

Part I

Six Technologies That Will Reshape Business in the Next Decade

1 Artificial Intelligence

Artificial intelligence (AI) is all around us. It underpins speech recognition, natural language processing (NLP), and machine vision. AI is behind the sophisticated spam filters that keep your inbox (mostly) free of junk mail. It flags unusual usage patterns on your credit card and assesses your credit score. It suggests tags for people in photos you post to Facebook. It offers intelligent suggestions when you use the search bar. AI is all around us, yet we have only scratched the surface of AI's future potential.

Like a femme fatale in a classic movie, AI seems simultaneously sexy and scary. While artificial intelligence may help us to unlock cures for diseases, discover new wonder materials, and predict the future, some people worry that it threatens human life as we know it, either economically or existentially.

What Is It and Why Is It Important?

Artificial intelligence is the umbrella term that describes the effort to mimic human skills and replicate human intelligence with machines. Various approaches have been used to build artificial intelligence over the years. In the 1980s and 1990s, “knowledge systems” were all the rage. Today, most modern AI uses a technique known as machine learning.

Machines learn from examples in the form of training data. Most machine learning systems are built with artificial neural networks (ANN), also known more simply as neural networks.

Electricity, Digital Computers, and Artificial Intelligence

About 120 years ago, electrification changed the world. Electricity refrigerates our food, washes our clothes, lights our homes, and powers our factories. Electricity transformed every industrial sector and now powers modern life.

Roughly 60 years ago, the first digital computers were built. Initially limited in their capability, computers evolved into powerful machines that brought us word processing, spreadsheets, the internet, video games, social media, streaming media, and smartphones. Like electricity before them, computers transformed business and changed our lives.

Artificial intelligence will have an impact as profound as both electrification and the digital computer. AI luminary Andrew Ng is the former chief scientist at Baidu, former lead of the Google Brain project, and now runs Landing.ai, a company that solves manufacturing problems using AI. In 2017, Ng observed, “Just as electricity transformed almost everything 100 years ago, today I actually have a hard time thinking of an industry that I don’t think AI will transform in the next several years.”

Artificial intelligence is a huge deal, worthy of the first and longest chapter of this book. Sundar Pichai, CEO at Google, once said, “AI is one of the most important things humanity is working on. It is more profound than, I dunno, electricity or fire.” While hyperbolic, this statement from the head of one of the world’s most powerful tech companies should make us all sit up and listen.

AI’s hype is largely justified. Just as digital technology was a vital component of any successful business strategy in the 1990s and 2000s, so AI must be central to strategic plans of the 2020s.

The Next Era of Computing: Traditional Digital versus Artificial Intelligence

Artificial intelligence and traditional digital computers are complementary. Importantly, AI solves problems that are costly or impossible for traditional computers to solve. The two technologies will coexist and work side-by-side, each solving a different set of problems. For a language translation app on a smartphone, traditional computing presents an attractive interface, while AI handles voice recognition and language translation functions.

AI solves problems using a radically different approach from traditional computers. Traditional computers are programmed to solve problems. Programs apply a set of rules to data and compute results. Said another way, they take rules and data as input, and output results. AI solves problems without using preprogrammed rules. Machine learning, a popular form of AI, takes data and results as its input and infers rules as its output. Through a complex training process, AI finds patterns and associations between data and results and divines its own rules for how they connect. This trait lets us solve entirely new problems with AI. It's why AI can seem so magical: It solves problems we don't know how to solve ourselves.

To find complex associations in data and build rules, AIs must be trained with thousands or sometimes millions of examples. Today's AIs are not like human brains. While some of the organizing principles are the same, human children learning about the world don't have to see a million automobiles before they can recognize one and adorably yell the word "car."

A 1950s Concept and 1980s Algorithms Meet Modern Computing Horsepower

The term "artificial intelligence" was first coined in the 1950s. The core algorithms behind today's AIs were first proposed in the 1970s and popularized in the mid-1980s. But it was 2012 before the recent

crop of AI breakthroughs began to appear. Why the quarter-century delay? Older computers lacked the performance to run AI applications. High-end graphics processors, GPUs from companies like Nvidia, eventually provided the computing horsepower needed. Their parallel number-crunching architectures, designed to create realistic video games, turn out to be pretty good for training an AI. As well as fast computers, AIs need training data to learn from. As digital storage costs fell and broadband speeds increased, data flooded in from many sources: billions of industrial sensors, millions of smart cameras, billions of people sharing trillions of photos and billions of videos, and trillions of clicks on social media. Users upload 500 hours of video to YouTube every minute and more than 1.2 billion photos to Google Photos every day (Source: Wikipedia).

With cheap, powerful computing, an avalanche of training data, and a small army of AI-savvy researchers and developers, artificial intelligence is now poised to solve myriad problems and create many exciting new capabilities.

What Can You Do with It?

Artificial intelligence can solve a wide variety of problems. Considering all the possible applications of AI can be overwhelming. I've found it helpful to cluster AI applications into eight broad categories:

1. Machine vision
2. Natural language processing (NLP) and voice platforms
3. Exploration and discovery
4. Better-informed decision-making
5. Predicting the future
6. Seeing the world through a new lens with super sensors
7. Solving complex problems by learning from experience
8. Creating and co-creating content

In all of these eight application categories, AI is used to find patterns and associations in data and to make statistical predictions. Each application uses this fundamental feature in different ways. Apply the associative capabilities of machine learning to images and you get machine vision; apply it to historical data and you get predictions; apply it to cursive text and you get handwriting recognition. In voice platforms, AIs trained on human speech determine what words you are saying. AIs trained on historical weather data are used to make predictions that inform the weather forecast.

Artificial intelligence finds important associations that we may not previously have discovered: the complex association between the molecular structure of a chemical compound and its physical properties, or the complex set of circumstances that lead to an outbreak of disease. This characteristic enables AI to solve problems that we don't yet know how to solve ourselves.

As we review each of the eight main applications of AI, think about how each one might impact your business, your life, and society at large.

Machine Vision: Computers Open Their Eyes

With the advent of artificial intelligence, machines have evolved eyes and ears. Computers can now see, hear, and “understand” something about the world that they inhabit. This understanding is still rudimentary. Computers can recognize an image of an apple and correctly categorize it with the five letters a-p-p-l-e, but they don't understand what an apple is, that it grew on a tree, or what it tastes like.

Machine vision has many interesting applications across a wide range of industries. Stocktaking robots audit shelf contents in grocery stores. Facial recognition algorithms turn faces into passwords. Agricultural robots spot-spray herbicide on weeds. Quality assurance AIs perform visual inspection on the manufacturing line. Autonomous machines—robots, drones, and self-driving vehicles—all rely on machine vision, too.

EarthNow, an ambitious startup funded by Bill Gates, Airbus, Softbank, and others, is a splendid example of machine vision's future

potential. EarthNow will operate a constellation of satellites, each containing four powerful, high-definition cameras. The company's goal is essentially to create a real-time version of Google Earth, though with an important twist: Artificial intelligence, built into the satellites, runs applications to interpret camera images and add intelligent insight. These applications will reveal important details about activity on our planet.

A lightning strike in a remote location can start a devastating forest fire that rages out of control. Detection, rapid response, and early containment may save millions of acres of land from incineration, prevent damage to structures, and save lives. With EarthNow, global fire detection is just another application that runs on the satellites. The network becomes an eye in the sky that watches for fire starts and automatically alerts local authorities, 24/7/365.

EarthNow has proposed a range of other exciting applications for their satellite constellation. The system could provide real-time traffic information to city managers, real-time crop health information to farmers, and alert law enforcement or government agencies to illegal fishing, mining, and logging activity. Marine biologists will be able to track whale migration and volcanologists will monitor volcanic activity. Global asset tracking applications include tracking ships at sea, trucks on the road, planes in the air, and shipping containers in transit. Other applications include improved weather forecasting, law enforcement, and news coverage. EarthNow is a powerful platform that raises serious privacy concerns. This is why EarthNow chose to build machine learning capabilities into the satellites themselves. Users have very limited access to data from applications—real-time traffic data, the GPS locations of whales—but not the image data itself. The machine vision capabilities of the EarthNow platform present an exciting new set of opportunities for research scientists, public safety professionals, local governments, and a wide range of businesses.

Natural Language Processing and Voice Platforms

Human language is complex, powerful, and full of rich nuance. It's a tough task to build a computer that can "hear" spoken natural human

language, break it down into its component parts, infer meaning and intent, and then act on that information to do something useful. Artificial intelligence is behind many incredible applications of natural language processing. It can find spelling and grammatical errors, translate from one language to another, review legal contracts, and summarize the important points of a report.

Natural language processing (NLP) is an umbrella term that describes a machine's ability to understand, process, and communicate using natural human language. NLP consists of two pieces: natural language understanding (NLU) and natural language generation (NLG). You can think of one as being language input and the other as language output. In more technical terms, NLU converts unstructured human language data to structured data that a computer can understand while NLG converts structured data to unstructured data in the form of human language. Artificial intelligence is essential to both NLU and NLG. Natural language processing sits at the heart of voice interfaces, language translation services, email sentiment analysis, and many other applications that involve human language.

Natural Language Understanding (NLU)

Computers that understand human language perform many useful business tasks. Sentiment detection assesses text to determine if it conveys positive or negative emotion, for example, to highlight positive online product reviews or surface negative customer emails that require a swift response. NLU is also used to detect profanity, hate speech, threats, abuse, and other conversation that may be deemed inappropriate.

Natural language understanding is used to analyze documents and provide decision support. Scriptbook, a startup from Antwerp, Belgium, reviews movie screenplays to predict their likely box office failure or success. The software helps studios to make greenlight decisions on scripts. Scriptbook analyzed the scripts of 62 movies released in 2015 and 2016. Thirty of these movies were box office successes; 32 were failures and lost money. Scriptbook's AI correctly predicted all 30 of the box office hits and correctly called 22 of the movies that were duds. With 52 correct

calls made on 62 movies, the AI scored far better than the Hollywood moguls had. Scriptbook also uses NLU to predict a film's likely MPAA rating, the likeability of characters, and the countries where a movie will find most success, all based on its script.

LawGeex, an Israeli company, uses natural language processing to automate the review of legal contracts and nondisclosure agreements. LawGeex challenged 20 U.S.-trained lawyers to identify legal issues in five real-life nondisclosure agreements (NDAs), faster than their AI. The test was overseen by an independent lawyer and performed with input from legal experts and law professors. The lawyers took an average of 92 minutes to review all five NDAs and achieved an average accuracy rate of 85%. The LawGeex AI scored an accuracy of 94%, equal to the best lawyer's score, and completed the entire task in just 26 seconds. Lawyers that I've told about this AI are generally delighted. Reviewing NDAs is not a favorite part of their work and they are excited to offload routine tasks and focus more time on higher value, higher revenue work.

Natural language understanding speeds data entry. AIs identify email addresses, physical addresses, dates of birth, and phone numbers from nonstandard forms, automatically. This technology, sometimes combined with handwriting recognition (thanks again, AI!), makes short work of data entry.

Voice recognition combines speech-to-text capabilities with NLU. This technology automatically creates subtitles for videos and presentations. Microsoft now includes this capability with some versions of its PowerPoint application.

Natural Language Generation (NLG)

Computers with the ability to write or speak in natural human language are a huge breakthrough. AIs generate language either from source data or from source text. For example, an auto-generated weather report is created from weather forecast data, while a translation from one language to another is performed based on source text.

Natural language generation (NLG) has many valuable business applications. Language translation is an obvious one. Another is the automatic creation of summaries and abstracts. These might be summaries of financial reports, legal documents, operations reports, performance reviews, news articles, or medical records. NLG is combined with image recognition to automatically caption or describe images. This is a valuable function for people with visual impairment and also improves the quality of image searches.

In March 2014, the *Los Angeles Times* published a short story on an earthquake in Beverly Hills, California. The article described the location, time, and strength of the earthquake (4.4 on the Richter scale, if you're interested) and posted it to the *L.A. Times* website within three minutes of the earthquake. The article was written by a simple piece of software that sourced seismic data from the U.S. Geological Survey. This simple automation populated data into a prewritten template but demonstrates the power of automation. In a world of 24-hour news cycles and shrinking ad revenue, NLG frees human reporters to focus on higher-value stories and investigative journalism. More advanced automation now writes weather and traffic reports, summarizes business results, and covers sports events. Wordsmith, an automated reporting platform used by the Associated Press (AP), uses NLG to generate stories about minor league baseball games, college basketball, and quarterly corporate earnings reports. The AP claims that Wordsmith produces more than 4,400 corporate recaps each quarter, more than 15 times the number it could previously handle using a human writing staff. Wordsmith offers a glimpse of the more sophisticated NLG capabilities coming in the near future.

Simplish uses AI to convert complex text, with a vocabulary of more than 100,000 words, into simpler text with a vocabulary of less than 2,000 words. NLG will make complex language accessible to broader audiences with lower levels of education, for example, an academic text translated for lay people or children. Future NLP applications might provide comprehensive copyediting services or disrupt the businesses of companies like CliffsNotes.

Quillionz uses AI to automatically generate questions, quizzes, and assessments from a body of text. This provides an amazing aid for teachers, who can guide the AI to focus on particular keyword topics and automatically create multiple-choice questions.

There is a dark side to NLG technology. If we think that online fake news is a problem now, just wait until next-generation NLG is unleashed by the Russian troll farms. A human troll farm might create a hundred fake “news” articles in a day; a weaponized AI could pump out a million pieces of disinformation in a single hour. OpenAI claims to have created NLG software that writes such high-quality text that they have chosen to withhold its release, based on such concerns.

A 2015 survey of 352 leading AI experts by the University of Oxford predicted that AI will translate languages better than human translators by 2024, will write high-school-level essays by 2026, but won’t write a best-selling novel until 2049. It’s a giant leap from teenage essays to Tolkein, but the direction is clear. While we are a long way from the first Pulitzer Prize–winning novel written by AI, sophisticated NLG technology will soon create complex documents that rival human output.

Voice Agents

Voice agents like Apple Siri, Microsoft Cortana, Amazon Alexa, Samsung Bixby, and Google Assistant improve significantly every year. Some add new capabilities almost every week. Google claims their voice assistant is now available on more than a billion devices. These “conversational computing” or “digital dialogue” platforms will reshape the way we work and become an increasingly important part of our lives. Paradoxically, voice agents may become both our managers, and our subordinates, guiding our actions and performing our errands.

Voice interfaces are an important component of hands-free computing and a natural complement to augmented reality. Hands-free technology offers the prospect of “computing for the rest of us” and voice interfaces are valuable for people with impaired vision, for use in sterile clinical environments where physical interfaces aren’t appropriate, and

for the 80% of people who either work with their hands or in highly mobile environments.

Artificial intelligence—in the form of speech-to-text, NLU, NLG, and speech synthesis—underpins the operation of voice platforms. Continued advances in AI will make conversational computing sound more human. Future voice agents will engage in complex back-and-forth conversations, speak realistically, use human idioms, and even add fake breathing sounds and natural hesitations and pauses to make them sound more human. Google's Duplex technology and Microsoft Cortana have already made great strides and we should expect *major* breakthroughs in this area in the coming years.

As voice agents become more sophisticated, they will become ever-present in our lives and help us navigate our days. We will use them to make reservations, manage our calendars, run errands, place orders, troubleshoot problems, give advice, and even provide emotional support. Ultimately, a transactional conversation with a digital voice agent will become indistinguishable from one held with a human. This prospect has profound implications for those that work in customer service.

Exploration and Discovery

The New York Police Department (NYPD) uses PATTERNIZR, an AI-based discovery tool, to spot crime patterns. With more than 68,000 robberies, larcenies, and burglaries occurring in New York in 2018 alone, the NYPD will take all the help they can get. The NYPD is split into 77 separate precincts. PATTERNIZR, rolled out in December 2016 but only revealed to the public in early 2019, looks for crime patterns that span precincts. PATTERNIZR frees up human analysts to focus on more complex analysis tasks. Just like a human analyst, PATTERNIZR compares factors including method of entry, the type of items stolen, the distance between crimes, and so on. To eliminate racial bias, the system is not given the race of suspects. PATTERNIZR has already proven useful. For example, the AI found a link between crimes that spanned precincts that had not previously been flagged as connected. In two

cases, a man used the threat of a syringe to steal a drill. The AI identified two other instances where a syringe was used as a threat in robberies. The NYPD used the information to locate the suspect and arrest him. He pleaded guilty to larceny and assault.

AI's predictive capabilities are a powerful tool for researchers. Material scientists use AI to predict the structures of materials that may have a desired set of physical properties. New alloys and compounds may be discovered as a result. AI-guided research could lead to the discovery and synthesis of new wonder materials such as room-temperature superconductors and high-efficiency battery electrolytes that would transform the energy sector and help to address the climate challenge. As we review in Chapter 10, pharmaceutical companies use a similar approach to help them discover new drugs. Predictive AI may help us discover therapeutic drugs that transform human health. How might AI boost your company's research efforts?

Better-Informed Decision-Making

Business analytics refine data to offer all manner of insights. Some use statistical techniques, some use heuristics, and many now leverage the power of AI to find associations in data, extract insights, and make recommendations. Intelligent decision support systems use such analytics to support data-driven decision-making. Credit agencies like Experian and credit card companies including American Express use machine learning to boost the speed and accuracy of credit approvals by crunching terabytes of consumer data. Some Customer Relationship Management (CRM) platforms use AI to intelligently prioritize leads. AI is used to guide recruitment decisions, spending decisions, investment decisions, purchase decisions, marketing decisions, design decisions, engineering decisions, and much more.

Mortgage companies use AI to assess loan risk and guide underwriting decisions. Underwriters make the near-Shakespearean decision of "to loan, or not to loan" with risk models built on historical data. An

underwriter assesses risk using 10 to 15 data points about a prospective borrower: salary, credit score, debt-to-earnings ratio, and so on. Based on this assessment, the underwriter either gives the thumbs up or thumbs down. Limited data doesn't provide a full picture of a person's ability to repay a loan. There's more to a person than 10 to 15 data points. Underwriters use this narrow data set to limit complexity and manage their workload. An underwriting AI considers hundreds of data points about a person and finds complex associations that create a more nuanced picture of a prospective borrower. Zest Finance and Underwriter.ai claim their underwriting AIs find low-risk loan candidates who don't qualify with a traditional underwriting approach. The upside for mortgage companies: they sell more loans without increasing their risk.

Predicting the Future

AI offers us a crystal ball. When AI is applied to historical data, it finds patterns and complex associations that allow it to make high-quality predictions about what might happen next. AIs are used to predict disease outbreaks, assess actuarial risk, and predict future demand on the electric grid. Atidot, Quantemplate, and Analyze Re use AI to predict insurance risk.

Law enforcement uses AI to predict crime. PredPol is a collaboration between the Los Angeles Police Department (LAPD) and the University of California at Los Angeles (UCLA). PredPol predicts where and when serious crimes are most likely to happen. PredPol scientists claim the system has double the accuracy of human analysts. Importantly, the system only predicts the locations of future crimes, not the identities of people predicted to commit those crimes. We are still a long way from the precrime concept described in Philip K. Dick's *Minority Report*.

Most businesses need to make forecasts or predictions, and AI is just starting to scratch the surface in this space. AIs will power all manner of business planning systems: demand forecasts, risk analysis, design trends, and more.

Seeing the World through a New Lens with Super Sensors

Sensors, turbocharged by AI to create “super sensors,” will lift the veil from the world, extend our existing five senses, and allow us to perceive the world more fully. Like the microscope before it, AI gives us a new lens to view the world through; a lens to experience the world in all its complexity and beauty. Super sensors will oversee business operations, monitor the operation of equipment, create new products, and provide us a more holistic view of people.

An exciting early example of super-sensing capabilities comes from the work of Dr. Dina Katabi, a professor and research lead at MIT’s Computer Science and Artificial Intelligence Lab (CSAIL). Katabi’s team has created a super sensor. What adds to the delight of this particular story is that it begins with *Star Wars* and ends with *Star Trek*.

As a child, Katabi was fascinated by the notion of “The Force” in the *Star Wars* movies. Her fascination persisted into adulthood. In *Star Wars*, Obi Wan Kenobi describes a mystical force that “surrounds us, penetrates us, and binds the galaxy together.” As she considered the fictional notion of The Force, Katabi realized that a real force surrounds us all—electromagnetic energy. If you flap your arms up and down, you literally create a disturbance in that force. Katabi wondered if sensing electromagnetic energy would let her see the world in a new way.

Katabi’s research team built a simple wall-mounted sensor and fitted it inside a room. The sensor emits and receives radio frequency (RF) signals, much like a Wi-Fi hotspot. RF signals tend to pass through walls but bounce back from people. The sensor picks up reflected RF signals and feeds data into a neural network that makes sense of the reflected signal data.

To train the AI, Dr. Katabi’s team captured video of people moving around inside the room. The AI was fed the video and RF sensor data as parallel inputs. The AI found complex associations between the RF sensor data and the video images and eventually correlated the two. With this insight, the AI can determine what is happening inside the room from just the RF sensor data and can register when a person is standing, sitting, or lying down. Since RF signals pass through most walls, the

AI can “see” through them and also “see” in the dark. Amazing. Katabi plans to use the sensor to monitor elderly patients under care. The sensor instantly detects when patients suffer a fall and calls for help. Incredibly, the sensor also detects vital signs—a patient’s breathing and heart rate—and their sleep state—awake, light sleep, deep sleep, and REM sleep. If you understand how well a person sleeps, you can tell a lot about their health. Disturbances in deep sleep can indicate depression or anxiety. Disturbances to REM sleep can indicate early-onset Alzheimer’s. Alzheimer’s is also indicated by repetitive patterns of movement or motion, which the sensor also detects. Changes in a person’s gait can indicate the development of Parkinson’s disease.

By applying AI to a low-cost, wall-mounted RF sensor, we can monitor vital signs, sleep states, falls, and provide clinical insight on the development of conditions that include Alzheimer’s, COPD, Parkinson’s, and depression. All without a single wire being attached to the patient. The system operates without using a camera, limiting privacy concerns. Patients are monitored with their full consent, comfortable in the knowledge that no video images are ever being gathered. The sensor enables clinicians to “see” a wealth of information, all gathered wirelessly. We started this story with *Star Wars* and ended up with the sick bay beds from *Star Trek*.

Super-sensing turbocharges simple sensors with AI to reveal more about our world. Google’s “Project Soli” technology uses short-range radar signals and machine learning to detect fine-grain finger gestures. This approach will transform sliders, buttons, and twistable crowns on mobile and wearable devices into virtual controls. More super sensors are detailed in Chapter 10 on the future of healthcare. What other super sensors will we create? What super sensors will your business create?

Solving Complex Problems by Learning from Experience

Some challenges—optimizing a system with many variables or programming a robot to walk on two legs—are too difficult, too complicated, or too laborious to tackle with traditional computers. AI solves some of these tricky problems using a technique called reinforcement learning.

Reinforcement learning is a branch of machine learning that uses a system of digital rewards and punishments as part of its training process. Reinforcement learning systems solve previously intractable problems through an iterative process of experimentation. The AI tries a range of strategies and learns the best way to approach a problem through an intelligent form of trial and error. It's like harnessing digital evolution.

Reinforcement learning teaches computers to perform complex optimizations, control complex equipment, and to play games really, really well. In 2018, researchers trained an AI to play the classic Sega console game, *Sonic the Hedgehog*. *Sonic* has two simple controls: run and jump. An AI was trained with the video game display as input and the game controls as the output. In reinforcement learning, AIs have an additional input known as a reward function. As the AI trains it tries to optimize the reward function. Game points increase the reward, and the reward decreases substantially if Sonic loses a life. At first, the AI plays terribly. Over time, the AI optimizes its model to run and jump at just the right moment, score maximum points, and keep the adorable blue hedgehog alive. The AI does not learn based on simple timing; it learns from what is happening on the screen, so it can succeed on game levels it has not seen before.

The most regularly cited example of reinforcement learning is DeepMind's AlphaGo system. DeepMind, a subsidiary of Alphabet, built AlphaGo to play the ancient Chinese game of Go. Winning strategies for Go are opaque; even grand masters can't always describe why they choose some of the moves they make—they say the move just “feels right.” There are more possible configurations for pieces on a Go game board than there are atoms in the universe. To build a machine that understands the nuances and subtle strategies of this complex game is a monumental challenge.

AlphaGo was not taught game strategies. It developed its own strategies through observation of many human versus human games. In March 2016, AlphaGo played 18-times world champion Lee Sedol, the best (human) Go player in the world. AlphaGo beat the legendary player, four games to one. To win, AlphaGo deployed several new strategies that went against hundreds of years of received wisdom among expert players. By observing AlphaGo's approach, human players have improved their play. This story offers an important lesson. Rather than consider AI

a threat to our unique humanity and our value within the workplace, we might instead think of AI as a sophisticated partner, one that boosts our skills and that ultimately elevates our humanity.

In 2017, DeepMind's next machine, named AlphaGo Zero, became a master Go player by playing millions of games against itself inside a simulation. It developed game strategies through practice rather than by observing human play. AlphaGo Zero now thrashes the original AlphaGo machine and is unassailable by all human grand masters.

Reinforcement learning isn't just used to play games. Researchers at Warsaw University used reinforcement learning to train bipedal robots to walk more efficiently. The AI that controls the robots varies the combinations of movements made by the robot's motors and experiments with different walking strategies. The robot's AI gains a small electronic reward for strategies that speed the overall efficiency and pace of the walk. With this approach, roboticists achieved more efficient and natural-looking walking motions for their robots. One robot learned to walk almost twice as fast as it could using the best initial walking strategy programmed by its human creator.

AI's ability to learn from experience is used to solve many business problems, including complex optimizations. AIs optimize traffic control systems, industrial chemical reactions, advertising bids, industrial automation, supply chain flow, product design, warehouse operations, inventory levels, yields, trading strategies, wind turbine controls, medication doses, smart grids, and commercial HVAC systems. Reinforcement learning also teaches AIs to drive. Like humans, AIs learn to drive by practicing. They drive real cars in real-world conditions but also drive millions of miles inside realistic software simulations. In part, *Tesla* AIs learn to drive from sensor data gathered while owners are driving.

Creating and Co-Creating Content

Computers now have the capacity for imagination. Artificial intelligence can compose music, paint pictures, and even write poetry. AI also co-creates content with humans: a partnership of digital intelligence with artists, designers, and engineers.

Most AIs create content using a relatively new AI approach known as Generative Adversarial Networks (GANs). The first GAN was built in 2014 by Ian Goodfellow, then a researcher at the Université de Montréal and now part of the Google Brain project. GANs link two AIs together in a sophisticated version of reinforcement learning where each AI trains the other. The AIs operate in an adversarial context: Each tries to catch the other out. Here's how it works. One AI is trained to create content and the other AI is trained to spot fake/generated content. Think of these AIs as a forger and an art detective. The detective, known as the discriminator AI, is initially trained with many real-world examples of "good" data (real content). The forger, known as a generator AI, tries to fool the discriminator by creating high-quality, realistic content. At first, the content created by the generator AI is terrible and the discriminator easily spots it as "fake." But the generator steadily improves until eventually the discriminator finds it harder to spot generated content. Over time, each AI gets better at its assigned task. While the relationship between the AIs is adversarial, the net result is that the discriminator acts as a coach to the generator AI, and the generator acts as a coach for the discriminator. Eventually, the generator AI creates incredible, high-quality content: images, videos, speech, music, prose, legal contracts, and engineering designs. GANs have even been used to automate the design of crowns for teeth.

GANs have incredible potential. They will fundamentally change the way we work. Within the next decade, many of you reading this book may partner with GAN-based AIs to co-create content and collaborate on business tasks.

While there are some worrying downsides to generative AI—most notably deepfakes—many applications offer incredible potential. The most powerful example of machines with an imagination today is generative design, a field pioneered by engineering design company Autodesk. Generative design is used extensively in engineering and architectural design. We discuss these applications further in Chapters 14 and 15. Generative design uses the imagination and content creation capabilities of GANs to generate hundreds or thousands of alternative design options from a single original design. Designers specify constraints—size,

weight, cost, and so on—and the tool generates options and evaluates each one using simulation tools. The designer simply picks the option that best meets their needs—perhaps the one that’s cheapest to make, easiest to manufacture, or has the lightest weight.

Generative AI will help to design lighter-weight aircraft, cars that are more resilient to crashes, and stronger, lighter robots. Generative architecture will improve the structural integrity and design of new buildings.

Researchers at the University of California, Berkeley, working in partnership with Glidewell Dental Lab, use GANs to design dental crowns. The AI uses digital x-rays of a patient’s upper and lower jaw to design a crown that perfectly fills the gap in the patient’s tooth line, optimizes bite contact, and looks aesthetically pleasing. Researchers claim that AI-generated crowns outperform those designed by humans. The approach should speed crown production, reduce costs, and free dentists to spend more time generating revenue by working in patients’ mouths, rather than designing crowns on a CAD machine in a back office.

Generative AI is an example of a broader category of AI that I refer to as “collaborative AI.” Collaborative AIs operate in partnership with humans in a creative process. Humanity’s use of tools distinguishes us from most other species. Traditional tools are subordinate—we wield a hammer, drive a car, and program a computer. Collaborative AI changes our relationship with tools. They are no longer subordinate, they now co-create with us. Collaborative AIs aren’t just tools, they’re partners.

Collaborative AI will co-create visuals for presentations, advertisements, and marketing brochures. Collaborative email software will auto-compose responses. Collaborative management software will co-create plans for complex projects. Many job functions will benefit from collaborative AI in the coming years.

Future Uses of AI

In a fast-moving field, new use categories, beyond the eight listed earlier, are bound to emerge. GANs, a technique that’s central to several of these application categories, were invented relatively recently. As research

embraces techniques beyond deep learning—cause and effect AI, common sense AI, capsules, and others—artificial intelligence will solve even more business problems than it can today.

AI is a big deal. Every leader should pay close attention. Every organization must understand how AI will shape product development, business operations, customer service, and workforce management.

How AI Works

You don't have to understand how AI works to use it. But such insight can help you to understand the capabilities and limitations of today's technology. While the following description is designed to be accessible to nontechnical types, feel free to skip to the next section if it gets too far into the weeds for you.

Neural Networks, Training, and Models

Neural networks underpin most of today's artificial intelligence. They operate quite differently from traditional digital computers. Traditional computers are glorified adding machines. Neural nets are organized more like the highly interconnected structures found in our brains.

Neural nets are made up of connected “nodes,” which act like neurons. Each node holds a numerical value. Unlike binary computers that work with zeros and ones, each node can have a range of values; the range depends on the application. Nodes are arranged into layers. The first layer is known as the input layer and the final layer is known as the output layer. All of the layers in between are known as hidden layers (see Figure 1.1).

Typically, the more layers there are, and the more nodes in each layer, the more capable the neural network. Neural networks with many layers are known as “deep” neural networks. This is where the term *deep learning* comes from.

Every node in the hidden layers has both inputs and outputs. Each node is connected to every node in the previous layer and every node in




Figure 1.1 A simple neural network.

the next layer. The value of each node is influenced by the values of all the nodes it is connected to in the previous layer. Here's the tricky bit: some nodes have a stronger influence on the value of proceeding nodes than others; their influence is weighted. The value of each node is the weighted sum of the values of the previous nodes. These weightings are determined during the training phase and collectively make up what is known as "the model." The model determines the functionality of the neural network: different weightings, different functionality. Information passes across the network from the input layer to the output layer via this complex web of weighted interconnections.

Neural networks are trained with a process known as backpropagation, or "backprop" as it's known in the business. The details of how backprop works is beyond the scope of this book. At a high level, backprop is a computationally intensive statistical approach that compares the desired output of a neural network with the actual output and then tweaks the weightings in the network to improve the accuracy of results. When the right result is given, the weightings of all the pathways through the neural network that lead to the correct result are

strengthened. If the result was incorrect, the pathways that lead to the incorrect result are weakened. Over time, with exposure to more and more data, the model becomes increasingly accurate. The network “learns” the correct complex associations between inputs and outputs.

Example: A Radiology AI

To train a neural network to read radiology charts and look for tumors, you would expose it to many example charts (the input), each tagged with a radiologist’s diagnosis—tumor or no tumor (the desired output). The output of the network is a single number, the probability that an image contains a tumor. Each time the neural net is exposed to a new image, the output of the network is compared with the correct result. If an image of a tumor is presented, the result should be close to 100%. If there’s no tumor the result should be close to zero. The backprop process is used to tweak the network’s model (the weightings of the connections between the nodes), strengthening the weightings of links that lead to the correct result, and weakening those that don’t. Once trained with enough data, the neural network will predict the right diagnosis with impressive accuracy. A more complex network might have several outputs. One could be the percentage chance of a tumor, another the probability of an embolism, another the probability of a broken bone, and so on.

If this all seems too difficult to understand, that’s okay. The key thing to understand is that neural nets can infer how to perform tasks from examples, *without* the need of a domain expert to supply explicit rules on how to perform that task.

Radiologists train for many years to read x-rays, computer tomography (CT), magnetic resonance imaging (MRI), and positron emission tomography (PET) images. After medical school, radiologists do additional training, often involving a four-year residency. Some do additional specialization training after that. Reading images to look for tumors and other ailments uses all of the radiologist’s skill, experience, and training. Yet this task is within the reach of a neural network. Given enough training data, an AI can be built with diagnostic abilities similar to those of a human radiologist, a person with about a decade of intense education

behind them. As we train the neural network, we are essentially codifying the collective knowledge, and several decades of professional experience, from hundreds of thousands of radiologists. Their experience and diagnostic insight are captured in the model that's generated.

Some radiologists already use AI-based tools to offer a “second opinion” as they read charts. As the accuracy of these tools surpasses that of human radiologists on routine charts, radiologists will be able to focus their attention on more complex, higher-value, and more patient-centered tasks and procedures. The progress made in radiology portends the future for other branches of medicine. Machine learning will be applied to many other fields of medical diagnosis and pathology in the coming decade.

While training an AI requires serious amounts of computing performance to create a model, using that model requires significantly less performance. The process of using a model is known as inference. Often, training occurs on workstations or in the cloud, while inference occurs on devices. Most future computer chips will include inference engines, silicon accelerators optimized to run AI models with relative ease.

Pattern Recognition

Pattern recognition is a core capability of many AI systems, including the radiology example we just discussed. Pattern recognition has many applications and comes in a range of different flavors. It's not important that you remember all these different approaches. They are listed here only to illustrate some of the fundamental capabilities of machine learning. As you read through them, think about how such a capability might be used to solve business problems in your organization.

- **Classification.** AI can classify data into similar types. For example, the radiology AI classifies images as positive or negative. A similar approach might be used to do visual inspection and quality assurance in a manufacturing plant, or to identify spoiled or underripe fruit at a fruit-packing plant.
- **Clustering.** Marketing professionals use clustering algorithms to partition consumers into market segments that share similar

characteristics—buying habits, affluence level, and needs or desires. Recommendation engines use clustering, too. *Spotify* recommends songs that you might enjoy by analyzing historical listening habits. A clustering algorithm finds the complex relationships between songs and listeners. The clustering algorithm might see that I like songs A, B, C, and D, and that you like songs B, C, D, and E. It may conclude that it's probable you will enjoy song A and I might enjoy song E. Clustering is useful to deliver personalized experiences.

- **Regression analysis** finds patterns that describe relationships between pieces of data. For example, regression analysis might observe that if Event A happens, most of the time Event B follows. More complex relationships are found, too, such as “if Datapoint A is below a certain threshold, and Event B and Event C are not happening, then Event D is 46% more likely to occur.” This approach is used to make predictions about the future with predictive analytics tools. Regression analysis is used by Walmart to predict how sales of certain food items are influenced by specific weather conditions.
- **Sequence labeling** is a pattern-recognition approach used in speech recognition, handwriting recognition, and gesture recognition. Sequence labeling is used to break sentences down into constituent words and phrases and to label them in a way that captures their context. For example, sequence labeling identifies which words are nouns, verbs, and proper names. Words are best interpreted in the broader context of a sentence. Sequence-labeling algorithms classify words within a sentence, or cursive letters within a handwritten word, by examining the broader context surrounding them.
- **Time-series prediction** is used in weather forecasting, stock market prediction, and to predict disasters. These algorithms analyze a set of historical data points and use that to project which data points might come next in a sequence.

These pattern-matching algorithms use complex mathematics to work their voodoo. You don't need to understand how pattern matching works, or even recall the names of all the various techniques listed

above. What you do need to understand is that AI makes it easier for computers to understand the physical world, to make predictions, and to find complex relationships hidden inside data. These tasks are at the root of solving many business problems.

Beyond Deep Learning: The Future of Artificial Intelligence

Most of the AI breakthroughs in the 2010s were built on deep learning technology and neural networks. Dramatic advances in machine vision, natural language processing, prediction, and content generation resulted. And yet, industry luminaries debate whether AI is about to enter a golden age of rapid technological advancement, or rapidly stagnate.

Stagnation or Golden Age?

The argument for stagnation is that deep learning has severe limitations—training needs too many examples and takes too long, and while these AIs pull off some amazing tricks, they have no true understanding of the world. Deep learning is built on algorithms from the mid-1980s and neural network architectures developed in the 1960s. Once we have perfected the implementation of deep learning technology and solved all the problems that we can with it, there are no viable technologies in the pipeline to keep things rolling. The current era of AI deployment will grind to a halt. So goes the stagnation argument.

On the other side of the debate are those who point to promising research that could take AI in new directions and solve a new set of problems.

Capsule Networks

Capsule networks are the brainchild of Geoff Hinton, the creator of backprop and one of the fathers of deep learning. Capsules overcome some of deep learning's shortcomings. The difference between capsule networks and traditional convolutional neural networks is beyond the

scope of this book, but capsules capture some level of understanding about the relationship between features in images, which makes image recognition engines more resilient and better at recognizing objects from many different angles.

Common Sense

AIs are trained to understand something about the world. Typical AIs operate within a bubble. They have no understanding of the way the world works. A lack of common sense limits their abilities. A household robot, on a search to find my reading glasses, should know that my desk and nightstand are good places to look first, and not inside the freezer.

Several organizations are trying to build AIs with common sense. They are building vast databases of the commonsense notions humans use to help them make high-quality decisions. For example, oranges are sweet, but lemons are sour. A tiger won't fit in a shoe box. Water is wet. Oil is viscous. If you overfeed a hamster, it will get fat. We often take this context for granted, but to an AI these notions are not obvious.

Researchers at the Allen Institute crowdsource commonsense insight using Amazon's Mechanical Turk platform. They use machine learning and statistical analysis to extract additional insights and understand the spatial, physical, emotional, and other relationships between things. For example, from a commonsense notion that "A girl ate a cookie," the system deduces that a cookie is a type of food and that a girl is bigger than a cookie. Allen Institute researchers estimate they need about a million human-sourced pieces of common sense to train their AIs.

The Cyc project, the world's longest-running AI project, takes a different approach. Since 1984, Doug Lenat and his team have hand coded more than 25 million pieces of commonsense knowledge in machine-usable form. Cyc knows things like "Every tree is a plant" and "Every plant dies eventually." From those pieces of information, it can deduce that every tree will die. Cycorp Company, Cyc's current developers, claims that half of the top 15 companies in the world use Cyc under license. Cyc is used in financial services, healthcare, energy, customer experience, the military, and intelligence.

As they mature, commonsense knowledge systems may help future AIs to answer more complex questions and assist humans in more meaningful ways.

Causal AIs

Causal AIs understand cause and effect, while deep learning systems work by finding correlations inside data. To reason, deep learning AIs find complex associations within data and assess the probabilities of association. Reasoning by association has proven adequate for today's simple AI solutions, but correlation does not imply causation. To create an AI with human-level intelligence, researchers will need far more capable machines. Some AI researchers, most notably Dr. Judea Pearl, believe that the best path forward for AI development is to design AIs that understand cause and effect. This would allow AIs to reason based on an understanding of causation. A deep learning AI associates events (A and B happen together) while a causal AI understands that one event caused the another (A caused B), and not the other way around (B caused A). The sophisticated machines needed to solve big problems like climate change will need to understand all of the causational relationships involved in highly complex systems. Causal AI will rely on the commonsense knowledge mentioned in the previous section to give it the vital context it needs for sound reasoning.

Neuromorphic Computers

Neuromorphic computers are inspired by the way brains work. Today's neural network designs are based on an understanding of neuroscience from the 1960s. Half a century later, we finally have the computing horsepower needed to implement these archaic machine models. The next time you ask Google Assistant or Alexa to play some Beatles music, remember that the foundation of the AI you're using was designed at the time when the Stones and The Beatles were first vying to be top of the charts. Neuromorphic computers, also referred to as cognitive

computers, are based on a much more recent understanding of how the brain operates. Nodes are hyperconnected, their connections can change over time (in the same way that the human brain exhibits plasticity), and there is no separation between memory and processing functions.

Major research projects—the Human Brain Project in the European Union and the BRAIN (Brain Research through Advanced Innovative Neurotechnologies) initiative in the United States—seek to advance our understanding of the human brain, mapping and understanding brain function. These efforts, and others like them, push the boundaries of our understanding and offer new frameworks for the design of future neuromorphic computers. New computer chips, inspired by neuromorphic insights, may accelerate AI functions, reduced power consumption for AI tasks, and enable exciting new capabilities.

Whether it's capsule networks, neuromorphic computing, or common sense and causal AI, there are plenty of avenues of research that should fuel future advances in AI in the coming decades.

Narrow, General, and Super Intelligence

All of today's AI is considered to be “narrow” AI. The holy grail of AI research is the development of “general” and “super” AIs. Let's quickly review these three categories.

Artificial Narrow Intelligence (ANI)

Artificial Narrow Intelligence (ANI), also known as weak AI or vertical AI, refers to any AI that solves problems or performs tasks with a level of intelligence equivalent to, or higher than, a human, but only within a narrowly defined domain. Every AI available today, and every AI described in this book, is an example of narrow AI. Narrow AI is only good at the task it was designed for and useless at others. A chess-playing AI can't filter the spam from your email, and your spam filter can't play chess.

Artificial General Intelligence (AGI)

Artificial General Intelligence (AGI), also known as strong AI, describes AIs with the equivalent intelligence of a human being in any area of human expertise that you might imagine. An AGI can perform any intellectual task that a human can. Researchers continue to chip away at the problem of creating an AGI but currently have no clear plan for how such a feat might be achieved. AGIs aren't created by bolting together enough ANIs. It doesn't work that way. Never say never, but AGI, if we achieve it, is likely several decades in the future.

Artificial Super Intelligence (ASI)

Artificial Super Intelligence (ASI) is where things get really exciting, or really scary, depending on your viewpoint. ASI defines an intelligence that surpasses the intellectual ability, knowledge, creativity, wisdom, and social skills of the very best human brains in almost every field. An ASI could be just 1% smarter or a million times smarter than the smartest human. In theory, an ASI machine that we build could design future, more powerful machines that operate in a way we aren't able to comprehend. Once ASI is achieved, rapid, runaway advances could follow. If that idea makes you uncomfortable, you're not alone. Exemplars of the high-tech and scientific world—including Elon Musk, Bill Gates, and Stephen Hawking—have all spoken out on the dangers of artificial super intelligence. It's hard not to think of the movie *Terminator* when you think about the professed dangers of AI becoming self-aware and designing better versions of itself. How engineers develop narrow AIs today may inform the design of super AIs in the future. This is why research and engineering focused on transparent and unbiased AIs is so vital.

Strategies to Get You Started

Artificial intelligence holds huge potential and will drive successive waves of innovation throughout every industrial sector. As your organization starts to build out a comprehensive, multiyear AI strategy, here are a few places to start your thinking.

Make Predictions about Customers, Operations, and Trends

Use AI's predictive capabilities to forecast demand, streamline operations, target marketing efforts, predict trends, and influence the design of future products. Energy companies use AI to forecast demand for electricity. Fashion designers use AI to suggest fall colors for a couple of years from now. Equipment makers use AI to predict when their machines will fail so they can schedule preventative maintenance. What could AI predict for your business?

See More, Understand More, Make Better Decisions

Future super sensors may improve the safety of autonomous vehicles, become “the next CT-scan” in medicine, or give us all a sixth and seventh sense. Super sensors will smooth business operations, revolutionize human interfaces, and transform the products and services of the future.

Every company should lead a strategic discussion around the possibilities super sensors present. Determine how your company can use super sensors to get eyes on your business, gain insight on operations in real time, and use those insights to make high-quality, data-driven business decisions. Inspired by Google Soli technology, consider how you could build super sensors into your products to improve the human-machine interface. How might you see your customers in new ways with super sensors? What other insights might super sensors reveal about your business?

Build a Comprehensive Digital Voice Strategy

Every business should have a digital voice strategy. Voice platforms offer customers a new way to interact with your brand, will complement customer support capabilities, and can be helpful for employees who use their hands to perform their work or don't work in traditional office settings.

First, determine the goals for voice platforms within your organization. Find areas where information delivered via voice interfaces improves worker efficiency or boosts their collaboration. Ask what conversational capabilities your brand should create for existing customers, and

how voice might be used to attract new customers. Consider the role of interactive, talking online ads, e.g., “Click to ask me questions about our latest model!” Explore how digital voice agents can reduce support costs and open new sales channels.

Customers always want more choice. When brands give us more choice, they honor our desire to be recognized as individuals. Consumers expect to interact with brands through many channels: face-to-face, phone, web, mobile app, social media, virtual reality, interactive augmented reality objects, chatbots, voice agents, and more. Each brand will offer as many of these choices as makes sense for them. A voice interface is a new front door to your business, and voice agents are ambassadors of your company’s brand. You must secure that front door the same way you secure your website. And you need to teach your ambassador to always put your best foot forward.

Businesses will need to train their chatbots to interact with a “tone” that is consistent with their brand voice. Marketing departments will need to define appropriate design parameters: Should the voice agent strike a casual or a formal tone? Calm and efficient, or excited to serve? Serious, or with a sense of humor? Rated G for everyone, or more PG-13? Voice agents for Disney, Harley Davidson, and Moët & Chandon will sound very different from one another.

Decide where to develop your voice agents: in-house, using third-party frameworks, or on existing voice platforms. Google, Amazon, Apple, and other tech giants enable developers to create apps for their voice services. Google calls these apps “actions” while Amazon calls them “skills.” The make-versus-buy decision will vary by brand. Most will build a presence on the major tech platforms. For competitive or security reasons, some brands may elect to build a dedicated voice agent that is embedded inside their apps and websites. Banks may prefer their own voice agents for security reasons. Bank of America’s Erica virtual assistant is a good example.

The future of voice may become more open. Today’s main voice platforms, Google Assistant and Amazon Alexa, resemble AOL and CompuServe, the walled garden online services of old. The rise of the internet,

with open standard web pages and browsers, broke down the walls of the online services and opened up the online world to anyone who wanted to participate. Efforts like the Open Voice Network seek to build standards for online voice akin to HTML for webpages and enable brands to build their own voice capabilities. Time will tell if this will fell the initial duopoly of the tech giants.

Understand and Mitigate for Bias

When AIs are trained with data that represents examples of human behavior, they learn unwanted human biases that are reflected in that data.

Amazon receives a small mountain of job applications each week. To triage resumes, Amazon built an experimental AI to surface candidates worthy of an interview. There was only one problem. According to a report by Reuters, it quickly became clear that Amazon's AI was biased against women. The problem with the AI had nothing to do with the way it was coded. The AI reflected historical bias in Amazon's hiring process. The unconscious bias of hiring managers had been codified in the recommendations of the AI, which learned to favor male candidates. Amazon canceled the project before it was ever deployed and reviewed its hiring practices. Sometimes AI holds a mirror up to our own humanity.

We should hold our AIs to even higher standards than we hold ourselves, and should strive to create AIs that don't reflect negative human bias. The state of New Jersey built an AI as part of their plan to do away with the bail bonds system. Cash bail has been a part of the U.S. court system for centuries but is criticized for penalizing poorer defendants. According to a 2013 study by the Drug Policy Alliance of New Jersey, 75% of the New Jersey jail population were people awaiting trial. The average wait time for trial was 314 days. Almost 40% of the people awaiting trial were there because they couldn't afford \$2,500 or less in bail. Under New Jersey law, everyone must be given bail, no matter the crime, so rich people always avoided jail and bail bondsmen made good money. In January 2017, the state of New Jersey replaced their bail system with an AI that created a public safety assessment (PSA) for each defendant. The PSA was used as a guide by judges. The PSA predicts

the chance defendants will commit crime while awaiting trial and whether they will appear for their court date. The new system reduced the number of people in New Jersey jails awaiting trial by 30%. The AI was trained with information about 1.5 million previous defendants in 300 jurisdictions. Race and gender information was deliberately removed from the training data. Developers also removed all information on the name, education level, employment, income, and home address of defendants. Any of these data can be a proxy for race and gender and thus might disadvantage some demographic groups. The way we train our AIs matters. We should expect our AIs to operate ethically, fairly, and without bias. Leaders must ensure that their teams work hard to strip bias from all the AIs they build.

Gather Data Today to Feed Your AIs of the Future

AIs have a voracious appetite for data. Build data pipelines to source the data today that you will need to train the AIs of the future. Your AI strategy informs your data strategy. Your data strategy informs your sensor strategy, your partnership strategy, your hiring strategy, and your IT strategy. Figure out what data you will need and where you will get it from. Think of this as another strategic “make versus buy” decision.

Data storage is relatively cheap. You never know what operational, customer, market, and other insights tomorrow’s AI might find hidden inside today’s data. It goes without saying that every business should adhere to all relevant data privacy and data retention laws. Wherever possible, businesses should gather and archive as much data as possible, particularly data that comes from machine operations, sensors, and market research. Today’s data is tomorrow’s AI fuel.

Develop New Products and Services

Use AI to spark innovation across your entire product line. Enhance and simplify human interfaces using voice and super sensors. Create new value and build products that intelligently manage and maintain themselves. Use generative design to accelerate research, push the boundaries

of product design, and speed service delivery. Marry human talent with collaborative AIs to build new service offerings at new price points.

Final Thoughts

At the beginning of the twentieth century, every business had to grapple with electrification. In the late twentieth century, every business had to understand and embrace digital technology. To remain competitive, every business must now build a comprehensive strategy that embraces artificial intelligence in every aspect of their operations. No exceptions. No delay.