

The global AI agenda



Preface

“The global AI agenda: Promise, reality, and a future of data sharing” is an MIT Technology Review Insights report produced in partnership with Genesys and Philips. It was developed through a global survey of over 1,000 executives across 11 different sectors and a series of interviews with experts having specific responsibility for or knowledge of AI. Denis McCauley wrote the report; Claire Beatty was the editor, and Nicola Crepaldi was the producer. The research is editorially independent, and the views expressed are those of MIT Technology Review Insights.

The survey

The survey was conducted in January and February 2020. The respondents are evenly distributed globally, with 20% based in each of North America, Europe, Asia, Latin America, and the Middle East and Africa. The sample is senior, 26% of respondents being C-suite executives, 30% directors, 16% heads of AI, and 10% heads of data or analytics. Over half (55%) of the organizations they represent are large, earning annual revenue of \$1 billion or more; nearly one-third (32%) have revenue of \$5 billion or more.

Of the 11 sectors represented, the largest contingents come from manufacturing (15%), IT and telecommunications (14%), consumer goods and retail (13%), financial services (11%), and pharma and health care (10%). The other sectors in the survey are professional services, energy and utilities, transport and logistics, travel and hospitality, media and marketing, and government.

Expert interviews

We would like to thank the following individuals for providing their time and insights:

George Bailey, Managing Director, Digital Supply Chain Institute

Tony Bates, CEO, Genesys

Kiran Bhat, Co-Founder and CTO, Loom.ai

Adi Chhabra, Head of Product Innovation, Vodafone UK

Kay Firth-Butterfield, Head of Artificial Intelligence and Machine Learning, World Economic Forum

Mike Hanrahan, CEO, Intelligent Research Lab, Walmart

Karen Hao, Senior AI Reporter, MIT Technology Review

Dirk Jungnickel, Senior Vice President, Enterprise Analytics, Emirates Group

Céline Le Cottonnec, Chief Data Innovation Officer, Bank of Singapore

Gary Marcus, Founder and CEO, Robust.AI, and Professor Emeritus, New York University

Hossein Rahnama, Founder and CEO, Flybits

Daniel Schreiber, CEO, Lemonade Insurance Company

Jeroen Tas, Chief Innovation and Strategy Officer, Philips



CONTENTS

1. Executive summary	4
2. Making business sense of AI	5
One step at a time	5
3. Use cases today and tomorrow	8
Use cases that deliver	9
If at first you don't succeed.....	11
Personal AI.....	11
4. The challenges of scaling	13
Process, culture, and data	14
5. Unlocking opportunities for data sharing	16
Value for data.....	18
From vision to reality	18
6. Conclusion	20

01 Executive summary

Artificial intelligence technologies are no longer the preserve of the big tech and digital platform players of this world. From manufacturing to energy, health care to government, our research shows organizations from all industries and sectors are experimenting with a suite of AI solutions across numerous use cases.

Among the organizations surveyed for this report, 72% had begun deploying AI by 2018, and 87% by 2019. Yet much remains unknown about AI's real, as opposed to potential, impact. Companies are developing use cases, but far from all are yet bearing fruit. How, business leaders ask, can we scale promising use cases to multiple parts of the enterprise? How can we leverage data, talent, and other resources to exploit AI to the fullest? And how can we do so ethically and within the bounds of regulation?

MIT Technology Review Insights surveyed 1,004 senior executives in different sectors and regions of the world to understand how organizations are using AI today and planning to do so in the future. Following are the key findings of this research:


- **AI deployment is widespread but will take time to scale.** AI is being deployed widely across sectors, but its reach within enterprises is likely to expand slowly. Most survey respondents (60%) expect AI to be used in anywhere from 11% to 30% of their business processes in three years' time, exercising an important, though not dominant influence in their operations. Financial services providers, manufacturers, and technology companies have the highest expectations of AI penetration.
- **Change management and data challenges do most to hinder scaling of AI.** Taking AI use cases beyond pilot stage is far from straightforward for any organization. The surveyed firms struggle most with the change management involved in modifying business processes to leverage AI, a challenge cited by 51% of respondents. Nearly as difficult are data challenges—cited by 48%—chief among which are difficulties integrating unstructured data and in interfacing with open-data platforms (problems reported by 57% and 53% of executives, respectively). Respondents' emphasis on the latter suggests a desire to access external data to feed their AI models.
- **The top AI use cases today are in the areas of quality control, customer care, and cybersecurity.** Around six out of 10 manufacturers and pharma companies are using AI to improve product quality. Nearly half of consumer goods and retail firms (47%) are using it in customer care. Over half (51%) of energy firms are leveraging AI for monitoring and diagnostics, 58% of financial services providers for fraud detection, and 52% of tech firms to strengthen cybersecurity.
- **Currently nascent, data sharing can magnify the impact of AI.** Two-thirds (66%) of surveyed firms are willing to share internal data externally to help develop new AI-enabled efficiencies, products, or even value chains. Manufacturers, consumer goods firms, retailers, and health sector organizations envision benefits to supply chain speed and visibility and reduced time to market of new products. Technology and financial services firms see gains to customer service, cybersecurity, and fraud detection, among other uses. Businesses are still cautious, however, and more clarity is needed in privacy regulation and industry standards, say 64% and 58% of respondents, respectively, before data sharing takes root.
- **Early AI adopters are benefitting the most, but also have war stories.** Firms with the longest experience of using AI are learning by trial and error and ultimately benefitting. Surveyed organizations that first deployed in 2015 are more likely than those deploying later to have seen their AI projects underperform in terms of return on investment (ROI). But the early adopters are also more likely than others to say that ROI has exceeded expectations.

02 Making business sense of AI

Few emerging fields of technology have generated as much excitement and debate in recent years as AI. Most of the excitement thus far (and much of the debate—see “Ethics first”) has centered around the technology industry, which in the US, China, and, to some extent, Europe has been investing billions in developing its AI capabilities.¹ Organizations in other sectors may not be spending on AI with the same abandon, but the survey suggests that most are at least testing the waters.

By the end of 2020, 97% of the large companies surveyed for this report will be deploying AI. The earliest adopters have been IT and telecoms firms, with 81% using AI by 2018—just ahead of financial services firms (78%) and consumer goods and retail (75%). The public sector is also finding numerous use cases for AI: by the end of 2019, 94% of government respondents said they had deployed AI.

Figure 1: In what year did you, or will you, first deploy AI in your organization?
(% of respondents)



Source: MIT Technology Review Insights survey, 2020

“For us, AI is a foundational technology that in the next couple of years will be found in the vast majority of our propositions.”


Jeroen Tas
Chief Innovation and Strategy Officer
Philips

AI will naturally play different roles in different types of businesses. For some, its capabilities will help to improve aspects of operational efficiency. Others expect it to be game changing. “For us,” says [Jeroen Tas](#), chief innovation and strategy officer at consumer and health-care equipment provider **Philips**, “AI is a foundational technology that in the next couple of years will be found in the vast majority of our propositions.” And for some companies born in the online environment, the success of the business model rests on it. One is Lemonade, a New York-based online provider of property and casualty insurance that is described as a “disrupter” of the established insurance industry.² Its chief executive [Daniel Schreiber](#) believes AI-powered bots such as Lemonade’s are “the future of insurance.”

One step at a time

Deployed widely though it is, AI is not about to conquer the enterprise. Very few surveyed executives (4%) believe it will be used in more than half of their business processes in

Figure 2: In three years, approximately what percentage of your business processes will use AI? (% of respondents)



Source: MIT Technology Review Insights survey, 2020

three years' time. Less than one-third (30%) expect it to be used in between 31% and 50% of processes. The majority, 60% of respondents, believe AI will find a place in anywhere from 11% to 30% of their processes—a considerable but not necessarily dominant influence on how most businesses operate. On this measure, AI will play an especially large role in the operations of financial services providers, manufacturers, and technology firms.

Unsure of its actual (as opposed to assumed) capabilities, and cognizant of the challenges it poses, businesses have been taking an iterative approach to AI's deployment. According to [Dirk Jungnickel](#), senior vice president, Enterprise Analytics, at **Emirates Group**, early attempts taken by his company in 2015 and 2016 to build AI capabilities often did not make it beyond the proof-of-concept or pilot stage. That began to change in mid-2018 when, he says, the firm began to “operationalize” its AI development efforts and “industrialize data science.”

[Céline Le Cotonnec](#), chief data innovation officer at **Bank of Singapore**, notes that many corporate organizations want to deliver AI use cases before having first set the right foundations in terms of IT architecture, AI capabilities, target operating model, or data governance. “My response is ‘first things first,’” she says. “We need to set up the right

foundation for data before moving ahead.” This includes having “analytics translators”—people who understand AI as well as the business to scope projects, manage the delivery, measure return on investment, and understand how feasible it will be to scale up.

For large companies setting aside substantial budgets to develop AI capabilities, developing and prioritizing use cases is a considerable challenge. According to [Mike Hanrahan](#), CEO of **Walmart's Intelligent Research Lab (IRL)**: “The first thing we had to invest time in was deciding where we should focus our resources.” His team identified over 250 different use cases and then filtered them down to a handful. “The filtering process was pretty complex in deciding what we should work on,” says Hanrahan. “It came down to deciding which cases were the most practical to scale.”

For Walmart, the priority cases related to management of inventory, which represents a large chunk of the multinational's cost base, and where even small improvements in efficiency generate significant savings. Across sectors, however, businesses are pursuing a variety of use cases, in the hope that at least a couple will generate early returns and serve to build confidence and excitement in AI throughout in the organization.

Ethics first

There are many examples of algorithms producing unintended consequences. Amazon, a company at the forefront of AI development is just one company that has attracted unflattering press for its AI development in recent years.

Amazon is not alone. Technology companies have come under heavy criticism in the past two years from policymakers, non-governmental organizations, and the media for perceived inadequate efforts to apply ethical guidelines to their development of AI techniques and technologies.³ Companies in other industries should not consider themselves immune to the same scrutiny and should draw lessons from the difficulties the tech leaders have encountered.

According to [Karen Hao](#), senior AI reporter at **MIT Technology Review**, few experts working in organizations know how to uphold ethics while developing an actual AI model. “The technology community,” she says, “can no longer deny [that inattention to ethics in AI model development] is a problem, and that it needs to be addressed early in any model’s development.”

The most important ethics challenge today, says Hao, is establishing algorithmic accountability. Software engineering provides a model for how AI engineering could develop: “It has strict practices for how to debug code, ensure that everything is running smoothly, that there are no unintended consequences, and that the results align with whatever principles have been established.”

Hao notes that some AI experts, in academia as well as the technology industry, recommend that model builders write down all the steps being taken to build a model, and systematically check them against the organization’s ethical principles. “Any company can do that,” says Hao. “The most important thing is having teeth in the process: if you’re finding misalignment, you should not be able to deploy a model until you’ve found a solution that aligns with your principles.”

Ultimate accountability for AI ethics naturally rests with boards and C-suites, but too few currently appreciate what is involved, says [Kay Firth-Butterfield](#), head of artificial intelligence and machine learning at the **World Economic Forum** (WEF). This is why the WEF has recently published a toolkit to help boards better understand AI and provide guidance on establishing oversight over AI development and use.⁴ Among the options it outlines for an AI oversight structure is the establishment of an ethics board or panel. She acknowledges the mixed record of such panels to date⁵ but stresses “they remain a good idea provided the right lessons are taken.”

[Gary Marcus](#), founder and CEO of **Robust.AI** and professor emeritus at New York University, argues, “Every company probably ought to have an ethics board, and not just for its use of AI. But we can’t leave everything to self-regulation, since the interests of companies are rarely fully aligned with the interests of society as a whole.” Even for companies with long experience of AI, he argues, human oversight should remain: “AI just isn’t smart enough yet to be fully trusted.”



03 Use cases today and tomorrow

Returns are visible from companies' early AI deployments. Among the survey respondents, an overwhelming majority say the return on investment (ROI) from AI projects is meeting (59%) or exceeding (37%) their expectations.

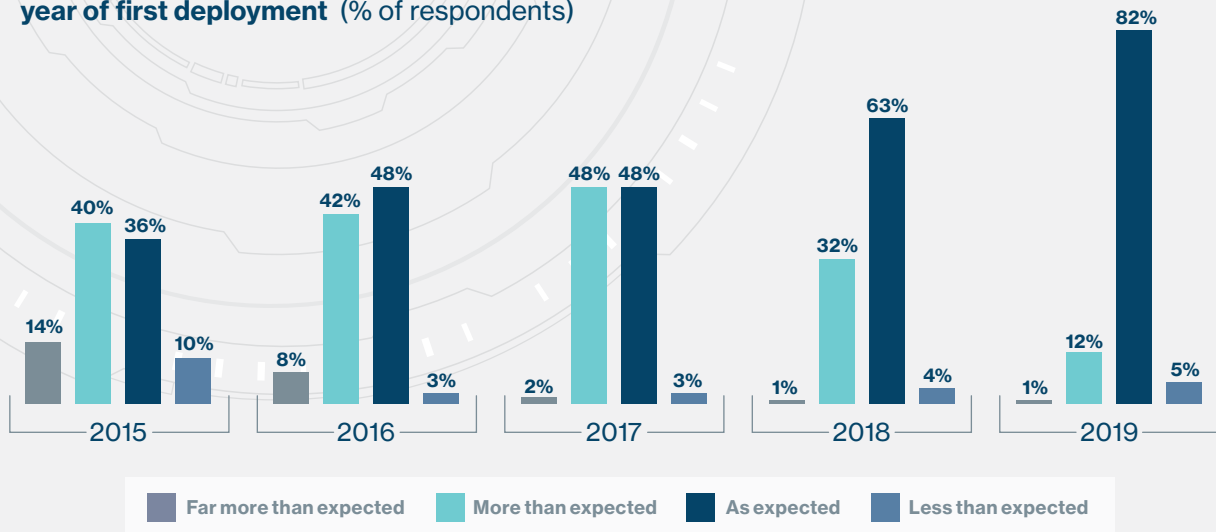
This holds true across all the sectors represented in the survey, although media and marketing firms are three times more likely than the rest to say that ROI has failed to meet expectations (12% versus 4% for the overall sample).

Experience, not surprisingly, appears to breed success. The earlier that surveyed firms first deployed AI, the more likely that their returns have surpassed expectations. Of


those firms introducing AI in 2015, over half—54%—say ROI has exceeded expectations, compared with 50% of those deploying in 2016 or 2017, and 33% of those doing so in 2018.

At the same time, early adopters are also more likely than the others to have been disappointed with ROI (10% versus no more than 5% of firms deploying in any other year). This suggests that the more experienced users of AI have learned well from the difficulties that often come with being a first mover. It also suggests that those introducing AI later are learning from the early adopters' mistakes.

Figure 3: Extent to which AI has delivered returns on investments, by year of first deployment (% of respondents)



Source: MIT Technology Review Insights survey, 2020



“In the telecoms industry, customer service eats up a lot of costs. Integrating AI with IVR (interactive voice response) almost instantly removes cost from these operations and it leads to faster decisions being taken to address customer issues.”

Adi Chhabra
Head of Product Innovation, Vodafone UK

Use cases that deliver

The returns that AI adopters are seeing most often take the form of improvements in operational efficiency and increased cost savings, cited by 51% of survey respondents. Better management decision-making and improved customer experience (cited by 44% and 41% respectively) are other oft-mentioned benefits from investment in AI.

Emirates provides an example of savings generated from a specific AI use case: premium-class meal catering. Dirk Jungnickel explains that AI algorithms predict, for each individual flight on each day, the volume of food that will be consumed by its business class passengers. “In premium classes, customers’ first choice of hot meals must be met, which ordinarily requires significantly over-stocking meals per passenger. The associated costs are not only in meals but in the additional fuel consumed in supporting the weight lift. The predictions that the algorithms generate help us to prevent over-catering, food and fuel waste, and reduce costs. On the scale that we operate, that translates into significant savings.”

For mobile operator Vodafone, the most successful use cases are found in customer service, according to [Adi Chhabra](#), head of product innovation at **Vodafone UK**. The benefits, he says, accrue both to cost efficiency and


to the customer experience. “In the telecoms industry, customer service eats up a lot of costs. Integrating AI with IVR (interactive voice response) almost instantly removes cost from these operations and it leads to faster decisions being taken to address customer issues.”

Among the survey sample as a whole, quality control, customer care, and fraud detection are currently the top AI use cases. A more detailed view, however, reveals considerable variety in the main use cases pursued by different sectors. For example, over half of financial services firms (58%) and government organizations (55%) cite fraud detection as their top AI use case. Céline Le Cotonnec says that in her previous role at a global insurance firm, fraud detection use cases were the first that were implemented, since the benefit promised to be sizeable: According to a Reinsurance Group of America (RGA) survey from 2017, 3-4% of all global claims are fraudulent, with the highest incidence (4.16%) in Asia.⁶


Cybersecurity use cases are being pursued by over half of surveyed firms in the technology (52% of respondents), professional services (59%), and government (58%) sectors. Around six in 10 manufacturers and pharma/health-care providers, meanwhile, are developing quality control use cases for AI. Customer care is a major AI use

Figure 4: Leading use cases by industry sector (% of respondents)


Manufacturing



Financial services



IT and telecoms



Pharma and health care



Transport and logistics



for 62% of firms in travel and hospitality, 48% in the technology sector, and 47% in consumer goods and retail.

Energy firms and utilities, according to 51% of respondents from that sector, use AI to monitor the state of their networks. Monitoring and diagnosis are also a natural AI focus for organizations in the health-care industry, and 42% of respondents from these organizations are pursuing such uses. According to Tas at Philips, AI has completely changed the way the industry diagnoses cancer, and its capabilities now extend further, to selecting treatments: “Once diagnosed, AI algorithms now help us to select the right therapy. That’s complex because the options could include surgery, ablation, chemotherapy, immunotherapy, or radiation, or a combination of these. Selecting the right therapies and pathways is becoming an insights-driven, AI-enabled exercise.”

If at first you don’t succeed ...

Cancer therapy is an example of an AI use case that would not have been viable a couple of years ago for most health-care organizations due to technology limitations, a lack of relevant data, high costs of development, or other factors. Its current viability for Philips highlights the fact that what are deemed unrealistic use cases today may very well become feasible within a couple of years. For Tas, one use case that will soon reach viability is personalized, preventive risk screening. “This means much more personalized pathways [than we have today], where we screen for diseases or conditions based on a detailed understanding of risk and deterioration. Then we can intervene with individuals at the right time before things truly deteriorate.”

Walmart’s Mike Hanrahan believes the growth of edge technology—in which computing takes place near the device and source of data, and away from the cloud—will increase the scalability of real-time video analysis techniques the company uses in its Intelligent Research Lab (see “Walmart’s AI factory”). The ability to run this AI-driven analysis at the edge, he says, will reduce the need for powerful but expensive servers. The sensors required are also becoming cheaper, he adds, and the networks less expensive to run.

Personal AI

The shift of computing power to the edge of networks is also likely to bring life to personal AI. The term is often used today to describe highly tailored personal assistants and chatbots, as well as digital replicas of humans (see “Getting personal with AI avatars”). It could eventually come to mean much more, as individuals learn how to

build or customize health, entertainment, safety, education, and other apps that use AI to deliver instant analysis and act on it.

AI is certain to take personalization to new levels in other ways. The marketing industry, for example, uses “hyper-personalization” to refer to the gathering and real-time analysis of customers’ behavioral data so that companies can tailor products and services that suit each individual customer. According to [Tony Bates](#), CEO of **Genesys**, a software company, this means personalized experiences as well as products and services, and it is likely to change how companies have traditionally thought about customer service and engagement.

“The use of algorithms is helping companies to understand the channels that customers want to use in interactions, and precisely when they want to use them,” he says. “AI makes it possible to learn about customer preferences near the ‘front door’—when and where interactions happen. It will challenge companies’ previous thinking about how they use traditional CRM [customer relationship management] systems and databases.”

“The use of algorithms is helping companies to understand the channels that customers want to use in interactions, and precisely when they want to use them. AI makes it possible to learn about customer preferences when and where interactions happen, challenging companies in how they use traditional CRM systems and databases.”

Tony Bates, CEO, Genesys

Getting personal with AI avatars

Personal AI can take many different forms. Among the most familiar today are apps that generate instant insights on an individual's health, wealth, and dating prospects, among other things. A more recent manifestation is the personalized avatar. This is not the static or animated depiction that often portrays social media users. It is rather an interactive avatar that is a digital recreation not just of a user's likeness but also of his or her expressions. A San Francisco startup, **Loom.ai**, founded by veterans of Hollywood animation studios such as Dreamworks and Lucasworld, uses deep learning techniques to study an individual's facial mannerisms and mirror his or her facial expressions and movements on avatars in different types of conversations and situations.



Virtual assistants, or chatbots, used by firms to interact online with customers and other stakeholders, will eventually incorporate personalized avatars.

Loom.ai's product is a platform that allows software developers to create and personalize avatars for use in enterprise settings. The company's co-founder, [Kiran Bhat](#), expects videoconferencing to be the most common use of such avatars as its platform gains adherents. Videocalls are a typical application of the technology, he says, as employees simply substitute their personalized avatar for their real persona while on the call. The voice remains theirs, but the avatar's expressions are the product of algorithms. Although basic, this use case has the potential to generate considerable savings for companies, says Bhat, who maintains that a call using avatars costs a fraction of one using streamed video. Given the hundreds or more conference calls that a multinational may hold over the course of a year, the savings could be sizeable, he says.

Virtual assistants, or chatbots, used by firms to interact online with customers and other stakeholders, will eventually incorporate personalized avatars, says Bhat. AI techniques imbue today's chatbots with a considerable degree of intelligence, but what's missing, according to Bhat, is an expressive face as well as a personalized voice. "Avatars in this setting must be emotive and highly expressive," he says, "but it's also very important to be able to personalize voice." Existing capabilities to do the latter remain somewhat primitive, but they will be perfected before too long, he believes.

While personalized avatars are likely to find consumer uses as well, such as in gaming in tandem with virtual reality, a considerable amount of market education will be needed for avatars to widely take hold, believes Bhat. Employees learning how to personalize their at-work personas will help breed familiarity and could bring that knowledge into their home lives. For now, though, Bhat and his colleagues believe the enterprise market will keep them busy for a long time to come.

04 The challenges of scaling

It is clear from the managers we've interviewed that scaling AI use cases is proving difficult.⁷ Existing technology limitations can hinder wider adoption, as suggested by Hanrahan in the case of real-time video analysis. When it comes to forms of personal AI, advances need to be made in natural language processing before chatbots, for example, become truly sophisticated, says Chhabra.

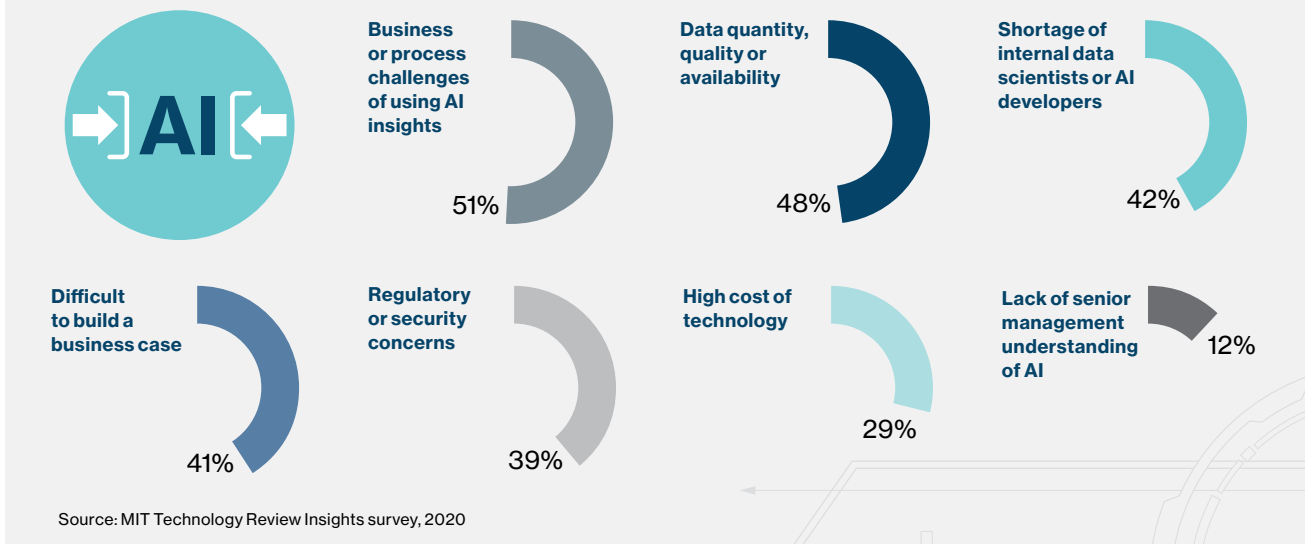
In some industries, regulation hinders the wider application of AI-enabled innovations. A case in point is algorithm-based insurance pricing in the United States. According to **Lemonade's** [Daniel Schreiber](#), this is currently allowed in only a handful of US states. "In the US, the world's largest insurance market, the regulatory environment has not yet made allowances for these next-gen technologies," he says.

Shortages of AI-related talent and skills are a frequent lament of CIOs and CTOs, and 42% of the respondents to our survey say a shortage of internal data scientists and related experts is a major constraint on their use of AI. The deficit is felt keenly among manufacturing and technology industry respondents, cited as an AI constraint by 48% and 47% of them, respectively.

Companies building advanced AI models internally are often constrained by a talent deficit, says Chhabra, as these require very sophisticated skills. Outside the technology industry, most firms are more likely to tap the capabilities of AI start-ups for such purposes, he says: "These are already working on some of the problems you want to address, and you can use their services or easily integrate their tools within your architecture."

"Implementing AI or machine learning in a business at scale is a huge change-management challenge. Everyone wants to do it independently, and there often is little understanding of the value of integrating data from different parts of the business in one platform, on top of which you can train your AI models."

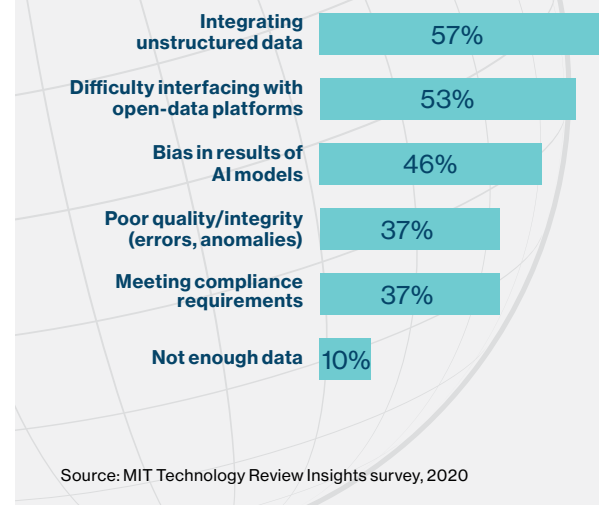
Dirk Jungnickel
Senior Vice President, Enterprise Analytics
Emirates Group

Figure 5: What are the greatest constraints to your company's use of AI? (% of respondents)

Process, culture and data

Organizations trying to take AI use cases beyond the pilot stage struggle above all with the extent of process and culture change that is required. Over half of the surveyed executives—51%—cite the business or process challenges among their chief constraints to using AI. “Implementing AI or machine learning in a business at scale is a huge change-management challenge,” says Jungnickel. It requires not only changing business processes, but breaking down organizational siloes. “Everyone wants to do it independently,” he says, “and there often is little understanding of the value of integrating data from different parts of the business in one platform, on top of which you can train your AI models.”

Data, indeed, is another major constraint, ranked a close second to process issues by the survey respondents (cited by 48% of them). AI models need data: the more a model ingests, the more accurate its analysis and the more likely that the decisions it prompts will hit the mark. The problem is less with the overall availability of data, however. Just 10% of respondents say they struggle with this. Difficulties integrating data in different formats, especially unstructured data, are the bigger problem, according to Le Cottonnec. Insurers should start to leverage on their unstructured data (call logs, pictures, e-mail), she says, in order to improve claims process efficiency or enhance customer experience. “We have enough data, but our ability to extract value from

Figure 6: What are the biggest data challenges that your company has encountered in the process of developing AI? (% of respondents)

unstructured data is limited at the moment.” Her industry is not alone: it is an issue experienced by 57% of the firms in the survey (and by considerably more professional services, transport, and government organizations).

Over half of the respondents—53%—also cite difficulties in interfacing with open-source platforms. This is less of a

constraint when organizations rely primarily on their own data to feed their models. But the 53% figure is a clear indicator that many businesses are looking farther afield

for data. They may believe that the power and accuracy of their models will increasingly come to rely on the ability to integrate data from a variety of internal and external sources.

Walmart's AI factory

Shoppers are unlikely to find a data center in their average neighborhood supermarket. They will, however, if they visit Walmart's store in Levittown, New York. Since April 2019 the 50,000-square-foot store has been the home of the company's Intelligent Retail Lab (IRL), set up as a test bed for emerging in-store technologies. Mike Hanrahan, the IRL's CEO, calls it Walmart's "artificial intelligence factory." Unlike other companies that develop and test AI-based systems in discrete R&D labs, Walmart chose to set it up in a fully operating store,



Using an operational store as an AI research lab allows Walmart to educate its associates and customers about how AI works and about Walmart's view of it in retail.

one of its busiest Neighborhood Markets in the country. "We toyed with building a lab where our engineering capacity is located," says Hanrahan. "But we felt that to properly 'productionize' AI, we needed to put the IRL in a living store that has hundreds of associates, 35,000 products, and all the interactions and nuance that you would expect in a real-world Walmart."

The data center (which customers can view) provides the computing power to analyze large volumes of data generated by intelligent cameras and sensors deployed in the store. This equipment monitors shelves in order to track inventory levels. Algorithm-based tools then alert staff when shelves need re-stocking, or alternatively if perishable food products are no longer fresh and need to be removed.

Beyond saving employees' time and ensuring customers can always find what they need, implementing AI use cases such as this in an operational store provides another advantage, says Hanrahan: "It allows us to educate our associates and customers about how AI works and about Walmart's view of it in retail. It also allows us to have conversations about some critical issues raised by AI, such as privacy, bias in algorithms, and the future of work."

The Levittown model, says Hanrahan, has been "over-engineered to enable us to develop whatever use cases we can possibly imagine," he says, which cannot easily be done at other stores, due to bandwidth and computing power constraints. In this regard, the IRL is playing an important role by experimenting with use cases. "We determine which cases can run efficiently at scale, and then push them out to the rest of the retail chain."


Unlocking opportunities for data sharing

In June 2019, 10 large pharmaceutical producers formed a consortium for the express purpose of sharing drug research data that each can use to train their AI algorithms.⁸ It is not the first time that drug majors have engaged in R&D collaboration. But it is the first instance in this industry of research collaboration in which AI is both enabler and beneficiary. Consortium members trawl each other’s data using “federated learning” techniques, which are a decentralized form of machine learning. This allows the data being searched to remain in each company’s servers rather than being pooled in a central repository. The companies can safeguard what they consider proprietary, while the use of

blockchain ensures full traceability of the data. The ultimate objective: simplifying and accelerating drug discovery and development, resulting in new and less-costly drugs and treatments reaching the market.

This example highlights how the use of AI, in combination with other technologies, can facilitate efficient and secure data sharing between companies, and the benefits that could result from powerful AI models built on shared data. The benefits could take the form of new efficiencies, new products and services, or even new value chains that form around data-sharing arrangements (see “Banking on a reputation of trust”).

Figure 7: How willing would you be for your company to share internal data with third parties for the purpose of building new value chains, products, or services? (% of respondents)




Source: MIT Technology Review Insights survey, 2020

Two-thirds (66%) of our survey respondents express a willingness to share their own data for these purposes. Manufacturers, consumer goods and retail firms, and technology companies are the most enthusiastic, although all sectors exhibit a large degree of readiness. How are they looking to gain from it? Manufacturers see their chief wins in the forms of greater supply chain speed and visibility, more efficient production operations, and faster and more innovative product development. Respondents from the consumer goods and retail and pharma and

health-care sectors cite the same supply chain and product development gains. IT and telecoms executives see the benefits chiefly in enhanced customer service experiences and stronger cybersecurity and fraud prevention, the same benefits as are top-of-mind for financial industry respondents.


Sharing data across value chains and between different providers in a customer journey can be a huge benefit to consumers, says Tony Bates at Genesys. “We generate

Figure 8: What would be the greatest benefits of sharing data with companies in your own or adjacent industries? (% of respondents)



Source: MIT Technology Review Insights survey, 2020

Figure 9: The greatest perceived benefits of data sharing, by industry



	New business models	Faster and more innovative product development	New or enhanced customer services and experiences	More efficient or innovative manufacturing	Greater speed and visibility across supply chains	Cybersecurity or prevention of fraud
Consumer goods and retail		X			X	
Financial services	X					X
Travel and hospitality		X			X	
IT and telecoms			X			X
Manufacturing		X		X	X	
Pharma and healthcare		X			X	
Transport and logistics		X			X	
Energy and utilities	X				X	
Media and marketing		X	X			X
Professional services					X	X
Government		X			X	

Source: MIT Technology Review Insights survey, 2020

billions of interactions that we can build common data models around. But we also need other forms of data. There is a need to share data across the different large technology stacks that we use and manage. If we can do this in an anonymized and secure way, we as an industry can ensure that our customers get much more personalized experiences than they are getting now.”

The pharma industry example cited above illustrates how advances in AI and other technologies create the potential to make data sharing more secure and attractive—even among industry rivals. For business-to-consumer companies in particular, the chief hindrance to date has been the difficulty of ensuring anonymity of their customers’ data. Chhabra at Vodafone says that complete anonymity has been elusive: “As much as you want to anonymize, users of shared data can find ways to track the customers the data refers to.” He believes that blockchain and newer developments such as Ocean Protocol, which works similarly to federated learning in allowing the decentralized searching of data, will eventually ensure that AI models can be run on fully anonymized data.

Value for data

Creating the ability to share data securely is one piece of the puzzle. The other, says Chhabra, is creating an incentive for companies to go ahead and share it. Those behind Ocean Protocol, a nonprofit platform developed by a Singapore-based foundation, see its technology underpinning the formation of “data marketplaces” in which companies, consumers, and other parties share or trade data.⁹ The benefits to companies can extend beyond those highlighted in Figures 8 and 9. Some will be able to monetize their data on such platforms, whether through charging fees or exchanging theirs for other data of similar value.

In whatever forms of sharing companies take part, they must learn how to value the data they hold and the data

Some 64% of survey respondents say that regulation needs to change or be clarified before companies will widely begin sharing data with third parties.

they need, says [George Bailey](#), managing director of the **Digital Supply Chain Institute**, a research organization. It may mean assigning a monetary price to different types of data, or identifying what data it is willing to exchange in return for that of other parties. He cites the examples of sporting goods producers that see a benefit in obtaining data from companies that support fitness apps, and consumer durables producers seeking data on end-user preferences and characteristics from retailers. Bailey also advocates data valuation and trading within companies as a means of overcoming entrenched data siloes.

From vision to reality

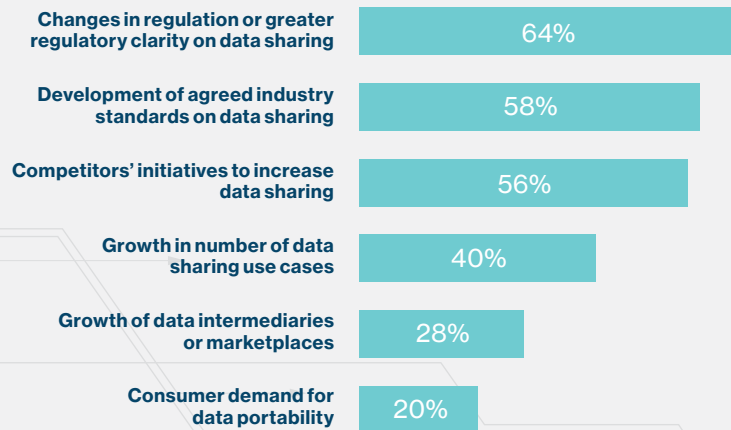
Data sharing on a scale that leads to new AI-enabled efficiencies, products, and value chains is a vision to be realized rather than a current reality. All those we interviewed believe it will come to pass, but that it will take time before board and C-suite misgivings about the security and privacy risks of doing so are eased. That view is reflected in the survey, where 64% of respondents say that regulation needs to change or be clarified, and 58% that industry standards need to be developed, before their firms will embrace data sharing widely.

Stringent data privacy regulations, such as the European Union’s General Data Protection Regulation (GDPR), in force since 2018, understandably give executives pause in light of the penalties they could face for failure to meet its requirements. Blockchain technology may be part of the solution, but Chhabra also highlights what is known as the “GDPR-blockchain paradox”—the former requires the ability of data erasure at an individual’s request, while non-erasure—full traceability—is inherent in the latter.

How such issues are resolved remains to be seen, but some organizations advocate the establishment of data trusts to facilitate sharing. One is the Open Data Institute (ODI), a UK-based nonprofit organization: it maintains that there is a strong appetite among UK businesses for such trusts set up as independent institutions to act as stewards of the data being shared.¹⁰

Philips, says Jeroen Tas, is a big proponent of creating a non-commercial trust to manage health-care data, and is in talks with the European Commission to support such initiatives. “There is a consensus in this industry that says, ‘Maybe we shouldn’t just be investing in brick-and-mortar hospitals and other facilities. Maybe the future of health care is a very solid health-data infrastructure that’s secure.’”

Figure 10: Which developments would be most likely to lead your company to engage more actively in data sharing with third parties?
(% of respondents)



Source: MIT Technology Review Insights survey, 2020

Banking on a reputation of trust

Most respondents to our survey believe that data sharing in alliances or ecosystems with a view to building new AI-based products, services, and even value chains is a good idea. How would such alliances work, and who could be the entities at the center of them? [Hossein Rahnama](#), founder and CEO of **Flybits**, a Toronto, Ontario-based tech start-up, believes that commercial banks are well placed to serve as hubs for future data alliances.

Rahnama describes Flybits as a platform-builder. Its technology enables companies to build AI-powered tools that can systematically upsell customers with “hyper-personalized” offers of related products or services at scale. Financial institutions are currently the startup’s main market because, says Rahnama, of the enormous underused “vaults” of data that they hold. Those vaults form the foundation of an ambitious role he sees banks playing in the data-sharing alliances that will take shape in the future—a role supported by Flybits’ platform technology. When banks partner with other companies, he believes, they’ll be able to develop insights about their shared customers that can be used to tailor meaningful experiences for individual consumers.

Large and small companies alike see the virtue of data sharing, according to Rahnama, and are willing to make their own internal data available to facilitate it.

“They realize that innovating with AI requires access to data that few organizations have at their sole disposal.” The biggest constraint, he says, is the reservations that firms have about data portability—namely how to encrypt and tokenize it to guarantee anonymity of customer and other sensitive data.

Rahnama believes that banks, with the help of technologies such as blockchain and edge computing, as well as the unique forms of trust they have built up over decades as secure repositories, can be the intermediaries behind data alliances. “Banks’ biggest assets are their trust protocols and the trust networks that they have created for financial assets. From an infrastructure and data perspective, they can leverage the same trust networks to become the data vault of their customers as well as the hubs of data alliances.”

With banks’ data vaults at the center, ecosystems of energy companies, telecoms operators, retail chains, and many other players can, according to Rahnama, share data in a secure manner and use AI algorithms to predict the needs of their customers while protecting the privacy of their customers’ data. And like any ecosystem, the value will build along with the membership: “The more members in that data alliance, the more data they share, the value of those insights will grow exponentially.”

06 Conclusion

This report, “The global AI agenda: Promise, reality, and a future of data sharing,” has explored the AI use cases that businesses are developing, the challenges they face in scaling AI, and their readiness to share data to enhance their AI capabilities. Five lessons for companies can be drawn from the research:

- **Adopt a portfolio management approach to AI development.** Even well-resourced giants such as Walmart need to be ruthless in prioritizing the use cases they can invest in. Projects that can demonstrate early returns in the form of cost efficiencies or improved customer satisfaction can build internal credibility and confidence. Hold onto projects that are shelved, as advances in technology and other developments can make today’s unlikely projects viable tomorrow.
- **Don’t wait to set governance and ethics guidelines.** The risks of an AI project failing or causing harm are arguably higher when organizations are only just learning about AI’s workings and so much remains unknown. Governance and risk-management regimes should be updated now to capture AI-specific challenges and risks. Ethical transgressions due to algorithm bias, for example, even if unintended, can bring reputation damage, as technology firms have discovered to their cost. Ethical AI use deserves more than lip service.
- **AI will not deliver without process change.** Developing an AI model and hoping for the best is a recipe for failure. Only process changes on the back of an AI deployment—for example, in how customer queries are handled, in how store shelves are re-stocked, in how patients are monitored—will deliver significant returns. Companies find bringing about such change the most difficult among AI-related challenges, but they can’t do without it.
- **Share your data to help AI work for everyone.** Hoarding data will eventually prove self-defeating, as companies come to find that the gains to AI model performance from data sharing outweigh any risks involved. Beyond commercial gains, societies will also benefit when data-rich AI models lead to accelerated vaccine discovery, for example, safer roads, or more-reliable public transport. Much needs to be done to build confidence in data sharing, but technology advances in areas such as blockchain and federated learning are likely to make it a safer proposition.
- **Establish the value of your existing and needed data.** Sharing data doesn’t mean giving it away for nothing. Before taking part in data alliances of any type, establish an exchange value of the data that you have to share, and also what you’re willing to pay or exchange for the types of data that you need. “Trading” data based on agreed values can also help to unlock internal siloes that hold data needed by the organization’s AI models.

About MIT Technology Review Insights

MIT Technology Review Insights is the custom publishing division of *MIT Technology Review*, the world's longest-running technology magazine, backed by the world's foremost technology institution – producing live events and research on the leading technology and business challenges of the day. Insights conducts qualitative and quantitative research and analysis in the US and abroad and publishes a wide variety of content, including articles, reports, infographics, videos, and podcasts. And through its growing MIT Technology Review Global Panel, Insights has unparalleled access to senior-level executives, innovators, and thought leaders worldwide for surveys and in-depth interviews.


Endnotes

1. In 2017, McKinsey estimated that tech giants such as Google and Baidu spent between \$30 and \$40 billion on AI-related R&D and acquisitions in the previous year alone. "Artificial Intelligence: The Next Digital Frontier?" McKinsey Global Institute, June 2017
2. "First, Fire All The Brokers: How Lemonade, A Millennial-Loved Fintech Unicorn, Is Disrupting The Insurance Business," Forbes, May 31, 2019; "Disrupting the Insurance Industry with AI," Harvard Business Review podcast, August 14, 2019
3. See, for example, "The Problem with AI Ethics: Is Big Tech's embrace of AI ethics boards actually helping anyone?" The Verge, 3 April 2019; and "In 2020, let's stop AI ethics-washing and actually do something," MIT Technology Review, 27 December 2019
4. Empowering AI Leadership: An Oversight Toolkit for Boards of Directors, World Economic Forum, 2020
5. "Google scraps AI ethics council after backlash: 'Back to the drawing board'," The Guardian, 5 April 2019
6. RGA 2017 Global Claims Fraud Survey, Reinsurance Group America, 16 February 2018
7. The consultancy McKinsey found in 2018 that no more than 8% of firms were pursuing practices that enabled the adoption of AI widely in the business. "Getting to scale with artificial intelligence," (podcast transcript), McKinsey & Company, October 2019
8. "Pharma groups combine to promote drug discovery with AI," Financial Times, 4 June 2019
9. Why we are investing in Ocean Protocol," Outlier Venture website, December 2017
10. Data trusts: lessons from three pilot, Open Data Institute, April 2019. The ODI was established in 2012 by Sir Tim Berners-Lee, the founder of the internet, and Sir Nigel Shadbolt, a leading UK expert on AI

Illustrations

Page 7, 12 illustration by Michele Singh with elements from frimufilms/Adobe stock. Page 9 illustration by Michele Singh. Page 15 illustration by Michele Singh with elements from webdesignhot.com and BSGStudio

While every effort has been taken to verify the accuracy of this information, MIT Technology Review Insights cannot accept any responsibility or liability for reliance by any person in this report or any of the information, opinions, or conclusions set out in this report.



MIT Technology Review Insights

 www.technologyreview.com

 @techreview @mit_insights

 insights@technologyreview.com