



Harvard Business Review

REPRINT R1804J
PUBLISHED IN HBR
JULY-AUGUST 2018

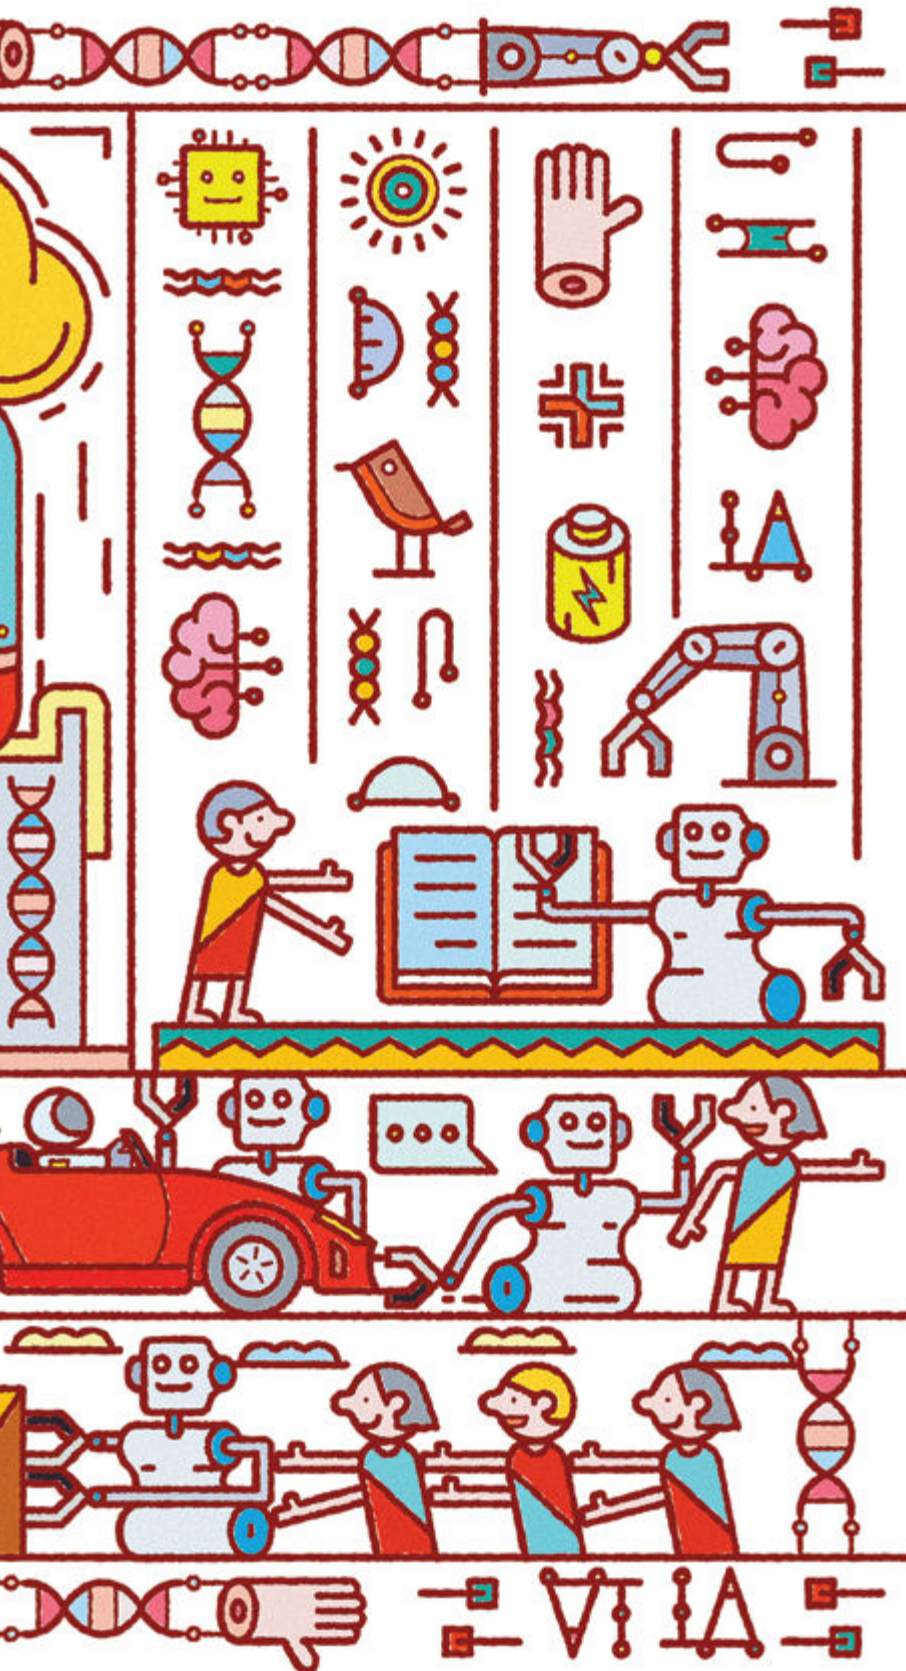
ARTICLE **TECHNOLOGY**

Collaborative Intelligence: Humans and AI Are Joining Forces

by H. James Wilson and Paul R. Daugherty



ILLUSTRATION BY KUO CHENG LIAO




COLLABORATIVE
INTELLIGENCE:

Humans and AI Are Joining Forces



BY H. JAMES WILSON
AND PAUL R. DAUGHERTY



IN BRIEF

THE OUTLOOK

Artificial intelligence is transforming business—and having the most significant impact when it augments human workers instead of replacing them.

THE DETAILS

Companies see the biggest performance gains when humans and smart machines collaborate. People are needed to train machines, explain their outputs, and ensure their responsible use. AI, in turn, can enhance humans' cognitive skills and creativity, free workers from low-level tasks, and extend their physical capabilities.

THE PRESCRIPTION

Companies should reimagine their business processes, focusing on using AI to achieve more operational flexibility or speed, greater scale, better decision making, or increased personalization of products and services.

ARTIFICIAL INTELLIGENCE IS BECOMING GOOD AT MANY “HUMAN” JOBS—DIAGNOSING DISEASE, TRANSLATING LANGUAGES, PROVIDING CUSTOMER SERVICE—AND IT’S IMPROVING FAST. THIS IS RAISING REASONABLE FEARS THAT AI WILL ULTIMATELY REPLACE HUMAN WORKERS THROUGHOUT THE ECONOMY. BUT THAT’S NOT THE INEVITABLE, OR EVEN MOST LIKELY, OUTCOME. NEVER BEFORE HAVE DIGITAL TOOLS BEEN SO RESPONSIVE TO US, NOR WE TO OUR TOOLS. WHILE AI WILL RADICALLY ALTER HOW WORK GETS DONE AND WHO DOES IT, THE TECHNOLOGY’S LARGER IMPACT WILL BE IN COMPLEMENTING AND AUGMENTING HUMAN CAPABILITIES, NOT REPLACING THEM.

Certainly, many companies have used AI to automate processes, but those that deploy it mainly to displace employees will see only short-term productivity gains. In our research involving 1,500 companies, we found that firms achieve the most significant performance improvements when humans and machines work together (see the exhibit “The Value of Collaboration”). Through such collaborative intelligence, humans and AI actively enhance each other’s complementary strengths: the leadership, teamwork, creativity, and social skills of the former, and the speed, scalability, and quantitative capabilities of the latter. What comes naturally to people (making a joke, for example) can be tricky for machines, and what’s straightforward for machines (analyzing gigabytes of data) remains virtually impossible for humans. Business requires both kinds of capabilities.

To take full advantage of this collaboration, companies must understand how humans can most effectively augment machines, how machines can enhance what humans do best, and how to redesign business processes to support the partnership. Through our research and work in the field, we have developed guidelines to help companies achieve this and put the power of collaborative intelligence to work.

HUMANS ASSISTING MACHINES

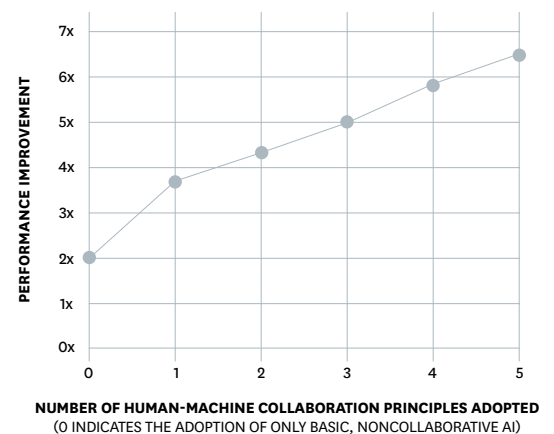
Humans need to perform three crucial roles. They must *train* machines to perform certain tasks; *explain* the outcomes of those tasks, especially when the results are counterintuitive or controversial; and *sustain* the responsible use of machines (by, for example, preventing robots from harming humans).

Training. Machine-learning algorithms must be taught how to perform the work they’re designed to do. In that effort, huge training data sets are amassed to teach machine-translation apps to handle idiomatic expressions, medical apps to detect disease, and recommendation engines to support financial decision making. In addition, AI systems must be trained how best to interact with humans. While organizations across sectors are now in the early stages of filling trainer roles, leading tech companies and research groups already have mature training staffs and expertise.

Consider Microsoft’s AI assistant, Cortana. The bot required extensive training to develop just the right personality: confident, caring, and helpful but not bossy. Instilling those qualities took countless hours of attention by a team that included a poet, a novelist, and a playwright. Similarly, human trainers were needed to develop the personalities of Apple’s Siri and Amazon’s Alexa to ensure that they accurately reflected their companies’ brands. Siri, for example, has just a touch of sassiness, as consumers might expect from Apple.

THE VALUE OF COLLABORATION

Companies benefit from optimizing collaboration between humans and artificial intelligence. Five principles can help them do so: Reimagine business processes; embrace experimentation/employee involvement; actively direct AI strategy; responsibly collect data; and redesign work to incorporate AI and cultivate related employee skills. A survey of 1,075 companies in 12 industries found that the more of these principles companies adopted, the better their AI initiatives performed in terms of speed, cost savings, revenues, or other operational measures.



AI assistants are now being trained to display even more complex and subtle human traits, such as sympathy. The start-up Koko, an offshoot of the MIT Media Lab, has developed technology that can help AI assistants seem to commiserate. For instance, if a user is having a bad day, the Koko system doesn’t reply with a canned response such as “I’m sorry to hear that.” Instead it may ask for more information and then offer advice to help the person see his issues in a different light. If he were feeling stressed, for instance, Koko might recommend thinking of that tension as a positive emotion that could be channeled into action.

Explaining. As AIs increasingly reach conclusions through processes that are opaque (the so-called black-box problem), they require human experts in the field to explain their behavior to nonexpert users. These “explainers” are particularly important in evidence-based industries, such as law and medicine, where a practitioner needs to understand how an AI weighed inputs into, say, a sentencing or medical recommendation. Explainers are similarly important in helping insurers and law enforcement understand why an autonomous car took actions that led to an accident—or failed to avoid one. And explainers are becoming integral in regulated industries—indeed, in any consumer-facing industry where a machine’s

output could be challenged as unfair, illegal, or just plain wrong. For instance, the European Union's new General Data Protection Regulation (GDPR) gives consumers the right to receive an explanation for any algorithm-based decision, such as the rate offer on a credit card or mortgage. This is one area where AI will contribute to *increased* employment: Experts estimate that companies will have to create about 75,000 new jobs to administer the GDPR requirements.

Sustaining. In addition to having people who can explain AI outcomes, companies need “sustainers”—employees who continually work to ensure that AI systems are functioning properly, safely, and responsibly.

For example, an array of experts sometimes referred to as safety engineers focus on anticipating and trying to prevent harm by AIs. The developers of industrial robots that work alongside people have paid careful attention to ensuring that they recognize humans nearby and don't endanger them. These experts

Human-machine collaboration enables companies to interact with employees and customers in novel, more effective ways.

may also review analysis from explainers when AIs do cause harm, as when a self-driving car is involved in a fatal accident.

Other groups of sustainers make sure that AI systems uphold ethical norms. If an AI system for credit approval, for example, is found to be discriminating against people in certain groups (as has happened), these ethics managers are responsible for investigating and addressing the problem. Playing a similar role, data compliance officers try to ensure that the data that is feeding AI systems complies with the GDPR and other consumer-protection regulations. A related data-use role involves ensuring that AIs manage information responsibly. Like many tech companies, Apple uses AI to collect personal details about users as they engage with the company's devices and software. The aim is to improve the user experience, but unconstrained data gathering can compromise privacy, anger customers, and run

afoul of the law. The company's “differential privacy team” works to make sure that while the AI seeks to learn as much as possible about a group of users in a statistical sense, it is protecting the privacy of individual users.

MACHINES ASSISTING HUMANS


Smart machines are helping humans expand their abilities in three ways. They can *amplify* our cognitive strengths; *interact* with customers and employees to free us for higher-level tasks; and *embody* human skills to extend our physical capabilities.

Amplifying. Artificial intelligence can boost our analytic and decision-making abilities by providing the right information at the right time. But it can also heighten creativity. Consider how Autodesk's Dreamcatcher AI enhances the imagination of even exceptional designers. A designer provides Dreamcatcher with criteria about the desired product—for example, a chair able to support up to 300 pounds, with a seat 18 inches off the ground, made of materials costing less than \$75, and so on. She can also supply information about other chairs that she finds attractive. Dreamcatcher then churns out thousands of designs that match those criteria, often sparking ideas that the designer might not have initially considered. She can then guide the software, telling it which chairs she likes or doesn't, leading to a new round of designs.

Throughout the iterative process, Dreamcatcher performs the myriad calculations needed to ensure that each proposed design meets the specified criteria. This frees the designer to concentrate on deploying uniquely human strengths: professional judgment and aesthetic sensibilities.

Interacting. Human-machine collaboration enables companies to interact with employees and customers in novel, more effective ways. AI agents like Cortana, for example, can facilitate communications between people or on behalf of people, such as by transcribing a meeting and distributing a voice-searchable version to those who couldn't attend. Such applications are inherently scalable—a single chatbot, for instance, can provide routine customer service to large numbers of people simultaneously, wherever they may be.

SEB, a major Swedish bank, now uses a virtual assistant called Aida to interact with millions of customers. Able to handle natural-language conversations, Aida has access to vast stores of data and can answer many frequently asked questions, such as how to open an account or make cross-border payments. She can also ask callers follow-up questions to solve their problems, and she's able to analyze a caller's tone of voice (frustrated versus appreciative, for instance) and use that information to provide better service later. Whenever the system can't resolve an



issue—which happens in about 30% of cases—it turns the caller over to a human customer-service representative and then monitors that interaction to learn how to resolve similar problems in the future. With Aida handling basic requests, human reps can concentrate on addressing more-complex issues, especially those from unhappy callers who might require extra hand-holding.

Embodying. Many AIs, like Aida and Cortana, exist principally as digital entities, but in other applications the intelligence is embodied in a robot that augments a human worker. With their sophisticated sensors, motors, and actuators, AI-enabled machines can now recognize people and objects and

work safely alongside humans in factories, warehouses, and laboratories.

In manufacturing, for example, robots are evolving from potentially dangerous and “dumb” industrial machines into smart, context-aware “cobots.” A cobot arm might, for example, handle repetitive actions that require heavy lifting, while a person performs complementary tasks that require dexterity and human judgment, such as assembling a gear motor.

Hyundai is extending the cobot concept with exoskeletons. These wearable robotic devices, which adapt to the user and location in real time, will enable industrial workers to perform their jobs with superhuman endurance and strength.

REIMAGINING YOUR BUSINESS

In order to get the most value from AI, operations need to be redesigned. To do this, companies must first discover and describe an operational area that can be improved. It might be a balky internal process (such as HR's slowness to fill staff positions), or it could be a previously intractable problem that can now be addressed using AI (such as quickly identifying adverse drug reactions across patient populations). Moreover, a number of new AI and advanced analytic techniques can help surface previously invisible problems that are amenable to AI solutions (see the sidebar "Revealing Invisible Problems").

Next, companies must develop a solution through co-creation—having stakeholders envision how they might collaborate with AI systems to improve a process. Consider the case of a large agricultural company that wanted to deploy AI technology to help farmers.

At Mercedes-Benz, cobot arms guided by human workers pick up and place heavy parts, becoming an extension of the worker's body.

An enormous amount of data was available about soil properties, weather patterns, historical harvests, and so forth, and the initial plan was to build an AI application that would more accurately predict future crop yields. But in discussions with farmers, the company learned of a more pressing need. What farmers really wanted was a system that could provide real-time recommendations on how to increase productivity—which crops to plant, where to grow them, how much nitrogen to use in the soil, and so on. The company developed an AI system to provide such advice, and the initial outcomes were promising; farmers were happy about the crop yields obtained with the AI's guidance. Results from that initial test were then fed back into the system to refine the algorithms used. As with the discovery step, new AI and analytic techniques can assist in co-creation by suggesting novel approaches to improving processes.

The third step for companies is to scale and then sustain the proposed solution. SEB, for example, originally deployed a version of Aida internally to assist

15,000 bank employees but thereafter rolled out the chatbot to its one million customers.

Through our work with hundreds of companies, we have identified five characteristics of business processes that companies typically want to improve: flexibility, speed, scale, decision making, and personalization. When reimagining a business process, determine which of these characteristics is central to the desired transformation, how intelligent collaboration could be harnessed to address it, and what alignments and trade-offs with other process characteristics will be necessary.

Flexibility. For Mercedes-Benz executives, inflexible processes presented a growing challenge. Increasingly, the company's most profitable customers had been demanding individualized S-class sedans, but the automaker's assembly systems couldn't deliver the customization people wanted.

Traditionally, car manufacturing has been a rigid process with automated steps executed by "dumb" robots. To improve flexibility, Mercedes replaced some of those robots with AI-enabled cobots and redesigned its processes around human-machine collaborations. At the company's plant near Stuttgart, Germany, cobot arms guided by human workers pick up and place heavy parts, becoming an extension of the worker's body. This system puts the worker in control of the build of each car, doing less manual labor and more of a "piloting" job with the robot.

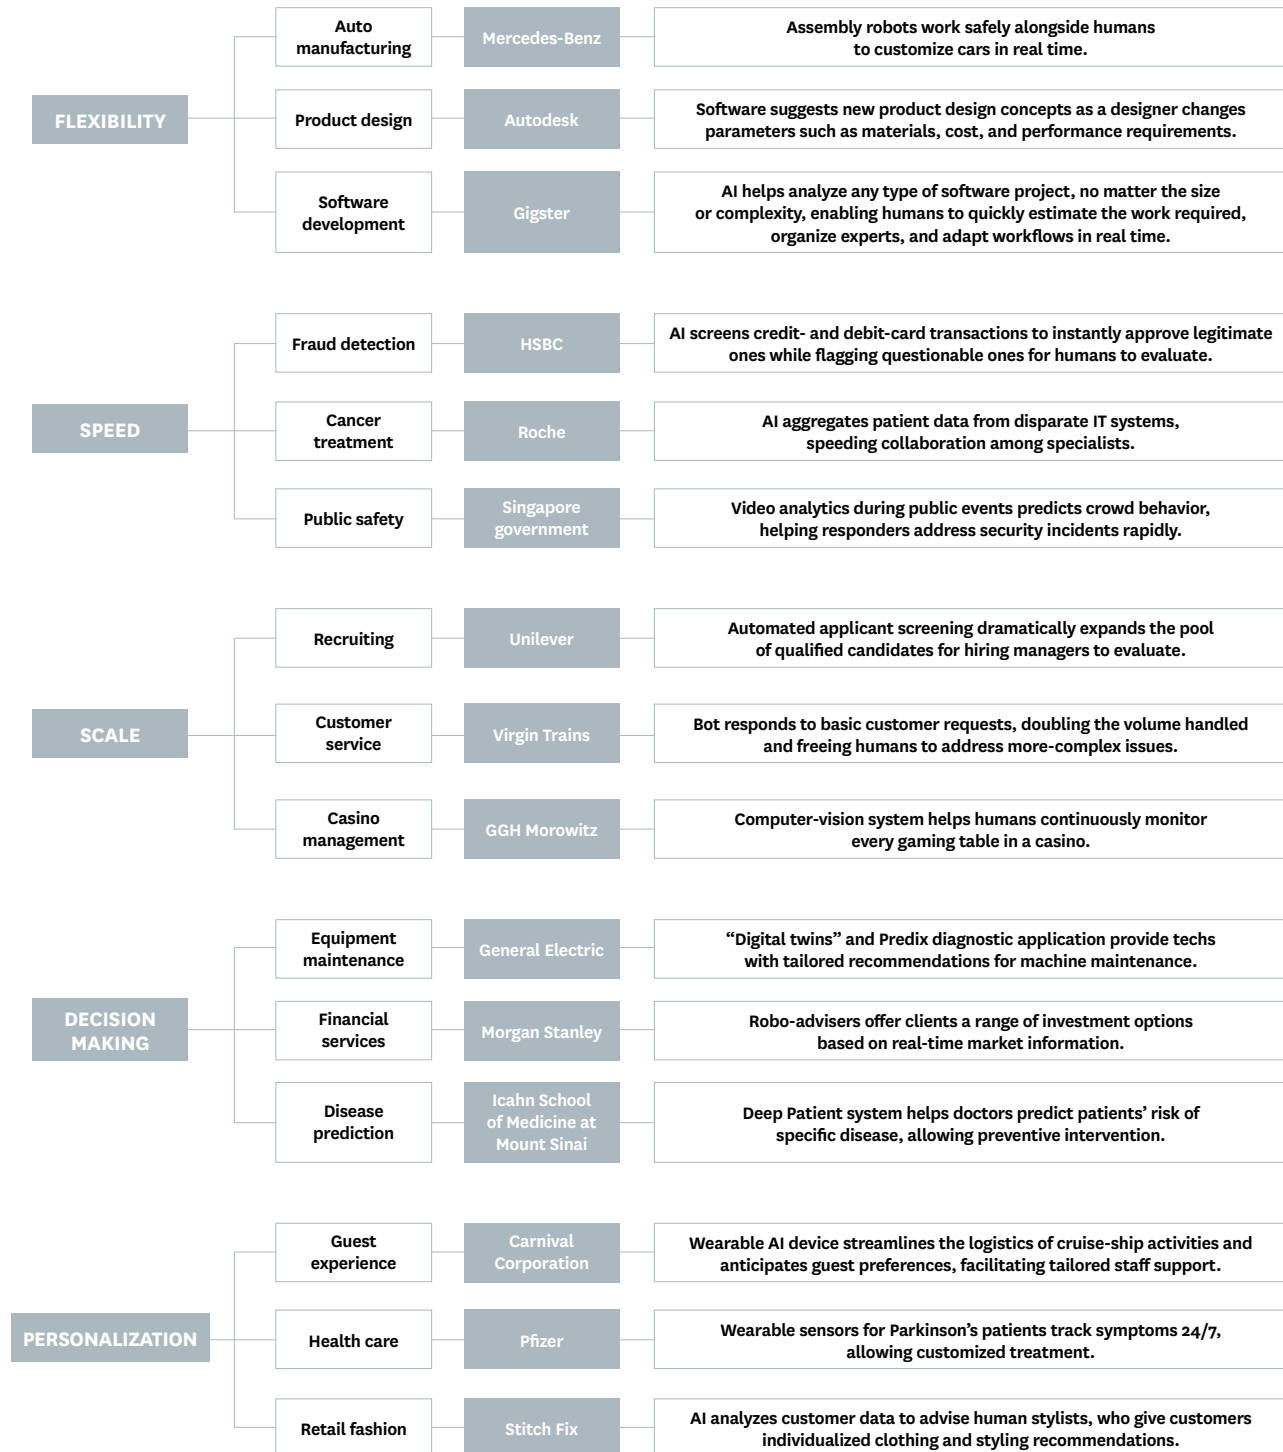
The company's human-machine teams can adapt on the fly. In the plant, the cobots can be reprogrammed easily with a tablet, allowing them to handle different tasks depending on changes in the workflow. Such agility has enabled the manufacturer to achieve unprecedented levels of customization. Mercedes can individualize vehicle production according to the real-time choices consumers make at dealerships, changing everything from a vehicle's dashboard components to the seat leather to the tire valve caps. As a result, no two cars rolling off the assembly line at the Stuttgart plant are the same.

Speed. For some business activities, the premium is on speed. One such operation is the detection of credit-card fraud. Companies have just seconds to determine whether they should approve a given transaction. If it's fraudulent, they will most likely have to eat that loss. But if they deny a legitimate transaction, they lose the fee from that purchase and anger the customer.

Like most major banks, HSBC has developed an AI-based solution that improves the speed and accuracy of fraud detection. The AI monitors and scores millions of transactions daily, using data on purchase location and customer behavior, IP addresses, and other information to identify subtle patterns that signal possible fraud. HSBC first implemented the system in the United States, significantly reducing the rate of undetected fraud and false positives,

ENHANCING PERFORMANCE

At organizations in all kinds of industries, humans and AI are collaborating to improve five elements of business processes.



REVEALING INVISIBLE PROBLEMS

Former U.S. Defense Secretary Donald Rumsfeld once famously distinguished among “known knowns,” “known unknowns,” and “unknown unknowns”—things you’re not even aware you don’t know. Some companies are now using AI to uncover unknown unknowns in their businesses. Case in point: GNS Healthcare applies machine-learning software to find overlooked relationships among data in patients’ health records and elsewhere. After identifying a relationship, the software churns out numerous hypotheses to explain it and then suggests which of those are the most likely. This approach enabled GNS to uncover a new drug interaction hidden in unstructured patient notes. CEO Colin Hill points out that this is not garden-variety data mining to find associations. “Our machine-learning platform is not just about seeing patterns and correlations in data,” he says. “It’s about actually discovering causal links.”

and then rolled it out in the UK and Asia. A different AI system used by Danske Bank improved its fraud-detection rate by 50% and decreased false positives by 60%. The reduction in the number of false positives frees investigators to concentrate their efforts on equivocal transactions the AI has flagged, where human judgment is needed.

The fight against financial fraud is like an arms race: Better detection leads to more-devilish criminals, which leads to better detection, which continues the cycle. Thus the algorithms and scoring models for combating fraud have a very short shelf life and require continual updating. In addition, different countries and regions use different models. For these reasons, legions of data analysts, IT professionals, and experts in financial fraud are needed at the interface between humans and machines to keep the software a step ahead of the criminals.

Scale. For many business processes, poor scalability is the primary obstacle to improvement. That’s particularly true of processes that depend on intensive human labor with minimal machine assistance. Consider, for instance, the employee recruitment process at Unilever. The consumer goods giant was looking for a way to diversify its 170,000-person

workforce. HR determined that it needed to focus on entry-level hires and then fast-track the best into management. But the company’s existing processes weren’t able to evaluate potential recruits in sufficient numbers—while giving each applicant individual attention—to ensure a diverse population of exceptional talent.

Here’s how Unilever combined human and AI capabilities to scale individualized hiring: In the first round of the application process, candidates are asked to play online games that help assess traits such as risk aversion. These games have no right or wrong answers, but they help Unilever’s AI figure out which individuals might be best suited for a particular position. In the next round, applicants are asked to submit a video in which they answer questions designed for the specific position they’re interested in. Their responses are analyzed by an AI system that considers not just what they say but also their body language and tone. The best candidates from that round, as judged by the AI, are then invited to Unilever for in-person interviews, after which humans make the final hiring decisions.

It’s too early to tell whether the new recruiting process has resulted in better employees. The company has been closely tracking the success of those hires, but more data is still needed. It is clear, however, that the new system has greatly broadened the scale of Unilever’s recruiting. In part because job seekers can easily access the system by smartphone, the number of applicants doubled to 30,000 within a year, the number of universities represented surged from 840 to 2,600, and the socioeconomic diversity of new hires increased. Furthermore, the average time from application to hiring decision has dropped from four months to just four weeks, while the time that recruiters spend reviewing applications has fallen by 75%.

Decision making. By providing employees with tailored information and guidance, AI can help them reach better decisions. This can be especially valuable for workers in the trenches, where making the right call can have a huge impact on the bottom line.

Consider the way in which equipment maintenance is being improved with the use of “digital twins”—virtual models of physical equipment. General Electric builds such software models of its turbines and other industrial products and continually updates them with operating data streaming from the equipment. By collecting readings from large numbers of machines in the field, GE has amassed a wealth of information on normal and aberrant performance. Its Predix application, which uses machine-learning algorithms, can now predict when a specific part in an individual machine might fail.

This technology has fundamentally changed the decision-intensive process of maintaining industrial equipment. Predix might, for example, identify some unexpected rotor wear and tear in a turbine, check the

turbine's operational history, report that the damage has increased fourfold over the past few months, and warn that if nothing is done, the rotor will lose an estimated 70% of its useful life. The system can then suggest appropriate actions, taking into account the machine's current condition, the operating environment, and aggregated data about similar damage and repairs to other machines. Along with its recommendations, Predix can generate information about their costs and financial benefits and provide a confidence level (say, 95%) for the assumptions used in its analysis.

Without Predix, workers would be lucky to catch the rotor damage on a routine maintenance check. It's possible that it would go undetected until the rotor failed, resulting in a costly shutdown. With Predix, maintenance workers are alerted to potential problems before they become serious, and they have the needed information at their fingertips to make good decisions—ones that can sometimes save GE millions of dollars.

Personalization. Providing customers with individually tailored brand experiences is the holy grail of marketing. With AI, such personalization can now be achieved with previously unimaginable precision and at vast scale. Think of the way the music streaming service Pandora uses AI algorithms to generate personalized playlists for each of its millions of users according to their preferences in songs, artists, and genres. Or consider Starbucks, which, with customers' permission, uses AI to recognize their mobile devices and call up their ordering history to help baristas make serving recommendations. The AI technology does what it does best, sifting through and processing copious amounts of data to recommend certain offerings or actions, and humans do what they do best, exercising their intuition and judgment to make a recommendation or select the best fit from a set of choices.

The Carnival Corporation is applying AI to personalize the cruise experience for millions of vacationers through a wearable device called the Ocean Medallion and a network that allows smart devices to connect. Machine learning dynamically processes the data flowing from the medallion and from sensors and systems throughout the ship to help guests get the most out of their vacations. The medallion streamlines the boarding and debarking processes, tracks the guests' activities, simplifies purchasing by connecting their credit cards to the device, and acts as a room key. It also connects to a system that anticipates guests' preferences, helping crew members deliver personalized service to each guest by suggesting tailored itineraries of activities and dining experiences.

THE NEED FOR NEW ROLES AND TALENT


Reimagining a business process involves more than the implementation of AI technology; it also requires

a significant commitment to developing employees with what we call “fusion skills”—those that enable them to work effectively at the human-machine interface. To start, people must learn to delegate tasks to the new technology, as when physicians trust computers to help read X-rays and MRIs. Employees should also know how to combine their distinctive human skills with those of a smart machine to get a better outcome than either could achieve alone, as in robot-assisted surgery. Workers must be able to teach intelligent agents new skills and undergo training to work well within AI-enhanced processes. For example, they must know how best to put questions to an AI agent to get the information they need. And there must be employees, like those on Apple's differential privacy team, who ensure that their companies' AI systems are used responsibly and not for illegal or unethical purposes.

We expect that in the future, company roles will be redesigned around the desired outcomes of reimagined processes, and corporations will increasingly be organized around different types of skills rather than around rigid job titles. AT&T has already begun that transition as it shifts from landline telephone services to mobile networks and starts to retrain 100,000 employees for new positions. As part of that effort, the company has completely overhauled its organizational chart: Approximately 2,000 job titles have been streamlined into a much smaller number of broad categories encompassing similar skills. Some of those skills are what one might expect (for example, proficiency in data science and data wrangling), while others are less obvious (for instance, the ability to use simple machine-learning tools to cross-sell services).

MOST ACTIVITIES AT the human-machine interface require people to *do new and different things* (such as train a chatbot) and to *do things differently* (use that chatbot to provide better customer service). So far, however, only a small number of the companies we've surveyed have begun to reimagine their business processes to optimize collaborative intelligence. But the lesson is clear: Organizations that use machines merely to displace workers through automation will miss the full potential of AI. Such a strategy is misguided from the get-go. Tomorrow's leaders will instead be those that embrace collaborative intelligence, transforming their operations, their markets, their industries, and—no less important—their workforces. 🗣️

HBR Reprint R1804J

 **H. JAMES WILSON** is a managing director of information technology and business research at Accenture Research. **PAUL R. DAUGHERTY** is Accenture's chief technology and innovation officer. They are the coauthors of *Human + Machine: Reimagining Work in the Age of AI* (Harvard Business Review Press, 2018).