"*Human + Machine* is a richly detailed guidebook that leaders need to capture the opportunities of AI and the fourth industrial revolution. If we heed Daugherty and Wilson's call to action on reimagining how we work and preparing people for fusion skills, we can ensure a brighter future for all of us."

— **KLAUS SCHWAB**, founder and Executive Chairman, World Economic Forum; author, *The Fourth Industrial Revolution*

"In the post–Information Age, every sector of the global economy is driven by technology, the source of disruption and new wealth everywhere. For those wishing to participate, you'll need a copy of *Human + Machine*."

— **MARK ANDERSON**, founder and CEO, Strategic News Service

"In *Human + Machine*, Daugherty and Wilson brilliantly illuminate with real-world examples how companies across every industry are reconceptualizing their businesses and organizations for the age of AI. This is just the beginning of the greatest business transformation in history, with humans and machines working together in ways never before imaginable. As the authors point out, we must invest in training millions of people for the jobs of tomorrow and establish guardrails to ensure that as AI evolves, the benefits accrue to all of humanity. *Human + Machine* is a roadmap to the future—read it if you're serious about understanding the impact of AI and how it is driving growth."

— **MARC BENIOFF**, Chairman and CEO, Salesforce

"*Human + Machine* is a well-written and well-researched book that addresses the missing middle of AI: how humans and machines can collaborate so as to augment—not replace—human skills. From the

# CONTENTS

PART TWO

# The Missing Middle

## *Reimagining Processes with AI*

*Those who can imagine anything, can create the impossible.*

—Alan Turing

*See, the world is full of things more powerful than us. But if you know how to catch a ride, you can go places.*

—Neal Stephenson, *Snow Crash*

# INTRODUCTION

# What's Our Role in the Age of AI?

In one corner of the BMW assembly plant in Dingolfing, Germany, a worker and robot are collaborating to build a transmission. The worker prepares a gear casing, while a lightweight robot arm, sensitive to and aware of its surroundings, picks up a twelve-pound gear. The worker moves on to her next task, while the robot precisely puts the gear inside the casing and turns away to pick up another.

In another part of the plant, where LP's song "Lost on You" hums across the floor, a different lightweight robot arm evenly applies a thick black adhesive to the edge of small car windows. Between applications, a worker walks over to wipe off the glue nozzle, pop in new glass, and carry away the finished windows, as if robot and human are part of a well-choreographed dance.[1]

Thanks to recent advances in artificial intelligence (AI), we are now at the cusp of a major transformation in business. It's a new era in which the fundamental rules by which we run our organizations are being rewritten daily. AI systems are not just automating many processes, making them more efficient; they are now enabling people and machines to work collaboratively in novel ways. In doing so, they are changing the very nature of work, requiring us to manage our operations and employees in

dramatically different ways.

For decades, robots have typically been large pieces of machinery, usually sectioned off from human workers, that would perform a dedicated task—unloading a stamping press, for example. That specific task was part of a rigid, fixed chain of work that would generally include humans doing other predefined tasks —for instance, inspecting the stamped metal parts in order to discard defects.

Contrast that traditional assembly line with a factory where robots are much smaller and more flexible, able to work alongside humans. A factory where those robots and other types of machinery are using embedded sensors and sophisticated AI algorithms. Unlike earlier generations of industrial robotics—which were typically bulky, unintelligent, and somewhat dangerous pieces of machinery—these new types of collaborative robots are equipped with the ability to sense their environment, comprehend, act, and learn, thanks to machine-learning software and other related AI technologies. All this then enables the work processes to be self-adapting, with fixed assembly lines giving way to flexible human-machine teams that can be put together on the fly. Now, in order to fulfill customized orders and handle fluctuations in demand, employees can partner with robots to perform new tasks without having to manually overhaul any processes or manufacturing steps. Those changes are baked into the system and are performed automatically.

The advances are not just in manufacturing. AI systems are being integrated across all departments, everything from sales and marketing to customer service to product R&D.

Take, for instance, a designer at Autodesk who decides to build a drone. Instead of modifying preexisting concepts and adjusting for various constraints like weight and propulsion, she inputs these parameters into the company's AI-enabled software. The software's genetic algorithm produces a vast and dizzying array of

new designs that no one has ever seen. Some are more bizarre than others, but all fit the initial constraints. The designer chooses one that will distinguish her drone from the rest and further tweaks the design to fit her aesthetic and engineering goals.

## From the Mechanistic to the Organic

The potential power of AI to transform businesses is unprecedented, and yet there is an urgent and growing challenge. Companies are now reaching a crossroad in their use of AI, which we define *as systems that extend human capability by sensing, comprehending, acting, and learning*. As businesses deploy such systems—spanning from machine learning to computer vision to deep learning—some firms will continue to see modest productivity gains over the short run, but those results will eventually stall out. Other companies will be able to attain breakthrough improvements in performance, often by developing game-changing innovations. What accounts for the difference?

It has to do with understanding the true nature of AI's impact. In the past, executives focused on using machines to automate specific workflow processes. Traditionally, these processes were linear, stepwise, sequential, standardized, repeatable, and measurable, and over the years they've been optimized through various time-and-motion analyses (think of those manufacturing assembly lines). But performance gains from that approach have recently been leveling off, as companies wring the last bits of efficiencies from mechanistic automation.

Now, to continue exploiting the full potential of AI technologies, many leading companies have begun to embrace a new view of business processes as more fluid and adaptive. In essence, they are moving beyond rigid assembly lines toward the idea of organic

teams that partner humans with advanced AI systems. This collaboration between workers and smart machines is leading to the reinvention of many traditional processes. As BMW and Mercedes-Benz have experienced, rigid assembly lines are giving way to flexible teams of employees working closely alongside robots. Moreover, these novel types of teams can continuously adapt on the fly to new data and market conditions. They are enabling companies to actually *reimagine* various work processes.

## The Third Wave

The key to understanding AI's current and future impact is its transformation of *business processes*.

A widespread misconception is that AI systems, including advanced robotics and digital bots, will gradually replace humans in one industry after another. Self-driving vehicles, for example, will one day replace taxi, delivery, and truck drivers. That may be true for certain jobs, but what we've found in our research is that, although AI can be deployed to automate certain functions, the technology's greater power is in complementing and *augmenting* human capabilities. In claims processing, for example, AI isn't replacing the need for humans; instead, it's doing the tedious grunt work, collecting data and doing a preliminary analysis, freeing human claims processors to focus on resolving complex cases. In essence, machines are doing what they do best: performing repetitive tasks, analyzing huge data sets, and handling routine cases. And humans are doing what they do best: resolving ambiguous information, exercising judgment in difficult cases, and dealing with dissatisfied customers. This kind of emerging symbiosis between man and machine is unlocking what we have called *the third wave of business transformation*.

To see how we got here, it helps to understand some historical context. The first wave of business transformation involved *standardized* processes. This era was ushered in by Henry Ford, who deconstructed the manufacture of automobiles so that they could be made on an assembly line. Each of those steps in that overall process could then be measured, optimized, and standardized to achieve considerable gains in efficiencies.

The second wave consisted of *automated* processes. This era emerged in the 1970s and peaked in the 1990s with the business process reengineering (BPR) movement, thanks to advances in information technology (IT): desktop computers, large databases, and software that automated various back-office tasks. Among other companies, retailers like Walmart rode this wave to become worldwide powerhouses. Other firms were able to reinvent themselves. UPS, for example, transformed itself from a package-delivery service to a global logistics company.

Now, the third wave involves *adaptive* processes. This third era, which builds on the previous two, will be more dramatic than the earlier revolutions enabled by assembly lines and digital computers, and will usher in entirely new, innovative ways of doing business. As we will see throughout this book, the leading firms in many industries are now reimagining their processes to be more flexible, faster, and adaptable to the behaviors, preferences, and needs of their workers at a given moment. This adaptive capability is being driven by real-time data rather than by an a priori sequence of steps. The paradox is that although these processes are not standardized or routine, they can repeatedly deliver better outcomes. In fact, leading organizations have been able to bring to market individualized products and services (as opposed to the mass-produced goods of the past), as well as deliver profitable outcomes.

## Think Like Waze

To help illustrate the profound difference between the old process thinking and the new, consider the history of GPS navigation. The first online maps were largely just a digital version of their paper counterparts. But soon, GPS navigation devices changed how we used maps, giving us directions after entering a destination. Even that, though, was still a fairly static process. Now, mobile map apps like Waze are taking advantage of real-time user data—about drivers' locations and speeds as well as crowd-sourced information about traffic jams, accidents, and other obstructions to create the perfect map in real time. All that data enables the system to update directions en route so that, if necessary, it can reroute drivers midcourse to minimize any possible delays. Whereas the old approach with GPS simply digitized a static paper-map route, Waze has combined AI algorithms and real-time data to create living, dynamic, optimized maps that can get people to their destinations as quickly as possible. Business approaches that use AI merely for automating existing static processes are like the early GPS navigation devices, whereas the current era of symbiotic collaborations between humans and machines are like Waze in that those traditional processes are being completely reimagined.

## Filling the "Missing Middle"

Unfortunately, popular culture has long promoted a man-versus-machine view—think of movies such as *2001: A Space Odyssey* and the *Terminator* series. The idea of intelligent machines as a potential threat to mankind has a long history and has resulted in many executives adopting a somewhat similar perspective, thinking exclusively of machines as threatening to replace

humans. But that view is not only woefully misguided; it's also perniciously shortsighted.

The simple truth is that machines are not taking over the world, nor are they obviating the need for humans in the workplace. In this current era of business process transformation, AI systems are not wholesale replacing us; rather, they are amplifying our skills and collaborating with us to achieve productivity gains that have previously not been possible.

As you shall see in this book, the third wave has created a huge, dynamic, and diverse space in which humans and machines collaborate to attain orders-of-magnitude increases in business performance. We call this the "missing middle"—"missing" because almost no one talks about it, and only a small fraction of companies are working to fill this crucial gap (see figure I-1).

**FIGURE I-1**

**The missing middle**

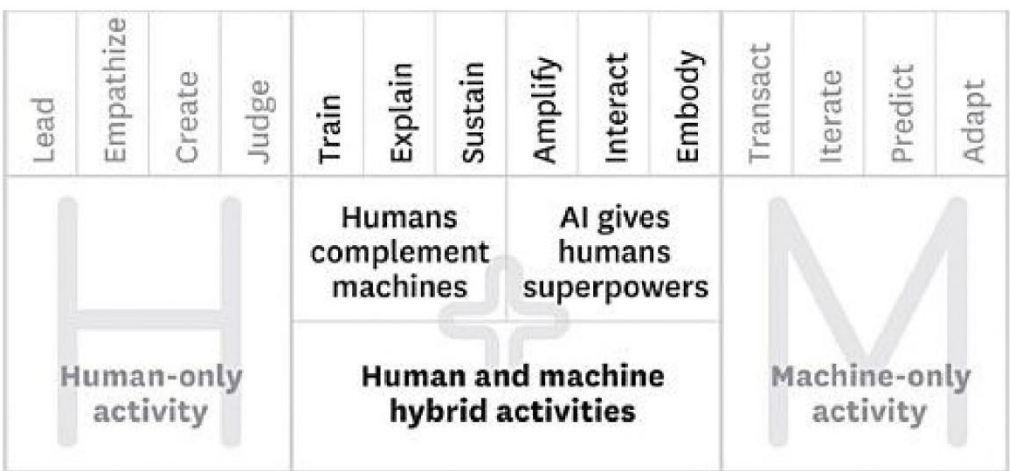


In the missing middle, humans work with smart machines to exploit what each party does best. Humans, for example, are

needed to develop, train, and manage various AI applications. In doing so, they are *enabling* those systems to function as true collaborative partners. For their part, machines in the missing middle are helping people to punch above their weight, providing them with superhuman capabilities, such as the ability to process and analyze copious amounts of data from myriad sources in real time. Machines are *augmenting* human capabilities.

In the missing middle, humans and machines aren't adversaries, fighting for each other's jobs. Instead, they are symbiotic partners, each pushing the other to higher levels of performance. Moreover, in the missing middle, companies can reimagine their business processes to take advantage of collaborative teams of humans working alongside machines. It's not just digital companies that are exploiting the missing middle. Rio Tinto, the global mining conglomerate, is a case in point. The company is using AI to manage its vast fleet of machinery—autonomous drills, excavators, earth movers, and so on—from a central control facility. This has freed human operators from working in dangerous mining conditions and also enabled Rio Tinto's team of data analysts to analyze information from sensors on the remote equipment to uncover valuable insights for helping the firm manage its fleet more efficiently and safely.[2]

## Separating Winners from Losers—and What's in This Book

As we mentioned earlier, in the current era of adaptive processes, the fundamental rules by which organizations are run are being rewritten daily. As leaders and managers in all kinds of enterprises begin to reimagine their business processes and rethink the basic relationship between employees and machines, they must

understand those rules and execute on them. That's why we wrote this book: to give people who are thinking about their organization, their team, or their career the knowledge they need that will separate winners from losers in the new age of AI.

In part one, we show and explain the current state of AI in companies. We begin on the shop floor, and in subsequent chapters, illustrate how organizations are currently deploying AI in different functions—back office, R&D, marketing, and sales. A key lesson here is that companies can't expect to benefit from human-machine collaborations without first laying the proper groundwork. Again, those companies that are using machines merely to replace humans will eventually stall, whereas those that think of innovative ways for machines to augment humans will become the leaders of their industries.

Chapter 1 describes how human-machine teams are transforming the factory floor, not just at BMW and Mercedes-Benz but also at many other large manufacturers. General Electric, for example, has been building "digital twins" of its deployed products, like the turbine blades on a jet engine. The company bases these virtual models on the current conditions of real machinery, enabling it to improve operations as well as predict breakdowns before they occur, thus fundamentally changing how it maintains commercial equipment.

The focus is back-office operations in chapter 2. Here, AI technology can help filter and analyze streams of information from a variety of sources and enable the automation of tedious, repetitive tasks as well as the augmentation of human skills and expertise. At a Canadian insurance and financial service provider, for example, an AI system processes unstructured financial data from news articles, reports, and emails to make specific recommendations, and the system can be trained to extract insights tailored to an individual analyst's preferences.

Chapter 3 describes how companies are using AI in research and

development. In each major step of the R&D process—observation, hypothesis generation, experiment design, and results analysis—AI technologies can lead to increased efficiencies and markedly improved outcomes. At GNS Healthcare, sophisticated machine-learning software finds patterns in patient medical health records and then automatically churns out hypotheses straight from that data. The system took just three months to recreate the results of a two-year study investigating adverse reactions between drug combinations in seniors on Medicare.

We turn to the marketing and sales function and describe how the impact of AI has been just as great—if not more so—as in other functions in chapter 4. Machine-learning technologies like Amazon's Alexa, Apple's Siri, and Microsoft's Cortana are increasingly becoming the digital embodiment of those companies' well-known brands. In other words, AI has become the brand.

In part two, we explore the "missing middle" and provide an executive guide for overhauling and "reimagining" the traditional notions of work. To exploit the full power of AI, companies must fill that gap by considering new employee roles, by establishing novel types of working relationships between humans and machines, by changing traditional concepts of management, and by overhauling their very concept of work itself.

Chapter 5 describes how machine learning, when integrated into processes, will lead to a variety of brand-new jobs. Specifically, employees will be needed to design and train algorithms, to explain the algorithms used, and to do so in a way that sustains the algorithms within a process. One such new role is that of *machine relations managers*, who will be similar to human resources managers except they will oversee AI systems instead of human workers. They will be responsible for regularly conducting performance reviews of a company's AI systems. They will promote those systems that perform well, replicating variants and deploying them to other parts of the organization. Those systems

with poor performance will be demoted and possibly decommissioned.

In chapter 6, we describe how people are achieving huge performance boosts by working with AI technologies that dramatically improve their human capabilities; they amplify, interact, and embody new human potential. (In a sense, this chapter is the flip side of chapter 5, in which we discussed how humans are helping machines to extend and amplify their capabilities.) These new types of human-machine relationships are helping people to "punch above their weight" by offloading tedious tasks and by enabling them to perform their work faster and more effectively through the expert guidance, advice, and support from AI systems.

Chapter 7 takes a hard look at managerial challenges introduced by AI that require different, new responses from management and leadership. A huge question here is, what steps must management take to facilitate reimagining processes? Specifically, management must support five crucial activities, including an emphasis on trial-and-error *experimentation,* building a *data supply chain* for AI, and others.

Finally, we explore the future of work itself in chapter 8. Specifically, as human-machine collaborations become increasingly prevalent, companies will need to hire for and develop eight new "fusion skills": *intelligent interrogation* (knowing how best to ask an AI agent questions, across levels of abstraction, to get the insights you need), *bot-based empowerment* (collaborating with intelligent agents to punch above your weight at work), *reciprocal apprenticing* (teaching AI agents new skills while also undergoing on-the-job training to work well within AI-enhanced processes), *holistic melding* (developing mental models of AI agents that improve collaborative outcomes), *rehumanizing time* (reimagining business processes to amplify the time available for distinctly human tasks and learning), *responsible normalizing*

(shaping the purpose and perception of human-machine collaborations as it relates to individuals, businesses, and society), *judgment integration* (choosing a course of action amid machine uncertainty), and *relentless reimagining* (thinking of novel ways to overhaul work, processes, and business models to obtain exponential increases in improvement).

## Five Crucial Principles

In our research, we have found that the leading companies in various industries—9 percent of our surveyed sample of more than fifteen hundred practitioners—are already riding the third wave. They have maximized automation and are now developing the next generation of processes and skills to capitalize on human-machine collaborations. They are thinking as Waze does, re-imagining processes as living and adaptive by using human and crowd inputs and real-time data. They are moving beyond the traditional thinking of simply digitalizing the old static maps.

How have these leading firms been accomplishing that? In our work, we found that they have succeeded by adopting five crucial principles having to do with their organizational mindset, experimentation, leadership, data, and skills (MELDS).

- *Mindset: assuming a radically different approach toward business by reimagining work around the missing middle, wherein people improve AI and, in turn, smart machines give humans superpowers.* Previously, the focus was on using machines to automate specific steps in a given work flow. Now, the potential collaboration between humans and machines is leading to the reinvention of many traditional processes. Rigid assembly lines are giving way to flexible teams of augmented humans and smart machines. Moreover, these teams are continuously

adapting on the fly to new data and to different human contributions. They are truly organic, with the associated work processes akin to living and breathing organisms. We foresee that AI technologies will be critical in helping companies bring work closer to the markets they serve, improving responsiveness to consumer demand. To achieve that, however, executives must embrace a distinctive, action-oriented mindset to reimagine their operations, as we will discuss throughout the book. That said, executives must also understand that they need to lay a foundation first instead of rushing to fill the missing middle. Specifically, they initially should focus on developing the full potential of their employees by applying automation to routine work; then they can proceed to concentrate on human-machine collaborations.

- *Experimentation: actively observing for spots in processes to test AI and to learn and scale a reimagined process from the perspective of the missing middle.* The age of standard business processes is coming to an end, and companies will no longer be able to rely on a strategy of replicating best-in-class practices from leading firms. And this is why experimentation is crucial. Executives must continually conduct tests to derive business processes that will work best for their unique set of conditions. A large part of that effort will require trial and error to determine what work should be done by humans, and what work would best be done by a collaboration between humans and machine (the missing middle).

- *Leadership: making a commitment to the responsible use of AI from the start.* Executives must always consider the ethical, moral, and legal implications of the AI technologies they deploy, and the systems must generate explainable results, promote algorithmic accountability, and eliminate biases. Firms also

need to pay close attention to ensure that the employees working with AI systems don't lose a sense of agency and that those individuals develop a heightened sense of empowerment in decision making. In addition, companies must provide the employee training and retraining required so that people will be prepared and ready to assume any new roles in the missing middle. In fact, investing in people must be a core part of any company's AI strategy.

- *Data: building a data "supply chain" to fuel intelligent systems.* AI requires extensive amounts of data, both in volume and variety. This includes "exhaust data"—data created as a byproduct of another process (for example, cookies from customer web browsing). Accumulating and preparing such information for use is one of the biggest challenges for organizations that deploy AI systems. Moreover, an organization's data should be able to flow freely, unencumbered by departmental silos. A company can then take full advantage of that information by applying it and other data to support, sustain, and improve AI and human performance in the missing middle.

- *Skills: actively developing the eight "fusion skills" necessary for reimagining processes in the missing middle.* The growing power of AI is fundamentally transforming the human-machine relationship. In the second wave, machines were generally being used to replace humans—think of how automation has decimated the ranks of factory workers, administrative assistants, bookkeepers, bank tellers, travel agents, and so on. But humans are needed now more than ever in the third wave. Humans are taking center stage in this current era of business process improvement. Specifically, the era of adaptive processes requires humans in the loop, not only to design,

develop, and train AI systems, but also to collaborate with them to fill the missing middle and achieve step-level increases in performance.

As you will see, the MELDS framework guides nearly all of the practical aspects of the book, and we will return to it many times throughout. In particular, chapter 7 will focus intently on "MELD," while chapter 8 will dive deeply into the "S".

---

The AI revolution is not coming; it is already here, and it is about reimagining your processes, across all functions of the company, to get the most benefit from this technology's power to augment human capability. This book is your road map for understanding and navigating the new landscape. Let's get started.

PART ONE

# Imagining a Future of Humans + Machines . . . Today

# 1

# The Self-Aware Factory Floor

*AI in Production, Supply Chain, and Distribution*

For centuries, factories have been the paragon of automation. And the people who work in factories, as a result, have often been measured by the standards of machines. So it's no surprise that the relationship between people and machines in industry has been fraught, with human workers feeling as if they've been dealt a losing hand. There's ample reason for that feeling. Since 2000, the United States has lost five million manufacturing jobs, roughly half of that through efficiency gains and automation.[1]

But things aren't as clear cut as they might first seem. As we discussed in the previous chapter, the second wave of business transformation was all about *automating existing processes*, and it was during this era that many humans were replaced by machines. In contrast, the third wave relies on *adaptive* processes that are reimagined from scratch, and the goal here is humans + machines. In this current era, thanks to AI, factories are acquiring a little

more humanity: jobs on manufacturing lines, for instance, have changed in nature and are increasing in number. And it's not just manufacturing. AI is boosting the value of engineers and managers, too. The emergence of AI is also creating brand-new roles and new opportunities for people up and down the industrial value chain.

In this era of reimagining processes with AI, the great irony is that some of the most-automated environments—the factory and other industrial settings—are experiencing a renaissance of human labor. From the assembly-line worker and maintenance specialist to robot engineer and operations manager, AI is rejiggering the concept of what it means to work in an industrial setting. In many cases, AI is freeing up time, creativity, and human capital, essentially letting people work more like humans and less like robots. One implication of the fact that people can work differently and better with the help of AI is that companies are gaining efficiencies and saving money. But perhaps more importantly in the long term is that companies are also starting to rethink their business processes. And as they do, they uncover the need for new kinds of jobs for people, and wholly new ways of doing business, which is our focus in part two of this book.

But let's not get ahead of ourselves. This is a complex journey. (For some historical perspective, see the sidebar "A Brief History of AI" at the end of this chapter.) Before we rewrite business processes, job descriptions, and business models, we need to answer these questions: what tasks do humans do best, and what do machines do best? There are jobs and tasks that will continue to shift to robots based on their comparative advantages in handling repetition and data processing power. But as we'll see, the transfer of jobs is not simply one way. In this chapter, we survey a number of companies that have already answered the human-machine question for manufacturing, maintenance, warehouses, and agriculture. These early movers have put people and AI-enhanced

machines into play, in the roles that they're best suited to, and they are reaping the benefits.

## The Arm That Learns

The third shift in a Tokyo factory is showtime for an emerging class of robotic arms that can learn new tricks overnight. Coupled with a camera and machine-learning software, these hinged and swiveling appendages can, on their own, figure out the most efficient ways to pick up parts and pieces and set them down somewhere else. No explicit programming is necessary.[2]

Robotic arms are used in factories to apply hot glue to widgets, to install windshields, and to smooth jagged metal edges, among other tasks. But, traditionally, engineers have preprogrammed them. Then, when robots' jobs change, engineers must reprogram them. In contrast, the new robotic arms, developed by Fanuc in partnership with software-maker Preferred Networks (both based in Japan), adapt on their own. They do it with an AI technique called deep reinforcement learning, in which the robot is given a picture of the successful outcome and then uses trial and error to figure out its own solution.

According to Shohei Hido, chief research officer at Preferred Networks, the arms take eight hours to become at least 90 percent accurate for this kind of task. This is virtually the same time and accuracy of an expert programming it, but because the arm is now autodidactic, the human expert is now freed to do other more complex tasks, especially those that require human judgment. What's more, once one robot learns a task, it can share its knowledge with other robots in the network. This means that eight arms working together for one hour can learn as much as one working on a problem for eight hours. Hido, who calls this process

"distributed learning," says, "You can imagine hundreds of factory robots sharing information."[3]

Now, imagine people working alongside these robots. Mechanical arms are great for highly repetitive tasks and heavy lifting, but in any factory, there will always be a subset of tasks that are just too complex to hand off to a robot—tasks like positioning numerous small wires or handling awkward or dynamic objects. A human is still needed in the loop.

So how do robot arms and humans work together? Historically, not so well. Robots, with their fast, decisive movements have been helpful and efficient, but also dangerous to people. They've often been cordoned off behind protective barriers. But that standard segregation is beginning to change. So-called collaborative robots from companies like Rethink Robotics, founded by robotics and AI pioneer Rodney Brooks, come equipped with sensors that allow them to recognize a range of objects and avoid knocking people around. When robots aren't so clumsy, they can work well with people. Factories that use Rethink Robotics products often divide the work between the robot and the human worker, working side by side, performing tasks best suited to their abilities. (For further examples of embodied AI, see the sidebar "AI in the Factory.")

## *AI in the Factory*

For a century, factory floors have been at the leading edge in robotic automation. From conveyor belts to robotic arms to AI-infused operations systems, the factory is getting smarter every day.

- Hitachi is using AI to analyze big data and workers'

routines to inform its robots, which deliver instructions to employees to meet real-time fluctuating demand and on-site kaizen objectives. In a pilot, the company observed an 8 percent productivity improvement in logistics tasks.[a]

- At Siemens, armies of spider-styled 3-D printed robots use AI to communicate and collaborate to build things in the company's Princeton, New Jersey, lab. Each bot is equipped with vision sensors and laser scanners. In aggregate, they join forces to manufacture on the go.[b]

- At Inertia Switch, robotic intelligence and sensor fusion enable robot-human collaboration. The manufacturing firm uses Universal Robotics' robots, which can learn tasks on the go and can flexibly move between tasks, making them handy helpers to humans on the factory floor.[c]

a. Dave Gershgorn, "Hitachi Hires Artificially Intelligent Bosses for Their Warehouses," *Popular Science*, September 8, 2015, www.popsci.com/hitachi-hires-artificial-intelligence-bosses-for-their-warehouses.

b. Mike Murphy, "Siemens is building a swarm of robot spiders to 3D-print objects together," *Quartz*, April 29, 2016, https://qz.com/672708/siemens-is-building-a-swarm-of-robot-spiders-to-3d-print-objects-together/.

c. Robotiq, "Inertia Switch Case Study – Robotiq 2-Finger Adaptive Gripper – ROBOTIQ," YouTube video, 1:32 minutes, posted July 28, 2014, https://www.youtube.com/watch?v=iJftrfiGyfs.

## Kindler, Gentler Robots

During the second AI "winter," Rodney Brooks challenged one of the fundamental ideas that had driven previous AI research—namely, the reliance on predetermined symbols and relationships between symbols to help computers make sense of the world (see the sidebar "Two AI Winters"). He claimed a much more robust approach: instead of cataloging the world in advance and representing it with symbols, why not survey it with sensors instead? "The world is its own best model," he wrote in a famous 1990 paper called "Elephants Don't Play Chess." (Brooks would later found iRobot, maker of the robotic vacuum Roomba, as well as Rethink Robotics. To date, iRobot has deployed the largest fleet of autonomous robots in the world: between 2002 and 2013, more than 10 million were sold.[4])

Now, Brooks's AI philosophy is alive and well in both research and industry. Rethink Robotics, in particular, demonstrates the power of an arm equipped with embedded sensors and algorithms for motion control that allow it to "feel" its way and adjust as it goes. The arm features elastic actuators and back-drivable joints, which means it can flex on contact to absorb energy. Consequently, even if it does knock into something (or someone), it wouldn't have nearly the force of a traditional robotic arm.

What's possible when robot arms can learn on their own, as with Fanuc's products? Or when an arm operates in a kinder, gentler way, as with Rethink's products?

On the assembly line, workers can collaborate with a self-aware robot arm. Say a worker is putting together a car and needs to put an interior panel on one of its doors. The robot can lift the panel and position it into place, while the worker performs fine adjustments and fastening without fear that a clunky machine will clock him in the head. AI helps both robots and people play to

their strengths, and in the process, the assembly line changes shape.

## *Two AI Winters*

The path to human-machine collaboration—a hallmark of the third wave of process improvement—was far from smooth. AI was initially greeted with considerable enthusiasm, only to be followed by results that didn't live up to the initial hype, and then more progress, leading to a second wave of hype then disappointment. Those down periods have become known as AI's two "winters."

The field of AI began in the 1950s, and during the decades that followed any research progress came only in fits and starts. By the 1970s, funding had dissipated so much that the era became known as the first AI winter. Then, during a few years in the 1980s, some researchers made progress in so-called expert systems—computer systems loaded with code that allowed a machine to perform a kind of rudimentary reasoning using "if-then" rules rather than following a strict, predetermined algorithm. But the desktop computer revolution was under way, and attention was diverted toward personal computers as they became increasingly affordable and practical for the average person. Again, money for AI dried up, and the second AI winter descended. It wasn't until the 2000s that AI began to draw major investment again.

One way that assembly lines can be reconfigured is through AI

itself. Engineers at the Fraunhofer Institute of Material Flow and Logistics (IML) have been testing embedded sensors to create self-adapting assembly lines in car plants. Essentially, the line itself can modify the steps in its process to fit the demands of various features and add-ons for highly customizable cars. Thus, instead of engineers designing an assembly line to make one kind of car at a time, these lines can adapt as needed. What's more, says Andreas Nettsträter, who coordinates strategic initiatives at IML, "If one station has a failure or is broken down, the others could also do what should have been done in this assembly station."[5]

This means that assembly-line workers are doing tasks that are less robotic (saving those tasks for the robot) and more nuanced, while process engineers don't need to reconfigure the line every time there's a change in demand or breakdown of a machine. They can spend their time working on more creative tasks to eke out further efficiencies, for instance.

## Follow the Data

What starts with smart arms can extend to an entire factory line and beyond: AI-enabled processes throughout manufacturing and industrial environments are freeing up human potential in a variety of contexts. Maintenance work, for instance, has been forever upended by AI. Sophisticated AI systems predict machine breakdowns *before* they occur, which means that maintenance workers can spend less time running routine checks and diagnostics and more time fixing a company's assets. (For other applications, see the sidebars "AI for Faster Machine Onboarding" and "AI in the Field—Unmanned Vehicles.")

## *AI for Faster Machine Onboarding*

Sight Machine, a startup in San Francisco, uses machine-learning analytics to enable its customers to reduce down-time when adding new machines to a factory floor. In one case, the technology was able to reduce a customer's downtime, usually inherent in breaking in new robotic systems, by 50 percent. In addition, the net gain was a 25 percent increase in performance when all the assets were up and running. Furthermore, not only does the technology help improve factory efficiency, but it also allows engineers and maintenance workers to spend more time tackling other, higher-value tasks.[a]

a. "Jump Capital, GE Ventures, and Two Roads Join $13.5 Million Series B Investment in Sight Machine," Sight Machine, March 22, 2016, http://sightmachine.com/resources/analytics-news-and-press/jump-capital-ge-ventures-and-two-roads-join-13-5-million-series-b-investment-in-sight-machine/.

General Electric, for example, keeps track of its products in the field using its AI-enabled system called Predix. The system relies on a "digital twin" concept in which all assets of a factory and beyond—from bolt to conveyor belt to turbine blade—are monitored and modeled on a computer. Predix collects and manages a lot of data, and it can be put to use to reimagine business processes in three fundamental ways:

## *AI in the Field—Unmanned Vehicles*

Acting as an extra pair of eyes in the sky or under the sea, drones, powered by AI, can keep human workers out of harm's way by allowing teams to remotely explore potentially dangerous terrain.

- Fortescue Metals Group, which operates the Cloudbreak iron ore mine, uses drones there to gather spatial information. Its fleet of flying robots has significantly reduced the safety risk to operators in high-risk areas.[a]

- At BHP Billiton Ltd., unmanned aerial vehicles equipped with infrared sensors and telescopic zoom can flag problems with safety beams and roads under construction. They also check blast zones to make sure they're clear of people before detonation.[b]

- Boeing's "Echo Voyager" is an unmanned deep-sea robot that inspects underwater infrastructures, takes water samples, creates maps of the ocean floor, and aids in oil and gas exploration.[c]

a. Allie Coyne, "Fortescue deploys survey drones at Cloudbreak mine," *IT News*, August 31, 2015, https://www.itnews.com.au/news/fortescue-deploys-survey-drones-at-cloudbreak-mine-408550.

b. Rhiannon Hoyle, "Drones, Robots Offer Vision of Mining's Future," *Wall Street Journal*, July 28, 2016, http://www.wsj.com/articles/drones-robots-offer-vision-of-minings-future-1469757666.

c. "Boeing's Monstrous Underwater Robot Can Wander the Ocean for

6 Months," *Wired*, March 21, 2016, https://www.wired.com/2016/03/boeings-monstrous-underwater-robot-can-wander-ocean-6-months/.

- *Reimagined maintenance*. GE keeps stats from the point of installation from a large number of customers and uses machine-learning technology to predict when certain parts might fail based on their current conditions. Previously, maintenance professionals would have a fixed schedule for checking or replacing certain parts—something like switching out a car's spark plugs every seventy-five thousand miles; checks and replacements can now happen on an as-needed basis. AI-enhanced prediction saves money and time, and it has the potential to keep maintenance workers more engaged in their jobs.[6]

- *Reimagined product development.* More data leads to improved R&D. GE is now attaching sensors to some of the hottest parts of turbines so that it can monitor physical changes. The sensors literally burn off at operating temperatures, but by that time, data from the turbine's cold-to-hot transition has been collected. This information can then help engineers better understand the thermodynamics of the materials used in the turbines and potentially improve that product's operating conditions. Engineers, thanks to AI, now have more data than ever to understand the operations of their systems.[7]

- *Reimagined operations.* The field data collected also enables GE to build digital twins of its deployed products, like its jet engines. Engineers can then test virtual flights in which the plane experiences cold, heat, dust, rain, and even a flock of birds.[8] The company is also monitoring ten thousand wind

turbines, and their digital twins are helping the turbines to adapt in real time. One valuable insight from an analysis of that data is that, depending on the direction of the wind, it might be best to have the leading turbine run slower than engineers might expect. When the front turbine absorbs less energy, the ones behind it can operate at close to their optimal levels, increasing energy generation overall. This application shows that digital twin technology can be applied beyond a single product to holistically optimize an entire wind farm's activity. According to GE, digital twins could increase wind-farm output by 20 percent and provide $100 million of value over the lifetime of a 100-megawatt wind farm.[9]

All three of these uses of Predix are freeing up human workers to do less routine work and more engaging work. The maintenance worker gets to spend more time on tricky fixes and less time on routine monitoring. The engineer has more data with which to see system successes and failures, which can lead to more creative solutions down the line. And finally, the digital twin models are providing an experimental space that's vastly larger than that in which most engineers play. These models allow engineers to be more creative in the questions they ask and allow for the emergence of previously hidden inefficiencies—with the potential for significant savings of time and money.

## The Warehouse That's Packed

Today, it's not uncommon to walk into a modern warehouse or distribution center and see robots rolling along the floor. (For a small sampling of these smarter supply chain and warehouse bots, see the sidebar "AI in Warehouse and Logistics.")

## *AI in Warehouse and Logistics*

Artificial intelligence is taking on the task of warehouse navigation and inventory and changing the way people think about warehouse design.

- When Amazon acquired Kiva Robots in 2012, it signaled that the mobile bots zipping around Amazon's warehouses were a key to their fulfillment advantage. Not only do the robots help lift and stack plastic bins filled with different products, they also do the legwork of autonomously transporting items around the facility to human "pickers," who then select the right products to fulfill different orders. Thanks to such increased efficiencies, the company has been able to offer same-day shipping for customers.[a]

- L'Oreal uses radio-frequency identification (RFID) technology and machine learning to help prevent forklift accidents in the company's warehouse in Italy. The tracking system warns forklift operators and pedestrians about other nearby vehicles, cutting down on collisions.[b]

a. Nick Wingfield, "As Amazon Pushes Forward with Robots, Workers Find New Roles," *New York Times*, September 10, 2017, https://www.nytimes.com/2017/09/10/technology/amazon-robots-workers.html.

b. Claire Swedberg, "L'Oréal Italia Prevents Warehouse Collisions via RTLS," *RFID Journal*, August 18, 2014, http://www.rfidjournal.com/articles/view?12083/2.

These robots are often sophisticated enough to see where they're going and understand what they're doing. But they have their limitations. Say a case of Cheerios is damaged, making it bulkier on one side. Most robots can't adapt. They'd need to skip it and move along to the next case. But robots from a company called Symbotic have the advantage of machine vision algorithms that allow them to assess an oddly shaped package and pick it up anyway. Even better, the robots can quickly measure shelf space to confirm that a box will fit. If it won't, the robot alerts a central control system, which automatically redirects that box to a shelf where it will fit. The bots zip around the warehouse floor at twenty-five miles per hour, carrying, sensing, and adapting as they go.

The difference between a traditional warehouse and one with Symbotic's robots is stark. Usually, trucks unload pallets of products at the dock; there's an area where pallets are stored until people can unpack them, and conveyor belts move cases of goods to various parts of the warehouse. But because the Symbotic robots immediately remove products from pallets and put them on shelves, there's no need to reserve space for storing pallets. And there's no need for conveyor belts either. Thus, a Symbotic-equipped warehouse can reclaim space for shelves. The ramifications are significant: in the best-case scenarios, either a warehouse can store twice as many goods as before, says Joe Caracappa, Symbotic vice president of business development, or it can operate in an area about half the size. Moreover, smaller warehouses can more easily fit into existing neighborhoods, and perishable items can be stored closer to their point of sale.

Because the only human interaction with the goods stored at a warehouse is when they're loaded on and off the trucks, we must ask the question: What happens to the human workers at the warehouse? Caracappa says Symbotic currently retrains many of them. Those who performed maintenance on conveyor belts, for

instance, are trained to fix robots. And there are new roles, too. Caracappa says system operators monitor the entire flow of robots. "Those roles are typically not in the warehouse before automation comes in," he explains, "but we'll hire them locally and the client will be part of the process."[10] (In part two of this book, we will explore these new types of jobs in depth when we discuss the missing middle in detail.)

## Supply Chains That Think

Smarter warehouses are just the beginning. AI technologies are now enabling entire supply chains to become increasingly intelligent, similar to the kinds of advances they've enabled on the factory floor. Of course, companies want to minimize any upstream disruptions to their supply chains, which can come from a number of sources—manufacturing quality problems at a supplier, political instability of a region, labor strikes, adverse weather events, and so on. To that end, AI can help collect and analyze data about suppliers, provide a better understanding of the variables in a supply chain, anticipate future scenarios, and so on. And firms also want to minimize downstream uncertainties. Here, AI can enable companies to optimize their demand planning, forecast more accurately, and better control their inventories. The result is more-agile supply chains capable of anticipating and dealing with the ups and downs of dynamic business conditions.

Consider just one part of the process: demand planning. Getting demand planning right is a pain point for many companies, but the use of neural networks, machine-learning algorithms, and other AI technologies can help lessen that pain. A leading health-food company, for example, leveraged machine-learning capabilities to analyze its demand variations and trends during promotions. The

analysis led to a reliable, detailed model that could highlight the expected results from a trade promotion. The gains included a 20 percent reduction in forecast error and a 30 percent reduction in lost sales.

Those improvements are the types sought by consumer goods giant Procter & Gamble, whose CEO recently stated his goal of cutting supply-chain costs by a whopping $1 billion a year. Part of those savings will come from near-term efforts like the use of AI and the internet of things (IoT) technologies to automate warehouses and distribution centers. And other savings will come from longer-term projects, including the customized automation of product deliveries of up to seven thousand different stock-keeping units. Whether these and other initiatives will enable P&G to save the company $1 billion annually in supply-chain costs remains to be seen, but it's safe to say that AI will be playing a significant role in those efforts.

## The Farms That Feed

AI technology is not only having a large impact on supply chains and the manufacturing of consumer goods and industrial machinery, but is also playing a big role in the production of food. The need for improved efficiency in the agricultural industry is huge. According to various statistics, 795 million people don't have enough food and, to keep pace with population growth, more food will be needed in the next fifty years than has been produced in the past ten thousand years combined. Both fresh water and arable soil are resources that have historically been difficult to acquire or maintain for agriculture. Precision agriculture—which leverages AI and fine-grain data about the state of crops—promises to significantly improve yield, reduce the waste of resources like

water and fertilizer, and increase overall efficiency.

To be effective, precision agriculture requires a vast network of diverse IoT sensors to collect granular data. This information might include aerial images captured by satellites or drones (to detect crop distress before it's visible at ground level), environmental sensors in the field (to monitor the chemical composition of the soil, for instance), sensors mounted on farm equipment, weather forecast data, and soil databases.

To help make sense of these various data streams, Accenture has developed a new service—the Precision Agriculture Service—that deploys AI to enable better decision making with respect to pest control, fertilizer usage, and so on. The idea is to process IoT-sensor data with a machine-learning engine to provide feedback that can then be used in one of two ways. It can either go directly to a farmer, who can then implement a solution. Or it can be routed directly to a farm's digital work-management system that will then deploy the recommendations automatically. In the system, a feedback loop that incorporates up-to-date sensor data and real-time analytics can help establish a kind of self-healing farm. Farmers can be part of the loop when they approve the system's recommendations, but over time, as the system becomes more reliable, they can spend their time managing other tasks that aren't so easily automated.

AI is also enabling entirely new agricultural models, such as the "vertical farm," in which plants can be grown in thirty-foot-high stacks of trays in urban settings such as city warehouses. At one such facility in Newark, New Jersey, run by AeroFarms, data is continuously collected on temperature, humidity, carbon-dioxide levels, and other variables, and machine-learning software analyzes that information in real time to grow the crops (including kale, arugula, and mizuna plants) as efficiently as possible. According to the company, the Newark facility is expected to use 95 percent less water and 50 percent less fertilizer than traditional

farms and, because the crops are grown indoors, pesticides aren't needed. AeroFarms is predicting that 2 million pounds of produce can be grown annually at the Newark vertical farm, only about fifteen miles from Manhattan.[11]

Precision agriculture is not yet widespread, but some of its technologies—the analysis of satellite data, for instance—have been used for years. The difference now is the pervasiveness of the IoT, which allows sensor data to talk to apps, and apps to talk to machine-learning-enabled systems. The ultimate goal with precision agriculture is that disparate systems can come together to produce recommendations that farmers can then act on in real time. The result is agricultural processes that produce less waste and higher yields. It's no wonder that precision-agriculture services are expected to grow to $4.55 billion by 2020.[12] As the technology becomes more widely used, the land will benefit, the farmer will benefit, and the hundreds of millions of people who need access to healthy affordable food will benefit. (See the sidebar "AI for Good: Akshaya Patra.")

## *AI for Good: Akshaya Patra*

Akshaya Patra, an India nonprofit with the vision that "no child in India shall be deprived of education because of hunger," combines the power of AI with blockchain (a digital, decentralized, public ledger) and IoT technologies. To achieve its vision, the company's midday meal program provides one wholesome lunchtime meal to keep children sufficiently motivated and nourished to pursue their education. Since 2000, when it began by feeding 1,500 children, its operations have expanded to 1.6 million

children per year in 2017; it commemorated its two-billionth meal served in 2016. Thus far, the nonprofit has demonstrated a 20 percent efficiency improvement in selected kitchens. Now feedback is digitized where once it was manually input, and blockchain is driving efficiencies in audit, attendance recording, and invoice processing. AI is used to accurately forecast demand, and IoT sensors monitor and sequence cooking processes to minimize waste and ensure consistent food quality. AI in combination with these other technologies will help Akshaya Patra expand its operations efficiently, meaning more children fed and kept in school.[a]

a. "About Us," Akshaya Patra, https://www.akshayapatra.org/about-us, accessed October 23, 2017.

## The Third Wave in Manufacturing

In this chapter, we're beginning to see how artificial intelligence can change the nature of business processes. Factories and industrial settings will continue to be highly automated environments for a variety of reasons; safety and efficiency are two primary drivers. And while new automation technologies will replace some human workers, there's still plenty of space for people, as long as executives look beyond jobs displacement and begin to think differently about work. This is the *leadership* part of our MELDS framework, which we detailed in the introduction, and it calls for executives to focus on reimagining processes and new roles for employees working in the missing middle (as we'll discuss in detail in part two). And, as we've seen in this chapter, some skills are becoming more in demand, and entirely new categories

of skills will be required. For example, as we will see in chapter 8, GE and the buyers of its equipment will always need maintenance workers, and they'll need those workers to be able to work well with new systems that can fuse their skills with advanced technologies in novel ways. This is the *skills* part of MELDS. Workers in these jobs will need to do what people do well: adapt to new situations and find novel, creative solutions to challenges that arise. Let the machines do the heavy lifting, the monitoring, and the monotonous tasks.

In the case of the researchers, engineers, farmers, and others, the data and analysis that AI systems provide can act as a third eye. And this is why the *data* part of MELDS is so important. Suddenly, very complex industrial or ecological systems become knowable. Engineers and managers can eliminate previously invisible inefficiencies, and they can make changes to certain aspects of a process with confidence. When you honestly assess the strengths of human and machine workers, and what they do well when they collaborate, a whole new world of possibilities emerges for running a business and designing your processes—that is, the important *mindset* part of MELDS. And by exploring those possibilities, companies can often develop novel businesses, like vertical farms. Indeed, it's through the *experimentation* part of MELDS that executives will be able to discover game-changing innovations that could potentially transform their company, if not their entire industry.

In the next chapter, we take artificial intelligence into the back office. It's where "second wave" automation is entrenched, and "third wave" AI has come as a welcome relief to many who have been working with awkward IT tools or inefficient processes. Here, too, we will see how AI and people's imaginations have been transforming seemingly mundane processes, opening up new possibilities for both efficiency and growth through human-machine collaborations.

## *A Brief History of AI*

The driving technology behind the current era of adaptive processes is AI, which has been evolving over decades. A brief history of the technology gives some context to its current state of advanced capabilities.

The field of artificial intelligence was officially born in 1956, when a small group of computer and research scientists, organized by John McCarthy, including Claude Shannon, Marvin Minsky, and others, gathered at Dartmouth College for the first-ever conference to debate the possibility that machine intelligence could imitate human smarts.[a]

The conference, essentially an extended brainstorming session, was based on the assumption that every aspect of learning and creativity could be described so precisely that it could be mathematically modeled and therefore replicated by machines. The goals were lofty; from the event proposal: “An attempt will be made to find how to make machines use language, form, abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves.” Of course, this was just the beginning.

The conference succeeded almost immediately in defining the field and unifying many of the mathematical ideas swirling around the concept of artificial intelligence. It also inspired entirely new areas of research in the decades that followed. For instance, Minsky, with Seymour Papert, wrote what was considered the foundational book on

scope and limitations of neural networks, a kind of AI that uses biological neurons as its model. Other ideas like expert systems—wherein a computer contained deep stores of "knowledge" for specific domains like architecture or medical diagnosis—and natural language processing, computer vision, and mobile robotics can also be traced back to the event.

One conference participant was Arthur Samuel, an engineer at IBM who was building a computer program to play checkers. His program would assess the current state of a checkers board and calculate the probability that a given position could lead to a win. In 1959, Samuel coined the term "machine learning": the field of study that gives computers the ability to learn without being explicitly programmed. In 1961, his learning program was used to defeat the fourth-ranked checkers player in the United States. But because Samuel was modest and didn't enjoy the politics of self-promotion, it was not until after his retirement from IBM in 1966 that the significance of his machine-learning work became more widely known.[b]

In the decades after the conference, machine learning remained obscure, as other kinds of AI took center stage. In particular, researchers in the 1970s and 1980s focused on a concept of intelligence that was based on physical symbols and manipulated by logical rules. These symbolic systems, however, found no practical success at the time, and their failures led to a period known as an "AI winter."

In the 1990s, however, machine learning began to flourish as its practitioners integrated statistics and probability theory into their approaches. At the same time, the personal computing revolution began. Over the next decade, digital systems, sensors, the internet, and mobile

phones would become common, providing all kinds of data for machine-learning experts to use when training these adaptive systems.

Today we think of a machine-learning application as one that builds models based on data sets that engineers or specialists use to train the system. It's a sharp contrast with traditional computer programming. Standard algorithms would follow a predetermined path set into motion by programmers' static instructions or code. A machine-learning system, conversely, can learn as it goes. With each new data set, it updates its models and the way it "sees" the world. In an era in which machines can learn and change based on their experiences and data, programmers have become less like rule makers and dictators, and more like teachers and trainers.

Now AI systems that deploy machine learning are everywhere. Banks use them for fraud detection; dating websites use them to suggest potential matches; marketers use them to try to predict who will respond favorably to an ad; and photo-sharing sites use them for automatic face recognition. We've come a long way since checkers. In 2016, Google's AlphaGo demonstrated a significant machine-learning advance. For the first time, a computer beat a human champion of Go, a game far more complex than checkers or chess. In a sign of the times, AlphaGo exhibited moves that were so unexpected that some observers deemed them to actually be creative and even "beautiful."[c]

The growth of AI and machine learning has been intermittent over the decades, but the way that they've crept into products and business operations in recent years shows that they're more than ready for prime time.

According to Danny Lange, former head of machine learning at Uber, the technology has finally broken out of the research lab and is fast becoming "the cornerstone of business disruption."[d]

a. "Artificial Intelligence and Life in 2030," Stanford One Hundred Year Study on Artificial Intelligence (AI100), September 2016, https://ai100.stanford.edu/sites/default/files/ai_100_report_0831fnl.pdf.

b. John McCarthy and Ed Feigenbaum, "Arthur Samuel: Pioneer in Machine Learning," Stanford Infolab, http://infolab.stanford.edu/pub/voy/museum/samuel.html, accessed October 23, 2017.

c. Cade Metz, "How Google's AI Viewed the Move No Human Could Understand," *Wired*, March 14, 2016, https://www.wired.com/2016/03/googles-ai-viewed-move-no-human-understand/.

d. Daniel Lange, "Making Uber Smarter with Machine Learning," presentation at Machine Learning Innovation Summit, San Francisco, June 8–9, 2016.

# 2

# Accounting for Robots

*AI in Corporate Functions*

Money laundering is a major concern for financial institutions, which can face heavy fines and stiff regulatory restrictions for any infractions. At one large global bank, up to ten thousand staffers were engaged in identifying suspicious transactions and accounts that might indicate money laundering, terrorist financing, and other illegal activities. Aggressive monitoring was required to meet the rigid expectations of the US Department of Justice, and the incurred costs were high, with an excessive number of false positives that the bank had to investigate.

In response, the bank implemented a full suite of advanced analytics tools for anti-money-laundering (AML) detection, including machine-learning algorithms to better segment transactions and accounts, and to set the optimal thresholds for alerting investigators to any suspicious activity. All this is done dynamically to incorporate the most recent data and latest results.

Moreover, the use of network analysis is helping to uncover valuable new patterns—for example, the closeness of a business relationship between two of the bank's customers can help determine the likelihood that, if one of them is involved in illicit activity, then the other might also be involved.

Thus far, the results have been impressive. The AML system has reduced false positive alerts by as much as 30 percent, allowing staffers more time to investigate those cases requiring human judgment and compliance expertise. The system has also helped reduce the time required to investigate each alert, resulting in a dramatic 40 percent decrease in costs.

## Allowing Humans to Be More Human

People rarely revel in performing repetitive or robotic tasks day in and day out. Talk to someone who's worked in a process with a lot of routine steps, and you'll learn how they relish an unusual situation that breaks up a workday or workweek. And if it gives them a chance to solve a hard problem, they feel as if they've made a difference to the organization or maybe even in someone's life. Research by Jordan Etkin of Duke University and Cassie Mogilner of the Wharton School suggests that some variety throughout a workday leads to increased happiness, likely tied to a greater sense of stimulation and productivity.[1] So, the question becomes, why continue to train people to work like robots? Why not let workers be more human? Or, as that global bank we discussed discovered, why not let staffers focus on higher-value tasks, requiring their judgment, experience, and expertise?

Our research has confirmed that in many cases, AI allows human workers to be more human. The rote nature of some administrative jobs like invoicing, bookkeeping, accounting,

complaints, forms processing, and scheduling arose initially from the use of standard IT technologies that required humans to adjust to the machine limitations of the 1990s and 2000s. Human resources, IT security, and banking compliance departments all use processes that are often made up of well-defined, repetitive tasks. This was the "second wave" of business process improvement.

This chapter examines innovative improvements in enterprise processes—a trend that's been building for years but, thanks to technological advances, has only recently become viable for most organizations. We give examples that address basic questions that anyone interested in deploying AI for enterprise-wide processes should be asking. What will such work look like in this new era of business process transformation? Which tasks are best suited for humans and which are best for machines? While it's true that many organizations can immediately see significant gains when they use AI in conjunction with their existing workforce, what happens if you completely rethink your processes around ultra-smart systems? What kind of growth, services, and products become possible?

## Your Office Robot

To answer those questions, let's start with a familiar process: categorizing and resolving complaints. In the past, much of the process work around sorting through customer complaints was done manually, and the tedium of many of those tasks detracted from people's job satisfaction. At Virgin Trains, a train-operating company in the UK, for example, a team of customer service reps would manually read, sort, and route complaints. These repetitive activities diverted employees' time and attention, and it fatigued

them more than other work they did, like directly talking to customers.

Because the read-sort-route process is clearly defined, it is in some ways an excellent example of a process ripe for automation. But because the incoming information is text-based and is considered "unstructured" in the eyes of software systems, parsing could have been difficult for a less advanced system. Enter AI. Virgin Trains has now installed a machine-learning platform, inSTREAM, with natural-language processing capabilities that can recognize patterns in unstructured data by analyzing a corpus of similar examples—in this case, complaints—and by tracking how customer service representatives interact with incoming text.

Now when a complaint arrives at Virgin Trains, it's automatically read, sorted, and packaged into a case ready file that an employee can quickly review and process. The most common complaints get appropriate, automated responses. If the software isn't fully confident in its assessment of a complaint, it flags it as an exception for a human worker to review; the worker's response effectively updates the software's model. Over time, this kind of feedback improves the algorithm's confidence for an increasing array of scenarios. The system can handle complaints that are terse or long-winded, generic or specific, or in English or other languages.

Thanks to this new technology, Virgin Train's complaints department has decreased its manual work by 85 percent. It also increased correspondence by 20 percent because the new capabilities prompted the company to further open itself to customer interactions. Previously, the company only accepted complaints through its website. Now, it can process inquiries of any type, including email, fax, snail mail, and social media.[2] (Virgin Trains is one of many companies adding some automated intelligence to the back office; see the sidebar "AI in Business

Processes" for more examples.)

## AI in Business Processes

For every company and business unit and division, there are a mass of behind-the-scenes activities. The introduction of AI can help to offset the burden of repetitive, low-visibility tasks so that employees can focus on higher-value tasks.

- At Goldman Sachs, AI studies up to a million different analyst reports to identify the top factors affecting share prices.[a]
- Woodside Petroleum uses IBM's Watson to extend the sharing of lessons learned through HR, legal, and exploration.[b]
- Along with human moderators, the *Huffington Post* uses AI to flag inappropriate comments, spam, and abusive language.[c]
- Arizona State University is now using an adaptive learning tool that deploys machine learning to provide a personalized tutor for students in introductory classes.[d]

a. Nathaniel Popper, "The Robots Are Coming for Wall Street," *New York Times*, February 25, 2016, https://www.nytimes.com/2016/02/28/magazine/the-robots-are-coming-for-wall-street.html.

b. Daniel Russo, "Hiring Heroes: How Woodside Energy Works with

IBM Watson," IBM Watson blog, September 11, 2017, https://www.ibm.com/blogs/watson/2017/09/hiring-heroes-woodside-energy-works-ibm-watson/.

c. Mike Masnick, "HuffPost Moderates Comments to Please Advertisers [Updated: Or Not]," Tech Dirt, October 30, 2012, https://www.techdirt.com/articles/20121022/12562620788/huffpost-moderates-comments-to-please-advertisers.shtml.

d. Seth Fletcher, "How Big Data Is Taking Teachers Out of the Lecturing Business," *Scientific American*, August 1, 2013, https://www.scientificamerican.com/article/how-big-data-taking-teachers-out-lecturing-business.

## Moving Well Beyond RPA

The Virgin Trains system is a relatively advanced form of back-office automation because it can analyze and adapt to unstructured data as well as the sudden influx of data. Such applications are called "robotic process automation" (RPA). Simply put, RPA is software that performs digital office tasks that are administrative, repetitive, and mostly transactional within a workflow. In other words, it *automates* existing processes. But in order to *reimagine* processes, firms must utilize more advanced technologies—namely, AI. (See the sidebar "AI Technologies and Applications: How Does This All Fit Together?" at the end of this chapter.)

Now we're talking about systems that deploy AI techniques such as computer vision, or machine-learning tools to analyze unstructured or complex information. It might be able to read various styles of invoices, contracts, or purchase orders, for instance. It can process these documents—no matter the format—and put the correct values into forms and databases for further

action. And then there are even more advanced systems that deploy sophisticated machine-learning algorithms not just to perform the tasks they've been programmed to do, but also to assess tasks and processes and adjust as needed. They can learn from observation by "watching over the shoulders" of human employees in order to improve their performance over time. In other words, they are exactly the kind of technology that is enabling the third wave of business process improvement—adaptable processes—that we discussed in the introduction chapter. These applications are more transformational and typically require human employees to actively participate, applying a kind of tacit knowledge or expertise that is difficult to explain or model. Think of the global bank's anti-money-laundering system that we discussed earlier. A complicated financial transaction is processed; an automated system flags it as being suspicious; and a human expert exercises the judgment to decide whether it warrants further investigation. This type of human-machine collaboration is also typical of the third wave of business process transformation.

Companies can deploy a range of these technologies, sometimes even for the same application. Case in point: Unilever's process for hiring employees. Say you're looking for a job and through LinkedIn you find a position at Unilever that might be suitable. For the first round of the application process, you'd be asked to play twelve online games based on tests from the field of cognitive neuroscience. The games help assess certain traits, like your risk aversion and ability to read emotional versus contextual cues. According to Unilever, there are no right or wrong answers to the games because, for example, an appetite for risk might be suitable for one type of job while an aversion to risk might be better for another type of position. For this round of the application process, advanced AI isn't required and a relatively basic technology like RPA would suffice.

But if you, as an applicant, made it to the next round, you'd then be asked to submit a video-recorded interview through your computer or smartphone in which you'd answer a set of questions designed for the specific position you were interested in. And here's where sophisticated AI technologies come into play: Your answers would then by analyzed by HireVue, an AI application that not only notes the words you use but also your body language and tone. The best candidates for the position are then invited to the company offices, where they can then be assessed by humans who would make the final hiring decision.

Not only does this Unilever example show how different technologies can be used for different parts of the same application; it also demonstrates the power of human-machine collaboration. Within 90 days of the new system going live, job applications had doubled to 30,000 from the same time period a year ago. In addition, the average time for someone to be hired plunged from four months to just four weeks, and the time that recruiters spent reviewing applications plummeted by 75 percent. Moreover, the company reports that, since installing the system, it has hired its most diverse class to date. For one thing, there was a dramatic increase in the universities represented, from 840 to 2,600.[3]

## How Do You Know Which Processes to Change?

Repetition. Replication. Redundancy. A well-outlined process. If these elements show up in your business operations, it's a clue that tasks or processes are ready to be changed.

Roger Dickey, a developer and founder of the fast-growing startup Gigster, recognized replication and redundancy in the code

of most software applications. At the same time, each new software project—regardless of how similar it was to others that had come before—was incredibly complex to build, complete with bugs and pitfalls that slowed production. Could AI be used to help reimagine the business processes required to build software?

The answer, as Gigster has found, is yes. The company typically uses AI to assess the needs of any given software project and automatically assembles an ad hoc team of crack developers to build it. If you are a small firm that needs an app or some other software product but don't have the time or resources to hire a team of developers yourself, you turn to Gigster. If you are a large corporation that doesn't want to divert resources away from established projects, you turn to Gigster.

Gigster effectively takes aim at multiple areas of the enterprise: HR (developer teams are assembled using AI), procurement (quotes are generated using AI), and IT (members of development work and are managed with AI-enabled assistance).

How does Gigster upend procurement and HR? Suppose you'd like to build an app that can help patients consolidate their medical records to share with their doctors. Where do you begin? First, you give Gigster a short document that explains the app's core function and how you envision a person will use it. On Gigster's end, the project description is cross-referenced with others in Gigster's portfolio of "data structures," which is essentially a catalog of software features. Dickey says his company has mapped the "software genome" and understands five hundred different features a product might have. Next, Gigster takes into account about twenty other customer requirements for how the user interface should look or how quickly the job needs to be completed, and so on. From the customer's mockup, description, and requirements, Gigster's AI quote generator leverages previous projects with similar constraints to estimate a price and a timeline.

If you agree to the price and timeline, the next set of Gigster's AI

features goes to work. The company deploys its "team builder" in which it matches the demands of your app to members of software development teams who can meet your needs. A typical team will consist of three to five people: a project manager, a designer or two, and a developer or two, all high-performing recruits who are monitored closely by Gigster's online system, which allows the company to guarantee quality and on-time products. This initial setup takes one to three days.

Because software developers work in a digital realm, everything they do can be recorded relatively easily and analyzed. "We believe work is measurable and that data has patterns and those patterns can be exploited to find new efficiencies in work," says Dickey. This means that Gigster knows what processes go into making a software project a success—based on hundreds of other projects just like it—and an AI tool can use that information to spot potential production hiccups before they spiral out of control. Moreover, whenever developers are having trouble with any particular pieces of code, an AI assistant can automatically put them in touch with someone who's recently solved or is grappling with a similar problem. It's "an AI assistant who knows where you are in the project," says Dickey, "and can match you with other people in the world doing the same thing."[4] That type of employee augmentation is one of the keys of the third wave of human-machine collaboration.

## How Do You Know How Much to Change?

By the very nature of its business—software—Gigster has been able to deploy AI to a range of IT and business processes. Other firms, however, might be better off applying AI to just a few processes. For those organizations, executives need to make judicious

decisions about how best to augment their existing employees. And they also must have a plan for scaling up their use of AI in processes.

Those issues were key concerns for SEB, a major Swedish bank, which has been busy installing a virtual assistant called Amelia. Built by IPsoft, Amelia (later renamed Aida in SEB's application) now interacts directly with SEB's 1 million customers. Within the first three weeks in this role, the software held more than four thousand conversations with seven hundred people and was able to resolve the majority of the issues, says Rasmus Järborg, SEB's chief strategy officer. The decision to move Aida to a customer-facing role came only after the bank had tested the software internally as a virtual IT help-desk agent, assisting its fifteen thousand employees.[5]

Aida is adept at handling natural-language conversations, and the technology is even able to monitor a caller's tone of voice as a way to offer better service in the future. The software adapts, learning new skills by monitoring human customer-service agents. This means its capabilities improve and increase over time; new tasks and processes within the customer service department can become automated with little direct effort by the people who work there.

SEB is the first bank to use Amelia for interactions with customers, and IPsoft has helped to set up an in-house talent pool of individuals who can mentor the software. These human mentors supervise learning and performance and identify new ways to apply the technology for customer service.[6] We discuss this type of human-machine collaboration in greater detail in chapter 5.

Aida is showing that automated natural-language customer communications are possible in large and complex business environments. As natural-language techniques improve and interfaces advance, they will continue spreading throughout

different business functions in various industries. In chapter 4 we'll discuss how various natural-language processing chatbots like Amazon's Alexa are becoming the new front-office faces of companies.

## Redefining an Entire Industry

As AI becomes increasingly capable of adding intelligence to middle- and back-office processes, the technology could potentially redefine entire industries. In IT security, for instance, a growing number of security firms are combining machine-learning approaches to build ultra-smart, continually evolving defenses against malicious software. These systems can unearth harmful viruses and malware before they cause damage, and they can predict vulnerabilities before they become liabilities that let hackers take over entire systems. In some cases, the IT security process is a closed, automated loop; humans can step away from the day-to-day controls and spend time researching threats or creating new simulations to further test and train bots. (See the sidebar "When Bots Collide.")

In traditional cybersecurity, a company might perform analytics on existing data, gathering signatures of threats and using them to protect against future threats. It's a static operation, and not able to adapt in real time. In contrast, AI-based approaches are able to recognize anomalous patterns as they arise. They do so by calibrating models based on network traffic behavior and scoring anomalies according to the extent to which they deviate from the norm. What's more, AI-based analytics improve as each alert is eventually resolved—either by machine or human—effectively incorporating each new insight into the system as it runs.