McKinsey Analytics

Entering a new decade of AI: The state of play

While organizations report measurable benefits from artificial intelligence, much work remains to scale impact, manage risks, and retrain the workforce.

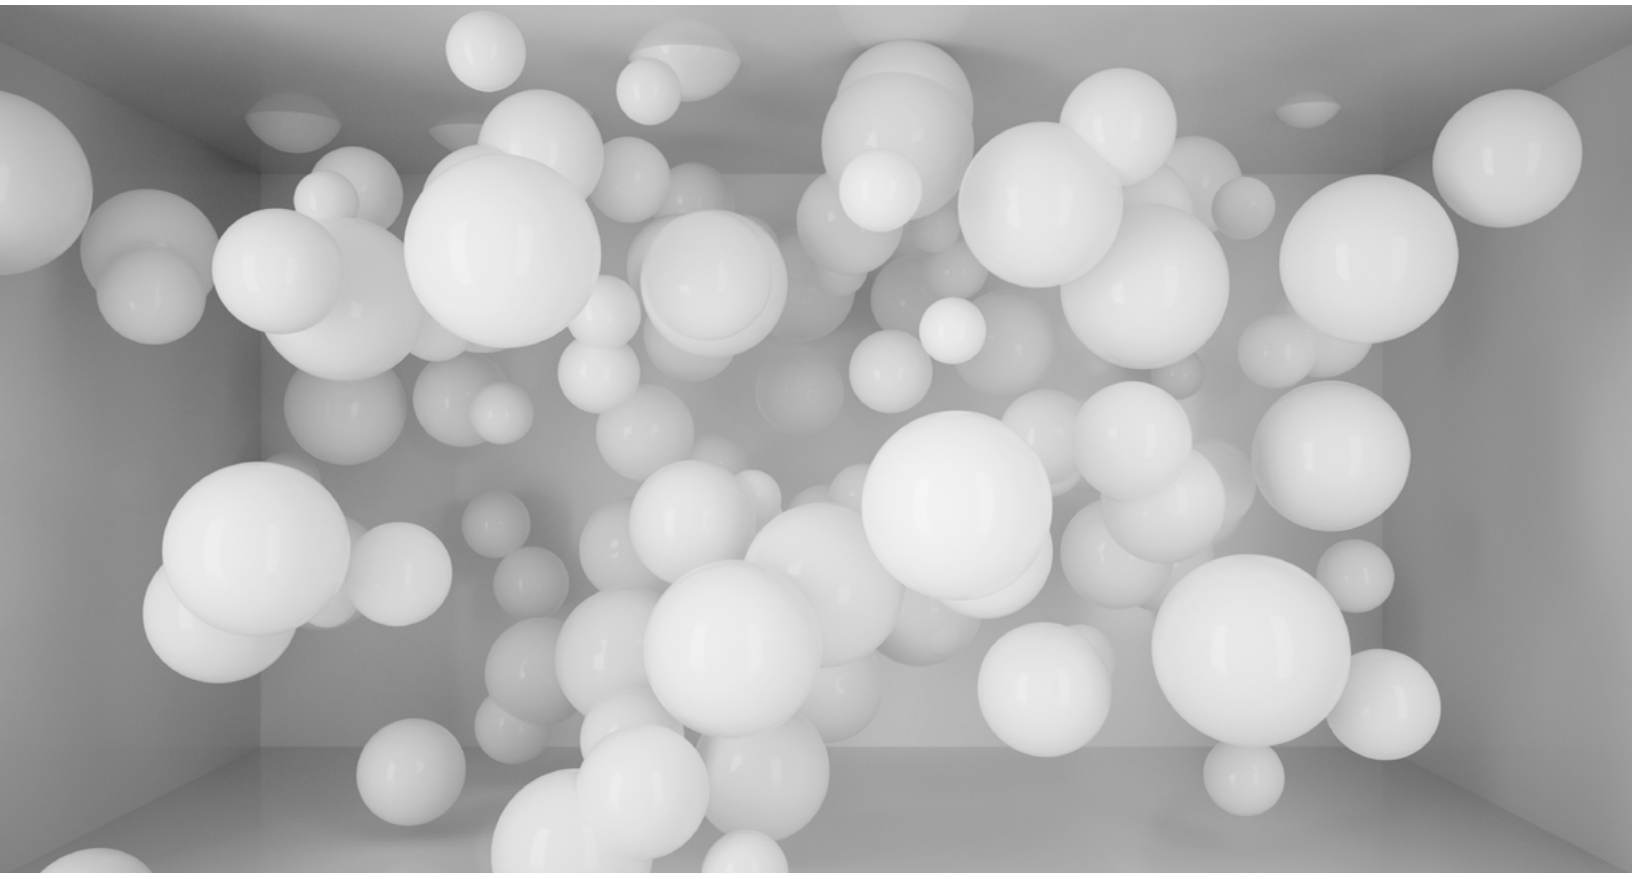

June 2020

In this episode of the *McKinsey on AI* podcast miniseries, McKinsey's David DeLallo speaks with McKinsey Global Institute partner Michael Chui and associate partner Bryce Hall about the latest trends in business adoption of artificial intelligence (AI). They discuss where the technology is being used most across industries, companies, and business functions; the keys to getting impact from AI investments and what lies ahead.

Podcast transcript

David DeLallo: Artificial intelligence. There's no shortage of predictions about how it could fundamentally change the way we live and work. Over the past few years, companies around the world have been figuring out exactly how AI technologies can improve their performance in a number of areas across their business.

But is AI actually delivering significant results? Moreover, what can we expect to see as we move into a new decade of AI use and development? I'm David DeLallo with McKinsey Publishing. To answer some of these questions today, I'm joined by Michael Chui, a McKinsey partner with the McKinsey Global Institute, who is based in our San Francisco office, and associate partner Bryce Hall from our Washington, DC, office.

Both have been doing plenty of work with businesses on AI and have also conducted a good amount of research in this space. Bryce and Michael, thanks so much for joining the podcast today.

Bryce Hall: Thank you.

Michael Chui: Thanks, David.

David DeLallo: To start off, it would be great to get a bit of a lay of the land here. Michael, can you start by answering the first question of where we are today overall in terms of adoption of AI across businesses?

Michael Chui: We've been fortunate to work with the *McKinsey Quarterly* to survey thousands of different executives around the world on their use of AI. In our most recent survey, we received responses from 2,300 executives globally.

And what we saw were significant increases in adoption levels of AI basically throughout the world, in North America, Europe, Asia–Pacific, and Latin America. And in Asia-Pacific and Latin America, the shares of respondents who say their companies have embedded AI across multiple functions or business units nearly doubled since the previous survey. When we look across all of these regions, even China, we see similar aggregate reported levels of adoption, even though there is a lot of variation within countries and across countries from individual respondents.

Bryce Hall: I would agree. Overall this research is showing significant growth in AI adoption. Across the board, we saw nearly a 25 percent year-over-year increase in the use of AI in standard business processes. We've moved beyond the phase of "Is AI a shiny object?" toward broader mainstream adoption and actual value creation.

David DeLallo: What are some differences that we see across companies, in terms of where they seem to be placing their AI investments?

Michael Chui: One of the things that we've discovered is the broad applicability of this technology across every sector and function. And that the companies that are able to drive the most value from certain use cases are in fact the most likely to be deploying AI within those use cases.

For instance, for companies that drive a lot of their value from top-line levers, such as marketing and sales, we see more investments, more experiments, and more value [from AI] being captured by those types of companies in marketing and sales, whether it's better segmentation, more personalized offerings, and other things of that sort.

Conversely, we also see companies that drive a lot of their value from bottom-line or operational levers investing in [AI in] those areas. So whether it's improving their logistics function, being able to reduce their inventory, increasing their inventory turns, increasing their OEE [overall equipment effectiveness], those are the types of use cases that those companies have started to invest in and have been capturing value on the cost side.

One of the other things that is interesting is that we're very early in this trend. As much as we're seeing this growth in adoption, less than a third of the companies that we surveyed have deployed AI in multiple businesses or functions. And so there's a lot of additional growth possible, given that AI can create value all across the business.

David DeLallo: So 30 percent of businesses have adopted AI across multiple areas in their company. Did you, Michael, find that number to be higher than you expected? Lower than you expected? Right on target?

Michael Chui: Well, it's funny. There's so much hype about AI that I think you could easily have expected people to think that there would be a lot more adoption. But we also know from our studies of doing work across many different technologies that it's hard work to deploy technology within an organization, not only because the technology problems are hard but also because the change management is really hard.

So in some ways it's not surprising that there's so much additional opportunity here. And we certainly know lots of companies that are going after that opportunity. But it is really hard work. There's a reason why this doesn't happen faster.

David DeLallo: So Bryce, I know you've done a lot of work with companies in helping them identify best practices to deploy this technology across business. And I believe we looked at some of that in the survey. Can you comment on some of the ways that you're finding companies are able to embed this more easily across their business?

Bryce Hall: Overall, based on this research and our work with companies in analytics, adoption of analytics, analytics transformation, and AI adoption, it's never just a technology challenge. One of the things that's been interesting about this research is the consistency in the core practices we've identified when we look at the AI high performers.

There is a set of things that they have learned to do particularly well, and they anchor on core elements of aligning an AI strategy to business goals, making sure that is linked to an overall business strategy.

There is a set of core competencies around AI talent, doubling down on investing in current employees and upskilling and training them by creating analytics and AI academies.

There is another set of core competencies around internal collaboration and having processes in place to bring business leaders, technologists, and data scientists together to identify where there are collaboration opportunities, making it far more likely that the AI technologies actually generate business value.

There is a set of core competencies around establishing a clear data strategy, around data governance, and around clear repeatable protocols and methodologies to be able to get from a concept to proving the concept and then scaling it across the organization.

So we're seeing companies investing tremendously in the whole set of specific core competencies around adoption and scaling.

Michael Chui: Those companies that are high performing on the core practices that Bryce mentioned, capturing the most top-line and bottom-line value from AI, are often two to four times more likely than those that aren't to do things such as aligning their AI strategy with their overall corporate strategy.

We know these practices are important because of how much more likely it is that those that have great results are actually doing them. It's just far and away a different level. That does suggest that, for all these other companies that have dipped their toe in the water or are capturing value from one use case or another, there are a set of things that they will need to do in order to be able to scale their impact from that individual use case to really having an organization-wide competitive advantage from AI.

David DeLallo: When we're talking about results at this stage of the game, what exactly are the returns that companies are seeing right now?

Michael Chui: We asked individual respondents how much value they are capturing from individual use cases, either in terms of increased top-line

revenue or decreased costs. And our overall result is that, for companies that have adopted an individual use case, they generally do see reportable improvements, either on the top line or the bottom line, depending on the lever or what the use case is meant to capture.

That said, much of the results are in either the 1 to 5 percent range or the 6 to 10 percent range on the top line, and sometimes slightly greater on the bottom line. But those are material differences when you scale them up to the size of a large company. So we are seeing real material impact from adopting AI in these individual use cases.

Bryce Hall: We can give a couple examples of companies we've worked with that have done this particularly well, where we see higher levels of impact. It has involved this transformation from a use-case mindset to one that looks across the entire business, starting with a domain-mapping view.

For example, we worked with a major airline recently on a project that started with looking across the main domains of its business: fleet optimization, pricing, cargo, and so on. And it started with the value at stake, mapping where there was potential value from AI and from analytics and technologies broadly.

And then the result of that becomes a portfolio of use cases and initiatives that can then target that value. And one by one, domain by domain, launching that portfolio of use cases ends up capturing more value because it's targeted—the team of people and the data are consistent across that domain of the business, and it's a quicker line of sight to actually capturing the value than an approach that targets individual use cases across the entire company.

David DeLallo: Another thing I want to dig into here are the specific technologies that enable AI. Did we see any differences between our research this past year and previous research, in terms of the types of technologies that companies are deploying? Are they becoming more sophisticated in the technologies that they're using? Is everybody using something similar? Or are there many differences?

Michael Chui: Again, much like we noticed in terms of which use cases are deployed, the actual specific technology capabilities that different companies are employing within different industries tend to be suited toward where the value is in those industries.

So take, for example, manufacturing industries—whether automotive, consumer packaged goods, or pharma—that's where we see more deployment of physical robotics. Conversely, in industries where lots of customer service is required, whether it's high-tech, telecom or financial services, we see more natural-language understanding, both text as well as speech, being deployed.

So the trends that we see are increased adoption overall and then a matching of investment or deployment with the types of activities that are most prominent within these industries.

Bryce Hall: And to build on that, not only do we see an increase in adoption of these AI capabilities overall across sectors, but we see increasingly sophisticated AI capabilities being adopted.

So natural-language text understanding, the ability to understand what someone has typed or written, is inherently less complicated than understanding the spoken word and speech understanding. It is less complicated than allowing computers to generate natural language for customer service and other applications. And so we increasingly see the broadest applications of natural-language text understanding.

But we are starting to see increased adoption more broadly of speech understanding, of actual language generation. And as this becomes more sophisticated, it will open up other opportunities, and we'll continue to see those adopted and scaled.

Michael Chui: There are a number of frontier technologies or frontier techniques in AI that are very interesting in the lab, but we don't see a lot of deployments for business value. For example, reinforcement learning is terrifically good at being able to train machines to play all kinds of games

better, but not as obvious is how much value that can create in the world. We do see some companies that are trying to use some of those techniques more in order to do optimization problems or robot learning problems.

So it's early days there. It's a similar situation for GANs—generative adversarial networks—which are very interesting, with abilities to create pictures of scenes, whether it's people or architecture, that never existed before. Some additional thought is being done there, so that's an interesting capability.

Can that actually be used in business? Can that be used for creating images for advertising that don't require you to license a stock image because it was just generated by a GAN?

So these are some of the things that companies are trying to figure out, how to get from the lab into business.

David DeLallo: One use I've seen of GANs, and I think, again, it's early days but it's beginning to be looked at more, is using it to create synthetic data. Is that something that we've seen happening yet?

Michael Chui: Yes. A huge problem, as folks in deep learning know, is having a sufficient and representative set of training data. Because GANs can be used to generate new data, you could create training sets using GANs.

Now, of course, you always have to be careful there because there can be bias within data sets from the "real world." GANs also can inherently create patterns within the data, which you might or might not want, so there is a lot of work to be done there. But, because GANs create new data, you could create synthetic data that can be used for training.

David DeLallo: That was another point I wanted to touch on here. As there's more deployment of AI happening and more sophisticated uses, we are seeing some unintended consequences coming out of those. Are we seeing companies becoming more aware of those risks and addressing those risks?

Bryce Hall: Well, this was interesting. In this research, we asked respondents to identify which risks are relevant and then which risks they have mitigated. And across a certain set of risks—such as cybersecurity, explainability, regulatory compliance, and others—among all respondents we still see a significant delta between respondents who say their company has identified a relevant risk but then have successfully been able to mitigate it.

There is a much lower difference in that number if we look at just the AI high performers. Typically, the AI high performers have identified cybersecurity and explainability risks and others and then been able to successfully mitigate that risk.

Michael Chui: For those of us in the field, there's a lot of discussion about a number of these risks, whether it's explainability, unintended bias, privacy, all those sorts of things. It would be easy to assume that, in fact, all the organizations that are deploying AI have this in hand, they all understand what the relevant risks are, and they're all working to mitigate them.

However, what was surprising to me when we looked at our responses was that, across all of these different risks that we identified—from cybersecurity all the way through equity and fairness and physical safety—for most of these risks, less than 50 percent of our respondents said it was even relevant. And, as Bryce said, even a smaller percentage of them had done anything to mitigate against those risks.

Now certainly there are some organizations for which physical safety might not be an AI-relevant risk. But if I look across all of these different risks that many people talk about, quite frankly the identification of them is relevant and the mitigation of these risks is especially important. Again, there is a lot more work to be done.

David DeLallo: Let's turn to an issue that is on everyone's minds in terms of how the use of AI could play out in the workplace. There's obviously great fear of job loss and other changes that AI

could bring about. Did we do any research around that? Or what are we seeing as the effects of AI in the workplace thus far? And any predictions going forward?

Michael Chui: This is a place that we found to be quite interesting in the survey research. We've published quite a bit in terms of what the potential is for jobs lost and jobs gained. But this survey research gives us a view as to what individual companies are actually doing.

We looked both in terms of what has happened in the past year as well as their expectations for the next three years, understanding, of course, that it's just an expectation. What's interesting is, in the past year, a plurality across different sectors said that, in fact, there are, so far, not huge amounts of changes.

However, we do see more respondents than in the past year saying that there potentially could be a decrease in the size of their workforce for an individual company as a result of the deployment of AI. Now that doesn't mean over the entire workforce that there'd be a decrease, because there might be other companies or other industries or other jobs where we'll see increases. But it is interesting that we see, broadly speaking, little change in the past year and people describing more of a likelihood of decreases going forward.

Bryce Hall: It's also interesting to see how that story varies across sectors and then across functions. So certain sectors—automotive and assembly, telecom, even consumer packaged goods, travel, transport, and logistics—those are the ones, this research shows, that can anticipate the largest workforce reductions over the next three years.

But then across functions, the story varies depending on the function. So there is an anticipated decrease in workforce levels in functions such as HR, manufacturing, supply-chain management, and service operations. But then there is an anticipated increase in the workforce in other functions, including marketing and sales and product development.

So this is not a one-size-fits-all technology. Based on this research, we anticipate it will impact different sectors and different functions in different ways.

David DeLallo: So given this variability in how AI can potentially impact the workforce and workplaces, what are things that the workforce can do broadly to prepare for these expected shifts we're seeing?

Michael Chui: One thing that you could do in terms of trying to prepare a workforce is retraining. What we do know is that what people do in their jobs will change as a result of AI—and not necessarily that everything will be automated but that the individual activities that someone does will change. So that means they'll need to do other activities. And that's been the history, and that's what we expect to see going forward.

We did ask our respondents about the degree to which they had actually deployed retraining programs as a result of their deployment of AI. Here what we found interesting is another big distinction between high performers, those who are capturing the most top-line and bottom-line value from AI, and the rest. We found that high performers have a much higher percentage of their workforce that has been retrained and they anticipate retraining in the next three years.

I think that's quite indicative of these companies that are really capturing the most value from AI, that they are also recognizing that they need to retrain their workforces in order to continue to capture that value and stay ahead of the curve.

Bryce Hall: I fully agree. Leading companies are doubling down on talent in various ways. We see, particularly over the past two to three years, an increased focus on developing analytics academies. We see retraining of current employees across a comprehensive variety of strategies, including in-person courses, online training, training for data science, translator training, developing role-based curriculum for roles across

the company, experiential on-the-job training and mentoring, and then comprehensive war-for-talent strategies to identify what future roles will be needed—and if they aren't available in the current company, ways to attract them and retrain them.

David DeLallo: Based on what we've seen in our research over the past few years and our work with companies around the world, when we look forward, we're entering a new decade here. It's 2020. What do you expect to see around AI's use in businesses going forward? What are some trends that we think we will see develop?

Michael Chui: I think as we look at this coming decade that we're going to see progress against all of these different dimensions that we asked our respondents about. We certainly would expect adoption to continue to increase, from the less than a third of companies that have adopted AI technologies across different businesses and functions today.

We'll also see increasing value capture. Companies that have been successful in individual use cases will add more use cases and capture more value. And also they will start to adopt these core practices that are prerequisites for capturing value at scale.

At the same time, they'll need to understand and mitigate more of the risks. Arguably, companies are far behind in terms of being able to understand the risks that are associated with AI.

And then as these implications for the workforce move forward as we anticipate, and people need to do different things, we'll see more and more retraining of workforces.

Bryce Hall: I fully agree. On the first point around broad adoption and scaling, if we think of the foundational trends that have led us to where we are today—increased processing power, reduction in cost of data storage, development of tools and platforms that enable the democratization of AI—we would expect all of those to continue. We see a high percentage of respondents say that they're actually, even today, able to capture value from increased revenues and cost reductions and increased anticipation of investment over the next three years. All of that leads us to say that the broad adoption and scaling trend will continue.

And then secondly, we've seen, as we've talked about, a proliferation of creative, innovative AI use cases. And each of those, in turn, opens up additional opportunities, adjacent use cases, including for other companies across different sectors, and more sophisticated AI capabilities that we anticipate will increasingly be adopted.

And then third, around the adoption and scaling point, we simply know more now about the core practices that need to be in place to actually land the change, to actually embed it across businesses, to manage change management. And so more and more companies will be investing in AI but then actually be able to capture value from that.

David DeLallo: This has been a great conversation. Michael and Bryce, thank you so much for joining us today, and we definitely look forward to any future research you're doing in this space.

Bryce Hall: Our pleasure. Thank you, David.

Michael Chui: Thanks, David. We're looking forward to it, too.

David DeLallo: And thanks as always to you, our listeners, for tuning into this episode of a McKinsey podcast. Please do visit us at McKinsey.com to download other podcasts and interesting articles on AI and other business topics. Have a great day.

Michael Chui is a partner of the McKinsey Global Institute (MGI), based in McKinsey's San Francisco office. **Bryce Hall** is an associate partner in the Washington, DC, office. **David DeLallo** is an executive editor in McKinsey Publishing, based in McKinsey's Stamford office.