ARTIFICIAL INTELLIGENCE AND LIFE IN 2030

ONE HUNDRED YEAR STUDY ON ARTIFICIAL INTELLIGENCE | REPORT OF THE 2015 STUDY PANEL | SEPTEMBER 2016

PREFACE

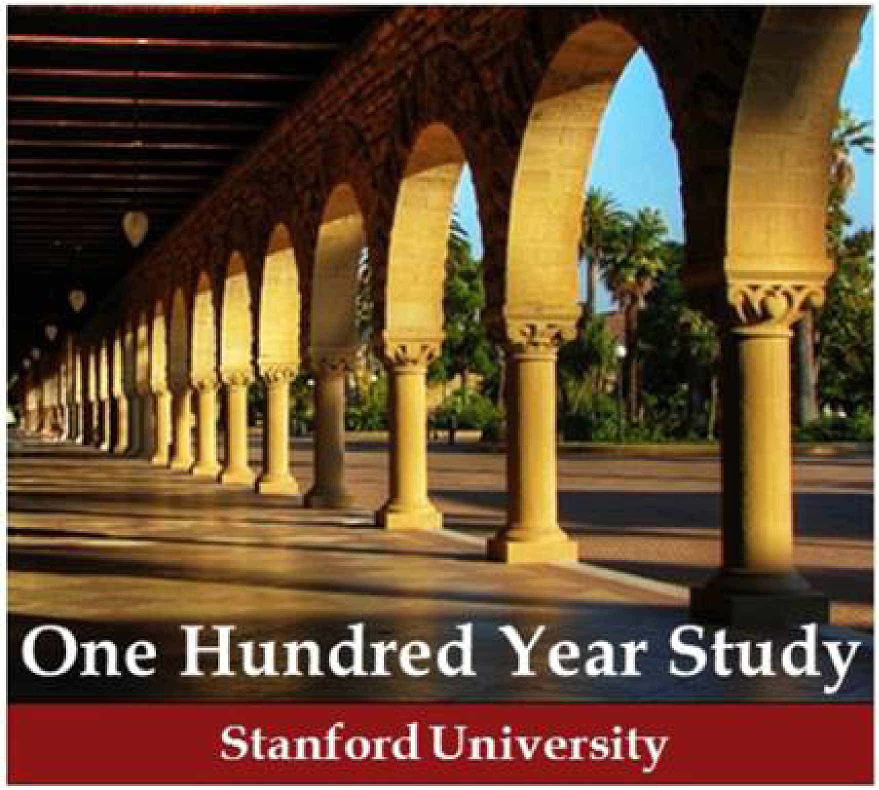

The One Hundred Year Study on Artificial Intelligence, launched in the fall of 2014, is a long-term investigation of the field of Artificial Intelligence (AI) and its influences on people, their communities, and society. It considers the science, engineering, and deployment of AI-enabled computing systems. As its core activity, the Standing Committee that oversees the One Hundred Year Study forms a Study Panel every five years to assess the current state of AI. The Study Panel reviews AI's progress in the years following the immediately prior report, envisions the potential advances that lie ahead, and describes the technical and societal challenges and opportunities these advances raise, including in such arenas as ethics, economics, and the design of systems compatible with human cognition. The overarching purpose of the One Hundred Year Study's periodic expert review is to provide a collected and connected set of reflections about AI and its influences as the field advances. The studies are expected to develop syntheses and assessments that provide expert-informed guidance for directions in AI research, development, and systems design, as well as programs and policies to help ensure that these systems broadly benefit individuals and society.[1]

The overarching purpose of the One Hundred Year Study's periodic expert review is to provide a collected and connected set of reflections about AI and its influences as the field advances.

The One Hundred Year Study is modeled on an earlier effort informally known as the "AAAI Asilomar Study." During 2008-2009, the then president of the Association for the Advancement of Artificial Intelligence (AAAI), Eric Horvitz, assembled a group of AI experts from multiple institutions and areas of the field, along with scholars of cognitive science, philosophy, and law. Working in distributed subgroups, the participants addressed near-term AI developments, long-term possibilities, and legal and ethical concerns, and then came together in a three-day meeting at Asilomar to share and discuss their findings. A short written report on the intensive meeting discussions, amplified by the participants' subsequent discussions with other colleagues, generated widespread interest and debate in the field and beyond.

The impact of the Asilomar meeting, and important advances in AI that included AI algorithms and technologies starting to enter daily life around the globe, spurred the idea of a long-term recurring study of AI and its influence on people and society. The One Hundred Year Study was subsequently endowed at a university to enable

1 "One Hundred Year Study on Artificial Intelligence (AI100)," Stanford University, accessed August 1, 2016, https://ai100.stanford.edu.

extended deep thought and cross-disciplinary scholarly investigations that could inspire innovation and provide intelligent advice to government agencies and industry.

This report is the first in the planned series of studies that will continue for at least a hundred years. The Standing Committee defined a Study Panel charge for the inaugural Study Panel in the summer of 2015 and recruited Professor Peter Stone, at the University of Texas at Austin, to chair the panel. The seventeen-member Study Panel, comprised of experts in AI from academia, corporate laboratories and industry, and AI-savvy scholars in law, political science, policy, and economics, was launched in mid-fall 2015. The participants represent diverse specialties and geographic regions, genders, and career stages.

The Standing Committee extensively discussed ways to frame the Study Panel charge to consider both recent advances in AI and potential societal impacts on jobs, the environment, transportation, public safety, healthcare, community engagement, and government. The committee considered various ways to focus the study, including surveying subfields and their status, examining a particular technology such as machine learning or natural language processing, and studying particular application areas such as healthcare or transportation. The committee ultimately chose a thematic focus on "AI and Life in 2030" to recognize that AI's various uses and impacts will not occur independently of one another, or of a multitude of other societal and technological developments. Acknowledging the central role cities have played throughout most of human experience, the focus was narrowed to the large urban areas where most people live. The Standing Committee further narrowed the focus to a typical North American city in recognition of the great variability of urban settings and cultures around the world, and limits on the first Study Panel's efforts. The Standing Committee expects that the projections, assessments, and proactive guidance stemming from the study will have broader global relevance and is making plans for future studies to expand the scope of the project internationally.

TABLE OF CONTENTS

As one consequence of the decision to focus on life in North American cities, military applications were deemed to be outside the scope of this initial report. This is not to minimize the importance of careful monitoring and deliberation about the implications of AI advances for defense and warfare, including potentially destabilizing developments and deployments.

The report is designed to address four intended audiences. For the general public, it aims to provide an accessible, scientifically and technologically accurate portrayal of the current state of AI and its potential. For industry, the report describes relevant technologies and legal and ethical challenges, and may help guide resource allocation. The report is also directed to local, national, and international governments to help them better plan for AI in governance. Finally, the report can help AI researchers, as well as their institutions and funders, to set priorities and consider the ethical and legal issues raised by AI research and its applications.

Given the unique nature of the One Hundred Year Study on AI, we expect that future generations of Standing Committees and Study Panels, as well as research scientists, policy experts, leaders in the private and public sectors, and the general public, will reflect on this assessment as they make new assessments of AI's future. We hope that this first effort in the series stretching out before us will be useful for both its failures and successes in accurately predicting the trajectory and influences of AI.

The Standing Committee is grateful to the members of the Study Panel for investing their expertise, perspectives, and significant time to the creation of this inaugural report. We especially thank Professor Peter Stone for agreeing to serve as chair of the study and for his wise, skillful, and dedicated leadership of the panel, its discussions, and creation of the report.

Standing Committee of the One Hundred Year Study of Artificial Intelligence

Barbara J. Grosz, *Chair* · Russ Altman · Eric Horvitz
Alan Mackworth · Tom Mitchell · Deidre Mulligan · Yoav Shoham

STUDY PANEL

Peter Stone, University of Texas at Austin, Chair
Rodney Brooks, Rethink Robotics
Erik Brynjolfsson, Massachussets Institute of Technology
Ryan Calo, University of Washington
Oren Etzioni, Allen Institute for AI
Greg Hager, Johns Hopkins University
Julia Hirschberg, Columbia University
Shivaram Kalyanakrishnan, Indian Institute of Technology Bombay
Ece Kamar, Microsoft Research
Sarit Kraus, Bar Ilan University
Kevin Leyton-Brown, University of British Columbia
David Parkes, Harvard University
William Press, University of Texas at Austin
AnnaLee (Anno) Saxenian, University of California, Berkeley
Julie Shah, Massachussets Institute of Technology
Milind Tambe, University of Southern California
Astro Teller, X

Acknowledgments: The members of the Study Panel gratefully acknowledge the support of and valuable input from the Standing Committee, especially the chair, Barbara Grosz, who handled with supreme grace the unenviable role of mediating between two large, very passionate committees. We also thank Kerry Tremain for his tireless and insightful input on the written product during the extensive editing and polishing process, which unquestionably strengthened the report considerably.

EXECUTIVE SUMMARY

Substantial increases in the future uses of AI applications, including more self-driving cars, healthcare diagnostics and targeted treatment, and physical assistance for elder care can be expected.

Artificial Intelligence (AI) is a science and a set of computational technologies that are inspired by—but typically operate quite differently from—the ways people use their nervous systems and bodies to sense, learn, reason, and take action. While the rate of progress in AI has been patchy and unpredictable, there have been significant advances since the field's inception sixty years ago. Once a mostly academic area of study, twenty-first century AI enables a constellation of mainstream technologies that are having a substantial impact on everyday lives. Computer vision and AI planning, for example, drive the video games that are now a bigger entertainment industry than Hollywood. Deep learning, a form of machine learning based on layered representations of variables referred to as neural networks, has made speech-understanding practical on our phones and in our kitchens, and its algorithms can be applied widely to an array of applications that rely on pattern recognition. Natural Language Processing (NLP) and knowledge representation and reasoning have enabled a machine to beat the Jeopardy champion and are bringing new power to Web searches.

While impressive, these technologies are highly tailored to particular tasks. Each application typically requires years of specialized research and careful, unique construction. In similarly targeted applications, substantial increases in the future uses of AI technologies, including more self-driving cars, healthcare diagnostics and targeted treatments, and physical assistance for elder care can be expected. AI and robotics will also be applied across the globe in industries struggling to attract younger workers, such as agriculture, food processing, fulfillment centers, and factories. They will facilitate delivery of online purchases through flying drones, self-driving trucks, or robots that can get up the stairs to the front door.

This report is the first in a series to be issued at regular intervals as a part of the One Hundred Year Study on Artificial Intelligence (AI100). Starting from a charge given by the AI100 Standing Committee to consider the likely influences of AI in a typical North American city by the year 2030, the 2015 Study Panel, comprising experts in AI and other relevant areas focused their attention on eight domains they considered most salient: transportation; service robots; healthcare; education; low-resource communities; public safety and security; employment and workplace; and entertainment. In each of these domains, the report both reflects on progress in the past fifteen years and anticipates developments in the coming fifteen years. Though drawing from a common source of research, each domain reflects different AI influences and challenges, such as the difficulty of creating safe and reliable hardware (transportation and service robots), the difficulty of smoothly interacting with human experts (healthcare and education), the challenge of gaining public trust (low-resource communities and public safety and security), the challenge of overcoming fears of marginalizing humans (employment and workplace), and the social and societal risk of diminishing interpersonal interactions (entertainment). The report begins with a reflection on what constitutes Artificial Intelligence, and concludes with recommendations concerning AI-related policy. These recommendations include accruing technical expertise about AI in government and devoting more resources—and removing impediments—to research on the fairness, security, privacy, and societal impacts of AI systems.

Contrary to the more fantastic predictions for AI in the popular press, the Study Panel found no cause for concern that AI is an imminent threat to humankind. No machines with self-sustaining long-term goals and intent have been developed, nor are they likely to be developed in the near future. Instead, increasingly useful applications of AI, with potentially profound positive impacts on our society and economy are likely to emerge between now and 2030, the period this report considers. At the same time, many of these developments will spur disruptions in how human labor is augmented or replaced by AI, creating new challenges for the economy and society more broadly. Application design and policy decisions made in the near term are likely to have long-lasting influences on the nature and directions of such developments, making it important for AI researchers, developers, social scientists, and policymakers to balance the imperative to innovate with mechanisms to ensure that AI's economic and social benefits are broadly shared across society. If society approaches these technologies primarily with fear and suspicion, missteps that slow AI's development or drive it underground will result, impeding important work on ensuring the safety and reliability of AI technologies. On the other hand, if society approaches AI with a more open mind, the technologies emerging from the field could profoundly transform society for the better in the coming decades.

While drawing on common research and technologies, AI systems are specialized to accomplish particular tasks. Each application requires years of focused research and a careful, unique construction.

Study Panel: Peter Stone, *Chair*, University of Texas at Austin, **Rodney Brooks**, Rethink Robotics, **Erik Brynjolfsson**, Massachussets Institute of Technology, **Ryan Calo**, University of Washington, **Oren Etzioni**, Allen Institute for AI, **Greg Hager**, Johns Hopkins University, **Julia Hirschberg**, Columbia University, **Shivaram Kalyanakrishnan**, Indian Institute of Technology Bombay, **Ece Kamar**, Microsoft Research, **Sarit Kraus**, Bar Ilan University. **Kevin Leyton-Brown**, University of British Columbia, **David Parkes**, Harvard University, **William Press**, University of Texas at Austin, **AnnaLee (Anno) Saxenian**, University of California, Berkeley, **Julie Shah**, Massachussets Institute of Technology, **Milind Tambe**, University of Southern California, **Astro Teller**, X

Standing Committee of the One Hundred Year Study of Artificial Intelligence: Barbara J. Grosz, *Chair*, Russ Altman, Eric Horvitz, Alan Mackworth, Tom Mitchell, Deidre Mulligan, Yoav Shoham

OVERVIEW

The frightening, futurist portrayals of Artificial Intelligence that dominate films and novels, and shape the popular imagination, are fictional. In reality, AI is already changing our daily lives, almost entirely in ways that improve human health, safety, and productivity. Unlike in the movies, there is no race of superhuman robots on the horizon or probably even possible. And while the potential to abuse AI technologies must be acknowledged and addressed, their greater potential is, among other things, to make driving safer, help children learn, and extend and enhance people's lives. In fact, beneficial AI applications in schools, homes, and hospitals are already growing at an accelerated pace. Major research universities devote departments to AI studies, and technology companies such as Apple, Facebook, Google, IBM, and Microsoft spend heavily to explore AI applications they regard as critical to their futures. Even Hollywood uses AI technologies to bring its dystopian AI fantasies to the screen.

Many have already grown accustomed to touching and talking to their smart phones. People's future relationships with machines will become ever more nuanced, fluid, and personalized.

Innovations relying on computer-based vision, speech recognition, and Natural Language Processing have driven these changes, as have concurrent scientific and technological advances in related fields. AI is also changing how people interact with technology. Many people have already grown accustomed to touching and talking to their smart phones. People's future relationships with machines will become ever more nuanced, fluid, and personalized as AI systems learn to adapt to individual personalities and goals. These AI applications will help monitor people's well-being, alert them to risks ahead, and deliver services when needed or wanted. For example, in a mere fifteen years in a typical North American city—the time frame and scope of this report—AI applications are likely to transform transportation toward self-driving vehicles with on-time pickup and delivery of people and packages. This alone will reconfigure the urban landscape, as traffic jams and parking challenges become obsolete.

This study's focus on a typical North American city is deliberate and meant to highlight specific changes affecting the everyday lives of the millions of people who inhabit them. The Study Panel further narrowed its inquiry to eight domains where AI is already having or is projected to have the greatest impact: transportation, healthcare, education, low-resource communities, public safety and security, employment and workplace, home/service robots, and entertainment.

Though drawing from a common source of research, AI technologies have influenced and will continue to influence these domains differently. Each domain faces varied AI-related challenges, including the difficulty of creating safe and reliable hardware for sensing and effecting (transportation and service robots), the difficulty of smoothly interacting with human experts (healthcare and education), the challenge of gaining public trust (low-resource communities and public safety and security), the challenge of overcoming fears of marginalizing humans (employment and workplace) and the risk of diminishing interpersonal interaction (entertainment). Some domains are primarily business sectors, such as transportation and healthcare, while others are more oriented to consumers, such as entertainment and home service robots. Some cut across sectors, such as employment/workplace and low-resource communities.

In each domain, even as AI continues to deliver important benefits, it also raises important ethical and social issues, including privacy concerns. Robots and other AI technologies have already begun to displace jobs in some sectors. As a society, we are now at a crucial juncture in determining how to deploy AI-based technologies in ways that promote, not hinder, democratic values such as freedom, equality, and transparency. For individuals, the quality of the lives we lead and how our contributions are valued are likely to shift gradually, but markedly. Over the next several years, AI research, systems development, and social and regulatory frameworks will shape how the benefits of AI are weighed against its costs and risks, and how broadly these benefits are spread.

An accurate and sophisticated picture of AI—one that competes with its popular portrayal—is hampered at the start by the difficulty of pinning down a precise definition of artificial intelligence. In the approaches the Study Panel considered, none suggest there is currently a "general purpose" AI. While drawing on common research and technologies, AI systems are specialized to accomplish particular tasks, and each application requires years of focused research and a careful, unique construction. As a result, progress is uneven within and among the eight domains.

A prime example is **Transportation**, where a few key technologies have catalyzed the widespread adoption of AI with astonishing speed. Autonomous transportation will soon be commonplace and, as most people's first experience with physically embodied AI systems, will strongly influence the public's perception of AI. As cars become better drivers than people, city-dwellers will own fewer cars, live further from work, and spend time differently, leading to an entirely new urban organization. In the typical North American city in 2030, physically embodied AI applications will not be limited to cars, but are likely to include trucks, flying vehicles, and personal robots. Improvements in safe and reliable hardware will spur innovation over the next fifteen years, as they will with **Home/Service Robots**, which have already entered people's houses, primarily in the form of vacuum cleaners. Better chips, low-cost 3D sensors, cloud-based machine learning, and advances in speech understanding will enhance future robots' services and their interactions with people. Special purpose robots will deliver packages, clean offices, and enhance security. But technical constraints and the high costs of reliable mechanical devices will continue to limit commercial opportunities to narrowly defined applications for the foreseeable future.

Society is now at a crucial juncture in determining how to deploy AI-based technologies in ways that promote rather than hinder democratic values such as freedom, equality, and transparency.

In **Healthcare**, there has been an immense forward leap in collecting useful data from personal monitoring devices and mobile apps, from electronic health records (EHR) in clinical settings and, to a lesser extent, from surgical robots designed to assist with medical procedures and service robots supporting hospital operations. AI-based applications could improve health outcomes and the quality of life for millions of people in the coming years. Though clinical applications have been slow to move from the computer science lab to the real-world, there are hopeful signs that the pace of innovation will improve. Advances in healthcare can be promoted via the development of incentives and mechanisms for sharing data and for removing overbearing policy, regulatory, and commercial obstacles. For many applications, AI systems will have to work closely with care providers and patients to gain their trust. Advances in how intelligent machines interact naturally with caregivers, patients, and patients' families are crucial.

Enabling more fluid interactions between people and promising AI technologies also remains a critical challenge in **Education**, which has seen considerable progress in the same period. Though quality education will always require active engagement by human teachers, AI promises to enhance education at all levels, especially by providing personalization at scale. Interactive machine tutors are now being matched to students for teaching science, math, language, and other disciplines. Natural Language Processing, machine learning, and crowdsourcing have boosted online learning and enabled teachers in higher education to multiply the size of their classrooms while addressing individual students' learning needs and styles. Over the next fifteen years in a typical North American city, the use of these technologies in the classroom and in the home is likely to expand significantly, provided they can be meaningfully integrated with face-to-face learning.

Beyond education, many opportunities exist for AI methods to assist **Low-resource Communities** by providing mitigations and solutions to a variety of social problems. Traditionally, funders have underinvested in AI research lacking commercial application. With targeted incentives and funding priorities, Longer term, AI may be thought of as a radically different mechanism for wealth creation in which everyone should be entitled to a portion of the world's AI-produced treasures.

AI technologies could help address the needs of low-resource communities, and budding efforts are promising. Using data mining and machine learning, for example, AI has been used to create predictive models to help government agencies address issues such as prevention of lead poisoning in at-risk children and distribution of food efficiently. These budding efforts suggest more could be done, particularly if agencies and organizations can engage and build trust with these communities. Gaining public trust is also a challenge for AI use by **Public Safety and Security** professionals. North American cities and federal agencies have already begun to deploy AI technologies in border administration and law enforcement. By 2030, they will rely heavily upon them, including improved cameras and drones for surveillance, algorithms to detect financial fraud, and predictive policing. The latter raises the specter of innocent people being unjustifiably monitored, and care must be taken to avoid systematizing human bias and to protect civil liberties. Well-deployed AI prediction tools have the potential to provide new kinds of transparency about data and inferences, and may be applied to detect, remove, or reduce human bias, rather than reinforcing it.

Social and political decisions are likewise at play in AI's influences on **Employment and Workplace** trends, such as the safety nets needed to protect people from structural changes in the economy. AI is poised to replace people in certain kinds of jobs, such as in the driving of taxis and trucks. However, in many realms, AI will likely replace tasks rather than jobs in the near term, and will also create new kinds of jobs. But the new jobs that will emerge are harder to imagine in advance than the existing jobs that will likely be lost. AI will also lower the cost of many goods and services, effectively making everyone better off. Longer term, AI may be thought of as a radically different mechanism for wealth creation in which everyone should be entitled to a portion of the world's AI-produced treasures. It is not too soon for social debate on how the economic fruits of AI technologies should be shared.

Entertainment has been transformed by social networks and other platforms for sharing and browsing blogs, videos, and photos, which rely on techniques actively developed in NLP, information retrieval, image processing, crowdsourcing, and machine learning. Some traditional sources of entertainment have also embraced AI to compose music, create stage performances, and even to generate 3D scenes from natural language text. The enthusiasm with which people have already responded to AI-driven entertainment has been surprising. As with many aspects of AI, there is ongoing debate about the extent to which the technology replaces or enhances sociability. AI will increasingly enable entertainment that is more interactive, personalized, and engaging. Research should be directed toward understanding how to leverage these attributes for individuals' and society's benefit.

What's next for AI research?

The research that fuels the AI revolution has also seen rapid changes. Foremost among them is the maturation of machine learning, stimulated in part by the rise of the digital economy, which both provides and leverages large amounts of data. Other factors include the rise of cloud computing resources and consumer demand for widespread access to services such as speech recognition and navigation support.

Machine learning has been propelled dramatically forward by impressive empirical successes of artificial neural networks, which can now be trained with huge data sets and large-scale computing. This approach has been come to be known as "deep learning." The leap in the performance of information processing algorithms has been accompanied by significant progress in hardware technology for basic operations such as sensing, perception, and object recognition. New platforms and markets for data-driven products, and the economic incentives to find new products and markets, have also stimulated research advances. Now, as it becomes a central force in society, the field of AI is shifting toward building intelligent systems that can collaborate effectively with people, and that are more generally *human-aware*, including creative ways to develop interactive and scalable ways for people to teach robots. These trends drive the currently "hot" areas of AI research into both fundamental methods and application areas:

Large-scale machine learning concerns the design of learning algorithms, as well as scaling existing algorithms, to work with extremely large data sets.

Deep learning, a class of learning procedures, has facilitated object recognition in images, video labeling, and activity recognition, and is making significant inroads into other areas of perception, such as audio, speech, and natural language processing.

Reinforcement learning is a framework that shifts the focus of machine learning from pattern recognition to experience-driven sequential decision-making. It promises to carry AI applications forward toward taking actions in the real world. While largely confined to academia over the past several decades, it is now seeing some practical, real-world successes.

Robotics is currently concerned with how to train a robot to interact with the world around it in generalizable and predictable ways, how to facilitate manipulation of objects in interactive environments, and how to interact with people. Advances in robotics will rely on commensurate advances to improve the reliability and generality of computer vision and other forms of machine perception.

Computer vision is currently the most prominent form of machine perception. It has been the sub-area of AI most transformed by the rise of deep learning. For the first time, computers are able to perform some vision tasks better than people. Much current research is focused on automatic image and video captioning.

Natural Language Processing, often coupled with automatic speech recognition, is quickly becoming a commodity for widely spoken languages with large data sets. Research is now shifting to develop refined and capable systems that are able to interact with people through dialog, not just react to stylized requests. Great strides have also been made in machine translation among different languages, with more real-time person-to-person exchanges on the near horizon.

Collaborative systems research investigates models and algorithms to help develop autonomous systems that can work collaboratively with other systems and with humans.

Crowdsourcing and human computation research investigates methods to augment computer systems by making automated calls to human expertise to solve problems that computers alone cannot solve well.

Algorithmic game theory and computational social choice draw attention to the economic and social computing dimensions of AI, such as how systems can handle potentially misaligned incentives, including self-interested human participants or firms and the automated AI-based agents representing them.

Internet of Things (IoT) research is devoted to the idea that a wide array of devices, including appliances, vehicles, buildings, and cameras, can be interconnected to collect and share their abundant sensory information to use for intelligent purposes.

Neuromorphic computing is a set of technologies that seek to mimic biological neural networks to improve the hardware efficiency and robustness of computing systems, often replacing an older emphasis on separate modules for input/output, instruction-processing, and memory.

The field of AI is shifting toward building intelligent systems that can collaborate effectively with people, including creative ways to develop interactive and scalable ways for people to teach robots.

AI policy, now and in the future

The measure of success for AI applications is the value they create for human lives. In that light, they should be designed to enable people to understand AI systems successfully, participate in their use, and build their trust. Public policies should help ease society's adaptation to AI applications, extend their benefits, and mitigate their inevitable errors and failures. Debate about how AI is deployed, including concerns about how privacy is protected and AI's benefits fairly shared, should be encouraged. Given the speed with which AI technologies are being realized, and concomitant concerns about their implications, the Study Panel recommends that all layers of government acquire technical expertise in AI. Further, research on the fairness, security, privacy, and societal implications of AI systems should be encouraged by removing impediments and increasing private and public spending to support it.

Misunderstandings about what AI is and is not could fuel opposition to technologies with the potential to benefit everyone. Poorly informed regulation that stifles innovation would be a tragic mistake.

Currently in the United States, at least sixteen separate agencies govern sectors of the economy related to AI technologies. Rapid advances in AI research and, especially, its applications require experts in these sectors to develop new concepts and metaphors for law and policy. Who is responsible when a self-driven car crashes or an intelligent medical device fails? How can AI applications be prevented from promulgating racial discrimination or financial cheating? Who should reap the gains of efficiencies enabled by AI technologies and what protections should be afforded to people whose skills are rendered obsolete? As people integrate AI more broadly and deeply into industrial processes and consumer products, best practices need to be spread, and regulatory regimes adapted.

While the Study Panel does not consider it likely that near-term AI systems will autonomously *choose* to inflict harm on people, it will be possible for people to *use* AI-based systems for harmful as well as helpful purposes. And though AI algorithms may be capable of making less biased decisions than a typical person, it remains a deep technical challenge to ensure that the data that inform AI-based decisions can be kept free from biases that could lead to discrimination based on race, sexual orientation, or other factors.

Faced with the profound changes that AI technologies can produce, pressure for "more" and "tougher" regulation is probably inevitable. Misunderstandings about what AI is and is not could fuel opposition to technologies with the potential to benefit everyone. Inappropriate regulatory activity would be a tragic mistake. Poorly informed regulation that stifles innovation, or relocates it to other jurisdictions, would be counterproductive.[2]

Fortunately, principles that guide successful regulation of current digital technologies provide a starting point. In privacy regulation, broad legal mandates coupled with tough transparency requirements and meaningful enforcement—rather than strict controls—encourage companies to develop processes and professional staff to enforce privacy controls, engage with outside stakeholders, and adapt their practices to technological advances. This in turn supports the development of professional trade associations and standards committees that spread best practices. In AI, too, regulators can strengthen a virtuous cycle of activity involving internal and external accountability, transparency, and professionalization, rather than narrow compliance.

A vigorous and informed debate about how to best steer AI in ways that enrich our lives and our society, while encouraging creativity in the field, is an urgent and vital need. AI technologies could widen existing inequalities of opportunity if access to them—along with the high-powered computation and large-scale data that fuel many of them—is unfairly distributed across society. These technologies will improve the abilities and efficiency of people who have access to them. Policies should be evaluated as to whether they foster democratic values and equitable sharing of AI's benefits, or concentrate power and benefits in the hands of a fortunate few.

As this report documents, significant AI-related advances have already had an impact on North American cities over the past fifteen years, and even more substantial developments will occur over the next fifteen. Recent advances are largely due to the growth and analysis of large data sets enabled by the internet, advances in sensory technologies and, more recently, applications of "deep learning." In the coming years, as the public encounters new AI applications in domains such as transportation and healthcare, they must be introduced in ways that build trust and understanding, and respect human and civil rights. While encouraging innovation, policies and processes should address ethical, privacy, and security implications, and should work to ensure that the benefits of AI technologies will be spread broadly and fairly. Doing so will be critical if Artificial Intelligence research and its applications are to exert a positive influence on North American urban life in 2030 and beyond.

2 Kate Crawford, "Artificial Intelligence's White Guy Problem," *The New York Times*, June 25, 2016, accessed August 1, 2016, http://www.nytimes.com/2016/06/26/opinion/sunday/artificial-intelligences-white-guy-problem.html.

An accurate and sophisticated picture of AI—one that competes with its popular portrayal—is hampered by the difficulty of pinning down a precise definition of artificial intelligence.

SECTION I: WHAT IS ARTIFICIAL INTELLIGENCE?

This section describes how researchers and practitioners define "Artificial Intelligence," and the areas of AI research and application that are currently thriving. It proffers definitions of what AI is and is not, and describes some of the currently "hot" areas of AI Research. This section lays the groundwork for Section II, which elaborates on AI's impacts and future in eight domains and Section III, which describes issues related to AI design and public policy and makes recommendations for encouraging AI innovation while protecting democratic values.

DEFINING AI

Curiously, the lack of a precise, universally accepted definition of AI probably has helped the field to grow, blossom, and advance at an ever-accelerating pace. Practitioners, researchers, and developers of AI are instead guided by a rough sense of direction and an imperative to "get on with it." Still, a definition remains important and Nils J. Nilsson has provided a useful one:

"Artificial intelligence is that activity devoted to making machines intelligent, and intelligence is that quality that enables an entity to function appropriately and with foresight in its environment."[3]

From this perspective, characterizing AI depends on the credit one is willing to give synthesized software and hardware for functioning "appropriately" and with "foresight." A simple electronic calculator performs calculations much faster than the human brain, and almost never makes a mistake.[4] Is a calculator intelligent? Like Nilsson, the Study Panel takes a broad view that intelligence lies on a multi-dimensional spectrum. According to this view, the difference between an arithmetic calculator and a human brain is not one of kind, but of scale, speed, degree of autonomy, and generality. The same factors can be used to evaluate every other instance of intelligence—speech recognition software, animal brains, cruise-control systems in cars, Go-playing programs, thermostats—and to place them at some appropriate location in the spectrum.

Although our broad interpretation places the calculator within the intelligence spectrum, such simple devices bear little resemblance to today's AI. The frontier of AI has moved far ahead and functions of the calculator are only one among the millions that today's smartphones can perform. AI developers now work on improving, generalizing, and scaling up the intelligence currently found on smartphones.

In fact, the field of AI is a continual endeavor to push forward the frontier of machine intelligence. Ironically, AI suffers the perennial fate of losing claim to its acquisitions, which eventually and inevitably get pulled inside the frontier, a repeating pattern known as the "AI effect" or the "odd paradox"—AI brings a new technology into the common fold, people become accustomed to this technology, it stops being considered AI, and newer technology emerges.[5] The same pattern will continue in the future. AI does not "deliver" a life-changing product as a bolt from the blue. Rather, AI technologies continue to get better in a continual, incremental way.

3 Nils J. Nilsson, *The Quest for Artificial Intelligence: A History of Ideas and Achievements* (Cambridge, UK: Cambridge University Press, 2010).
4 Wikimedia Images, accessed August 1, 2016, https://upload.wikimedia.org/wikipedia/commons/b/b6/SHARP_ELSIMATE_EL-W221.jpg.
5 Pamela McCorduck, *Machines Who Think: A Personal Inquiry into the History and Prospects of Artificial Intelligence*, 2nd ed. (Natick, MA: A. K. Peters, Ltd., 2004; San Francisco: W. H. Freeman, 1979), Citations are to the Peters edition.

The human measure

Notably, the characterization of intelligence as a spectrum grants no special status to the human brain. But to date human intelligence has no match in the biological and artificial worlds for sheer versatility, with the abilities "to reason, achieve goals, understand and generate language, perceive and respond to sensory inputs, prove mathematical theorems, play challenging games, synthesize and summarize information, create art and music, and even write histories."[6]

This makes human intelligence a natural choice for benchmarking the progress of AI. It may even be proposed, as a rule of thumb, that any activity computers are able to perform and people once performed should be counted as an instance of intelligence. But matching any human ability is only a sufficient condition, not a necessary one. There are already many systems that exceed human intelligence, at least in speed, such as scheduling the daily arrivals and departures of thousands of flights in an airport.

AI's long quest—and eventual success—to beat human players at the game of chess offered a high-profile instance for comparing human to machine intelligence. Chess has fascinated people for centuries. When the possibility of building computers became imminent, Alan Turing, who many consider the father of computer science, "mentioned the idea of computers showing intelligence with chess as a paradigm."[7] Without access to powerful computers, "Turing played a game in which he simulated the computer, taking about half an hour per move."

But it was only after a long line of improvements in the sixties and seventies—contributed by groups at Carnegie Mellon, Stanford, MIT, The Institute for Theoretical and Experimental Physics at Moscow, and Northwestern University—that chess-playing programs started gaining proficiency. The final push came through a long-running project at IBM, which culminated with the Deep Blue program beating Garry Kasparov, then the world chess champion, by a score of 3.5-2.5 in 1997. Curiously, no sooner had AI caught up with its elusive target than Deep Blue was portrayed as a collection of "brute force methods" that wasn't "real intelligence."[8] In fact, IBM's subsequent publication about Deep Blue, which gives extensive details about its search and evaluation procedures, doesn't mention the word "intelligent" even once![9] Was Deep Blue intelligent or not? Once again, the frontier had moved.

An operational definition

AI can also be defined by what AI researchers do. This report views AI primarily as a branch of computer science that studies the properties of intelligence by synthesizing intelligence.[10] Though the advent of AI has depended on the rapid progress of hardware computing resources, the focus here on software reflects a trend in the AI community. More recently, though, progress in building hardware tailored for neural-network-based computing[11] has created a

Intelligence lies on a multi-dimensional spectrum. According to this view, the difference between an arithmetic calculator and a human brain is not one of kind, but of scale, speed, degree of autonomy, and generality.

6 Nilsson, *The Question for Artificial Intelligence*.
7 Nilsson, *The Question for Artificial Intelligence*, 89.
8 McCorduck, *Machines Who Think*, 433.
9 Murray Campbell, A. Joseph Hoane Jr., and Feng-hsiung Hsu, "Deep Blue," *Artificial Intelligence* 134, nos. 1 and 2 (2002): 57–83.
10 Herbert A. Simon, "Artificial Intelligence: An Empirical Science," *Artificial Intelligence* 77, no. 2 (1995):95–127.
11 Paul Merolla John V. Arthur, Rodrigo Alvarez-Icaza, Andrew S. Cassidy, Jun Sawada, Filipp Akopyan, Bryan L. Jackson, Nabil Imam, Chen Guo, Yutaka Nakamura, Bernard Brezzo, Ivan Vo, Steven K. Esser, Rathinakumar Appuswamy, Brian Taba, Arnon Amir, Myron D. Flickner, William P. Risk, Rajit Manohar, and Dharmendra S. Modha, "A Million Spiking-Neuron Integrated Circuit with a Scalable Communication Network and Interface," accessed August 1, 2016, http://paulmerolla.com/merolla_main_som.pdf.

tighter coupling between hardware and software in advancing AI.

"Intelligence" remains a complex phenomenon whose varied aspects have attracted the attention of several different fields of study, including psychology, economics, neuroscience, biology, engineering, statistics, and linguistics. Naturally, the field of AI has benefited from the progress made by all of these allied fields. For example, the artificial neural network, which has been at the heart of several AI-based solutions[12 13] was originally inspired by thoughts about the flow of information in biological neurons.[14]

Human intelligence has no match in the biological and artificial worlds for sheer versatility, with the abilities "to reason, achieve goals, understand and generate language... create art and music, and even write histories."

AI RESEARCH TRENDS

Until the turn of the millennium, AI's appeal lay largely in its promise to deliver, but in the last fifteen years, much of that promise has been redeemed.[15] AI technologies already pervade our lives. As they becomes a central force in society, the field is shifting from simply building systems that are intelligent to building intelligent systems that are human-aware and trustworthy.

Several factors have fueled the AI revolution. Foremost among them is the maturing of machine learning, supported in part by cloud computing resources and wide-spread, web-based data gathering. Machine learning has been propelled dramatically forward by "deep learning," a form of adaptive artificial neural networks trained using a method called backpropagation.[16] This leap in the performance of information processing algorithms has been accompanied by significant progress in hardware technology for basic operations such as sensing, perception, and object recognition. New platforms and markets for data-driven products, and the economic incentives to find new products and markets, have also contributed to the advent of AI-driven technology.

All these trends drive the "hot" areas of research described below. This compilation is meant simply to reflect the areas that, by one metric or another, currently receive greater attention than others. They are not necessarily more important or valuable than other ones. Indeed, some of the currently "hot" areas were less popular in past years, and it is likely that other areas will similarly re-emerge in the future.

Large-scale machine learning

Many of the basic problems in machine learning (such as supervised and unsupervised learning) are well-understood. A major focus of current efforts is to scale existing algorithms to work with extremely large data sets. For example, whereas traditional methods could afford to make several passes over the data set, modern ones are designed to make only a single pass; in some cases, only sublinear methods (those that only look at a fraction of the data) can be admitted.

Deep learning

The ability to successfully train convolutional neural networks has most benefited the field of computer vision, with applications such as object recognition, video labeling, activity recognition, and several variants thereof. Deep learning is also making significant inroads into other areas of perception, such as audio, speech, and natural language processing.

Reinforcement learning

Whereas traditional machine learning has mostly focused on pattern mining, reinforcement learning shifts the focus to decision making, and is a technology that will help AI to advance more deeply into the realm of learning about and executing actions in the real world. It has existed for several decades as a framework for experience-driven sequential decision-making, but the methods have not found great success in practice, mainly owing to issues of representation and scaling. However, the advent of deep learning has provided reinforcement learning with a "shot in the arm." The recent success of AlphaGo, a computer program developed by Google Deepmind that beat the human Go champion in a five-game match, was due in large part to reinforcement learning. AlphaGo was trained by initializing an automated agent with a human expert database, but was subsequently refined by playing a large number of games against itself and applying reinforcement learning.

Robotics

Robotic navigation, at least in static environments, is largely solved. Current efforts consider how to train a robot to interact with the world around it in generalizable and predictable ways. A natural requirement that arises in interactive environments is *manipulation*, another topic of current interest. The deep learning revolution is only beginning to influence robotics, in large part because it is far more difficult to acquire the large labeled data sets that have driven other learning-based areas of AI. Reinforcement learning (see above), which obviates the requirement of labeled data, may help bridge this gap but requires systems to be able to safely explore a policy space without committing errors that harm the system itself or others. Advances in reliable machine perception, including computer vision, force, and tactile perception, much of which will be driven by machine learning, will continue to be key enablers to advancing the capabilities of robotics.

Computer vision

Computer vision is currently the most prominent form of machine perception. It has been the sub-area of AI most transformed by the rise of deep learning. Until just a few years ago, support vector machines were the method of choice for most visual classification tasks. But the confluence of large-scale computing, especially on GPUs, the availability of large datasets, especially via the internet, and refinements of neural network algorithms has led to dramatic improvements in performance on benchmark tasks (e.g., classification on ImageNet[17]). For the first time, computers are able to perform some (narrowly defined) visual classification tasks better than people. Much current research is focused on automatic image and video captioning.

Natural Language Processing

Often coupled with automatic speech recognition, Natural Language Processing is another very active area of machine perception. It is quickly becoming a commodity for mainstream languages with large data sets. Google announced that 20% of current mobile queries are done by voice,[18] and recent demonstrations have proven the possibility of real-time translation. Research is now shifting towards developing refined and capable systems that are able to interact with people through dialog, not just react to stylized requests.

AI technologies already pervade our lives. As they become a central force in society, the field is shifting from simply building systems that are intelligent to building intelligent systems that are human-aware and trustworthy.

---

12 Gerald Tesauro, "Practical Issues in Temporal Difference Learning," *Machine Learning*, no. 8 (1992): 257–77.

13 David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis, "Mastering the game of Go with deep neural networks and tree search," *Nature* 529 (2016): 484–489.

14 W. McCulloch and W. Pitts, W., "A logical calculus of the ideas immanent in nervous activity," *Bulletin of Mathematical Biophysics*, 5 (1943): 115–133.

15 Appendix I offers a short history of AI, including a description of some of the traditionally core areas of research, which have shifted over the past six decades.

16 Backpropogation is an abbreviation for "backward propagation of errors," a common method of training artificial neural networks used in conjunction with an optimization method such as gradient descent. The method calculates the gradient of a loss function with respect to all the weights in the network.

17 ImageNet, Stanford Vision Lab, Stanford University, Princeton University, 2016, accessed August 1, 2016, www.image-net.org/.

18 Greg Sterling, "Google says 20% of mobile queries are voice searches," *Search Engine Land*, May 18, 2016, accessed August 1, 2016, http://searchengineland.com/google-reveals-20-percent-queries-voice-queries-249917.

Natural Language Processing is a very active area of machine perception. Research is now shifting towards developing systems that are able to interact with people through dialog, not just react to stylized requests.

Collaborative systems
Research on collaborative systems investigates models and algorithms to help develop autonomous systems that can work collaboratively with other systems and with humans. This research relies on developing formal models of collaboration, and studies the capabilities needed for systems to become effective partners. There is growing interest in applications that can utilize the complementary strengths of humans and machines—for humans to help AI systems to overcome their limitations, and for agents to augment human abilities and activities.

Crowdsourcing and human computation
Since human abilities are superior to automated methods for accomplishing many tasks, research on crowdsourcing and human computation investigates methods to augment computer systems by utilizing human intelligence to solve problems that computers alone cannot solve well. Introduced only about fifteen years ago, this research now has an established presence in AI. The best-known example of crowdsourcing is Wikipedia, a knowledge repository that is maintained and updated by netizens and that far exceeds traditionally-compiled information sources, such as encyclopedias and dictionaries, in scale and depth. Crowdsourcing focuses on devising innovative ways to harness human intelligence. Citizen science platforms energize volunteers to solve scientific problems, while paid crowdsourcing platforms such as Amazon Mechanical Turk provide automated access to human intelligence on demand. Work in this area has facilitated advances in other subfields of AI, including computer vision and NLP, by enabling large amounts of labeled training data and/or human interaction data to be collected in a short amount of time. Current research efforts explore ideal divisions of tasks between humans and machines based on their differing capabilities and costs.

Algorithmic game theory and computational social choice
New attention is being drawn to the economic and social computing dimensions of AI, including incentive structures. Distributed AI and multi-agent systems have been studied since the early 1980s, gained prominence starting in the late 1990s, and were accelerated by the internet. A natural requirement is that systems handle potentially misaligned incentives, including self-interested human participants or firms, as well as automated AI-based agents representing them. Topics receiving attention include computational mechanism design (an economic theory of incentive design, seeking incentive-compatible systems where inputs are truthfully reported), computational social choice (a theory for how to aggregate rank orders on alternatives), incentive aligned information elicitation (prediction markets, scoring rules, peer prediction) and algorithmic game theory (the equilibria of markets, network games, and parlor games such as Poker—a game where significant advances have been made in recent years through abstraction techniques and no-regret learning).

Internet of Things (IoT)
A growing body of research is devoted to the idea that a wide array of devices can be interconnected to collect and share their sensory information. Such devices can include appliances, vehicles, buildings, cameras, and other things. While it's a matter of technology and wireless networking to connect the devices, AI can process and use the resulting huge amounts of data for intelligent and useful purposes. Currently, these devices use a bewildering array of incompatible communication protocols. AI could help tame this Tower of Babel.

Neuromorphic Computing
Traditional computers implement the von Neumann model of computing, which separates the modules for input/output, instruction-processing, and memory. With the success of deep neural networks on a wide array of tasks, manufacturers are actively pursuing alternative models of computing—especially those that are inspired by what is known about biological neural networks—with the aim of improving the hardware efficiency and robustness of computing systems. At the moment, such "neuromorphic" computers have not yet clearly demonstrated big wins, and are just beginning to become commercially viable. But it is possible that they will become commonplace (even if only as additions to their von Neumann cousins) in the near future. Deep neural networks have already created a splash in the application landscape. A larger wave may hit when these networks can be trained and executed on dedicated neuromorphic hardware, as opposed to simulated on standard von Neumann architectures, as they are today.

Overall trends and the future of AI research

The resounding success of the data-driven paradigm has displaced the traditional paradigms of AI. Procedures such as theorem proving and logic-based knowledge representation and reasoning are receiving reduced attention, in part because of the ongoing challenge of connecting with real-world groundings. Planning, which was a mainstay of AI research in the seventies and eighties, has also received less attention of late due in part to its strong reliance on modeling assumptions that are hard to satisfy in realistic applications. Model-based approaches—such as physics-based approaches to vision and traditional control and mapping in robotics—have by and large given way to data-driven approaches that close the loop with sensing the results of actions in the task at hand. Bayesian reasoning and graphical models, which were very popular even quite recently, also appear to be going out of favor, having been drowned by the deluge of data and the remarkable success of deep learning.

Over the next fifteen years, the Study Panel expects an increasing focus on developing systems that are human-aware, meaning that they specifically model, and are specifically designed for, the characteristics of the people with whom they are meant to interact. There is a lot of interest in trying to find new, creative ways to develop interactive and scalable ways to teach robots. Also, IoT-type systems—devices and the cloud—are becoming increasingly popular, as is thinking about social and economic dimensions of AI. In the coming years, new perception/object recognition capabilities and robotic platforms that are human-safe will grow, as will data-driven products and their markets.

The Study Panel also expects a reemergence of some of the traditional forms of AI as practitioners come to realize the inevitable limitations of purely end-to-end deep learning approaches. We encourage young researchers not to reinvent the wheel, but rather to maintain an awareness of the significant progress in many areas of AI during the first fifty years of the field, and in related fields such as control theory, cognitive science, and psychology.

A growing body of research is devoted to the idea that a wide array of devices can be interconnected to collect and share their sensory information. Such devices can include appliances, vehicles, buildings, cameras, and other things.

SECTION II: AI BY DOMAIN

Though different instances of AI research and practice share common technologies, such as machine learning, they also vary considerably in different sectors of the economy and society. We call these sectors "domains," and in this section describe the different states of AI research and implementation, as well as impacts and distinct challenges, in eight of them: transportation; home/service robotics; healthcare; education; low-resource communities; public safety and security; employment and workplace; and entertainment. Based on these analyses, we also predict trends in a typical North American city over the next fifteen years. Contrary to AI's typical depiction in popular culture, we seek to offer a balanced overview of the ways in which AI is already beginning to transform everyday life, and how those transformations are likely to grow by the year 2030.

Autonomous transportation will soon be commonplace and, as most people's first experience with physically embodied AI systems, will strongly influence the public's perception of AI.

TRANSPORTATION

Transportation is likely to be one of the first domains in which the general public will be asked to trust the reliability and safety of an AI system for a critical task. Autonomous transportation will soon be commonplace and, as most people's first experience with physically embodied AI systems, will strongly influence the public's perception of AI. Once the physical hardware is made sufficiently safe and robust, its introduction to daily life may happen so suddenly as to surprise the public, which will require time to adjust. As cars will become better drivers than people, city-dwellers will own fewer cars, live further from work, and spend time differently, leading to an entirely new urban organization. Further, in the typical North American city in 2030, changes won't be limited to cars and trucks, but are likely to include flying vehicles and personal robots, and will raise social, ethical and policy issues.

A few key technologies have already catalyzed the widespread adoption of AI in transportation. Compared to 2000, the scale and diversity of data about personal and population-level transportation available today—enabled by the adoption of smartphones and decreased costs and improved accuracies for variety of sensors—is astounding. Without the availability of this data and connectivity, applications such as real-time sensing and prediction of traffic, route calculations, peer-to-peer ridesharing and self-driving cars would not be possible.

Smarter cars

GPS was introduced to personal vehicles in 2001 with in-car navigation devices and has since become a fundamental part of the transportation infrastructure.[19] GPS assists drivers while providing large-scale information to technology companies and cities about transportation patterns. Widespread adoption of smartphones with GPS technology further increased connectivity and the amount of location data shared by individuals.

Current vehicles are also equipped with a wide range of sensing capabilities. An average automobile in the US is predicted to have seventy sensors including gyroscopes, accelerometers, ambient light sensors, and moisture sensors.[20] Sensors are not new to vehicles. Automobiles built before 2000 had sensors for the internal state of the vehicle such as its speed, acceleration, and wheel position.[21]

19 Mark Sullivan, "A brief history of GPS," *PCWorld*, August 9, 2012, accessed August 1, 2016, http://www.pcworld.com/article/2000276/a-brief-history-of-gps.html.
20 William J. Fleming, "New Automotive Sensors - A Review," *IEEE Sensors Journal* 8, *no 11*, (2008): 1900-1921.
21 Jean Jacques Meneu, ed., "Automotive Sensors: Now and in the Future," *Arrow*, September 24, 2015, accessed August 1, 2016, https://www.arrow.com/en/research-and-events/articles/automotive-sensors-now-and-in-the-future.

They already had a number of functionalities that combined real-time sensing with perception and decision-making such as Anti-lock Braking Systems (ABS), airbag control, Traction Control Systems (TCS), and Electronic Stability Control (ESC).[22] Automated capabilities have been introduced into commercial cars gradually since 2003 as summarized in the following table.

Context	Automated Functionality	Release Date
Parking	Intelligent Parking Assist System	Since 2003[23]
Parking	Summon	Since 2016[24]
Arterial & Highway	Lane departure system	Since 2004 in North America[25]
Arterial & Highway	Adaptive cruise control	Since 2005 in North America[26]
Highway	Blind spot monitoring	2007[27]
Highway	Lane changing	2015[28]

These functionalities assist drivers or completely take over well-defined activities for increased safety and comfort. Current cars can park themselves, perform adaptive cruise control on highways, steer themselves during stop-and-go traffic, and alert drivers about objects in blind spots during lane changes. Vision and radar technology were leveraged to develop pre-collision systems that let cars autonomously brake when risk of a collision is detected. Deep learning also has been applied to improve automobiles' capacity to detect objects in the environment and recognize sound.[29]

As cars will become better drivers than people, city-dwellers will own fewer cars, live further from work, and spend time differently, leading to an entirely new urban organization.

Self-driving vehicles

Since the 1930s, science fiction writers dreamed of a future with self-driving cars, and building them has been a challenge for the AI community since the 1960s. By the 2000s, the dream of autonomous vehicles became a reality in the sea and sky, and even on Mars, but self-driving cars existed only as research prototypes in labs. Driving in a city was considered to be a problem too complex for automation due to factors like pedestrians, heavy traffic, and the many unexpected events that can happen outside of the car's control. Although the technological components required to

22 Carl Liersch, "Vehicle Technology Timeline: From Automated to Driverless," Robert Bosch (Australia) Pty. Ltd., 2014, accessed August 1, 2016, http://dpti.sa.gov.au/__data/assets/pdf_file/0009/246807/Carl_Liersch_Presentation.pdf.
23 "Intelligent Parking Assist System," *Wikipedia*, last modified July 26, 2016, accessed August 1, 2016, https://en.wikipedia.org/wiki/Intelligent_Parking_Assist_System.
24 The Tesla Motors Team, "Summon Your Tesla from Your Phone," Tesla, January 10, 2016, accessed August 1, 2016, https://www.teslamotors.com/blog/summon-your-tesla-your-phone.
25 Lane departure warning system," *Wikipedia*, last modified July 24, 2016, accessed August 1, 2016, https://en.wikipedia.org/wiki/Lane_departure_warning_system.
26 "Autonomous cruise control system," *Wikipedia*, last modified July 30, 2016, accessed August 1, 2016, https://en.wikipedia.org/wiki/Autonomous_cruise_control_system.
27 "Blind spot monitor," *Wikipedia*, last modified April 20, 2016, accessed August 1, 2016, https://en.wikipedia.org/wiki/Blind_spot_monitor.
28 Dana Hull, "Tesla Starts Rolling Out Autopilot Features," *Boomberg Technology*, October 14, 2015, accessed August 1, 2016, http://www.bloomberg.com/news/articles/2015-10-14/tesla-software-upgrade-adds-automated-lane-changing-to-model-s.
29 Aaron Tilley, "New Qualcomm Chip Brings Deep Learning To Cars," *Forbes*, January 5, 2016, accessed August 1, 2016, http://www.forbes.com/sites/aarontilley/2016/01/05/along-with-nvidia-new-qualcomm-chip-brings-deep-learning-to-cars/#4cb4e9235357.