

MIT Sloan Management Review RESEARCH REPORT

In collaboration with



FINDINGS FROM THE 2020 ARTIFICIAL INTELLIGENCE
GLOBAL EXECUTIVE STUDY AND RESEARCH PROJECT

Expanding AI's Impact With Organizational Learning

By Sam Ransbotham, Shervin Khodabandeh, David Kiron,
François Candelon, Michael Chu, and Burt LaFountain

REPRINT NUMBER 62270

AUTHORS

SAM RANSBOTHAM (@ransbotham) is a professor in the information systems department at the Carroll School of Management at Boston College, as well as guest editor for *MIT Sloan Management Review's* Artificial Intelligence and Business Strategy Big Ideas initiative.

SHERVIN KHODABANDEH is a senior partner and managing director at BCG, and the co-leader of BCG GAMMA (BCG's AI practice) in North America. He can be contacted at shervin@bcg.com.

DAVID KIRON is the editorial director of *MIT Sloan Management Review*, where he directs the publication's Big Ideas program. He can be contacted at dkiron@mit.edu.

FRANÇOIS CANDELON is a senior partner and managing director at BCG, and the global director of the BCG Henderson Institute. He can be contacted at canelon.francois@bcg.com.

MICHAEL CHU is a partner and associate director at BCG, and a core member of BCG GAMMA. He can be reached at chu.michael@bcg.com.

BURT LAFOUNTAIN is a partner and managing director at BCG, and a core member of BCG GAMMA. He can be reached at lafountain.burt@bcg.com.

SPECIAL CONTRIBUTORS

Rodolphe Charme di Carlo and Allison Ryder

CONTRIBUTORS

Sylvain Durantou, Connie Gao, Carolyn Ann Geason, Tam Harbert, Michele Lee DiFilippo, Martin Reeves, Lauren Rosano, Leonor Saitkoulov, and Yunke Xiang

To cite this report, please use:

S. Ransbotham, S. Khodabandeh, D. Kiron, F. Canelon, M. Chu, and B. LaFountain, "Expanding AI's Impact With Organizational Learning," *MIT Sloan Management Review* and Boston Consulting Group, October 2020.

Collaborating organizations:



Copyright © MIT, 2020. All rights reserved.

Get more on artificial intelligence from *MIT Sloan Management Review*.

Read the report online at <https://sloanreview.mit.edu/ai2020>

Visit our site at <https://sloanreview.mit.edu/big-ideas/artificial-intelligence-business-strategy>

Get the free artificial intelligence enewsletter at <https://sloanreview.mit.edu/enews-artificial-intelligence>

Listen to our artificial intelligence podcast at <http://sloanreview.mit.edu/aipodcast>

Contact us to get permission to distribute or copy this report at smr-help@mit.edu or **617-253-7170**

- 1 / Executive Summary**
- 3 / Introduction: Continuous Learning With AI at Repsol**
- 4 / Few Organizations Get Significant Benefits With AI**
- 5 / Beyond Automation: Precision, Speed, and Learning**
- 6 / Learning to Learn With AI**
- 8 / Machines Must Learn, but So Must We**
- 10 / Better Together: How Humans and AI Interact**
- 12 / Learning With AI Requires Learning to Do Things Differently**
- 13 / Conclusion**

Expanding AI's Impact With Organizational Learning

Executive Summary

Only 10% of companies obtain significant financial benefits with artificial intelligence technologies. Why so few? Our research shows that these companies intentionally change processes, broadly and deeply, to facilitate organizational learning with AI. Better organizational learning enables them to act precisely when sensing opportunity and to adapt quickly when conditions change. Their strategic focus is organizational learning, not just machine learning.

Organizational learning with AI is demanding. It requires humans and machines to not only work together but also learn from each other — over time, in the right way, and in the appropriate contexts. This cycle of mutual learning makes humans and machines smarter, more relevant, and more effective. Mutual learning between human and machine is essential to success with AI. But it's difficult to achieve at scale.

Our research — based on a global survey of more than 3,000 managers, as well as interviews with executives and scholars — confirms that a majority of companies are developing AI capabilities but have yet to gain significant financial benefits from their efforts. More than half of all respondents affirm that their companies are piloting or deploying AI (57%), have an AI strategy (59%), and understand how AI can generate business value (70%). These numbers reflect statistically significant increases in adoption, strategy development, and understanding from four years ago. What's more, a growing number of companies recognize a business imperative to improve their AI competencies. Despite these trends, just 1 in 10 companies generates significant financial benefits with AI.

We analyzed responses to over 100 survey questions to better understand what really enables companies to generate significant financial benefits with AI. We found that getting the basics right — like having the right data, technology, and talent, organized around a corporate strategy — is far from

sufficient. Only 20% of companies achieve significant financial benefits with these fundamentals alone. Getting the basics right and building AI solutions that the business wants and can use improve the odds of obtaining significant financial benefits, but to just 39%.

Our key finding: Only when organizations add the ability to learn with AI do significant benefits become likely. With organizational learning, the odds of an organization reporting significant financial benefits increase to 73%.

Organizations that learn with AI have three essential characteristics:

1. They facilitate systematic and continuous learning between humans and machines. Organizational learning with AI isn't just machines learning autonomously. Or humans teaching machines. Or machines teaching humans. It's all three. Organizations that enable humans and machines to continuously learn from each other with all three methods are five times more likely to realize significant financial benefits than organizations that learn with a single method.

2. They develop multiple ways for humans and machines to interact. Humans and machines can and should interact in different ways depending on the context. Mutual learning with AI stems from these human-machine interactions. Deploying the appropriate interaction mode(s) in the appropriate context is critical. For example, some situations may require an AI system to make a recommendation and humans to decide whether to implement it. Some context-rich environments may require humans to generate solutions and AI to evaluate the quality of those solutions. We consider five ways to structure human-machine interactions. Organizations that effectively use all five modes of interaction are six times as likely to realize significant financial benefits compared with organizations effective at a single mode of interaction.

3. They change to learn, and learn to change. Structuring human and machine interactions to learn through multiple methods requires sig-

nificant, and sometimes uncomfortable, change. Organizations that make extensive changes to many processes are five times more likely to gain significant financial benefits compared with those that make only some changes to a few processes. These organizations don't just change processes to use AI; they change processes in response to what they learn with AI.

Organizational learning with AI demands, builds on, and leads to significant organizational change. This report offers a clear, evidence-based view about how to manage organizational learning with AI.

Introduction: Continuous Learning With AI at Repsol

In 2018, Repsol, a global energy and utility company based in Spain, launched an ambitious digital transformation effort. Up and down the company's value chain, Repsol executives identified and began implementing more than 190 digital transformation projects. Seventy percent of them now incorporate artificial intelligence in some way. These projects are not experiments: They make essential contributions to Repsol's business model, from upstream drilling operations to personalized offers at downstream retail service stations. "AI is now a cornerstone of an overwhelming majority of the most relevant digital cases transforming our business units," says Valero Marín, Repsol's CIO and chief digital officer.

At its upstream facilities, Repsol created an operations excellence center to use AI to optimize the process of drilling productive wells. Analyzing more than 100 million data points a day, AI solutions help reduce nonproductive time by 40%-50% across 30 drill sites. The AI system identifies inefficiencies and root causes, along with potential solutions, that subject-matter experts then evaluate before either taking action or feeding corrections back into the system for continuous improvement. Continuous refinement of the algorithm is an ongoing source of learning for both machine and management.

At its midstream refineries, Repsol prepares crude oil for commercial sale, a complex process that coordinates oil blending, storage, and delivery. Repsol uses AI to reduce more than 5 million scenarios to a small set of feasible production options for 30-day time frames. Human engineers then analyze the AI-generated options, incorporating difficult-to-quantify context, and feed their own analyses back to the AI system. Engineers and operations managers collectively changed how they work in order to work with AI.

Downstream, the company uses AI to prepare personalized offers for 8 million customers across its international network of more than 5,000 service stations. The system generates between 200,000

and 400,000 offers each day. After machine learning algorithms process the resulting outcomes data, managerial insights improve the AI systems that prepare the offers. The personalized promotions have led to a large sales increase — the equivalent of having 3%-4.5% more service stations. That is an important measure for Repsol, because regulations limit the company's ability to acquire or build new service stations. The results reflect an achievement that neither humans nor machines could achieve alone.

Notably, Repsol does more than teach machines how to drill, blend, and serve. In effect, Repsol changed

ABOUT THE RESEARCH

This report presents findings from the fourth annual research effort between *MIT Sloan Management Review* and Boston Consulting Group. We fielded a global survey in spring 2020, attracting over 3,000 total respondents representing 29 industries and 112 countries. We then interviewed 17 academics and executives researching or leading AI initiatives in large organizations in a broad range of industries, including insurance, software, banking, manufacturing, health care, retail, and mining.

Our assessment of "significant financial benefits" uses a threshold that varies based on organization size. For the largest organizations in our sample, with revenues of more than \$10 billion, passing this threshold requires more than \$100 million in revenue and/or cost improvements annually from the use of AI. For smaller organizations the thresholds were lower: \$20 million in improvements for organizations with revenues between \$500 million and \$10 billion, \$10 million in improvements for organizations with revenues between \$100 million and \$500 million (or nonprofits), and \$5 million for organizations with less than \$100 million in revenues. (Sensitivity analysis indicates that our conclusions do not qualitatively depend on these specific thresholds.) We used machine learning (clustering and tree-based models) to distill responses to over 100 survey questions down to seven areas of organizational AI effort. We then used logistic regression to quantify the relationship between the seven areas and significant financial benefits.

its processes to continuously learn with AI. Process improvements beget new behaviors and new human knowledge, which is then fed back to machines. These dynamics play out continuously, deliberately, and systematically across the organization. Repsol's ability to learn *with* AI is fundamental — as our research shows — to obtaining significant benefits with AI.

Based on over 3,000 survey responses from organizations in 29 industries around the world, as well as interviews with academics and executives, we offer insights into cross-industry adoption of AI, along with

a quantitative analysis of those companies achieving significant financial benefits with AI.¹ The bulk of this report offers an evidence-based view of how organizations learn with AI — both the scope of its challenge and its opportunity to transform the organization.


Few Organizations Get Significant Benefits With AI

Industry adoption of AI continues to increase as more companies perceive that AI drives both strategic opportunity and risk. As of 2020, 57% of companies report having AI pilots or deployed AI solutions, a large increase from 46% in 2017. Increasingly, AI-specific strategies guide these efforts. More than half of respondents (59%) say they have an AI strategy, up from 39% in 2017. (See Figure 1.)


FIGURE 1: FEW ORGANIZATIONS SEE SIGNIFICANT BENEFITS WITH AI

Despite increasing AI understanding, strategy development, and use, few organizations report significant financial benefits with AI.


We understand how AI will change the ways our organization generates business value



We have a strategy for AI



We are piloting or have deployed AI solutions



... but only
11%
of organizations report significant financial benefits in 2020

(Percentage of respondents who agree or strongly agree)

These AI strategies increasingly intersect with organizational strategy and essential operations. Indian telecommunications company Bharti Airtel uses AI to analyze satellite imagery to optimize the deployment of cell towers. With AI, Airtel quickly knows exactly where to put infrastructure “within one meter to get maximum revenue and provide the best service,” says CIO and head of cloud and security business Harmeen Mehta. “AI is not a sexy, shiny tool; it’s part of the company’s fundamentals. It’s a necessity for us.” Airtel hasn’t just found an application that fits AI. It has found that AI is necessary to compete.

Overall, a growing number of executives recognize that competitors are using AI, and they believe that AI will generate significant strategic benefits for their organization. Competitive dynamics, ecosystem forces, and financial incentives all motivate the increasing dependence on AI in business strategy. (See Figure 2, page 5.)

A deeper understanding of AI technologies and their business potential supports this belief: More than 70% of our survey respondents say they un-

Leaders share one outstanding feature: They intend to become more adept learners with AI.

derstand how AI will change the way they generate business value, a substantial increase from 57% in 2017. Despite these trends, only about 1 in 10 companies — a group we call Leaders — report obtaining “significant” financial benefits with AI. (We provide a detailed explanation of how we define and calculate these benefits in the “About the Research” section.) What counts as significant depends on company revenues. In our 2019 report,² we discussed the fact that only 3 in 10 companies reported any impact with AI; this year we delve deeper into the subset that obtain significant financial benefits.

Our analysis of Leaders’ survey responses highlights several useful insights. First, realizing significant financial benefits with AI requires far more than a foundation in data, infrastructure, and talent. Second, even embedding AI in business processes is not enough. Third, the returns on AI investments extend beyond immediate financial results. Financial benefits are a consequence of learning with AI, rather than the only goal.

Beyond Automation: Precision, Speed, and Learning

Our survey analysis demonstrates that Leaders share one outstanding feature: They intend to become more adept learners with AI. Organizations that sense and respond quickly and appropriately to changing conditions, such as a new competitor or a worldwide pandemic, are more likely to take advantage of those disruptions. They view AI as more than a tool for cost cutting and automation.


For example, German sports car manufacturer Porsche uses AI to make complicated region-specific production decisions to match inventory with local demand in cities around the world. Porsche CIO Matthias Ulbrich says, “Based on the AI, we continuously learn how to better tailor the precise mix of configurations of cars, out of millions of potential options, that we deliver to each market.” In a topsy-turvy global

economy, shifts in market demand and regulatory environments intensify the need for accurate, and continually adjusted, predictions. For Porsche, improving its ability to allocate the right products to the right market is an ongoing motivation to learn with AI.

But sensing changing conditions isn’t enough. The benefits of precision often require speed. Ninety percent of Leaders use AI to increase how quickly they

FIGURE 2: COMPETITIVE PRESSURES DRIVE AI STRATEGIES

Competitive dynamics and financial pressures encourage integrating AI with business strategy.



(Percentage of respondents who agree or strongly agree with each statement)

sense and respond to the business environment. Baidu, the China-based search company, places AI at every layer of its technology stack. “We made AI an early strategic imperative,” says Yanjun Ma, director of the deep learning platform at Baidu Research, “and it has had a direct impact on our precision, speed, and efficiency.” Speed helps Leaders capitalize on precision.

Precision, speed, and learning are not a function of financial investment alone; they require large-scale organizational shifts in mindsets, processes, and behaviors. “As more and more of the core of a company is built around software and data, the nature of the organization changes,” says Marco Iansiti, the David Sarnoff Professor of Business Administration at Harvard Business School. Rather than just applying AI to specific cases, a corporate architecture based on algorithms and data enables an organization to not just use AI, or even extensively automate with AI, but to learn with AI. This volume of organizational change takes time and effort. “It’s an architectural transition that takes a lot of time for a traditional organization,” cautions Iansiti. “It’s a massive change.”

Learning to Learn With AI


Our research affirms that learning with AI requires significant organizational effort. To learn more about how Leaders make these changes successfully when others struggle, we applied machine learning to our survey responses. These machine learning models distilled responses from over 100 questions and identified seven activities that affect the likelihood of becoming a Leader. (See Figure 3.)

Organizations that develop competencies in these seven activities can raise their probability of becoming a Leader from 2% to 73%. These increases won’t come from a series of targeted applications of AI, even if those applications benefit from the organization’s growing experience with the technology. Instead, organizations increase their potential benefits with AI with each activity. We group these activities into four categories:

Discovering AI. AI implementations in targeted areas can (even in isolation) occasionally yield surprisingly sizable returns. For example, a sophisticated

FIGURE 3: ACHIEVING LEARNING WITH AI

These seven activities dramatically increase the likelihood that an organization will achieve significant financial benefits with AI.



(Percentage is the likelihood of achieving significant financial benefits with AI)

CROSS-FUNCTIONAL SOLUTION TEAMS, NOT AI TEAMS

Leaders set up a “business problem-solving team” rather than an “AI team.” Our survey results show that nearly 80% of Leaders assemble teams from multiple functional areas to capitalize on their respective experiences and strengths (76%), gather feedback from end users when prototyping the AI system (78%), and use this feedback to evolve solutions from the original design during subsequent development (80%). As a result, Leaders scale AI solutions only after they prove their positive financial impact (74%). Their focus isn’t on how to use AI; it’s on solving business problems.

We observe this business focus at both tech companies and nontech companies alike. For example, Arun Narayanan, chief data officer of global mining company Anglo American, mentions that his company’s team is not an AI team but rather “like a product team in a software company, with a focus on services, communications, marketing, customer engagement, sustainability, etc. Together, they diagnose the business challenge and then define holistic solutions to solve that problem, with or without AI.”

Walmart vice president of machine learning Prakhar Mehrotra gradually built a relationship between data science and managers after arriving from Uber. He spent time learning retail literally from the ground up, walking the store aisles. He hosted meetings to demonstrate to managers what AI can and cannot do. He assembled teams of merchandisers and supply chain managers as well as data scientists and machine learning engineers. “It’s a full-stack team that works on a problem. They brainstorm ideas. Usually, the managers are very, very good at correcting the algorithm before it ever hits the business,” he says. “I want to test the algorithm with their experience. I want to know all the mistakes it can make. Managers and data scientists have a shared goal to make the algorithm better.”

model that reduces customer churn can, by itself, create substantial value. Organizations that are still developing their infrastructure, talent, and strategy to embrace AI have some small potential (2%) of realizing more than incremental value. While 2% certainly isn’t a high probability, it attests to the potential of AI even in targeted (but important) applications.

Building AI. To advance from targeted applications to organizational use of AI, companies invest in data, technology, and algorithms to build AI capabilities.³ They marshal the technical skills necessary to capitalize on those investments.⁴ Their overall business strategy inherently incorporates the capabilities of AI, without relegating AI to a side strategy.⁵ Organizational efforts in these areas increase their probability of becoming Leaders from 2% to 21%. But even this tenfold increase doesn’t garner more than an absolute chance of 1 out of 5 for significant financial benefits.

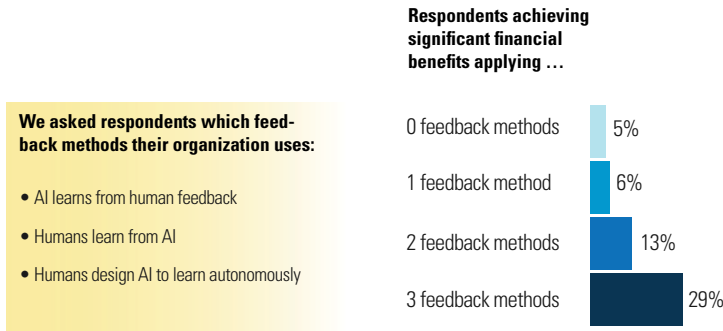
Scaling AI. Adding the ability to embed AI into processes and solutions improves the likelihood of significant benefits dramatically, but only to 39%. Several factors contribute to this improvement. With many technology options, effectively scaling AI overcomes the difficulty of matching the right technology to a specific problem. (See the sidebar “Cross-Functional Solution Teams, Not AI Teams.”) From the production side, embedding AI deeply and broadly within processes demonstrates that an organization can solve problems and execute projects effectively. From the consumption side, managers throughout the organization learn more about how to use AI, how AI models work, and how much (or how little) confidence to put in them. Scaling AI is clearly about more than enterprise-wide process automation. Organizations that effectively produce and consume AI create some business value.⁶ However, the absolute chance of obtaining significant financial benefits with discovering, building, and scaling AI is still much less than 50%.

Organizational learning with AI. Some organizations almost double again their prospects of obtaining significant financial benefits with AI (to 73%). How? They create opportunities for mutual learning between humans and AI. They haven’t just learned to teach machines what they, as humans, know; they’ve learned

how to use human-machine interactions to refine processes quickly as circumstances change. They haven’t uncovered a single prescriptive structure for human and machine roles in processes; they’ve learned how to adapt human and machine roles to each situation. Growing collective organizational knowledge, in the form of both digital data and human experience, supports ever-improving decision-making. Organizational learning with AI becomes a systematic, continuous process of improvement. Both aspects — ensuring that people and AI learn from each other, and structuring effective ways to work together — require significant, complex changes to an organization. No wonder so few of them have learned how to learn with AI.

FIGURE 4: ENABLING LEARNING THROUGH FEEDBACK

Combining multiple feedback methods increases potential success with AI.



(Percentage is the likelihood of achieving significant financial benefits with AI)

Machines Must Learn, but So Must We

Organizational learning with AI isn't just machines learning autonomously. Or just humans teaching machines. Or just machines teaching humans. Organizations that use only one of these methods have only about a 6% likelihood of reporting significant financial benefits.

More successful organizations combine multiple teaching and learning methods to take advantage of humans' and machines' distinctive strengths and weaknesses. For example, humans work well with broad contexts and new situations, while machines retain and process vast data well. Enabling humans and machines to learn from each other often produces better results than either could achieve alone. Mutual learning enables humans and machines to understand the other's logic and reasoning: It is not simply having a human check the output of the AI and vice versa. For example, ensuring that the AI knows why the human is overriding a decision or favoring a certain outcome improves the algorithm. Ensuring that the human knows why the AI is making a decision ensures that humans can trust the outcome and can learn from the AI's decision-making.

ing. Organizations that effectively apply multiple feedback methods are six times more likely to realize significant financial benefits than organizations with a single feedback method. (See Figure 4.)

We surveyed respondents on three human-AI learning methods: AI learns from human feedback, humans learn from AI, and humans design AI to learn autonomously. The survey results, in combination with our executive interviews, show the potential benefits of using each feedback method and integrating all methods in a single solution.

AI learns from human feedback. Like any new worker, many AI algorithms go through an adjustment period in a new business environment. They learn from human mentors, drawing on human business experience to improve their effectiveness over time. At DHL, a machine-vision-based tool helps workers inspect shipping pallets to determine how stackable they are. This helps the company better utilize space in cargo planes while saving time and increasing revenue. The AI is trained to spot markers and other signs of troublesome pallets, and it integrates expertise from workers who have years of experience identifying non-stackable pallets. But, as Gina Chung, vice president of innovation at DHL says, "The worst day for AI is the first day, because the algorithm gets better over time." Built-in learning is intentional, says Chung, "especially with the early versions of AI that we release into operations. There's always a loop to ensure that we capture feedback." This feedback is essential to help the AI identify "a pallet that can't be stacked for some reason, such as the fragility of its contents," she explains. A new worker, even a machine, can learn from experienced coworkers.

Humans learn from AI. Managers typically understand what affects outcomes in their businesses. But, in some cases, AI can offer alternatives that humans haven't considered. Early on at Lyft, engineers designed an algorithm to maximize revenue by matching driver supply and customer demand. "It looked at all the possible combinations of riders and drivers and picked the combination that — based on the ride being requested, where the driver was located, all of the system dynamics — would maximize revenue," says Elizabeth Stone, former vice president

of science at Lyft. Then, as data scientists began testing other objectives, something interesting emerged. One AI solution discovered that optimizing conversion rates — the percentage of times a user actually ordered a ride after opening the app — would deliver more ride requests in the future. More ride requests ultimately means more revenues — measurably more than the company’s former focus on matching supply and demand produced. Having humans in the loop who could think through and test possible objectives for the machine learning algorithms was critical. Combining human knowledge about the business with AI’s computation power enabled Lyft to improve a key strategic metric. The change influenced an array of business activities, including operations, revenue targets, performance management, and marketing.

Humans design AI to learn autonomously. AI systems can learn autonomously. However, even if AI systems do most of their learning without human intervention, humans still design AI learning strategies. Porsche’s Ulbrich describes how a coffee maker helped the sports car manufacturer improve a quality assurance and corporate process. One engineer realized that the coffee maker sounded different depending on how well it was working — a watery cappuccino sounded different from a delicious espresso. After discussion, that observation led Porsche’s technology unit in Berlin to realize that it could similarly detect differences in the quality of vehicle components, such as car doors, based on the sounds they made. By feeding quality assurance results back to the acoustic anomaly detector, the AI system learned how to hear potential defects. So although a human had the idea for an acoustic anomaly system, implementation required AI because, as Ulbrich observes, “humans cannot listen 24 hours a day the way an AI algorithm can.” Gradually, the system taught itself to recognize potential defects in the production process. Humans can design AI to learn autonomously.

The three feedback methods are, individually, capable of yielding significant business value. But using all three methods in a single solution can deliver even more value.

We spoke with the head of AI research at one financial institution, where her team had deployed AI as a decision support tool for traders. In this scenario, human

THREE ARE BETTER THAN ONE

At one financial institution, high-frequency traders and AI algorithms learn a lot from one another and by themselves. First, the AI learns by observing traders and the screens and charts they are monitoring. It learns which specific data points to focus on — out of the vast volume of data that financial markets create each second. This observation significantly narrows down the information that the AI system must process, allowing it to ignore noise.

Second, the AI system learns autonomously from these observations, identifying intricate relationships between data points and trades made. In effect, the AI uncovers the underlying decision-making logic among traders on the bank’s platform.

Third, the trader learns from the information the AI shares. Because the AI observes the history of multiple traders, not just one, it facilitates learning between traders.

And the cycle of learning begins again as the AI observes how traders use the information it provides. The AI system observes that the traders who use this system have collectively improved the speed and effectiveness of their decision-making. Mutual learning is more than providing specific corrective feedback in specific situations. It is an ongoing partnership that improves human and machine decision-making alike through multiple modes of interaction between AI and humans.

traders had the edge over autonomous AI algorithms. Even so, the team found a way to use AI to improve traders’ effectiveness by combining all three feedback methods. Initially, the algorithm learned from human traders which market information — among billions of possible data points — was most salient by observing which market indicators the traders spent the most time monitoring. The algorithm then learned autonomously to identify patterns among salient information and shared these insights with traders. Subsequently, traders learned of new information from the AI system that directly improved their decision-making. “AI support systems bring relevant information to traders’ attention,” says the executive. “Even though AI can’t make the trades directly, it adds value by increasing the efficiency of the trader in processing information.” (See the sidebar “Three Are Better Than One.”)

Mutual learning is not just about providing specific corrective feedback in a specific situation. Rather, it is an ongoing process that improves both human and machine

decision-making. Implementing this process requires multiple modes of interaction between AI and humans.

Better Together: How Humans and AI Interact


Matching the right mode of human-machine interaction to the situation affects whether feedback advances or limits organizational learning with AI.

In some processes, managers design AI solutions to largely work alone. For example, at toolmaker Stanley Black & Decker, image-processing algorithms monitor the quality of its tape-measure manufacturing. Cameras capture images as tape measures pass through various points of manufacture and flag defects in real time before the company wastes additional resources on defective tapes. These AI systems work independently in real time because waiting for human input would slow the process. But humans still have a role, because as Carl March, Industry 4.0 director of analytics at Stanley Black & Decker, says, “Defects, at times, still warrant some actions to give additional validation, as there are sometimes gray areas.” The process still involves human effort in exceptional cases, but it doesn’t have to slow the main process flow.

Other processes may require a much greater role for human discernment and input. At Walmart, employees not only have extensive operational experience managing in-store product assortment but also an ability to understand context that even extensive historical data lacks — such as sudden, radical shifts due to the COVID-19 pandemic. Machine learning depends on historical data being relevant to current and future states. “But when faced with COVID,” Walmart vice president of machine learning Prakhkar Mehrotra says, “the world completely changed and we could no longer predict the future from the past.” For assortment management, Walmart designs AI solutions to present recommendations to managers. Managers can agree, disagree, and comment to improve both the current recommendation and future recommendations. In a COVID-19 world, management disagreement with AI solutions was a critical source of new machine learning. Walmart ensures that decisions reflect the entirety of knowledge available — in databases as well as in people’s heads. Well-designed roles support mutual learning.

FIGURE 5: LEADERS SUCCESSFULLY DEPLOY MULTIPLE INTERACTION MODES

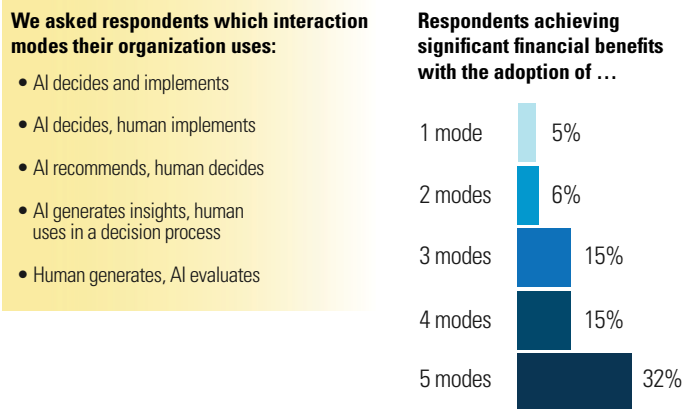
Leaders successfully integrate humans and AI.



(Percentage of Leaders reporting success with each mode)

FIGURE 6: MORE INTERACTION MODES LEAD TO BETTER RESULTS

Organizations adopting multiple modes of human-AI interaction are more likely to achieve significant financial benefits with AI.



(Percentage of Leaders reporting success with each mode)

Our survey results reveal five ways humans and AI work together. Leaders use each mode well but report particular proficiency integrating human and AI roles. (See Figure 5, page 10.)

Leaders haven't found just one way to structure and refine human-AI interactions. Instead, they deploy multiple modes of human-AI interaction. Organizations that successfully use all five modes are six times more likely to attain significant financial benefits than those able to use just one or two. Furthermore, companies gain the most when they increase their expertise from four to five modes. (See Figure 6, page 10.) Broader competencies allow organizations to fit a wider variety of interaction modes to a wider variety of situations.

Repsol illustrates how a single organization can use these five distinct modes of human-AI interaction, deploying each as situations require.

AI decides and implements: In this mode, AI has nearly all the context and can quickly make decisions. Human involvement would only slow down an otherwise fast process. Repsol embeds AI in its customer relationship management system to deliver real-time personalized offers, like discounts and free car washes, to consumers at its 5,000 retail service stations, with humans providing only a light layer of oversight and supervision and maintaining compliance with local regulations. (See the sidebar "Responsible AI.")

AI decides, human implements: AI can capture the context well and make decisions, but humans — rather than software or robotics, for instance — implement the solutions. Repsol uses this mode for AI predictive maintenance in offshore production facilities. AI identifies parts at risk of failure and schedules a maintenance review. Post-review, human operators then schedule the replacement, taking into account part availability and scheduled maintenance.

AI recommends, human decides: This mode is appropriate when organizations must make a large number of decisions repeatedly and the AI can incorporate most but not all of the business context. Repsol's AI for crude-oil blending integrates and analyzes millions of factors, including the type of

RESPONSIBLE AI

The term *responsible AI* doesn't yet have a universally accepted definition. For our survey, we defined responsible AI in terms of improving fairness in algorithms and reducing biases in decision-making; promoting inclusivity and a diversity of perspectives; providing model interpretability and explainability to AI end users; ensuring data privacy and security in AI; complying with legal requirements; and monitoring the social impact and ethics of AI. Among Leaders, 90% say they are developing or already have developed responsible AI strategies. But action lags intent — only 57% of Leaders have specific roles and processes to enforce responsible AI. A smaller percentage have an appointed leader for responsible AI (52%), and an even smaller percentage offer training on responsible AI processes and practices (43%).

That doesn't surprise Alice Xiang, head of fairness, transparency, and accountability research at the Partnership on AI, a research institute focused on the intersection of AI and society. "There is a constant struggle," she says. "Organizations often don't have a business incentive to engage and invest in this space unless they are subject to legal or reputational risk."

Compliance and risk management need not be the only motivations for adopting responsible AI practices. Of the organizations that already have in place responsible AI strategies, 72% find that these strategies actually increase the financial benefits of AI, and 62% report that they decrease operational risk. If organizations view responsible AI practices only from a compliance perspective, they may miss out on financial benefits. Responsible AI doesn't have to be altruistic; it can make financial sense.

Mutual learning is an ongoing process that improves both human and machine decision-making.

crude and the operating conditions of the refinery, to recommend blending schedules for the next 30 days. Humans then decide which blending schedule to use, depending on expected global market conditions.

AI generates insights, human uses them in a decision process: In this mode, inherently creative work requires human thought, but AI insights can inform the process. For workforce planning, algo-

"Algorithms don't know an org chart ... They cut across the organization."

—Prakhar Mehrotra, vice president, machine learning, Walmart

rithms at Repsol combine forecasts from machine learning models with human experience and insight to determine future workforce needs. Human resource managers use the result as input to hiring and training plans.

Human generates, AI evaluates: Humans generate many hypothetical situations but rely on AI to tediously assess many complex dependencies. For example, Repsol uses digital twins of physical assets, such as wells, to simulate the consequences of possible operational changes and validate hypotheses. Using multiple engineering and operational efficiency models, managers can simulate consequences before actually changing a physical well.

Furthermore, organizations not only need to use these multiple modes of human-machine interaction but also must be able to switch between these interaction modes as changing contexts demand. For example, the sudden and massive changes in behaviors driven by 2020 COVID-19 pandemic lockdowns showed that certain algorithms required more human oversight. At Walmart, for example, Mehrotra describes the early days of the pandemic as “a type of cold-start problem” and says the question was, “How quickly can AI agents learn what has recently happened?” Initially, processes required a larger human role, but as data accrued, the AI system could make better recommendations. Even well-designed processes may change as situations demand, and Leaders are prepared to change defined roles as required.

Our research shows that organizational learning with AI depends on success with three aspects of human-machine interactions: breadth (using as many of the five modes as possible), fit (selecting the appropriate mode for each context), and agility (switching between modes as needed).⁷

Learning With AI Requires Learning to Do Things Differently

Organizational learning with AI demands a complex set of organizational behaviors. Learning how to apply one or more modes of human-machine interaction in a given use case is one thing. Learning how and when humans and machines should give or receive feedback is another. Adapting to what the organization is learning with AI invites new organizational architectures, processes, behaviors, and attitudes. Learning with AI entails significant, and sometimes uncomfortable, change.

In the tape inspection example, Stanley Black & Decker didn't just swap a machine inspector for a human inspector in a quality control process. The real-time capabilities of AI enabled it to move inspection processes further upstream to reduce additional processing and catch systemic production errors before large numbers of defects created significant waste. This early-detection effort then helped production improve overall. In this situation, Stanley Black & Decker didn't automate a routine process; it learned how to produce better. What's more, the change wasn't local to a quality control function — it affected many parts of the production process.

As Walmart's Mehrotra puts it, “Algorithms don't know an org chart. They don't know the definition of merchandising versus supply chain functions. They cut across the organization.” Algorithmic indifference to organizational norms, roles, and behaviors demands a managerial response. In short, management needs to adapt to process changes that result from mutual learning with AI. Such adaptation is not easy to achieve.

“It's very hard to change managers,” notes Kartik Hosanagar, the John C. Hower Professor of Technology and Digital Business at the Wharton School. “AI is a huge shift. Companies need to take the long-term perspective, be willing to take a lot of hits in the short term, accept failure, and yet continue to

It is difficult to create significant value with AI because it is difficult to learn with AI.

invest.” This reflects another way that organizations need to learn with AI. They need to recognize, for example, the potential discomfort that AI may instigate. They need to learn when the improvements with AI offset the discomfort. They need to learn how to work with AI, not just at an individual level but across the organization.

It is difficult to create significant value with AI because it is difficult to learn with AI. The good news is that this degree of change and discomfort can be worth it. Respondents in the survey who reported “making extensive changes to many processes” at their company were five times more likely to report significant financial benefits than those who made only “small changes to a few processes.” (See Figure 7.)

It isn’t that change itself creates value or causes the organization to learn. The relationship isn’t so clearly defined. Instead, learning with AI simultaneously demands, builds on, and leads to significant organizational change. Still, the net effect we observe is change — substantial change. The resulting process changes are both extensive and widespread throughout the organization.

Conclusion

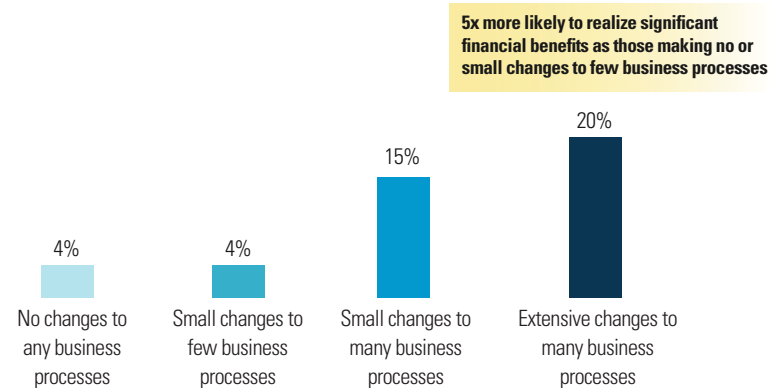
Adoption of AI continues to increase, and many organizations now use AI technologies to generate some business value.

But significant financial benefits are elusive, and few organizations achieve them. Many organizations struggle to build an AI foundation that rests on the right data, technology, and talent. Or they may have built this foundation, use it to churn out AI solution after solution, and yet wonder why the financial benefits are only incremental. Significant financial benefits are likely only when organizations define multiple, effective ways for humans and AI to work and learn together.

Although rare, these successes aren’t confined to a single industry or available only to digital natives or large companies. Instead, these most successful

FIGURE 7: BUSINESS PROCESS CHANGE CONNECTS TO AI OUTCOMES

Organizations that extensively change business processes when integrating AI solutions are five times as likely to realize significant financial benefits.



(Percentage of respondents who agree or strongly agree with each statement)

organizations distinguish themselves through effort and a commitment to learn with AI. They don’t just get good at working with machines; they get good at tailoring human and machine roles dynamically as situations change. They don’t facilitate machine learning; they facilitate mutual learning.

They don’t just use AI; they learn with AI.

Reprint 62270.

Copyright © Massachusetts Institute of Technology, 2020.

All rights reserved.

ACKNOWLEDGMENTS

We thank each of the following individuals, who were interviewed for this report:

Gina Chung, vice president of innovation and head of Americas Innovation Center, DHL

Gail Evans, global chief digital officer, Mercer

Kartik Hosanagar, John C. Hower Professor of Technology and Digital Business, the Wharton School, University of Pennsylvania

Marco Iansiti, David Sarnoff Professor of Business Administration, Harvard Business School

Yanjun Ma, director, deep learning platform, Baidu

Carl March, director of analytics, Industry 4.0, Stanley Black & Decker

Valero Marín, CIO and chief digital officer, Repsol

Prakhar Mehrotra, vice president of machine learning, Walmart

Harmeen Mehta, CIO and head of cloud and security business, Bharti Airtel

Kate Mulroney, senior director, advanced analytics and AI, Novo Nordisk

Arun Narayanan, chief data officer, Anglo American

Elizabeth Stone, former vice president of science, Lyft

Mattias Ulbrich, CIO, Porsche, and CEO, Porsche Digital

Manuela Veloso, managing director, head of AI research, J.P. Morgan

Alice Xiang, head of fairness, transparency, and accountability research, Partnership on AI

The research and analysis for this report was conducted under the direction of the authors as part of an *MIT Sloan Management Review* research initiative in collaboration with and sponsored by Boston Consulting Group.

REFERENCES

1. Our interpretation of “significant financial benefits” necessarily varies based on company revenue. See “About the Research” for details.
2. S. Ransbotham, S. Khodabandeh, R. Fehling, et al., “Winning With AI,” MIT Sloan Management Review and Boston Consulting Group, Oct. 15, 2019, <https://sloanreview.mit.edu>.
3. S. Ransbotham, D. Kiron, P. Gerbert, et al., “Reshaping Business With Artificial Intelligence,” MIT Sloan Management Review and Boston Consulting Group, Sept. 6, 2017, <https://sloanreview.mit.edu>.
4. S. Ransbotham, P. Gerbert, M. Reeves, et al., “Artificial Intelligence in Business Gets Real,” MIT Sloan Management Review and Boston Consulting Group, Sept. 17, 2018, <https://sloanreview.mit.edu>.
5. Ransbotham et al., “Winning With AI.”
6. Ransbotham et al., “Winning With AI.”
7. Our survey results indicate that smaller businesses — in contrast with larger organizations — capture most of their AI-related benefits from learning with the most autonomous mode (AI decides and implements). One reason may be that smaller companies typically have less organizational complexity, which can make it easier to realize significant financial benefits from automation than at larger companies.

MIT SLOAN MANAGEMENT REVIEW

At *MIT Sloan Management Review (MIT SMR)*, we share with our readers an excitement and curiosity about how the practice of management is transforming in the digital age. Our expert contributors help leaders explore the trends that are shaping how organizations operate, compete, and create value in a technology-fueled world. We deliver the kind of evidence-based analysis and practical insight that will inspire readers to do great work.

MIT SLOAN MANAGEMENT REVIEW BIG IDEAS

MIT Sloan Management Review's Big Ideas Initiatives develop innovative, original research on the issues transforming our fast-changing business environment. We conduct global surveys and in-depth interviews with front-line leaders working at a range of companies, from Silicon Valley startups to multinational organizations, to deepen our understanding of changing paradigms and their influence on how people work and lead. Working with collaborative organizations, we publish our findings in annual research reports, case studies, videos, articles, and interactive data visualizations. For over a decade, *MIT SMR's* Big Ideas Initiatives have explored topics such as artificial intelligence, digital leadership, sustainability, data and analytics, the internet of things, and performance management.

BOSTON CONSULTING GROUP

Boston Consulting Group partners with leaders in business and society to tackle their most important challenges and capture their greatest opportunities. BCG was the pioneer in business strategy when it was founded in 1963. Today, we help clients with total transformation — inspiring complex change, enabling organizations to grow, building competitive advantage, and driving bottom-line impact.

To succeed, organizations must blend digital and human capabilities. Our diverse, global teams bring deep industry and functional expertise and a range of perspectives to spark change. BCG delivers solutions through leading-edge management consulting along with technology and design, corporate and digital ventures — and business purpose. We work in a uniquely collaborative model across the firm and throughout all levels of the client organization, generating results that allow our clients to thrive.

BCG HENDERSON INSTITUTE

The BCG Henderson Institute is Boston Consulting Group's strategy think tank, dedicated to exploring and developing valuable new insights from business, technology, science, and economics by embracing the powerful technology of ideas. The Institute engages leaders in provocative discussion and experimentation to expand the boundaries of business theory and practice and to translate innovative ideas from within and beyond business. For more ideas and inspiration from the Institute, please visit bcghendersoninstitute.com.

BCG GAMMA

BCG GAMMA is BCG's global team dedicated to applying artificial intelligence and advanced analytics to critical business problems at leading companies and organizations. The team includes 900-plus data scientists and engineers who utilize AI and advanced analytics expertise (e.g., machine learning, deep learning, optimization, simulation, text and image analytics, etc.) to build solutions that transform business performance. BCG GAMMA's approach builds value and competitive advantage at the intersection of data science, technology, people, business processes, and ways of working. For more information, please visit our webpage.

MIT Sloan

Management Review

PDFs • Reprints • Permission to Copy • Back Issues

Articles published in *MIT Sloan Management Review* are copyrighted by the Massachusetts Institute of Technology unless otherwise specified.

MIT Sloan Management Review articles, permissions, and back issues can be purchased on our website, **shop.sloanreview.mit.edu**, or you may order through our Business Service Center (9 a.m. - 5 p.m. ET) at the phone number listed below.

Reproducing or distributing one or more *MIT Sloan Management Review* articles **requires written permission.**

To request permission, use our website **shop.sloanreview.mit.edu/store/faq**, email **smr-help@mit.edu**, or call 617-253-7170.