MIT SMR CONNECTIONS

CUSTOM RESEARCH REPORT

How AI Changes the Rules:

New Imperatives for the Intelligent Organization

ON BEHALF OF: SAS

CONTENTS

MIT SMR Connections develops content in collaboration with our sponsors.
It operates independently of the *MIT Sloan Management Review* editorial group.

Executive Summary

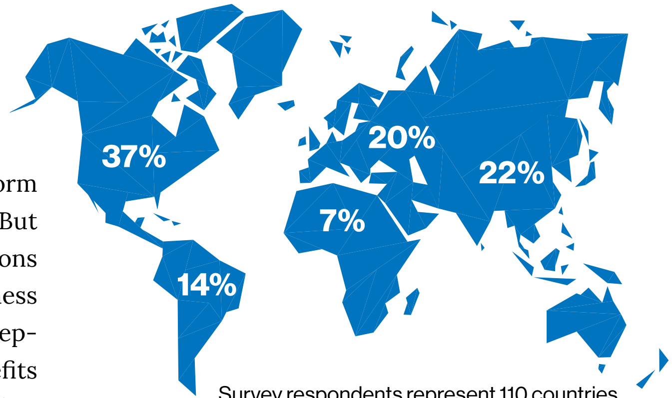

Survey respondents represent 110 countries. Findings from each of the five global regions were consistent with the overall results.

Many leaders are excited about AI's potential to profoundly transform organizations by making them more innovative and productive. But implementing AI will also lead to significant changes in how organizations are managed, according to our recent survey of more than 2,200 business leaders, managers, and key contributors. Those survey respondents, representing organizations across the globe, expect that reaping the benefits of AI will require changes in workplace structures, technology strategies, and technology governance.

The energy for exploring AI is widespread: Nearly two-thirds of survey respondents reported increased spending on AI technologies in the past year. However, for most, it's still too early to realize benefits at scale. Less than half of respondents reported active adoption, with just 1 in 20 indicating that they have implemented AI broadly, while 18% have implemented AI in a few processes and 19% are running pilot projects.

The responses from those who have implemented AI indicate that these initiatives have implications for general management and technology leaders in three significant ways:

1. AI will drive organizational change and ask more of top leaders. The majority of respondents to our survey expect that implementing AI will require more significant organizational change than other emerging technologies. AI demands more collaboration among people skilled in data management, analytics, IT infrastructure, and systems development, as well as business and operational experts. This means that organizational leaders need to ensure that traditional silos don't hinder AI efforts, and they must support the training required to build skills across their workforces.

2. AI will place new demands on the CIO and CTO. AI implementation will influence the choices CIOs and CTOs make in setting their broad technology agendas. They will need to prioritize developing foundational technology capabilities, from infrastructure and cybersecurity to data management and development processes — areas in which those with more advanced AI implementations are already taking the lead compared with other respondents. CIOs will also need to manage the significant changes to software development and deployment processes that most respondents expect from AI.

Many CIOs will also be charged with overseeing or supporting formal data governance efforts: CIOs and CTOs are more likely than other executives to be tasked with this, according to our survey. As leading practitioners note, AI requires both quality data and ongoing support to improve the efficacy of its results and to achieve strong ROI.

3. AI will require an increased focus on risk management and ethics. Our survey shows a broad awareness of the risks inherent in using AI, but few practitioners have taken action to create policies and processes to manage risks, including ethical, legal, reputational, and financial risks. Managing ethical risk is a particular area of opportunity. Those with more advanced AI practices are establishing processes and policies for technology governance and risk management, including providing ways to explain how their algorithms deliver results. They point out that understanding how AI systems reach their conclusions is both an emerging best practice and a necessity, in order to ensure that the human intelligence that feeds and nurtures AI systems keeps pace with the machines' advancements.

The report that follows explores these findings in depth. Read on to learn more about the changes that leaders must prepare for to successfully implement trusted AI.

Building the Intelligent Organization

Artificial intelligence (AI) is recognized as a technology of vital strategic importance to enterprise leaders. The majority of respondents to our global survey have begun the AI journey, reporting that they are in the planning stages or have already piloted or implemented the technology. And organizations are stepping up their financial investments, with 62% reporting increased spending on AI in the past year (see Figure 1, "Spending on Emerging Technologies").

Our survey of more than 2,200 leaders, managers, and contributors across a wide range of industries and geographies reveals how the drive to implement AI is reshaping organizational culture and processes and creating new mandates for CIOs and other technology leaders. It also provides insight into organizations' perceptions of AI risk and ethical issues, and how those attitudes are affecting technology governance.

While AI promises to transform business and create value, for most organizations we surveyed, it is still too early to realize those benefits at scale. Just 5% are implementing AI widely across the organization, while 18% have implemented it in a few processes, and 19% are running pilot projects (see Figure 2, "Most Are in the Early Stages of the AI Journey"). Although another 13% are planning AI adoption, the largest group, 27%, are still investigating it.

The CEO's Role in Driving AI

There's an opportunity to learn from the organizations now implementing AI, and this report will look at how they differ from the rest. One distinction we noted: A commitment from the top appears to be key in driving AI forward. We found that those organizations broadly implementing AI are significantly more likely than others to have top management involved in identifying use cases. They are also more likely to identify the CEO and board as the primary champions for the introduction of new strategic technologies.

Leadership requires both taking action and setting an example. At the Toronto headquarters of Sun Life, an insurance and asset management company, President and CEO Dean Connor has set the expectation that leaders in the business should be able to discuss architecture and technology, and he devoted half a day of an annual executive summit to a data analytics boot camp.

Figure 1: Spending on Emerging Technologies

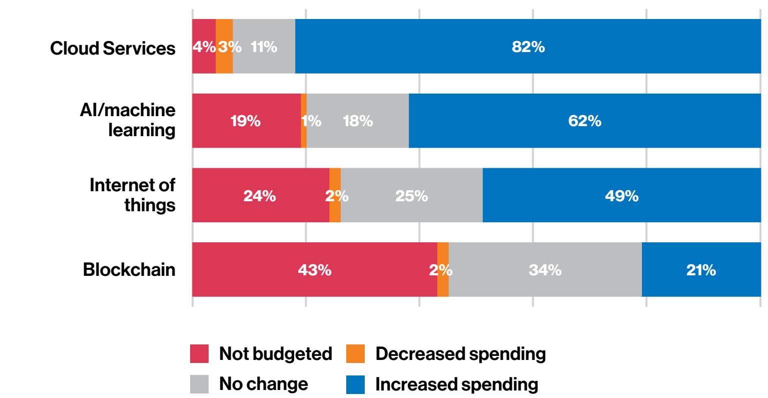

Spending on cloud services has increased for the largest number of survey respondents, followed by spending on AI/machine learning.

Figure 2: Most Are in the Early Stages of the AI Journey

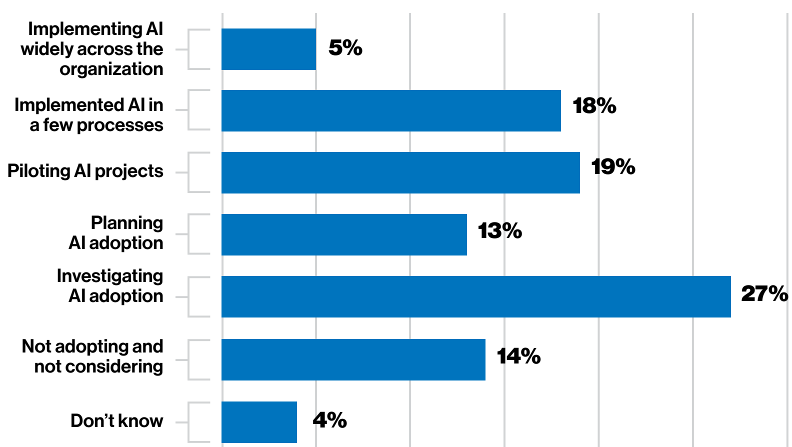

While 42% of respondents are actively engaging with AI in their organizations, only 5% of all respondents reported implementing it widely to date. Meanwhile, 40% are investigating or planning adoption.

> **"AI requires a long-term philosophy.... Companies that are going to succeed in AI are the ones that have a long-term view, that understand that this is a significant competitive edge and this is a long-term investment. Everything else is just lipstick on a pig."**
>
> RAY WANG, CONSTELLATION RESEARCH

"It sent the message, 'If the CEO is discussing topics like our API strategy, you should be thinking about them too,'" says Eric Monteiro, senior vice president and chief client experience officer at Sun Life.

Peter Guerra, North America chief data scientist at Accenture, cautions leaders not to see AI simply as the latest hot technology offering a quick payback. Instead, they should be thinking about the direction of their businesses and how AI can help solve problems. "It should be something that drives everything you do, meaning it underpins everything that you do," Guerra says. "You don't focus on 'I'm going to go do AI.' You focus on 'I'm going to do supply chain better, and I'm going to leverage AI to do that.'"

AI also requires a long-term mindset that can be difficult to square with the pressures CEOs face, says Ray Wang, principal analyst, founder, and chairman at Constellation Research. "AI requires a long-term philosophy. You have to understand that you're going to make this long-term investment with an exponential payoff toward the end. Leaders often don't understand how to do that, so they keep making short-term decisions for earnings per share instead of thinking about the long-term health of the company," he says. "Companies that are going to succeed in AI are the ones that have a long-term view, that understand that this is a significant competitive edge and this is a long-term investment. Everything else is just lipstick on a pig."

Supporting Grassroots AI Initiatives

Succeeding with AI is not only about decisions at the top. Guerra says that the best-case scenario for adopting AI and delivering a positive long-term impact involves balancing leadership support and grassroots enthusiasm.

"The best examples I've seen have been grassroots, bottom-up efforts combined with a supportive, top-down drive," Guerra says. "Where I haven't seen AI be successful is when it's been only top-down driven ... or when it's only grassroots, bottom-up. That leads to complete chaos, with nobody setting priorities."

The kind of effort that combines top-down and grassroots approaches is working at American Fidelity Assurance Company, says Shane Jason Mock, vice president for research and development at the financial services and insurance firm. One of its successful projects used machine learning and robotic process automation to interpret and route customer emails for faster, more effective service.

The company's R&D and innovation teams meet regularly with its executive team to discuss business needs that AI and machine learning can address. It's a two-way channel. Sometimes the executive team will

AI Creates New Opportunities for the Midmarket

We compared survey responses from organizations whose annual revenues range from $50 million to $499 million with those from larger enterprises and found that midmarket companies are less likely to be piloting or implementing AI, less likely to have increased their spending on AI and cloud, and less likely to rate their foundational technology capabilities as mature. What are the implications of these findings for midmarket organizations? One possibility is that midmarket players that have the commitment to innovate with AI have the opportunity to gain a competitive advantage over their peers. By moving sooner, they can stay ahead of the AI maturity curve.

What's Working: Top Leadership Sets the Tone at Sun Life

At Sun Life, President and CEO Dean Connor sets the expectation that executives in the business must understand information architecture and technology in order to engage as collaborators in the development of AI applications. He has even devoted a segment of an executive summit to a data analytics boot camp.

discuss problems, and Mock's group will suggest solutions; at other times, the R&D team will bring an idea for a new AI or machine learning use case to the attention of executives, Mock explains.

Creating an environment where the power of AI can be tapped throughout the organization is a cultural shift the U.S. Air Force is pursuing as it aims to democratize access to AI technologies, according to Capt. Michael J. Kanaan, cochair for AI at U.S. Air Force headquarters. Applying AI across a global enterprise employing 450,000 people means looking for ways to use it across typical business functions — warehousing, supply chain, and human resources — as well as military-specific missions like aircraft operations, intelligence, surveillance, reconnaissance, and cyber operations.

"Empowering everyone to execute their organization's missions with AI technologies," as Kanaan describes the goal, demands a leadership commitment to creating new, cross-functional structures. The Air Force has done this by bringing together people from 26 support and mission operations. "We have in one room people who represent all of these things so that they have a say in the matter and can represent what they know best," Kanaan says.

AI Demands Leadership Focus on Data

Leveraging AI for business advantage requires top leadership's commitment to managing data as a key asset, according to industry experts and AI practitioners we interviewed.

"Being able to treat data as an enterprise asset means that there is somebody whose primary job is to understand where the data is and how it can be used to further the enterprise mission or goals," says Melvin Greer, chief data scientist for the Americas at Intel.

The experience of working with AI and machine learning had led to a profound change in the way Sun Life works with data, says Monteiro.

"It's raised awareness of the importance of data in all our applications," he says. When a business unit team works to develop a new application, its members now think about the fact that other teams at Sun Life may be using the same data in the future. They know that the data has to be validated and its definitions must be clear. "They're worried now about the quality of data creation, even in the process of designing applications or changing applications, which just wasn't true in the past," Monteiro adds.

Treating data as an asset, and devoting resources to managing it, can represent a challenging shift in thinking for corporate leaders, says Astrid Undheim, vice president of analytics and AI at Telenor Group, an Oslo-based telecommunications company serving the Nordic countries and Asia. "If we think about the big barrier to really succeeding with AI, the big job that needs to be done that executives may not truly understand is the need for good quality and high volumes of data. AI and prediction models often need completely different sets of data than what we have had before," Undheim says.

"And that's a huge task," she adds. "It's not about the data that we have in our systems and how can we use AI to create value from that. The thinking is really, how can we change our processes? How can we change the way we work, and how can we collect data that we

"It was important to lead with asking, 'What is the objective, strategically, of our company? What questions do we want solved?' If you say, 'Let's just start gathering data,' you might not be going down the best path."

ROB STILLWELL, SEACOAST BANK

need to build models for solving our most important business problems? It's not like you just hire experts, AI people, and they will do magic. Most often, you have to go all the way back to what data are you collecting before these AI experts can actually do their job."

Seacoast Bank recognized the importance of working on its data early in the machine learning journey, but first it identified its business objective: to gain intelligence that would help it to better serve customers.

"It was important to lead with asking, 'What is the objective, strategically, of our company? What questions do we want solved?' If you say, 'Let's just start gathering data,' you might not be going down the best path. But if you start with, 'What's our objective?' then you can start getting the appropriate data to do it. In our case, only then did we develop the machine learning and everything on top of it," says Rob Stillwell, senior vice president and business analytics officer at Seacoast.

With its objective set, the 93-year-old financial institution serving southern Florida worked to reinvent the way it collected and analyzed data about its customers so it could serve them better. For example, instead of focusing separately on each of the various transactions the bank conducted, Seacoast created what Stillwell calls a holistic data set on the bank's customers. The bank now uses machine learning models on this data to analyze the lifetime value of a customer relationship and target services at increasing that value.

"We worked on data for about two years before utilizing machine learning, because we knew how important it was to get the data part right," Stillwell says, adding that the second element was determining how any data analytics or machine learning effort would match the bank's strategy.

How AI Drives Organizational Change

This expectation that AI requires a commitment to organizational change is supported by our survey respondents, who see AI as driving more change than other emerging technologies. Sixty-three percent of respondents overall said they expect AI to drive dramatic or significant organizational change (see Figure 3, "AI Is Expected to Drive Organizational Change"). Among respondents

Figure 3: AI Is Expected to Drive Organizational Change

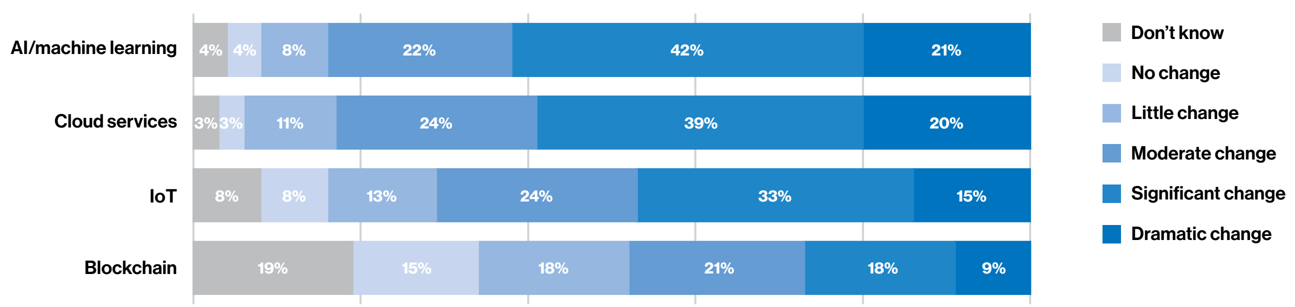

Sixty-three percent of respondents expect AI/machine learning to drive dramatic or significant change, slightly ahead of cloud and well ahead of IoT or blockchain. *Note: Some percentages do not total 100 due to rounding.*

who are widely implementing AI, 80% expect it to drive significant or dramatic change. Cloud, however, followed closely, with 59% seeing it as a significant or dramatic driver of change; it also emerged as a key enabling technology for AI implementation.

Large majorities of respondents see AI increasing activities that cross organizational boundaries: They expect it will increase the need for cross-functional collaboration, require more training across disciplines, and increase the number of multidisciplinary teams (see Figure 4, "AI Demands a More Connected Organization").

This drive for cross-functional teamwork is reflected in how survey respondents report that data science expertise is situated within their organizations. Data science teams are placed to enhance collaboration, working across both IT and the business units at 22% of organizations, or being embedded in the business units at 21%. Data science groups are less often centralized or placed within IT.

Meanwhile, business or operational experts are identifying AI use cases at 47% of respondent organizations overall (see Figure 5, "Who's Making the (Use) Case for AI?"); the likelihood that business experts are collaborating with data scientists in this way increases at organizations implementing AI.

While practitioners interviewed all support the need for collaboration between data scientists and business units, there are different philosophies and approaches to how they are organizing for AI and enabling collaboration.

Figure 4: AI Demands a More Connected Organization

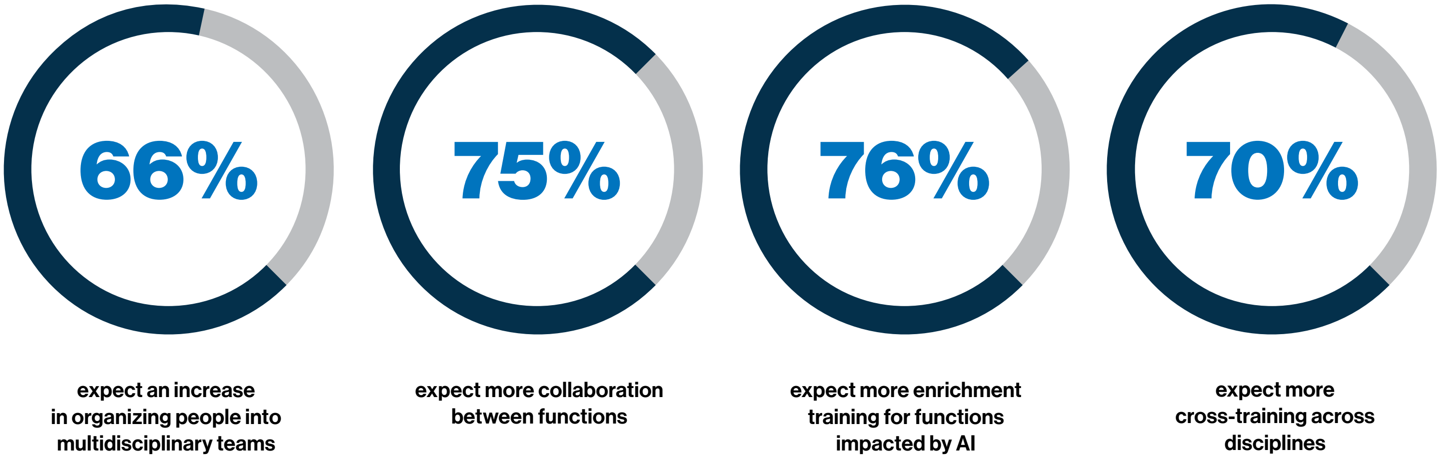

The majority of survey respondents expect that AI will require them to increase activities that bridge functions and disciplines across the organization.

What's Working: Seacoast Bank Put Its Data House in Order Early

Seacoast Bank knew that to accomplish its machine learning goal of improving service to customers, it would first need to collect the right data and get it into good shape. Before it began to build models to help it better understand and increase the lifetime value of the customer relationship, it worked on its data strategy for about two years.

Figure 5: Who's Making the (Use) Case for AI?

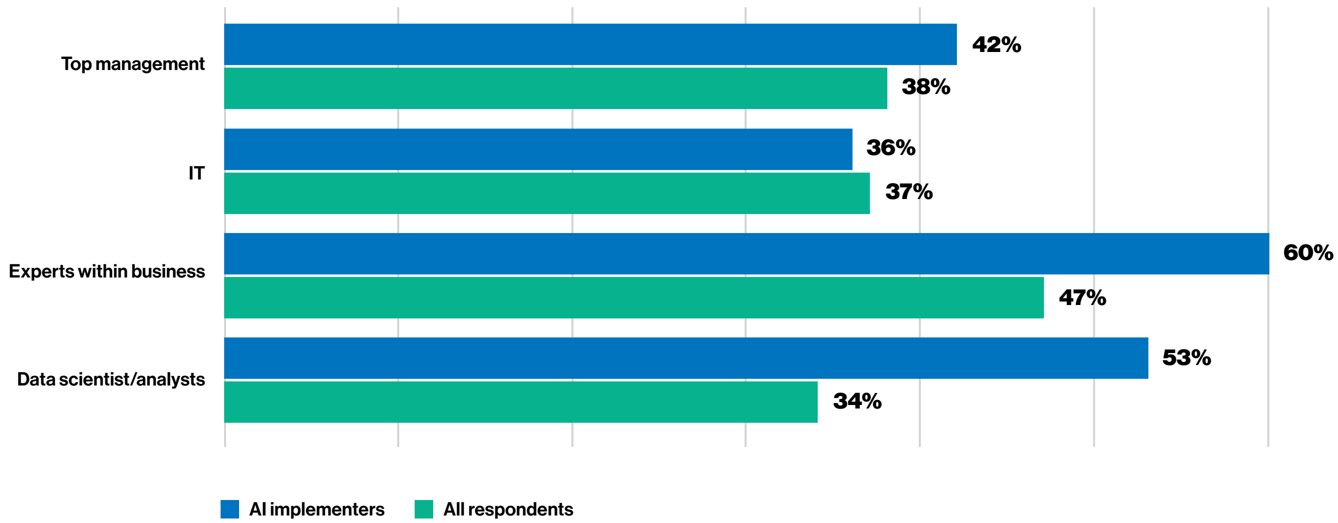

Business domain or operational experts are most likely to be involved in developing use cases for AI and are much more likely to be involved at organizations already implementing AI.

Sun Life, for example, has organized data science and AI expertise in what Monteiro calls "a federated model," in which these experts are managed centrally but embedded within various business units. (Sun Life also has a small number of data scientists available for consulting throughout the organization.) This helps the company ensure that its business strategy drives the development of use cases, such as matching clients to advisers who have experience dealing with a client's particular stage of life, predicting the next best action for a client's health or financial future, evaluating insurance policy risks, and predicting and preventing fraud.

"They are part of the management team's discussions every day," Monteiro says of the data and AI experts. Because they work in the business unit, they understand the relative importance and challenges of various use cases, and they know what the quality of the data is, he adds. "They know what the business outcome is. They understand the process that decision model is going to have to fit into way better than if they were off working in a centralized team."

Others see an advantage in centralizing their AI and data science capability. Telenor has been working with deep learning and neural networks for four years, following on earlier experience with machine learning. Its applications have tackled improving customer service (by automatically helping customers with queries), human resources management (by sensing which employees are at risk of leaving), network management optimization, and an internet of things project that uses sensors to predict air quality in cities.

Telenor's Undheim says that some of the business units of the multinational telco have centralized data science and AI expertise.

"There are many reasons for centralizing it. One reason is to create strong communities of experts. Another very good reason is to align on tools and platforms and really make sure that the company uses the same tools, because it's quite demanding to set up the required platforms and tools," Undheim says.

While organizations will make different choices about where to place data science expertise, many of the practitioners interviewed recognize, like Telenor, that there is a need for an enterprisewide approach to the data and technology infrastructure to support AI. Next, we'll explore these and other implications that AI has for the CIO's role and agenda.

Building Workforce Skills for AI

Finding workers with the right skills is the biggest challenge in adopting emerging technologies, according to leaders in our survey. And when it comes to implementing AI, this concern is particularly urgent. Most respondents assess their organizations' data science capabilities as developing, with only 17% considering them to be mature. But only a minority of organizations have training in place to address any deficits.

The survey illuminated a gap between expectations and action for many organizations. AI implementations will change the way people work: 75% of respondents expect that their organizations will need to provide more training for functions impacted by AI. At the same time, only 8% of respondents overall reported a broad range of training available in their organizations (see Figure 6, "Training for AI Skills Development Tracks Implementation"). Again, AI implementers are more likely to offer workforce training on AI.

An important part of the U.S. Air Force's AI initiative is developing relevant skills in the existing workforce, says Capt. Michael J. Kanaan, cochair for AI at U.S. Air Force headquarters. Identifying skills already present is the first element of this strategy.

Kanaan says his organization is launching a computer language initiative modeled on a Defense Department practice of recognizing recruits who speak a foreign language. "It's designed to treat computer programming languages as the functional, objective equivalent of human foreign language aptitude and proficiency," Kanaan says. "We are providing them incentives, and it drives us internally to create jobs of the future, to leverage those same skill sets and talent. We want to provide the entirety of our workforce multiple modalities for global access to coding and programming education."

Kanaan says this approach differs from organizations' typical efforts, which focus on recruiting only elite data scientists for their AI initiatives. "I deal with these conversations every day — 'Where do we find this top 1% of machine learning data scientists?' or, 'How do I compete with the likes of MIT or Stanford or any number of math-heavy institutions?'"

Instead, Kanaan wants to find talent in the broader workforce, by looking for those with interests in programming, video-game building, and related activities, and providing them with educational opportunities. "It's about tapping into your indigenous talent," he says. "I bet most managers don't walk down the hall and say, 'Hey, do you know how to program?' while they're going out and trying to hire programmers. It's the citizen coder that will make a big difference for you, the person you don't know down the hall, that's really important to us."

To pursue this goal, leaders released a self-assessment for all Air Force members and found that more than 2,900 employees have relevant skills, including developing user interfaces, data curation, and even machine learning. The message in conducting that assessment, Kanaan says, was, "If you know how to program or code, we want to value you, so tell us, and then we can start talking about opportunities."

The next step, he says, is to provide identified employees with 24-7 access to training materials so they can learn new skills when their schedules permit.

Figure 6: Training for AI Skills Development Tracks Implementation

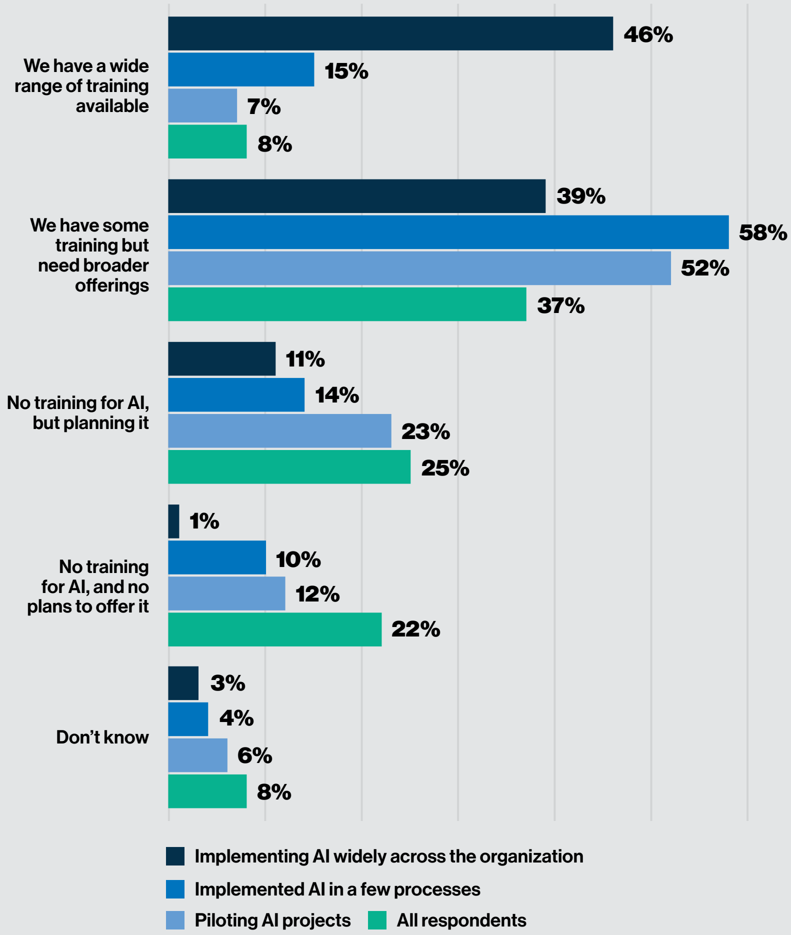

Even among the most advanced AI implementers, fewer than half reported that broad training for AI skills is available within their organizations.

AI's Implications for IT Leadership

AI implementation will have a significant impact on C-suite technology leaders, particularly CIOs, chief data officers, and chief analytics officers. Developing the capabilities required for production AI will influence IT road maps, software development and deployment processes, and how the organization treats data.

Sharpening Focus on Critical Technology Competencies

Most respondents reported that they are still developing foundational technology capabilities that are key to realizing the benefits of AI applications: cloud/data center infrastructure, cybersecurity, data management, and development processes and workflow. As organizations move up the AI implementation curve, they are more likely to report mature capabilities in these areas (see Figure 7, "AI Implementers Report Most Mature Technology Capabilities").

These are clear indications that CIOs and C-suite technology leaders must have a plan for increasing these capabilities.

"You have to work closely with your CIO," says Capt. Michael J. Kanaan, co-chair for artificial intelligence at U.S. Air Force headquarters. "We're talking about simple things here. If the cloud isn't there, you're not going to do this very well. If the data sets aren't there, you're not going to do this very well. If you're not making good choices on your DevSecOps" — as the Air Force refers to DevOps that keep security concerns in focus — "it's not going to work. AI is at the top of that big, big iceberg."

Cloud services in particular are a critical asset for AI, says Eric Monteiro, senior vice president and chief client experience officer at Sun Life. Paying for on-demand cloud computing resources is more cost-effective than buying and operating the computing infrastructure required by AI. It also offers the global financial services and insurance company more flexibility to serve different business units according to their individual needs and to access the latest technologies.

"The computing models and the solutions are all more flexible with the cloud," Monteiro says. "Having our infrastructure in the cloud allows us to provide for the different needs of each business without having to buy the whole thing. And the cloud enables us to use the latest and greatest tool kits. There's no way we could buy all those tool kits, and manage and certify them, in a way that's fast enough for the needs of our analytics community. Going to the cloud means someone else is doing that."

Figure 7: AI Implementers Report Most Mature Technology Capabilities

Advancing certain technology capabilities — cloud/data center, cybersecurity, data management, and development processes and workflows — may build a foundation for successful AI implementations.

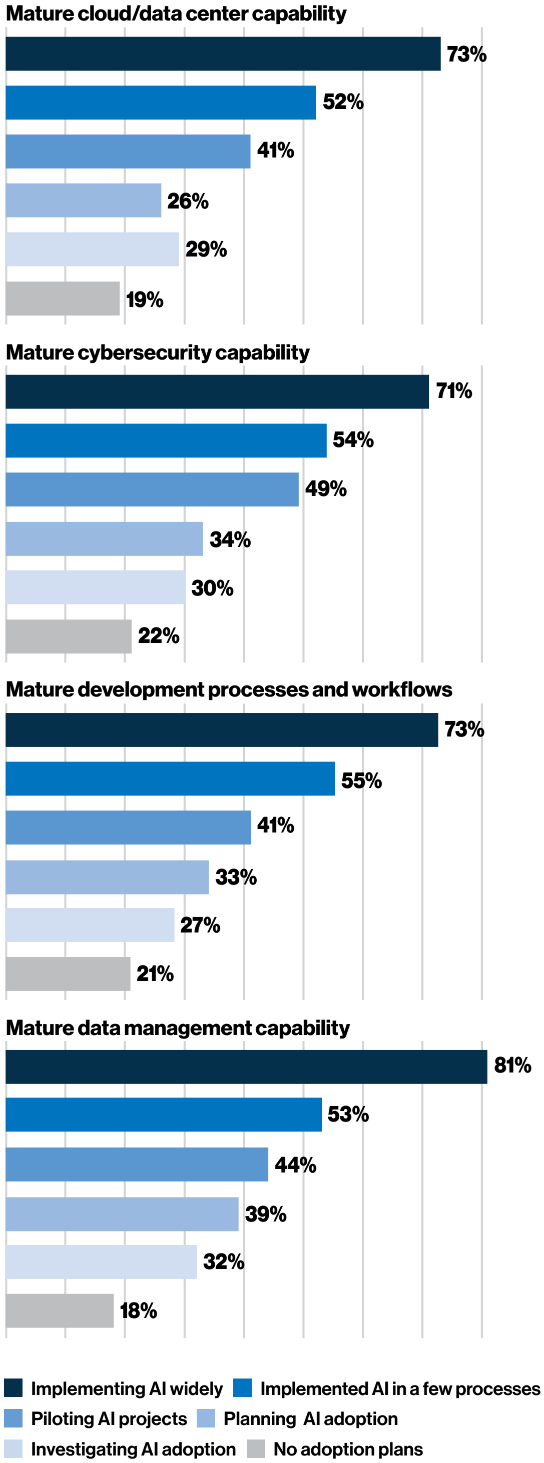

Whether it's computing resources, DevOps, or new technologies for experimenting with different algorithms and data, AI practitioners are pressing IT to respond, says Astrid Undheim, vice president of analytics and AI at Telenor Group, a Norway-based telecommunications company.

"From my perspective, I think that AI and machine learning experts are putting quite high demands on IT, more demand in terms of agility in IT," Undheim says. "The difference between IT that works and IT that doesn't work for machine learning is quite big. And for telcos that have traditionally outsourced a lot of IT, it is demanding to make the IT infrastructure agile enough."

"From my perspective, I think that AI and machine learning experts are putting quite high demands on IT, more demand in terms of agility in IT. The difference between IT that works and IT that doesn't work for machine learning is quite big."

ASTRID UNDHEIM, TELENOR

AI Changes How Software Is Developed and Deployed

Respondents to our survey strongly indicated that they see AI changing software development and deployment processes. Sixty-one percent said they see AI driving significant or dramatic changes to software development processes, with 57% expecting it to similarly influence software deployment processes (see Figure 8, "AI Drives Changes in Software Development and Deployment Processes"). And those who have already implemented AI are more likely to report a strong impact on both software development and deployment.

AI brings significant change compared with traditional technology implementations because the deployment process is dynamic, requiring continuous monitoring and retraining. Managing these systems requires ongoing management of the predictive AI and machine learning models a company develops, not just before but also after they have been deployed. It means being ready to make improvements and corrections to these models, says Ray Wang, principal analyst, founder, and chairman at Constellation Research.

Figure 8: AI Drives Changes in Software Development and Deployment Processes

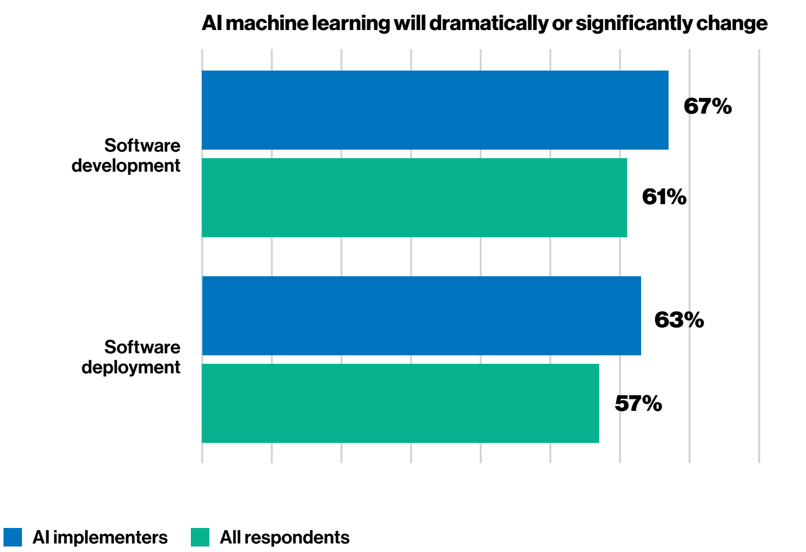

The majority of respondents expect AI to require changes in how they develop and deploy software. Those with experience implementing AI are even more likely to hold this view.

"You're always collecting data, and you're always refining the model; this isn't something that's static. You also have to make sure that once you train the system, you can also 'unlearn' the system," Wang says. "You have to be able to take corrective action. If a pattern that's assumed to be correct is incorrect, how do you retrain the system?" Leaders need an answer to this question, he says.

Machine learning and AI algorithms are designed to improve results as additional informative data goes into them. Tests can show that a predictive model works well, but even then, it still requires ongoing maintenance, says Linda Zeger, founder and principal consultant at the data analytics and system design consultancy Auroral LLC. "When you start putting [models] out there, things change, and over time they

> **"It's a little different than a typical software project in that, even once it's deployed, the human interaction with it can change. And so the models you originally have might become outdated."**
>
> LINDA ZEGER, AURORAL

may not work as well. This is how it is really different than a typical software development process. The training and testing processes should continue as the environment or operating conditions change," Zeger says.

One vivid reason for the need to constantly reevaluate AI and machine learning models is that the initial model is unlikely to incorporate all the variables in a situation in which it is taking actions or account for how users will interact with the AI system over time. Zeger points to self-driving cars as an example. Test drivers can be told that the car will automatically brake if it senses a pending collision but that if for any reason the car fails to behave as expected or is unsafe, the driver is expected to serve as a backup operator.

But even if the cars perform perfectly in a test environment, the results can change when a consumer gets behind the wheel, Zeger says. "Once somebody owns such a car, they might become complacent and accustomed to the car braking for them. They might start losing patience and they think, 'The car is doing pretty well. I don't really have to watch it so much.' So they react differently, and now the success of the self-driving car as a whole may not be as good. At this point, you would need to retrain and reevaluate the car/driver system."

This scenario, Zeger says, illustrates the human factor involved. "It's a little different than a typical software project in that, even once it's deployed, the human interaction with it can change. And so the models you originally have might become outdated," she says.

How Data Is Influencing Software Development Practices

The demands that AI places on an organization — the need to manage data holistically and proactively — has influenced software development practices and increased the need for business-IT collaboration on architectural strategies, experts and practitioners say. There is a greater awareness of how the data that flows throughout the organization applies to all software development, not just the development of machine learning and AI models, as enterprises plan for the future use of the data they collect and generate.

This shift in attitude and approach has also influenced IT staffers, Monteiro says. "The folks who design applications in IT, they don't just think about applications now; they also think about the data. They now see that the applications they're creating or designing create data that's going to be used later in the process. This just wasn't true 10 years ago. So if it's an application process or a website design or whatever, data quality is now a core part of that design principle."

AI development also requires more — and earlier — business and IT collaboration than traditional application development, Monteiro says. Choices about aspects of AI such as computing architecture, how the data will flow in a particular application, how the new AI system will change business processes in various parts of the company, how people will interact with the systems through user interfaces, and more are now part of early-stage talks.

"It's raised the bar on up-front alignment between the business and IT on architecture, and data and API strategies," Monteiro says. "It creates a lot of complexity early on in projects. In fact, one of the things we found is that projects slowed down in the beginning because there was a lot of alignment required. Lots of questions came up that never would have come up before. But then, when it's time to use the data, we've accelerated it because we answered a lot of the earlier questions when we had these discussions," he says.

This means that business experts need to understand more about how technology works, just as technology experts need to be smarter about business when discussing AI projects, Monteiro says. This cultural shift requires IT leaders to be comfortable with business colleagues' questions about details, down to the technology functions, data, and computing architecture that support them. "Historically, business leaders wouldn't be bothered about whether this is in that database or this database. But now they do, and it does matter," he says.

Sharing Responsibility for the Organization's Data

Data governance is an important component of an AI-ready data strategy, but our survey found formal data governance efforts among just 46% of respondents. AI implementers, in contrast, are far more likely to have a data governance program: 74% of those with broad implementations of AI and 62% of those with point implementations of AI reported engaging in these efforts. All respondents, including AI implementers, are still driven by security and regulatory compliance, while data access, data quality, and data ethics appear to be lesser concerns (see Figure 9, "Security and Compliance Still Top the Rankings of Issues That Drive Data Governance").

Nonetheless, the drive to implement new technologies is pushing more activity in quality, governance, and accessibility. A majority of respondents said they are doing more in data security (64%), data quality (61%), data privacy and accessibility (each 58%), and data governance (57%).

Who is responsible for data governance varies widely, according to our survey data. About 40% of respondents said the CIO or CTO is accountable for data governance, but responses from the other 60% revealed a range of other approaches. Even in organizations where there is shared responsibility or it is assigned to another C-suite executive, the structures and systems required for data governance will inevitably be part of the technology leader's agenda. Regardless of who is ultimately accountable, different aspects of governance may be assigned to different groups, centralized, or assigned to business leaders.

At global information services provider Equifax, data governance essentially operates with three lines of defense. At the first level of defense, business units, the data and analytics center of excellence, data stewards, and users

Figure 9: Security and Compliance Still Top the Rankings of Issues That Drive Data Governance

#1 Concern that data is kept secure

#2 Concern that data is managed in compliance with privacy or other regulations

#3 Concern that data quality is trusted by internal users

#4 Concern that data is used ethically

#5 Concern that data can be easily accessed for AI, analytics, or other applications

Despite the growing spotlight on issues such as data quality, data ethics, and data accessibility, all respondents, including those with more-advanced AI practices, remain most concerned about security and compliance when it comes to data governance.

share responsibility for data access and governance. At a second level, internal audit is responsible for ensuring that all data policies have been implemented and executed. The third line of defense is provided by the corporate risk management team along with security experts, to ensure that data use is meeting the data governance guidelines for the entire enterprise, says Vickey Chang, vice president for data and analytics at the U.S. Information Services unit at Equifax.

What's Working: DBS Bank Centralizes Governance

Governance over AI can be a sprawling effort, encompassing technology risk, compliance, privacy, and discussions of responsible, appropriate uses. DBS Bank has centralized all of those activities to streamline the review and approval process. Although that approach may not be practical for every enterprise, it is now easier for project leaders to gain the necessary sign-offs at DBS.

Chang works with a team of about 15 data scientists on predictive models that use neural networks that, for example, help financial institutions evaluate loan applications. Among her most recent projects: making neural network models, known as black boxes, explainable to regulators with oversight of the credit industry. This work involves building many statistical models, and Equifax's governance process is an internal check.

"If my team builds a model using multiple data sources, we will need to get all the proper data approvals to make sure we are using data in the right way to fulfill our customer requests," Chang says. A data stewardship and governance team assessment is conducted with each business unit that holds the data, to ensure that Chang's team is compliant with the data regulations (which can vary by country and industry) that apply to that business unit.

DBS Bank in Singapore established its Data First program two years ago as part of a strategic effort to drive data innovation and elevate data management and governance standards across the bank. It involves having a senior leader from each business unit serve as a data owner who is responsible for making sure that everyone is thinking about data strategically. All aspects of legal and regulatory requirements impacting data are also centrally managed under the center of excellence team for data headed by an executive director in the Legal, Compliance, and Secretariat group.

Centralizing the management of legal risk, regulatory compliance, data privacy, and oversight of appropriate use of data has also streamlined DBS's process of reviewing and approving AI and analytics use cases, says Lam Chee Kin, managing director and head of the Legal, Compliance, and Secretariat group. "You create one organization to handle the harmonization of data frameworks and to be a single point of advice," he says. In contrast, many other organizations set up separate units to oversee singular concerns such as data privacy, banking secrecy, competition law, outsourcing, legal risks, and so on, he says. "Because it's in different places, a person trying to run one data project ends up having to talk to multiple people. It can create a lot of operational complexity."

And as AI brings bigger questions about appropriate use of technology and data to the forefront, developing a formal yet streamlined process for evaluating use cases may become a strategic priority for leadership. ●

How AI Raises the Stakes for Trust and Risk

AI's potential to solve difficult problems excites those who see it ushering in a new era comparable to the Industrial Revolution of the 18th century. At the same time, the implication that humans are granting inanimate technology a kind of agency and autonomy — and even some degree of personality, as in the case of automated personal assistants who have names and speak to us — creates anxiety for many. While some of those worries belong in the realm of science fiction, using AI does carry real potential risks that organizations must manage.

Precisely because the technology can enable autonomous decision-making and actions by machines, questions about its reliability are more urgent than for other technologies. To establish a baseline for trust in AI, we asked our survey group to rate just how reliable they find the results and recommendations of AI-based systems in two different contexts: in their personal lives, as consumers; and at work, interacting with their organizations' AI-based systems. On a scale of 1 to 10, with 10 being most reliable, respondents' overall rating for personal AI technology was 7, while AI used in their organizations earned a 6.

While those ratings lean positive, they demonstrate that trust in AI is at best guarded. This cautious approach was confirmed when we measured people's degree of concern about eight commonly discussed AI risks and found that most were rated fairly highly — 7 on a 10-point scale (with 10 representing the highest concern). The six risks of high concern were that AI may:

- Deliver inadequate return on investment.
- Produce bad information.
- Be used unethically.
- Support biased, potentially illegal decisions.
- Produce results that humans cannot explain.
- Be too unpredictable to manage adequately.

The two risks eliciting less concern were that AI may disrupt workflows or productivity (rated 3) and that AI may deliver bad customer experiences (rated 4).

Figure 10: AI Risk Management Structures Are Emerging

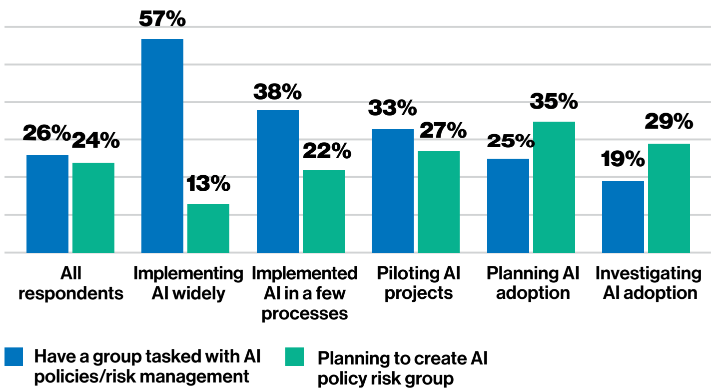

As AI becomes more widely implemented, it appears to be driving the creation of groups to develop policies around its use and manage associated risks.

New Risks Call For New Risk Management Structures

An important component of building trust in AI is managing the associated risks — particularly through oversight that seeks to understand and verify how models function, mitigate bias, and anticipate unintended consequences. How are respondents doing this? About half are acting to create organizational structures to manage AI risk: 26% have a group that sets policies and manages AI risk, and 24% plan to create one. Organizations that have implemented AI are much more likely to already have a group tasked with setting policies and managing AI risk: 57% of those that have implemented AI broadly in their enterprises do, as do 38% of the point implementers (see Figure 10, "AI Risk

Trust, Risk, and Culture

Our global survey included respondents from all regions, and we found almost no significant variance in response patterns by geography. The one exception was to the questions about AI trust and risk. Europeans surveyed expressed markedly less trust in AI systems used in their personal lives. And North Americans are much more concerned that AI may deliver bad customer experiences, rating that as highly as most of the other risks identified above.

> **"We need to be able to defend our models and how we made those decisions in front of a regulator, which happens often, actually. Therefore, we can't afford not to do this."**
>
> ERIC MONTEIRO, SUN LIFE

Management Structures Are Emerging"). Large companies of 5,000 or more employees are also more likely to act on AI risk: 34% have set up a group to do this.

According to our survey, the responsibility for managing AI risk is as likely to fall to the CIO, CTO, or CEO as it is to be shared; very few organizations appear to be assigning accountability along traditional lines of risk management, to legal or financial executives. Those implementing AI are more likely to say the responsibility is shared, which is perhaps indicative of their experience working more cross-functionally with this technology. Very few overall have any plans for remediating harm caused by AI applications (see Figure 11, "Who Has a Plan to Remediate Harm Caused by an AI Application?").

An important development among companies that are implementing AI is establishing a management-review mechanism, says Peter Guerra, North America chief data scientist at Accenture.

"The big thing that I'm seeing is a lot of clients, especially those that are highly regulated, are putting together boards that review anything that they want to operationalize from an AI perspective," Guerra says. These boards often have data scientists as well as business leaders, representatives from the legal department, and other relevant experts.

Eric Monteiro, senior vice president and chief client experience officer at Sun Life, says the global financial services and insurance company has established a three-layered approach to govern its use of AI and mitigate the risk of bias in its applications. First, Sun Life provides training to its model developers that includes guidelines and standards to test for bias. Second, a team of experts evaluates every model to examine its behavior after the testing phase, to determine whether the inclusion of additional data introduces bias. Third, Sun Life has another AI model-evaluation team set up specifically for what Monteiro calls "more critical models," such as those that impact the company's financial results, to provide an additional layer of scrutiny.

Monteiro says Sun Life takes model validation seriously because it's a regulated financial institution. "We need to be able to defend our models and how we made those decisions in front of a regulator, which happens often, actually. Therefore, we can't afford not to do this," he says.

Linda Zeger, founder and principal consultant at the data analytics and system design consultancy Auroral LLC, says that maintaining documentation about AI projects — including a model's scope, limitations, appropriate and inappropriate uses, and accuracy, as well as notes about users' interactions with it — is an important way to mitigate risks. A company can provide documentation not only internally, but also to partners, customers, and end users to help people understand what an algorithm can and can't do and how they should use it.

"I think one of the greatest risks — and this is why communication is so important — is that the algo-

Figure 11: Who Has a Plan to Remediate Harm Caused by an AI Application?

We have a plan for remediating harm that AI applications might cause

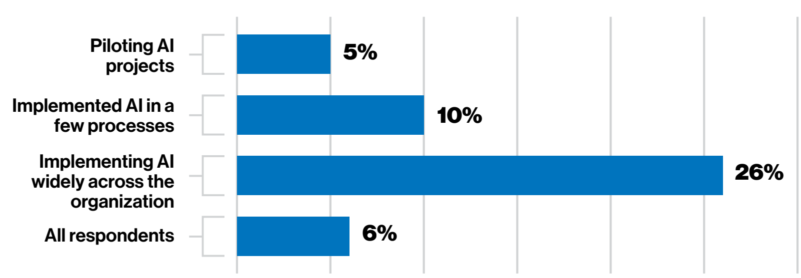

AI applications have the potential to cause harm if they malfunction or recommend poor decisions — but only 1 in 4 of the most advanced practitioners has a plan for remediating such harm.

rithm gets used where it's not applicable and where it wasn't designed for," Zeger says. She cites the example of a self-driving car that is tested in some conditions (such as typical weather conditions) but not others (such as unusually heavy smoke). "That is one reason I emphasize that it's really important that the designers of the system document the limitations and scope and accuracy and that the end users receive that documentation — to say, 'This system has only been designed for typical conditions.' It's almost like a black-box warning that's really obvious to the end user," she says.

Figure 12: Ability to Explain AI Results Gains Traction at the Leading Edge

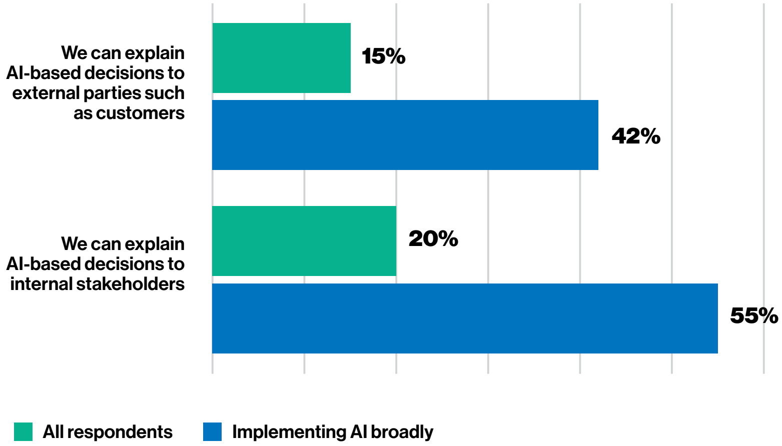

Understanding — and communicating — the logic behind how an AI-based system has produced a result is key to building trust in AI and ensuring that decisions aren't biased.

The Quest for Explainability

An issue that is new to many who consider technology risk management is adequate explainability — the ability to identify the key factors that a model used in producing its results, which may be recommendations in a decision-support system or actions in an automated process. In our survey, the most action on explainability of AI results is being taken by those with the broadest AI implementations. Among this group, 55% make explanations available to internal stakeholders, and 42% make explanations available to external parties such as customers (see Figure 12, "Ability to Explain AI Results Gains Traction at the Leading Edge").

Explainability is important for regulatory compliance in use cases such as hiring or granting credit, where decisions must be shown to be free of illegal bias. It is also enshrined in the European Union's General Data Protection Regulation (GDPR), which gives individuals the right to know how their data has been used to reach a decision. This policy is echoed in emerging data regulations in other jurisdictions.

In addition to GDPR, Melvin Greer, chief data scientist for the Americas at Intel, points to Canada's federal privacy laws and moves in 26 U.S. states to develop laws like GDPR. "I think it's safe to say that we are moving into a regulatory and legal environment that is going to put more emphasis on data privacy and protection," Greer adds.

Heavily regulated industries such as financial services and health care are well ahead in understanding and managing explainability.

Monteiro says that Sun Life invests in its process for examining and validating models, and it experiments with new tools to increase its ability to describe the work of algorithms.

"If somebody came in for an insurance application and didn't get approved, we need to be able to tell the regulators why," he says. "Regulators want to make sure that it's not creating systemic bias or societal problems."

Monteiro says that emerging software tools are getting better at describing correlations that occur within the algorithms. "You can say with this model, 'The key variables that mattered for this decision were X, Y, Z. And the model took these into account and got to these outcomes,'" he says.

Explainability also plays a role in building trust in AI internally, by providing visibility into how a model works and why it is making a particular recommendation to business stakeholders. Because credit agency Equifax serves financial services companies with its AI-driven scoring tools, explainability is particularly important. Thus, the company has invested heavily in innovating effective AI models that are also explainable.

Vickey Chang, vice president for data and analytics at Equifax's U.S. Information Services unit, says the company previously used a neural network model that she regarded as a black box because the underlying logic could not be seen. Equifax subsequently developed an in-house machine learning model that provides the reasons behind a recommendation, such as why a given credit application was approved or denied. She says this newer model, which uses technology that Equifax has patented, outperforms its neural network predecessor.

At Telenor Group, explainability has become a steeper challenge as the Norwegian telco's AI group has advanced from machine learning to deep learning and its larger, more complex models. Astrid Undheim, vice president of analytics and AI at Telenor, says deep learning models are black boxes that she describes as being "too big to be able to explain properly." That means the company takes a cautious approach.

"You need to test the system much more rigorously to be able to trust it," she says. "I think another thing is that these models are not mature enough to be used in critical domains because of this black-box problem. So in my view, we need to start using them for problems with low risk and then, in that same process, learn, and also put focus on advancements in deep learning that involves, for instance, explainability and interpretability of the model. These are unsolved problems at the moment."

AI Puts Ethics Conversations on the Table

Making the workings of AI models more transparent is also aligned with efforts to ensure that this powerful technology is used ethically and that developers consider how a model developed for one purpose might potentially be misused for another. Advanced practitioners such as Capt. Michael J. Kanaan, cochair for AI at U.S. Air Force headquarters, believes that "if you are in the business of developing AI solutions, building them, or providing them to someone, you have a moral obligation to talk about these topics and the dual-use purposes."

While it's possible that ethics concerns are being addressed via broader AI review groups that the majority of respondents are committed to, just 10% of respondents to our survey have set up an ethics-focused AI review board, and 14% are planning to establish one. Those with broad implementations of AI are ahead on this issue: 42% report that they have an AI ethics oversight group, and 40% have adopted an AI code of ethics (see Figure 13, "Focus on AI Ethics"). Hiring for the emerging job of AI ethicist lags, with about 1 in 5 of the most advanced implementers having done so, compared with 1 in 20 of respondents overall.

DBS Bank in Singapore is one enterprise that has established principles and a process for evaluating proposed data use cases to ensure responsible data use. The review process is a natural outgrowth of the bank's experience with developing its data governance practices. DBS also decided to add AI capabilities to its board and increase its focus on using data responsibly, launching a Responsible Data Use Committee that brings together a diverse array of viewpoints to vet use cases.

Jeffery Lee, executive director of legal and compliance at DBS, says the bank's approach to ensuring that data use cases are appropriate builds on foundational data-governance principles: The baseline for any use case

Figure 13: Focus on AI Ethics

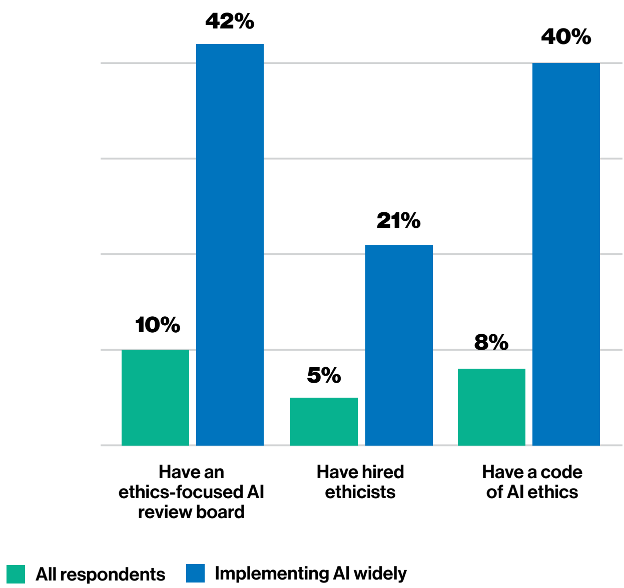

As with AI risk management overall, those with the most advanced AI practices have a much greater likelihood of having taken specific actions to apply ethical considerations to their use of the technology.

is that the relevant data is of good quality and will be used in a way that conforms to regulatory and security requirements. Then, proponents of a particular use case must be able to articulate how the proposed use of the data is purposeful, unsurprising (meaning that customers and employees would not be shocked to learn how data was being used), respectful to customers and employees, and explainable; this is where the Responsible Data Use Committee may be convened to consider the use case. Finally, there is a model governance process that looks at the development, testing, validation, documentation, deployment, communication, and ongoing review of a particular model.

This review process safeguards the bank's work in an emerging field by asking challenging questions, says Lam Chee Kin, DBS managing director and head of the bank's Legal, Compliance, and Secretariat group. "Why do we feel comfortable that this data use case is responsible? Can we hold our heads up high to society so that we're actually using data in a responsible way? This requires very, very broad and diverse societal perspectives on whether you should allow something or not," he says. "We've chosen to staff that committee along the lines of a broad cross section of perspectives. For example, our sustainability team is there, along with human resources, legal, and compliance. We have a country perspective as well, because what may be right for Singapore may not be right for India or China."

What's Working: Telenor Handles 'Black Boxes' With Care

As Telenor Group's AI practice has advanced to using deep learning with larger, more complex models that are difficult to explain, it has taken a cautious approach regarding which domains this technology is used in. Until model explainability and interpretability improve, Telenor believes deep learning is best applied to low-risk problems and advocates much more rigorous testing to build trust.

DBS's third stage of review, model governance, is a work in progress, says Lam. "How do you govern an AI [system], and do you in some situations have to disregard the AI? How do you determine whether an AI is fit for purpose? In what situations do you or do you not involve humans? All these questions are really important," he says.

"The reason for putting a ton of time and effort into this is that if we get it right, it is potentially a competitive advantage," Lam says. "If we get this correct, if you can develop a framework that can be applied broadly — when do you release an AI, how do you turn off the AI when it's going wrong, when do you involve humans, how do you deal with this properly? — if we get that part correct, that's a real differentiator."

Working to Mitigate Bias

Managing the risk of bias in AI applications is fundamental to sound data science, because inadequate data sets can skew a model that predicts a machine failure just as easily as they can distort a model with an impact on humans. However, the question of managing bias becomes more urgent with models that make recommendations affecting people, where the potential consequences of doing the wrong thing are greater. Identifying and mitigating bias risk requires organizations to scrutinize data sets for adequate diversity and to bring diverse points of view to the table as models are developed.

"No one at any of the companies that have had issues on this ethics piece got up in the morning and said, 'I'm going to make a discriminatory AI,'" says Kanaan of the U.S. Air Force. But such cases have "illuminated inherent biases that we have to address. Now we get to have that conversation."

Kanaan says efforts to mitigate bias start with understanding the purpose of any project and then evaluating whether there is appropriate data to build an unbiased model. He and other experts interviewed cite recruiting as a use case where it may be particularly difficult to assemble a data set that does not perpetuate bias. Training data based on past successful hires will almost inevitably reproduce biased decision-making, no matter how often a model is tweaked.

Frameworks for AI and Data Ethics

Organizations seeking to develop ethical guidelines to govern their use of AI and data can review a growing body of frameworks created by many organizations around the world working to promote ethical use of AI. The Berlin-based nonprofit Algorithm Watch, which is devoted to evaluating algorithmic decision-making processes designed to predict human behavior or to make automated decisions, maintains the "AI Ethics Guidelines Global Inventory," with links to dozens of examples from industry, academia, and governments.

Below are some prominent examples of frameworks or principles that organizations can draw on when creating their own guidelines:

The International Association of Privacy Professionals has published "Building Ethics Into Privacy Frameworks for Big Data and AI," a white paper that explains how organizations can implement data ethics in their operations. It draws on a meeting the organization held with United Nations Global Pulse.

The Singapore Personal Data Protection Commission worked with industry leaders, technology experts, and academic researchers to publish the "Model Artificial Intelligence Governance Framework," which it says "helps translate ethical principles into pragmatic measures that businesses can adopt." (DBS Bank was a participant.) The framework focuses on four areas: internal governance, decision-making models, operations management, and customer relationship management.

Data for Democracy's Global Data Ethics Project is a community-based effort that encourages data practitioners to abide by a set of 10 specific guidelines to ensure that data projects support fairness, openness, reliability, trust, and social benefit.

The European Commission published "Ethics Guidelines for Trustworthy AI" in April 2019. It advocates that uses of AI should be lawful and respectful of ethical principles and should take into account the social environment.

The Institute for Ethical AI & Machine Learning, a U.K.-based research center that supports the responsible development and operation of machine learning systems, has published "The Responsible Machine Learning Principles." Its themes include assessing the impact of incorrect predictions and their effect on people; monitoring bias; working to explain and justify the results of machine learning systems; and ensuring data privacy and the security of data and machine learning models.

The Australian government's Department of Industry, Innovation, and Science has released "Artificial Intelligence: Australia's Ethics Framework," which delineates core principles for AI, including regulatory compliance and explainability, while adding issues like accountability for intended and unintended impacts and creating a process by which people can challenge the results.

The Future of Life Institute, a nonprofit research organization, has developed the Asilomar AI Principles, which include a set of ethics and values that cover ensuring privacy and transparency into algorithms, along with a statement that AI should not be used for lethal weapons.

Intel's Greer says he sees a significant increase in practitioners' conversations around ethics, with bias being top of mind. "I think people are determined to try and get it right, which is really, really a positive," he says.

Mitigating bias is a constant challenge, Greer says. "It does require some diligence and some forethought in order to ensure that we aren't selecting data sets that simply reinforce our predefined conclusions." Likewise, he wants to ensure that data analysts aren't being drawn from a homogeneous population and that the tools they're using aren't inherently biased.

The work of reducing the risk of bias extends to the practice of one data science team adopting another's work. "The common process is to adopt, then adapt, other people's algorithms or training to models. However, in doing so, it can lead to this hereditary insertion of the same biases that were used to train the original models into this new analysis," Greer says.

Greer also advocates for diverse participation in efforts seeking to establish principles for the ethical use of data (see "Frameworks for AI and Data Ethics"). He is involved with Data for Democracy, a group run by volunteers that is actively building a diverse community of people that includes professional data scientists, citizen data scientists, ethicists, and sociologists. The ideal, he says, is to assemble "a cross-functional diverse team of people who are able to look at the use of data and project ways to ensure that the guidelines and policies that are created benefit all communities."

And ultimately, assembling similarly diverse data science teams, and ensuring that ethical guidelines are followed, will be essential to any organization if AI is to gain the trust needed to deliver on its promise. ●

About the Research

MIT SMR Connections conducted a global online survey in June and July 2019, drawing 2,280 survey respondents from *MIT Sloan Management Review* readers. Respondents represent a broad range of functions and industries, with more than 80% of respondents identifying themselves as holding C-suite, board, or management roles. The survey drew respondents from all regions of the world.

To provide a rich context for discussion of the quantitative research results, we interviewed analytics experts, including practitioners, consultants, and academics. These individuals provided insight into how the drive to implement AI is changing organizational culture, technology strategy, and technology governance.

Acknowledgments

MIT SMR Connections is grateful to the following AI practitioners and experts who shared their experiences and insights in interviews for this report.

Vickey Chang, vice president, data and analytics, U.S. Information Services, Equifax

Melvin Greer, chief data scientist, Americas, Intel

Peter Guerra, chief data scientist, North America, Accenture

Sameer Gupta, managing director, chief analytics officer, DBS Bank

Capt. Michael J. Kanaan, cochair for artificial intelligence, U.S. Air Force

Lam Chee Kin, managing director, head of the Legal, Compliance, and Secretariat group, DBS Bank

Jeff Lee, executive vice president and chief digital officer, Seacoast Bank

Jeffery Lee, executive director, legal and compliance, DBS Bank

Shane Jason Mock, vice president, research and development, American Fidelity Assurance Company

Eric Monteiro, senior vice president, chief client experience officer, Sun Life

Jimmy Ng, CIO, DBS Bank

Rob Stillwell, senior vice president, business analytics officer, Seacoast Bank

Astrid Undheim, vice president, analytics and AI, Telenor

Jan Van Haaren, chief product and technology officer, SciSports

Ray Wang, principal analyst, founder, and chairman, Constellation Research

Linda Zeger, founder and principal consultant, Auroral LLC

SPONSOR'S VIEWPOINT

Using AI and Analytics to Transform Your Business

Jennifer Chase,
Senior Vice President, Worldwide Marketing, SAS

As the world becomes more digital, artificial intelligence (AI) is increasingly critical to the way we do business. Leaders are no longer deciding whether they will implement AI — they are deciding how.

This year's *MIT SMR* Connections survey shows that organizations are rethinking the way they operate to gain value from AI. They are transforming their processes, encouraging collaboration across the enterprise, and finding new ways to use AI and analytics to get tangible results.

Still, many organizations are facing challenges in getting their AI programs off the ground: Less than half of survey respondents reported active adoption. Most are still in the early stages of executing AI strategies. The good news is that leaders are committed to transforming their businesses. They understand that making changes today leads to big benefits tomorrow.

SAS has long been an advocate of AI and its associated technologies. We've been a pioneer in analytics, including machine learning, for more than 40 years and have decades of experience in natural language processing. And we're embedding AI into our core solutions so that our customers automatically benefit from AI capabilities.

But it's important to remember that even if you have the most powerful analytics available and expand automation across ever more business processes, you still need humans to drive business strategy. Machines are not taking over the world. They don't understand strategy. They lack the vision required to truly drive change at a strategic level.

Strategic vision can only come from business leaders and their teams. The survey explains how commitment from an organization's C-suite is key to a successful AI effort. When the technology becomes part of the business culture, it's more accessible to everyone.

At SAS, we know that bringing AI into the day-to-day workplace — and making it more accessible — is integral to improving the world around us. Business leaders need to account for the convergence of people, processes, and technology. AI is one piece of the puzzle.

Fortunately, the survey results indicate that businesses are shifting in the right direction. I'm excited to see this evolution as more organizations make AI part of their everyday decisions. It's a big undertaking — transformation often is. But in the end, it's well worth the effort.

— **Jennifer Chase**, Senior Vice President, Worldwide Marketing, SAS