

**MIT SMR CONNECTIONS**

STRATEGY GUIDE



# Assessing AI Readiness: Planning for Today — and Tomorrow

Part 4 in a series on the journey to AI success.

---

ON BEHALF OF:



**Assessing AI Readiness: Planning for Today — and Tomorrow** ..... 1

- 1. Balance Current Needs and Future Goals
- 2. Embrace Risk; Accept Failure
- 3. Establish an Ethical Approach
- 4. Invest In Continuous Learning

**Pulling Ahead of the Pack: Dick’s Race to Retail AI Leadership** ..... 2

**Implementing an Ethical AI Practice** ..... 4

**Checklist: Plan for AI Now — and Down the Road** ..... 5

**Sponsor’s Viewpoint**..... 6

CONTENTS

This guide is the last in a series on AI readiness. Learn more at [sloanreview.mit.edu/SAS-AdoptingAI](https://sloanreview.mit.edu/SAS-AdoptingAI).

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

Copyright © Massachusetts Institute of Technology, 2021. All rights reserved.

# Assessing AI Readiness: Planning for Today — and Tomorrow

**Executives in just about every industry are talking about artificial intelligence (AI)** — and those conversations come with plenty of questions. This four-part series of Strategy Guides, produced by SAS and *MIT Sloan Management Review Connections*, provides answers, offers advice on overcoming challenges and avoiding pitfalls, and shares expert insights designed to help organizations find the best ways to harness the power of AI. This fourth Strategy Guide explores ways to balance current and future priorities, embrace risk and the possibility of failure, and incorporate AI into a larger digital transformation and analytics strategy. It also offers a snapshot of a well-known retailer's AI strategy, along with tips for building an ethical AI framework.

No question about it: Artificial intelligence is becoming an increasingly essential component for virtually every type of business — and its importance is likely to escalate in the near future. As Beena Ammanath, executive director of the Deloitte AI Institute, puts it: “Today, AI is no longer an afterthought. Tomorrow, it will be everywhere.”

As established earlier in this series, AI is impacting industries from health care to manufacturing to financial services to retail and beyond. Nearly two-thirds of the 2,700-plus executives who participated in Deloitte's most recent annual “State of AI” survey report that AI is already giving them a competitive edge. Nearly 75% expect to incorporate AI into every enterprise application with the next three years. And more than three-quarters agree that AI will transform their organizations.<sup>1</sup>

“In the next 10 years, we can expect that access to infrastructures and speed of computing will allow organizations to scale AI like never before,” Ammanath says. The result: a more connected world that might include self-managing supply chains, call centers that know what customers need before they call, extreme levels of individual product personalization, and even satellites that can predict climate-related crises to improve warning and response times. “This all sounds futuristic now, but with the pace of change we're seeing in the market, it could soon be reality,” Ammanath says.

## Balance Current Needs and Future Goals

Meanwhile, one massive AI challenge facing every organization is how to manage today's initiatives while also planning for what's down the road.

The AI team at Cleveland Clinic knows all about weighing the big-picture, extended view against the need to quickly solve smaller, more

immediate problems. “We have an opportunity right now to revolutionize the entire health care spectrum, from how we research disease and generate discovery to the entire care delivery process,” says Chris Donovan, executive director of enterprise analytics for the 6,500-bed health system, which employs more than 70,000 people in several U.S. and overseas locations. While continuously working toward that ambitious long-term goal, his team continues to tackle a variety of current projects.

One recent example is a model that allows care coordinators to determine which patients within a specific population are most at risk for future health problems or increased health care expenditures. “The question is, can we change from acute care — that is, making you better — to preventative care — that is, keeping you from getting sick?” Donovan says. The team built a model with algorithms that use predictive risk scores to identify such patients, including those who aren't showing symptoms of a particular condition but appear to be moving

**One often overlooked aspect of AI might be called “the expectation gap” — that is, the difference between what people hope a particular AI solution will do and how it actually operates. That's a consequence of the fact that, in many cases, AI teams are still figuring out how those solutions work even as they're being deployed.**

in that direction. “By using that score and the results, we can provide that information to our care coordinators. They can use that to inform their work and decide which patients to focus on and reach out to first,” he says. “That’s been really successful.”

Balancing short- and long-term priorities is the biggest single AI challenge for Dick’s Sporting Goods, says Steve Miller, senior vice president of strategy and analytics for the 800-store retail chain.

## Pulling Ahead of the Pack: Dick’s Race to Retail AI Leadership

At Dick’s Sporting Goods, every AI initiative is all about serving “athletes,” the venerable retail chain’s term for its customers.

“We think about it in two buckets,” says Steve Miller, senior vice president of strategy and analytics for the Pittsburgh-based company, which operates more than 800 retail stores nationwide. “Bucket No. 1 is: How do we use AI and machine learning [ML] to create amazing experiences for our athletes? And Bucket No. 2 is: How do we drive efficiency?”

Key to successes in initiatives in both those buckets is Dick’s ScoreCard loyalty program, which provides a wealth of data about its more than 20 million active users. “These users are the lion’s share of our sales — far more than 70% of our sales,” Miller says. The company draws on that foundation in, for instance, developing initiatives to predict what products or discounts will be most appealing to athletes, or to personalize their search results on the dicks.com website.

Because many orders ship directly from Dick’s stores, the company also uses predictive algorithms to balance the needs of in-store and online customers, making sure it has inventory available for both. “So rather than saying we won’t sell an item online if there are only two or three in a store, just to protect that inventory for the in-store shopper, we’re going to let the AI decide,” Miller says. “And the AI is going to know how not to let down that brick-and-mortar shopper but also maximize the inventory available for that online shopper.”

Ultimately, Miller says, Dick’s will continue to differentiate itself through great in-store experiences and human expertise. “It’s about offering easy ways to pick a product and get it the same day through our network of stores, or talking to an expert who really understands how to pick out a baseball bat or a golf club,” he says. “AI is super important, and it will absolutely change the world — but if that’s your only business strategy, it’s important to remember that a lot of good old-fashioned retail principles are not going away.” ●

“There’s so much opportunity, particularly in the short term, that it’s still very exciting,” he says. “But we need to make sure that we’re thinking about what our athletes” — the company’s official term for its customers — “will expect in five years.”

Dick’s addresses that challenge in large part through strategic change-management efforts, Miller says. For instance, “It’s getting business owners to buy into using an algorithm, using AI to guide decision-making — not to *make* decisions, but to *guide* them — as opposed to just relying on their own intuition and Excel modeling,” he says. It also involves getting other IT staffers to buy into the AI team’s efforts: “It’s something we’re all figuring out together.” (For more on the company’s AI-related initiatives, see “Pulling Ahead of the Pack: Dick’s Race to Retail AI Leadership.”)

### Embrace Risk; Accept Failure

One often overlooked aspect of AI might be called “the expectation gap” — that is, the difference between what people hope a particular AI solution will do and how it actually operates. That’s a consequence of the fact that, in many cases, AI teams are still figuring out how those solutions work even as they’re being deployed. That, of course, can increase the chances that a solution will perform poorly or fail.

“With AI, there are a lot of unknowns,” says Ritu Jyoti, program vice president for the Worldwide Artificial Intelligence and Automation Research Practice and global AI research lead at IDC. “So you have to have a culture where people are embracing innovation and there’s no penalty for any kind of risk-taking.”

As an example, she cites one major corporation whose C-level executives have sent an organization-wide message that failure in the pursuit of AI-related innovation is not only acceptable, but expected, especially in the early stages. “So people understand that it’s OK to do projects in small pieces and fail because they have that risk-taking culture,” Jyoti says.

That’s important in an environment in which — depending upon which research source you cite — 50% to 85% of AI projects fail in terms of not delivering as expected, having little or no impact, or generating new problems rather than solving existing ones.

Jose Murillo of Grupo Financiero Banorte is all too familiar with the odds. “Right now, I’m running three experiments per week,” says Murillo, chief analytics officer for the organization, which is one of Mexico’s largest banks. “I would be lying to you if I tell you that I am able to scale everything up or that they’re all successes.”

## It's no secret that organizations of all kinds and sizes, in every industry, are concerned about AI-related ethical issues ranging from ensuring fairness to preventing bias to protecting privacy — and if they aren't, they should be.

Any analytics officer can describe specific projects that just didn't pan out. Curtis Dudley, vice president of enterprise analytics and data services at Mercy, a St. Louis-based health network, recalls trying to develop a model for predicting the likelihood of patients failing to show up for clinic appointments. "It was good at predicting the *number* of no-shows that we were going to have — in fact, it was pretty darn close," he recalls. "But it was not good at predicting with certainty who would and wouldn't show up, or predicting the time slot, which was what we really needed to know if we're going to increase availability for patients." (For the record, he notes that many of Mercy's other AI projects, such as one calculating the likelihood of a patient's readmission after chemo treatment, have been successful and even won awards.)

### Establish an Ethical Approach

It's no secret that organizations of all kinds and sizes, in every industry, are concerned about AI-related ethical issues ranging from ensuring fairness to preventing bias to protecting privacy — and if they aren't, they should be.

"While the opportunity for benefits in this space is just huge, we really have to be careful to do this right," says Donovan, of Cleveland Clinic. "There's an equally large opportunity for harm if we don't develop and apply a thoughtful approach as to how best to use these tools and approaches in health care."

AI-savvy retailers are beginning to offer previously unparalleled levels of personalization for customers, says Brian Kilcourse, managing partner with Retail Systems Research and former senior vice president and CIO of Longs Drugs. But that opportunity comes with its own set of ethical concerns. "Consumers are willing to outsource buying decisions to trusted allies," Kilcourse says. "If I can trust in an entity to make sure that my prescriptions are available on time, that my basic food needs get replenished, that I get a reminder when it's time to replace my shoes, that's a decision I don't have to make." However, when trustworthiness comes into question, that cozy relationship evapo-

rates. "That's the Achilles heel of all of this automated decision-making and predictive capability," Kilcourse says. "It takes time to win consumers' trust, but only seconds to lose it. Companies that use modern AI-enabled capabilities to personalize the value proposition need to be very careful to avoid being 'creepy'."

In independent consultant Pritam Bhavnani's view, ethics is a broader issue extending far beyond the AI team. "The company culture is set by the leadership," says Bhavnani, formerly vice president of supply chain transformation and vice president of advanced manufacturing engineering at Honeywell Aerospace USA. "You set that tone at the top, and you reinforce it regularly, constantly. That's how you ensure you're ethical in everything you do, not just the AI element of it." In fact, writing in *MIT Sloan Management Review*, researchers Thomas H. Davenport of Babson College and Vivek Kaytal of Deloitte argue that the ultimate responsibility for ethical AI extends beyond the C-suite: "Since an AI ethical mishap can have a significant impact on a company's reputation and value, we contend that AI ethics is a board-level issue."<sup>2</sup>

Who should actually oversee AI ethics? Some companies have chief AI officers, whose roles obviously include overseeing AI ethics. Others have chief ethics or trust officers, whose roles may encompass AI oversight. And still others may have ethics committees or advisory boards, or include an ethicist on the AI team. Whatever route they choose must involve not only setting ethical standards, but enforcing them, investigating real violations, and planning for mitigation and recovery from the mishaps that almost certainly will occur.

Bottom line: "To reach the full potential of AI, we need to think about ethics upfront," says Ammanath, of Deloitte. There's progress on that front, with leaders increasingly recognizing that imperative and addressing it early on in conversations about AI, she adds. "It's no longer just about 'how do we create value from data and AI?' Now, it's also about 'What are the ethical implications we should consider?'"

But there's no one-size-fits-all answer to that question, notes Jyoti, of IDC. "It depends on the organization, the location, the industry, and the regulations," she says. "But they all need to have their own ethical charters, their own codes of ethics. It's extremely important." (For more on developing such guidelines, see "Implementing an Ethical AI Practice.")

### Invest in Continuous Learning

Ultimately, according to researchers and practitioners alike, it's important to address AI as just one part of a larger digital transformation

## Implementing an Ethical AI Practice

It's no secret that, along with great promise, AI brings the potential for generating ethical problems from bias and discrimination to privacy and security concerns and more. In turn, many organizations are developing their own guidelines for both preventing and responding to ethical mishaps.

However, given the sheer number of possible ethical issues that might arise, that may sound as if it's easier said than done. The good news: There are plenty of existing ethical frameworks available to leverage as examples or blueprints. AlgorithmWatch, a Germany-based initiative that conducts research on automated decision-making, lists and links to nearly 175 such ethical charters in its [AI Ethics Guidelines Global Inventory](#). The organization notes that while these individual frameworks differ significantly by industry, organization, and location, all emphasize transparency, equality, accountability, and safety — the foundation for any AI code of ethics.

Two particularly well-known frameworks are:

**Deloitte's Trustworthy AI:** This framework is designed to address AI ethics throughout an AI project's life cycle, with the ultimate goal of managing risk and improving returns on investment. It emphasizes [six key dimensions](#) for ethical AI:

- Fairness and impartiality
- Transparency and explainability
- Responsibility and accountability
- Robustness and reliability
- Respectfulness for data privacy
- Safety and security

In the Trustworthy AI model, all six overlap with both internal AI governance and external regulatory compliance.

**European Commission Ethics Guidelines for Trustworthy AI:** Released in 2019, the [EC's guidelines](#) state that “the foundations of trustworthy AI adhere to ethical principles based on fundamental rights, including respect for human autonomy, prevention of harm, fairness, and explicability [or explainability].”

Those foundations translate to seven requirements that, like Deloitte's similar six key dimensions, should be followed throughout the AI life cycle:

- Human agency and oversight
- Technical robustness and safety (including reliability, security, and resiliency)
- Privacy and data governance
- Transparency
- Diversity, nondiscrimination, and fairness
- Societal and environmental well-being
- Accountability

Then there's FATE. While not a formal ethical framework, FATE — a commonly used acronym for *fairness, accountability, transparency, and explainability* — offers a quick, high-level summary of key AI ethical principles. As Iain Brown, an SAS data science leader, notes in the Sponsor's Viewpoint for this Strategy Guide, some organizations have turned to the approach as a simple way to begin thinking about trustworthy AI. ●

and analytics strategy. As established throughout this series of Strategy Guides and webinars, the AI approaches most likely to succeed are those focused on establishing seamless, integrated capability with the ability to evolve and adapt over time, rather than on implementing stand-alone projects or applications.

As just one example, Cleveland Clinic's AI strategy encompasses the entire analytics spectrum, connecting what Donovan calls the four key pillars of data, people, process, and technology. “We want to deliver both data and AI as a service to our organization,” Donovan says. “We want to create a distributed capability that meets our diverse users where they are and enables them to meaningfully engage in the application of these new technologies.”

A major part of the organization's strategy for achieving that goal is intensive education — and not just for its data scientists. “We really want to look at our entire workforce,” Donovan says. Accordingly, the Clinic is training many employees in statistics, data science, AI, and an-

alytics via certificate programs developed in collaboration with partner companies and local universities, rather than through traditional degree-granting programs. “That's because things just move too fast in this field,” he says. “There are too many people that want to educate, and we want to focus that learning on the areas that are going to be most applicable to their roles.”

Among many other efforts, Cleveland Clinic also launched a data governance office to centralize its efforts to work on privacy, security, data, and related areas, and established a multidisciplinary AI team — including a patient — to review use cases and requests to use the organization's data. All are necessary for AI success, Donovan says: “There's got to be a huge investment across the whole organization to raise the bar for everybody in analytics capabilities. It doesn't do us any good to deliver a fantastic model or result or insight if the people that we're delivering it to don't have the capability or the understanding of how to use it to change a decision that they're making.” ●

### REFERENCES

1. Deloitte State of AI in the Enterprise Survey, 3rd edition, July 14, 2020.
2. T. Davenport, V. Katal, “Every Leader's Guide to the Ethics of AI,” MIT Sloan Management Review, Dec. 6, 2018.

# CHECKLIST:

## PLAN FOR AI NOW — AND DOWN THE ROAD

Use this quick checklist to ensure that you're taking the optimal approach for launching AI initiatives today and in the future.

- ✔ **Set short- and long-term priorities.** Continue developing high-value solutions to current business problems while working toward future AI capabilities.
- ✔ **Empower employees to innovate.** Invite them to suggest how AI might improve efficiency or productivity, or what functions might be strong candidates for automation. And, of course, acknowledge the risk: Don't punish people when their ideas don't pan out.
- ✔ **Create a culture of lifelong learning.** AI is dynamic. As it changes, so will the skill sets your initiatives need to succeed. Be prepared to continuously retool, reskill, and upskill your team members, either internally or by outsourcing the effort.
- ✔ **Learn to manage change, continuously.** Again, AI is dynamic, and as it evolves, your business goals and technical needs will as well. Having a standing strategy for addressing those changes will help you adapt more quickly.
- ✔ **Develop an AI ethics practice.** Set and enforce standards to ensure that your organization uses AI responsibly, fairly, transparently, and safely, in ways that protect both people and data. (See "Implementing an Ethical AI Practice." for examples.)
- ✔ **Establish accountability for ethical oversight.** As we've established, responsibility for ethical use of AI should start at the board and C-level. Beyond that, though, it's important to determine the point person (or people) who will actually monitor AI projects, enforce those ethical standards, and respond, quickly and appropriately, when problems arise.
- ✔ **Be prepared to investigate.** When something does go wrong — and it will — you'll need to undertake a robust root-cause analysis to find out what went wrong. You'll need to determine what data and algorithms were used, how users engaged with them, and what actions were taken in response. Having a transparent start-to-finish view of AI decision-making will make this process much easier.
- ✔ **Invest in upskilling.** Offer training, such as in-house or external certificate programs, to help existing employees come up to speed on the latest data, AI, and related skills.
- ✔ **Build a community of practice.** Help prevent your AI team from reinventing the wheel with each new project. Establish a center of excellence or other mechanism for AI team members to seek assistance and share what they know.

For a webinar based on this Strategy Guide, visit [http://bit.ly/MIT-SMR\\_SAS\\_AdoptAI\\_4](http://bit.ly/MIT-SMR_SAS_AdoptAI_4).



**Iain Brown** is head of data science for SAS, U.K. and Ireland, and adjunct professor of marketing analytics at the University of Southampton. Over the past decade, he has worked across all sectors, providing thought leadership on the topics of risk, artificial intelligence, and machine learning and their ethical applications. At SAS, he has delivered numerous projects and driven innovation in AI, machine learning, deep learning, and natural language understanding. An experienced public speaker and published author, he has presented at international conferences and conventions and has published papers in the *European Journal of Operational Research*, *International Journal of Forecasting*, and the international journal *Expert Systems with Applications*. He was included on DataIQ's 2020 list of the 100 most influential people in data.



## SPONSOR'S VIEWPOINT

### A Q&A With Iain Brown

In this Q&A, Iain Brown, SAS's head of data science for the United Kingdom and Ireland, discusses some key risks, ethical issues, and platform questions that organizations should consider before adopting AI and takes a quick look at current and emerging AI trends.

*This conversation has been edited for clarity, length, and editorial style.*

#### **Q: In your view, what are some of the biggest risks associated with the development of AI solutions?**

**Iain Brown:** In my view, the biggest risk we see today is how transparent the process leading up to an AI-derived decision is. It's all too easy for AI to be used in a way that starts to make decisions that a business — and, more important, its customers — may not be happy with. Now the risk here is typically a result of not having adequate safeguards around the process. By that, I mean not having clear transparent lineage to how those decisions have been made and how as a business they can cyclically update those decisions, monitor those decisions, validate those decisions — and make new decisions as more information comes online.

What I typically advise in situations where this arises is the adoption of a structured approach to the deployment of AI. As an example, I have seen several organizations adopt the “FATE” approach — which stands for fairness, accountability, transparency, and explainability — to good effect. If organizations don't have a structured approach, they won't necessarily understand what their models are doing and subsequently can't explain them. If there's no transparency, I would argue that businesses are no longer in control, as the decision-making is being done automatically, increasing the risk of bias.

So it's imperative that organizations considering AI should be thinking about the operationalization [AIOps] piece to generate value, but they should also be thinking about how this will work longer-term, when those models start making decisions. We need to keep a view of what those decisions are, so there needs to be a degree of human oversight in this process as well.

#### **Q: That leads right into a broader discussion about AI ethics.**

**Brown:** That's a very complex and thorny topic. If you're holding AI to a certain standard, how do you measure that standard? Where do you put it as a baseline for what should and shouldn't be done? It's imperative that organizations understand what they're developing, why they're developing it, and how these decisions will impact their customers.

Going back to the FATE analogy, I think that's core to having a good process and protocols in place. The ethical decisions are more than just the models. And it's all too easy to simply blame the data — that is, the data itself may not be fit for the purpose, and therefore, bad decisions are made. It goes much further than that. How algorithms are chosen and tuned needs

**“We are all consumers, and our expectations are increasing in terms of what organizations do with the data that we provide. I think that will only increase with the next generation of consumers.”**

to be considered. Although an algorithm in its own right is not biased, the way algorithms are developed, the way they're trained, the way they're tuned, and the way the parameters interact — ultimately, bias does feed into the model along with the data that's being fed into it.

In using AI, organizations need to be very careful with how they wield what is potentially a powerful tool in a business's toolbox. They need to fully understand what they're trying to develop, and they should be validating at every step that it's making the right decisions for their business and the right ones for their customers as well.

**Q: Will most organizations be looking toward a heterogeneous environment going forward, rather than a one-size-fits-all platform?**

**Brown:** I think they will. Having a single platform is great, but that doesn't mean you can't have multiple platforms that work together. From an SAS perspective, we have some amazing capabilities, some amazing technology, but we work within ecosystems that have a very wide mix of technology. And they need to be working hand in hand.

I do think the ecosystems will continue to grow in terms of what capabilities are there. But this goes back to a fundamental point: If you're just adding functionality for the sake of it and not looking at it from a strategic perspective, there's a risk that you'll over-complicate decisioning and generate inefficiencies. So I think that organizations will continue to extend and add to the platforms internally, but there still needs to be a joined-up view across those so that you can still get to the roots of the decision, and there should be as much transparency as possible when these functionalities are working together.

**Q: What are some AI trends you're seeing today, and what's on the horizon?**

**Brown:** We're seeing a trend towards composite AI adoption. By that, I mean the combination of different AI techniques to solve organizations' problems. I'm seeing a definite increase in the uptake

of AI techniques focused around unstructured data, such as natural language generation and natural language processing, to provide much more conversational platforms in terms of what organizations offer their customer bases. We're seeing that kind of growth trend toward the embedding of AI in systems.

Chatbots are one example. I still think they've got a long way to go to truly mimic humanlike interaction, and probably 90% to 95% of the chatbots we experience today are still really rules based — they have a manual list of answers to the key questions that most of us will ask, but there is no adaptive element to that.

In the retail world, we're seeing cases where there's an augmented approach to how you'll be buying clothes or buying produce in the future, and utilizing AI for much stronger recommendations, personalization, greater relevance to you personally. They're viewing customers not as a segment, but as individuals, deciding how to treat them based upon the data they provide.

But that goes both ways. We are all consumers, and our expectations are increasing in terms of what organizations do with the data that we provide. I think that will only increase with the next generation of consumers.

Organizations need to be conscious that if they're just harvesting information and there's no benefit in return, people eventually switch off. Ultimately, the organizations that will succeed are the ones that provide something worthwhile, some reward, in exchange for the data that individuals provide to them.

---

## ABOUT SAS

Conquer your analytics challenges, from experimental to mission critical, with faster decisions in the cloud. SAS Viya enables everyone — data scientists, business analysts, developers, and executives alike — to collaborate, scale, and operationalize insights everywhere.

**To learn more, visit [www.sas.com/viya](http://www.sas.com/viya).**