunteer, with Research Assistants Apenisa Matairavula (centre) and Saminsoni Weleilakeba from the Fiji Association for the Deaf, Fiji 2008*. January 1, 2008. Department of Foreign Affairs and Trade, https://www.flickr.com/photos/dfataustralianaid/10706279134.

Hernán Piñera. *Locked in His World*. February 12, 2013. Flickr, https://www.flickr.com/photos/hernanpc/12677922605/.

Agriculture

Jesse Allen. *Precision Farming in Minnesota*. September 18, 2009. NASA, https://en.wikipedia.org/wiki/Agriculture#/media/File:Precision_Farming_in_Minnesota_-_Natural_Colour.jpg.

Business Operations

Flickr user "David." *_MG_5801*. December 19, 2009. Flickr, https://www.flickr.com/photos/65193799@N00/4197816847.

Consumer Convenience

Wikimedia user "Tokumeigakarinoaoshima." *SoftBank Pepper*. July 18, 2014. Wikimedia, https://commons.wikimedia.org/wiki/File:SoftBank_pepper.JPG.

Disaster Prevention and Response

United Nations Development Programme. *Haiti Earthquake Cash for Work*. January 23, 2010. Flickr, https://www.flickr.com/photos/unitednationsdevelopmentprogramme/4298505043/.

Edward N. Johnson. *Camp Humphreys Firefighters 'Turn Up the Heat' with Live Fire Training in Korea*. Flickr, https://www.flickr.com/photos/usaghumphreys/8123607930/in/photostream/.

Education

Wikimedia user "Mrjohncummings." *10 years old Dipa and 12 years old Laboni study in class 2 at "Unique Child learning Center", Mirmur-Dhaka, Bangladesh*. December 5, 2009. Wikimedia, https://commons.wikimedia.org/wiki/File:10_years_old_Dipa_and_12_years_old_Laboni_study_in_class_2_at_%22Unique_Child_learning_Center%22,_Mirmur-Dhaka,_Bangladesh.jpg.

Tulane Public Relations. *Studying with a View*. April 20, 2011. Wikimedia, https://commons.wikimedia.org/wiki/File:Studying_with_a_view_(5653578171).jpg.

Energy

Hans Hillewaert. *Windmills D1-D4 (Thorton Bank)*. September 29, 2008, Wikimedia, https://commons.wikimedia.org/wiki/File:Windmills_D1-D4_(Thornton_Bank).jpg.

Environment

Pixabay user "MoikaP." *Elephant*. Pixabay, accessed October 7, 2016, https://pixabay.com/en/elephant-uganda-joy-water-animals-1535882/.

Health Care: Prevention and Screening

Wikimedia user "Nephron." *Maternal malaria placenta – high mad*. Wikimedia, accessed October 7, 2016, https://commons.wikimedia.org/wiki/File:Maternal_malaria_placenta_-_high_mag.jpg.

Health Care: Treatment and Monitoring

Nick Casberg. *Hand Holding Blood Draw Tube*. Freerange, accessed October 7, 2016, https://freerangestock.com/photos/5254/hand-holding-blood-draw-tube.html.

Industrial Operations

Steve Jurbetson. *Tesla Motors Assembly Line*. March 12, 2012. Flickr, https://www.flickr.com/photos/jurvetson/6858583426.

Public Safety

Photo courtesy of Qylur Intelligent Systems: http://www.qylur.com/.

Social Good

Flickr user "Freedom House." *Syrian Refugees*. September 9, 2015. Flickr, https://www.flickr.com/photos/syriafreedom/21264442335.

Transportation

Scania Lines Up for Platooning Trials. April 4, 2012. Scania, https://www.scania.com/group/en/scania-lines-up-for-platooning-trials/.

Flickr user "Smoothgroover22." *Google Self-Driving Car*. May 28, 2014. Flickr, https://www.flickr.com/photossmoothgroover22/15104006386.

Center for Data Innovation

The Center for Data Innovation is the leading global think tank studying the intersection of data, technology, and public policy. With staff in Washington, DC and Brussels, the Center formulates and promotes pragmatic public policies designed to maximize the benefits of data-driven innovation in the public and private sectors. It educates policymakers and the public about the opportunities and challenges associated with data, as well as technology trends such as predictive analytics, open data, cloud computing, and the Internet of Things. The Center is a nonprofit, nonpartisan research institute affiliated with the Information Technology and Innovation Foundation.

Daniel Castro

Daniel Castro is director of the Center for Data Innovation and vice president at the Information Technology and Innovation Foundation. His research interests include health IT, data privacy, e-commerce, e-government, electronic voting, information security, and accessibility. Previously, Castro worked as an IT analyst at the Government Accountability Office where he audited IT security and management controls at various government agencies. He has a B.S. in foreign service from Georgetown University and an M.S. in information security technology and management from Carnegie Mellon University.

Joshua New

Joshua New is a policy analyst at the Center for Data Innovation. He has a background in government affairs, policy, and communication. His research focuses on methods of promoting innovative and emerging technologies as a means of improving the economy and quality of life. New graduated from American University with degrees in C.L.E.G. (communication, legal institutions, economics, and government) and public communication.

datainnovation.org