

MIT SMR CONNECTIONS

STRATEGY GUIDE



Technical Readiness for AI:
What Your Organization Needs to Know

Part 3 in a series on the journey to AI success.

ON BEHALF OF:



Technical Readiness for AI: What Your Organization Needs to Know 1

- 1. Examine Your Ecosystem
- 2. Know Your Data
- 3. Pick the Appropriate Infrastructure
- 4. Make Decisions About People and Processes
- 5. Bring It Back to the Business

Demystifying ModelOps and More 2

AI Technical Readiness: A Five-Point Framework 4

Checklist: Gauge Your Technology Ecosystem's Readiness for AI 5

Sponsor's Viewpoint..... 6

CONTENTS

This guide is third in a series on AI readiness. Learn more at 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.

Technical Readiness for AI: What Your Organization Needs to Know

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. In this third Strategy Guide, you will learn to assess your organization's readiness, from a technology standpoint, to undertake specific AI initiatives, with real-life examples from leading retail and health care organizations.

Is your organization technically prepared to launch an AI initiative? As we established earlier in this series, answering that question begins well before you begin talking about tools and systems — first by pinpointing the problem you're trying to solve (Part 1) and next by assessing your business readiness to address that problem with AI (Part 2).

When it comes to technical readiness for AI, Beena Ammanath, executive director of the global Deloitte AI Institute, sums up the ultimate objective this way: “Technical readiness is all about getting to the solution for a business problem in the fastest, most optimized manner possible.” Reaching that end goal requires ensuring that your ecosystem, data, platform, your overall approach, and your people are all truly up to the task.

Examine Your Ecosystem

Before launching any AI project, it's important to understand your organization's ecosystem — that is, the ever-expanding network of technologies both within and beyond its walls. Ecosystems — which typically involve multiple layers of users, supply-chain partners, IT providers, financial organizations, end consumers, and any number of other parties — are dynamic, expanding, and changing

“Technical readiness for AI is all about getting to the solution for a business problem in the fastest, most optimized manner possible.”

Beena Ammanath, Executive Director, Deloitte AI Institute

in real time. “AI has shown us that there's a heavy demand for ecosystems,” says Ammanath. “Regardless of which industry you're in or what problem you're trying to solve, it's a guarantee that you're going to need a technology ecosystem to rely on beyond the cloud and data that you might have.”

An AI-ready ecosystem should emphasize:

Openness. Such ecosystems don't just evolve naturally. McKinsey Digital researchers stated recently that any company hoping to use an ecosystem successfully must develop “next-generation integration architecture to support it and enforce open standards that can be easily adopted by external parties.”

Flexibility. To function effectively, AI systems should be as flexible, modular, and easy to integrate as possible. Kevin Martelli, managing director of data and analytics for KPMG, often advises AI-focused clients to build their ecosystems in what he calls a “layer and open matter.” He describes it this way: “We all know that things change rapidly and new tools are coming into the marketplace all the time. So you need to build these platforms and technologies and ecosystems in a way that you can layer on new capabilities.”

Interchangeability. That's critical for any ecosystem, Martelli says: “How easily can you move out those pieces to put in new pieces to solve what you're looking to solve? It's all about plug-and-play.”

Know Your Data

Another important step in launching any AI project is establishing an overall view of the data involved, from knowing how it's collected and where it's stored to ensuring that it's appropriate to the initiative in question. Completing the following two data prerequisites can help ensure your organization's technical readiness for AI.

Locate your data — all of it. Consider doing a comprehensive overview of your data architecture and relationships to determine, for instance, where data might be siloed, scattered, inaccessible, or inadequately secured. Consultant and author Seth Earley, well known for observing that “there is no AI without IA” (information architecture), refers to the result of this process as an *ontology*, a holistic picture of every piece of business data. Without such a master blueprint, an AI system is likely to evolve in a fragmented way that inhibits its impact, says Earley, author of “The AI-Powered Enterprise: Harness the Power of Ontologies to Make Your Business Smarter, Faster, and More Profitable” (Lifetree, 2020). His recommendation for approaching the task, in a nutshell: Begin with use cases, as discussed previously in this series, then develop the details for organizing and categorizing the data for access by those who are using it.

Identify the right data — and the right amount. Obviously, it’s important to make sure you’ve got enough high-quality, reliable data for any AI project and that it’s relevant to the specific initiative as well. But data challenges and requirements will vary significantly by industry and use case. For example:

- Retailers are facing a major shift away from the purely transactional data they’ve traditionally used to run their businesses toward being able to ingest new, non-transactional data types that move them closer to real-time response to market conditions. That’s engendering a cultural shift as well. “Retail is very good at reacting, but it’s not very good at predicting,” says Brian Kilcourse, managing partner with Retail Systems Research, and former senior vice president and CIO of Longs Drugs. “But in today’s world, the internal data that you have from your operational systems is not sufficient to be able to react in time. So retailers are trying to anticipate a scenario and set up the pieces so that when something does happen, they can react almost instantaneously.”
- For health care organizations (as in other industries), unstructured data presents a major challenge. “We have a ton of unstructured data, we have clinical notes, we have lots of data across our EMR [electronic medical record system], not to mention all of our other systems,” notes Chris Donovan, executive director for enterprise analytics at Cleveland Clinic, a 6,500-bed health system with 70,800 employees in several U.S. and overseas locations. “So you need to have data for AI to work against, and you need an environment that can make that data available and accessible to AI.” That also requires being able to incorporate the insights generated back into decision-making processes by, for instance, getting the results into an EMR or a clinical or operational workflow.

Pick the Appropriate Infrastructure

One clear critical enabler for AI technical readiness is having sufficient computing power, whether that’s on-premises, in the cloud,

Demystifying ModelOps and More

You’ve probably heard of **DevOps**, a term coined in 2009 to describe a then-new system-oriented approach for speeding up delivery of high-value IT applications and services. Essentially, DevOps — which practitioners describe as a kind of mindset or environment, as opposed to a specific technology or standard — breaks down the traditional barriers between *development* and *operations*.

Now meet **ModelOps**. Short for “model operationalization,” the approach focuses on model life cycle and governance. Like DevOps, ModelOps is intended to expedite the journey from development to deployment — in this case, moving AI models from the data science lab to the IT organization as quickly and effectively as possible.

To clear up one common misconception: Although the terms are sometimes used interchangeably, ModelOps isn’t the same as **MLOps** or **AIOps**. ModelOps focuses on *all* predictive analytics models, while the other two approaches focus more narrowly on machine learning and AI operationalization, respectively. It’s also distinct from, but related to, **DataOps**, which IT research firm Gartner defines as “a collaborative data management practice focused on improving the communication, integration, and automation of data flows between data managers and data consumers across an organization.”

So why does ModelOps matter? Because without it, your AI projects are much more likely to fail completely or take longer than you’d like to launch. Research firm IDC estimates that only about half of all models ever make it to production, and of those that do, about 90% take three months or longer to deploy.

ModelOps greatly improves an AI project’s chances for success by building the gap between the analytics and production teams. The approach eliminates silos, streamlines handoffs, helps ensure quality, and, above all, ensures scalability. As Gartner researchers put it in a 2020 report: “ModelOps lies at the center of any organization’s enterprise AI strategy.” If you’re hoping to see high-quality AI projects cross the finish line as quickly as possible, it’s an approach well worth considering. ●

or based on a hybrid arrangement. While that determination will vary widely based on industry and use case, more and more organizations simply default to a “cloud-first” technology environment. “It’s really, really hard to do this if you’re not in the cloud,” says Steve Miller, senior vice president of strategy and analytics for Dick’s Sporting Goods. “That’s absolutely a requirement.”

But simply being in the cloud isn’t enough, Miller adds: “It’s also very important that you know how to scale.” In other words, your systems must be able to add resources as needed to accommodate a growing amount of work — a critical capability in fast-moving, constantly changing AI projects. “Real-time models need to be able to respond in milliseconds, because otherwise you lose the consumer’s attention,” Miller says, using a retail example. “And the rest of benefits you might get from that personalized experience or that algorithm — they fall by the wayside if the algorithm is taking too long.”

That’s often easier said than done. An Accenture survey of 1,500 executives found that 76% understood that being able to scale was critical to their success with AI but also believed they would struggle to make that happen.

Donovan concurs that cloud platform capabilities are typically required, but many organizations rely on in-house systems as well. “Basically, you need the ability to do the math, so to speak, so whether that’s cloud or on-prem, that’s obviously a core technology requirement.” The front end requires a digital technology infrastructure with the ability to capture and understand data, make it available, and transform it into structures or formats applicable for AI, he says.

Make Decisions About People and Processes

Taking the following four steps will also help move your organization down the path to technical readiness for AI.

Establish ownership: Specifically, answer this question: Who owns the AI software and hardware — the AI team or the IT team, or both? “This is where you get organizational boundaries that need to be clearly defined, clearly understood, and coordinated,” says independent consultant Pritam Bhavnani. “That’s a real challenge.” (For Bhavnani’s additional advice on assessing technology preparedness for AI, see “AI Technical Readiness: A Five-Point Framework.”)

Decide whether to buy or build: As with any IT implementation, an AI initiative initially raises one inevitable question: Do you buy the

technology you need — or do you build it in-house? For Fernando Lucini, managing director of Accenture Applied Intelligence, the answer is clear: Buy or outsource before you build. “Don’t try to build everything yourself,” cautions Lucini, who is also the consulting company’s global data science and machine learning engineering lead.

But whatever approach you adopt, be sure to maintain it, he adds: “Once it’s deployed, if you don’t keep on tweaking it and making sure the learning is right, it becomes stale very quickly. And then it seems like you’ve failed because it wasn’t good — but it wasn’t good because you didn’t maintain it, so it became irrelevant.”

Adopt ModelOps: This emerging approach — which is more of a culture or mindset than a specific set of tools or processes — focuses on effective operationalization of all types of AI and decision models. “ModelOps is critical because it’s a huge part of what makes an analytics life cycle repeatable and sustainable,” says Bhavnani, who was previously vice president of supply chain transformation and vice president of advanced manufacturing engineering at Honeywell Aerospace USA. “ModelOps and governance are very much overlapping.” But remember that ModelOps is part of a spectrum that includes DevOps and DataOps, both of which are also typically required for true technical readiness for AI. (For more information, see “Demystifying ModelOps and More.”)

Enlist data engineers: As discussed previously in this series, it’s important to build a team with the right mix of skills. From a technical perspective, that effort includes engaging not only data scientists, but also data-engineering specialists — that is, those with significant expertise in using analytics and business intelligence

“There are a lot of technology and AI capabilities that are built assuming that the business will come, and then the business doesn’t. They don’t see the value, or it doesn’t meet their business needs, or there’s some misunderstanding.”

Kevin Martelli, Managing Director of Data and Analytics, KPMG

tools, database software, and the SQL data language, as well as the ability to consistently produce clean, high-quality, ethical data. And, of course, they should be familiar with ModelOps and the related approaches as well.

Bring It Back to the Business

Ultimately, being technically prepared for AI maps back to connecting the desired business objectives and outcomes. As Martelli puts it, using a “Field of Dreams” analogy: “There are a lot of technology and AI capabilities that are built assuming that the business will come, and then the business doesn’t. They don’t see the value, or it doesn’t meet their business needs, or there’s some misunderstanding.”

A more effective approach is for the AI team to help their business counterparts understand AI’s capabilities and potential. “Have a better way of brainstorming with them or interacting with them to generate the ideas that are going to help the business, as opposed to steering them in a direction based on how we think they want it to work,” Martelli says. And that effort should go straight

to the top, says Lucini, of Accenture: “Make sure that the distance between the C-suite – the people who are setting up the strategy for the company – and the people who are doing the AI work is a very small distance.”

As the business team becomes better educated about how AI can solve specific problems, the more likely they are to see real value and become enthusiastic partners in the projects.

As noted in Parts 1 and 2 of this series, that’s even more likely to happen if you start with a project with smart, tangible ROI, says Ritu Jyoti, program vice president for the Worldwide Artificial Intelligence and Automation Research Practice and global AI research lead at IDC. As an example, she cites the case of a struggling startup that used a small machine learning algorithm to help regain control of its accounts receivable process. “It was humanly impossible to go and track so many records, but the use of this technology helped them do that,” she said. The result: “They were able to recover millions of dollars for the company. That’s a big win” – and one that the business side will clearly understand. ●

AI Technical Readiness: A Five-Point Framework

Independent consultant Pritam Bhavnani can count what he considers as the most important technical factors for successful AI implementation on one hand.

“There are five basic areas that organizations need to address to be technologically ready for AI,” says Bhavnani, formerly the vice president of supply chain transformation and vice president of advanced manufacturing engineering at Honeywell Aerospace USA. Those areas are:

1. Data gathering. This involves collecting data in everything from sensors and scanners to manual data entry to ERP systems to operational processes and beyond. “This is the foundation of an AI system,” Bhavnani says.

2. Data storage. This involves two types of data: time-based, which is related to operational processes, and transactional data, associated with a customer or product or part. They must be merged and stored, which, as Bhavnani notes, involves more than just dropping the data into a file. “It’s about organizing it in a manner that makes the next step, analysis, easy and executable.”

3. Data analysis. This involves obtaining insights and results that can ultimately deliver the improvements desired. “The analysis is done using multiple algorithms, which help identify correlations and issues that then drive actions,” Bhavnani says. For example, a particular credit card transaction might trigger certain parameters indicating potential fraud, such as transactions for more than a set amount or from international locations.

4. Insight-to-action communication. This involves making sure that the insights or results generated in the analysis get sent to the right place. For instance, in the credit card case, the flagged transaction should be communicated to the bank’s security team or to the cardholder or both, so that they can decide whether to approve or block the transaction.

5. Security. This, of course, involves ensuring that all systems are protected from unauthorized access. “With all that data in one place, it can be a tempting target for the wrong kind of people,” Bhavnani notes. “So setting up the IT and AI infrastructure with appropriate security is paramount for the success of any AI initiative.” ●

CHECKLIST:

GAUGE YOUR TECHNOLOGY ECOSYSTEM'S READINESS FOR AI

Use this quick checklist to ensure that you've addressed the key technology components necessary for adopting AI.

- [✓] Do a data inventory.** Know what data you've got, where it's stored, and how it fits into your overall ecosystem.
- [✓] Build strong connections between your business and AI teams.** Establish clear, direct lines of communication between the two groups. Ensure that they're well aligned in terms of the desired outcome for each initiative.
- [✓] Start with a pilot project.** Too often, AI initiatives languish at the proof-of-concept. Instead, collaborate with the business side on a pilot project that clearly demonstrates how AI can solve a key problem.
- [✓] Inventory and classify the data sources needed for your project.** Validate its availability for training and production. Tag and label data for future usage, even if you're not sure yet what that usage might be. Over time, you'll create an enterprise inventory that will help future projects run faster.
- [✓] Define scalability requirements.** As we've established, inability to scale as needed is among the single biggest reason AI projects fail. Take a proactive approach: Know when your initiative is likely to scale and prepare to accommodate that demand in a hurry.
- [✓] Determine your cloud strategy.** Will you go all in with one cloud service provider (CSP)? Or will you use different CSPs for different initiatives? Or will you take a hybrid approach, with some workloads running on-premises and some with a CSP? One consideration: Some major CSPs typically offer more than just scalability and storage space, such as providing tools and libraries to help build algorithms and assisting with deploying models into production.
- [✓] Enlist the data and machine learning engineers.** They're the all-important key link between your data scientists and your IT operations/production specialists.
- [✓] Embrace ModelOps.** A variant on the venerable DevOps approach, ModelOps streamlines AI projects by bridging the gap between analytics and production teams. (For more information, see "Demystifying ModelOps and More").

For a webinar based on this Strategy Guide, visit http://bit.ly/MIT-SMR_SAS_AdoptAI_3.



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 technical readiness for AI, customer adoption trends, IT's changing role, and mission-critical considerations for technology and talent.

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

Q: What does it mean, from both a technology and a cultural standpoint, to be ready for AI? What are the enablers?

Iain Brown: From a technology perspective, much of it comes down to maturity around data and capabilities and resources that exist internally. Many organizations want to do more with what they've got, but their ecosystems may not be ready to deal with the volume of data or the processing power needed to understand the data in the best way possible.

The technology enablers include the right data feeds coming in, having a consolidated view across the businesses, and having the right technological capabilities. There's both a software and a hardware aspect to this: making sure you have the right hardware, localized or cloud based, that can process the data that you need in a timely manner, and then the actual software, the analytical algorithms that you can utilize.

Culture is a harder piece to deal with. There's still a lot of pushback within organizations around AI adoption. That may have to do with a misunderstanding of what AI is and what it can achieve, and, in some regards, how much AI will take over the human aspect of the role. There needs to be buy-in from the business, across the business, for successful AI adoption.

There also needs to be a clearly defined reason for AI's use and a well-governed application of it, as well as a positive augmentation of existing human efforts. When organizations try to go to wholesale automation of a process, they typically struggle to, first, get buy-in internally, and second, deploy something that's robust, fair, governed, and interpretable in certain regards – and that has the right kind of human oversight.

Q: What are you seeing in terms of SAS customers and AI adoption?

Brown: Typically, organizations have focused on structured data, so they use data that already exists in well-defined databases. It's easy to manipulate. It's numeric. It's in a constant form. Where we've seen a big uptake in the last half-decade is in unstructured data usage. This could be as simple as textual documentation – think webchats, call center records – or more complex in the form of image, video, or audio.

The classic stat is that around 70% or 80% of all data is unstructured. But I'd say most organizations still aren't getting the value out of that level, so they're exploring how to generate

value from it. This usually means starting small with aspirations to grow.

We'll typically go in and have a look from a value perspective for where it's easy to start a project. That means the right level of data already exists; someone's got a good understanding of it; there's a business problem that's clearly defined.

That's often where you can grow utilization as well. If you can prove the value and show the business "this is how AI should be applied, and here's a use case in your organization of how it's working well," you'll see much more momentum behind the adoption of AI across other departments and applications.

Q: How is IT's role evolving in response to AI?

Brown: In terms of operationalizing, a lot falls back on IT. You have data scientists who are building great innovative things. But unless they can be deployed in the ecosystem or the infrastructure that exists – and typically that involves IT – there's no point in doing it.

IT's role has been elevated in terms of decision-making, and that's drawing IT and business closer together. So there needs to be a greater collaboration between the two. But I'd go further than that. I think the data science community and AI teams should be working very closely with IT and the business, being the conduit to join the two so there's a clear idea and definition of the problem that's being faced, a clear route to production. Without that, you're going to have disjointed processes and issues with value generation.

Q: In terms of AI-related talent, what's mission critical and what's "nice to have"?

Brown: Diversity of thought is extremely important. There's a lot of debate about how data scientists should grow as a community. At the moment, unfortunately, there's not a lot of diversity, and that causes all sorts of potential issues.

In terms of skills that are most applicable – this goes back to the STEM-type areas: science, technology, engineering, and mathematics. But I would say STEAM, including the arts, is also important. Having creative thinking alongside the very statistical and the technical knowledge can differentiate organizations in terms of the way they utilize AI.

You also need strong capabilities and skills from a delivery and IT perspective. You need to have DevOps- and Model-Ops-type personas sitting in there as well.

Q: Same question in terms of applications, technologies, tools: What's mission critical and what's "nice to have"?

Brown: You need to have a platform-type approach. Having a centralized set of capabilities that could be used across multiple projects and departments – that's where things really work well.

I've seen where you have different silos popping up within an organization, building their own wheels, as it were. That not only doesn't make financial sense for the organization, it also doesn't allow for efficiency or scalability gains when you start to tackle really difficult problems that touch multiple departments.

Having a platform of capabilities, whatever they may be, is essential to making sure that an organization generates value as quickly as possible and accesses the underlying data to generate that easily as well.

I do think a cloud-based environment is the way forward. You can't rely on localized versions of anything; that approach doesn't scale. Relying on third-party cloud providers – the Microsofts, Googles, and the AWSs of the world – is a good way to lower the barriers to entry. But, again, you need the right capabilities sitting alongside those environments to make the best use of them and to get the value out of the data.

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.