

O'REILLY®



The AI Ladder

Demystifying AI Challenges

Rob Thomas

REPORT

The AI Ladder

Demystifying AI Challenges

Rob Thomas
Foreword by Tim O'Reilly

Beijing • Boston • Farnham • Sebastopol • Tokyo

O'REILLY®

The AI Ladder

by Rob Thomas

Copyright © 2019 O'Reilly Media. All rights reserved.

Printed in the United States of America.

Published by O'Reilly Media, Inc., 1005 Gravenstein Highway North, Sebastopol, CA 95472.

O'Reilly books may be purchased for educational, business, or sales promotional use. Online editions are also available for most titles (<http://oreilly.com>). For more information, contact our corporate/institutional sales department: 800-998-9938 or corporate@oreilly.com.

Acquisitions Editor: Rachel Roumeliotis

Development Editor: Nicole Tache

Production Editor: Kristen Brown

Copyeditor: Octal Publishing, LLC

Interior Designer: David Futato

Cover Designer: Karen Montgomery

Illustrator: Rebecca Demarest

August 2019: First Edition

Revision History for the First Edition

2019-08-23: First Release

The O'Reilly logo is a registered trademark of O'Reilly Media, Inc. *The AI Ladder*, the cover image, and related trade dress are trademarks of O'Reilly Media, Inc.

The views expressed in this work are those of the author, and do not represent the publisher's views. While the publisher and the author have used good faith efforts to ensure that the information and instructions contained in this work are accurate, the publisher and the author disclaim all responsibility for errors or omissions, including without limitation responsibility for damages resulting from the use of or reliance on this work. Use of the information and instructions contained in this work is at your own risk. If any code samples or other technology this work contains or describes is subject to open source licenses or the intellectual property rights of others, it is your responsibility to ensure that your use thereof complies with such licenses and/or rights.

This work is part of a collaboration between O'Reilly and IBM. See our [statement of editorial independence](#).

978-1-492-07310-9

[LSI]

Table of Contents

Foreword.....	v
The AI Ladder.....	1
Introduction	1
AI Is the New Electricity	1
A Brief Definition of AI	2
AI Challenges	3
The AI Ladder	5
Modernize Your Architecture: Making Your Data Ready for an AI and Multicloud World	7
Collect: Make Data Simple and Accessible	8
Organize: Create a Business-Ready Analytics Foundation	10
Analyze: Build and Scale AI with Trust and Transparency	12
Infuse: Operationalize AI Throughout the Business	14
Conclusion	17

Foreword

Over the years, I've seen how the right name can help people understand a new and important technology concept. When Rob Thomas first told me about his idea for the "AI Ladder" at a crowded cocktail party during IBM's Think event, I thought "This is one of those ideas, like Open Source Software, Web 2.0, big data, and the Maker movement, that isn't just a label, but a map to help guide people into what had previously been Terra Incognita."

Everyone is talking about "AI" these days, but most companies have no real idea of how to put it to use in their own business. They just know that they want some of what sounds like magic. But no vendor can help pull AI out of your company like a rabbit out of a hat. The companies who have succeeded have ascended what Rob calls the "AI Ladder." First, that means understanding what business problem you are trying to solve. Next, you need to get your data in order. And that doesn't just mean your traditional sources of business and customer data. AI isn't just about your existing data; it might require creating or acquiring data from other sources that are relevant to your business problem, and that can be used to train the AI models that will then be used to understand and respond to that data in the real world. Think about the enormous amounts of data that needed to be collected to build today's speech and image recognition capabilities. You need a data architecture that can support disparate sources of data.

You also need developers and business people alike to gain new skills, and you need culture change to adapt to a new way of working. As Peter Norvig, the coauthor of the leading textbook on AI, put it in a talk at the first O'Reilly AI conference, "The workflow of

an AI developer is quite different from the workflow of the software developer.”

You also need to let go of many of your assumptions about the way things work. So often, the first iteration of new technology is just used to make a slightly better version of what went before—the “horseless carriage,” so to speak—and only later do companies realize how much they might need to change their business model to truly make use of the new capabilities. This is why each new technology revolution breeds new business leaders, whereas old stalwarts so often fall behind.

The processes for gathering, organizing, analyzing, and ultimately infusing AI throughout your organization can be thought of as a ladder. A ladder helps you climb far higher than you could go on your own by turning an impossible leap into a series of steps. Knowing what the ladder looks like helps you evaluate your organization’s readiness for AI. Just as you can’t leap to the top of the roof, neither can you get there with a ladder that has weak or missing steps. The analogy isn’t perfect, because unlike climbing a ladder, readying an organization for AI isn’t a linear journey. But knowing the steps of the ladder helps you identify where your organization is strong, and where the gaps are.

That seems like a simple idea. And, yes, it is. It’s simple like “open source,” or “mobile first,” and any another idea that is obvious in retrospect but only became so after a lot of struggle and many stumbles. The companies that advocated for those simple ideas while everyone else was resisting them got ahead; those that didn’t, fell behind.

This report introduces a roadmap that will help companies without the benefit of years of advanced AI research and hundreds of deep learning PhDs to take advantage of one of the next big steps forward in computing.

Paul Cohen, former DARPA program manager for AI and now dean of the new School of Information Science at the University of Pittsburgh has said, “The opportunity for AI is to help humans model and manage complex interacting systems.” Don’t let that opportunity pass you by.

— *Tim O’Reilly, Founder and CEO,
O’Reilly Media, Inc.*

The AI Ladder

Introduction

Artificial intelligence (AI) is one of the greatest opportunities of our time. Fueled by vast amounts of data and unprecedented advances in machine learning, it has the potential to add almost **16 trillion dollars to the global economy by 2030**. AI is poised to transform organizations across every industry and alter the very way people work. **According to Gartner, AI augmentation (a combination of human and artificial intelligence) will recover 6.2 billion hours of worker productivity in 2021**. Despite these promising forecasts, AI adoption has been slower than anticipated. It has been reported that **81% of business leaders** do not understand the data and infrastructure required for AI. This report aims to present executives and line-of-business professionals with a unified, prescriptive approach—*the AI Ladder*—for successful implementation of AI.

AI Is the New Electricity

AI is one of the greatest challenges and opportunities of our time. It is poised to change the way people work, how enterprises operate, and how entire industries transform. AI initiatives offer more than just cost savings—they actually help organizations predict and shape future outcomes; allow people to do higher-value work; automate decisions, processes, and experiences; and reimagine new business models. Ultimately, this means increased revenue.

AI, however, is often portrayed as some mystical thing; a magical black box that's being put to work, with little understanding of *how* it works. People view AI as something to be relegated to experts who

have mastered and dazzled us with it. In this environment, AI has taken on an air of mysticism with promises of grandeur and out of the reach of mere mortals. **AI is the new electricity.**

When electricity was first discovered, it was considered to be the domain of sorcerers—a magic power that left audiences puzzled about where it came from and how it was generated. All new innovations go through a similar evolution: discovery, exploration, application, and eventual ubiquity.

We find ourselves in a similar state today. We understand the power of AI, but we haven't fully discovered how to unleash its potential. The reality is: AI is not magic. Applying it is hard work. There is no wand to be waved at enterprise inefficiencies, and having the technology alone is not enough.

To set the context for this report, let's briefly define AI and discuss the common challenges in implementing AI at an enterprise level.

A Brief Definition of AI

AI is an umbrella term for a family of techniques that allow machines to learn from data and to act on what they have learned rather than simply following rote instructions created by a programmer. Machine learning is a branch of software engineering, and it is almost always a part of a larger system that also incorporates traditional software. As I've noted in my talks, if it's written in Python, we call it "machine learning." If it's written in PowerPoint, we call it "AI."

AI is behind advances in speech recognition, image recognition, and autonomous vehicles. It has enabled the creation of voice-activated assistants for your phone and your home as well as playing an important role in customer care, social media, and cybersecurity.

But for businesses, AI can be defined as a way of radically improving three things: predictions, automation, and optimization. First, AI is about *predictions*—organizations want to be able to forecast what's going to happen in their business, at both the macro and micro level. Next, there's tremendous value in *automating* critical yet time-consuming business processes that are often done manually, freeing employees to focus on higher-value, more creative work. Finally, AI is about *optimization*, whether that means optimizing routing and logistics, marketing spend, or the configuration of your cloud instal-

lation. AI is a tool for improving human decision making, at speed and scale, and has the potential to augment the work of every employee.

It's time to demystify the challenges of AI so that organizations can successfully bring the power of AI into their business.

AI Challenges

The reality is that companies of all sizes and across all industries are struggling to adopt AI. The challenges they face can be categorized into five key buckets.

The first is *lack of understanding*. Many organizations jump in and are looking to implement an “AI solution” because of its increasing popularity, and they assume that it will fix any business problem. At its core, AI represents a powerful new set of software and data engineering techniques for making sense out of vast masses of unstructured data. But it isn't a magic wand that can do anything, and it must be applied to problems that it is well suited to solve. That means the first step for organizations is to understand the business problems they are trying to solve, ask the right questions, and identify whether AI is the right approach to achieve their business goals.

The second problem organizations face is *getting a handle on their data*. Data is the foundation and fuel for AI. Good data is needed for training machine learning models, and then for the resulting AI-infused business processes to do their work. There are three different kinds of data problems:

Lack of data

Companies need to start by collecting their data, acquiring additional data from third parties, and making data accessible throughout their organization.

Too much data

Although a lack of data can hamper AI adoption, the same can be said for having too much data. When companies have too much data spread across different environments and databases, it quickly becomes a data engineering problem. In this instance, companies need to collect and organize their data to make it ready for AI.

Bad data

Garbage in, garbage out is as true in the days of AI as it was in the early days of programming when the phrase was first coined. The problem is that even though business leaders list improving the use of data as a top priority, **60% are challenged by managing data quality**. Much of the most important work in AI involves data cleaning.

The next problem organizations face is *lack of relevant skills*. AI requires even the most experienced software engineers to **relearn a lot of what they take for granted** about how to program. In current forms of software development, the programmer spells out what the computer executes. In AI-based projects, the programmer feeds training data to a machine learning algorithm, which learns from the data and constructs a mathematical model that represents the task to be performed. When presented with real-world data, the system is able to recognize the same patterns it saw in the training data, and its output is then incorporated into traditional procedural programs that act on what the model recognizes. This training process is time-consuming, and the regular software development workflows of Continuous Integration (CI) and Continuous Deployment (CD) don't easily apply.

The challenge is that AI skills are rare and therefore in high demand, so there's a shortage of skilled workers available to hire. This makes it even more important that the technology being built and used is more easily accessible to everyone within the business, regardless of skill level.

Next, there's the issue of *trust*. It is critical to ensure AI recommendations or decisions are fully traceable—enabling enterprises to audit the lineage of the models and the associated training data, along with the inputs and outputs for each AI recommendation. As more applications make use of AI, businesses need visibility into the recommendations made by their AI applications. In the case of certain industries like finance and health care, in which adherence to the European Union's General Data Protection Regulation (GDPR) and other comprehensive regulations present significant barriers to widespread AI adoption, applications must explain their outcomes in order to be used in production situations.

When companies put AI models into production, it is essential to break open the “black box” of AI and develop a strategy for

continuous monitoring of their output after deployment. Without this, organizations have zero visibility into what their AI is doing, how often it is being used, what the outcomes are, and what biases might have been revealed in the data used to train the model. Author Cathy O’Neil refers to the ways that models can encode bias as “weapons of math destruction.”

Finally, there’s the issue of *culture and business model change*, which is required to take advantage of the opportunity that new technology provides. Just as existing companies failed to embrace the internet and then the mobile revolution, they are often unwilling to do the deep rethinking of their business models and business workflows that will allow them to fully embrace the opportunities of AI. AI doesn’t just improve existing business processes. *It allows you to rethink the process altogether and do things that were previously impossible.*

AI lets your organization glean intelligence from unthinkable amounts of data. But that doesn’t mean you have old-school business analysts poring over the equivalent of smarter spreadsheets to make decisions. It means that you have new automated workers that can complete tasks reliably. The former business-analysts-turned-data-scientists need to become the managers of these new workers, evaluating their work and intervening when things go wrong, not necessarily by overriding their decision, but by going back and improving the data used to train them so that they can do the task successfully the next time. It’s also important to understand that AI is not about executing a single business project—it’s about transforming an entire business culture. It’s about creating a culture of iteration, experimentation.

The AI Ladder

As stated, AI is not magic. As companies look to harness the potential of AI, they need to use data from diverse sources, support best-in-class tools and frameworks, and run models across a variety of environments. However, **81% of business leaders** do not understand the data and infrastructure required for AI.

The vast majority of AI failures are due to failures in data preparation and organization, not the AI models themselves. Success with AI models is dependent on achieving success first with how you collect and organize data.

No amount of AI algorithmic sophistication will overcome a lack of data [architecture] ... bad data is simply paralyzing.

—MIT Sloan

Organizations need a thoughtful and well-designed approach, particularly in today's hybrid multicloud world. They need an approach that is modern and open by design, granting organizations the flexibility to work as one, with open source on any cloud. With a unified, prescriptive, and open information architecture, organizations can modernize their data architecture to make their data ready for an AI and multicloud world.

Simply put, there is no AI without an information architecture. An information architecture is the foundation on which data is organized and structured across a company. It allows organizations to eliminate data silos and avoid lock-in and run anywhere with agility. Also, with an information architecture designed for AI, businesses can automate and govern the data and AI life cycle with a unified approach so that they can ultimately operationalize AI with trust and transparency.

The AI Ladder (Figure 1-1) has been developed by IBM to provide organizations with an understanding of where they are in their AI journey as well as a framework for helping them determine where they need to focus. It is a guiding principle for organizations to transform their business by providing four key areas to consider: how they collect data, organize data, analyze data, and then ultimately infuse AI into their organization.

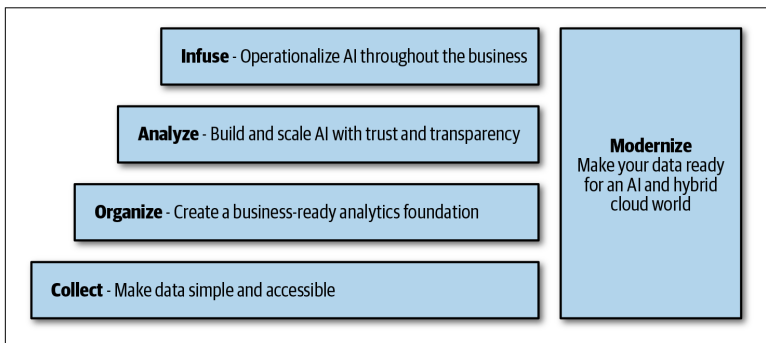


Figure 1-1. The AI Ladder, a guiding strategy for organizations to transform their business by connecting data and AI.

Breaking an AI strategy down into pieces—or rungs of a ladder—serves as a guiding principle for organizations, regardless of where they are on their journey. It allows them to simplify and automate how they turn data into insights by unifying the collection, organization and analysis of data, regardless of where it lives. By using the ladder to AI as a guiding framework, enterprises can build the foundation for a governed, efficient, agile, and future-proof approach to AI.

Let's examine the four steps that make up the AI Ladder in more detail:

1. *Collect*: Make data simple and accessible.
Collect data of every type regardless of where it resides, enabling flexibility in the face of ever-changing data sources.
2. *Organize*: Create a business-ready analytics foundation.
Organize all data into a trusted, business-ready foundation with built-in governance, protection, and compliance.
3. *Analyze*: Build and scale AI with trust and transparency.
Analyze data in smarter ways and benefit from AI models that empower organizations to gain new insights and make better, smarter decisions.
4. *Infuse*: Operationalize AI throughout the business.
Apply AI across the enterprise in multiple departments and within various processes—drawing on predictions, automation, and optimization.

No matter where you are on your journey, it's still difficult work. Having the technology alone is not enough. That is why the ladder—every step—is so critical. Each of the steps on the AI Ladder are covered in detail in the following sections. However, before you take your first step and begin collecting the data, you must modernize your data architecture.

Modernize Your Architecture: Making Your Data Ready for an AI and Multicloud World

“Modernize,” in this context, means building an information architecture for AI that provides choice and flexibility across the organization. To meet today's demands and to stay competitive tomorrow,

organizations need an efficient, agile architecture for their data that they can reliably build on. Enter hybrid, multicloud platforms.

Hybrid, multicloud platforms are the future of data architecture. A hybrid, multicloud platform allows organizations to take advantage of their data and applications across any cloud (public, private, on-premises) via containers. As enterprises modernize for an AI and multicloud world, they will find there is less “assembly required” in expanding the impact of AI across the organization.

In most organizations, data is spread across a variety of environments: public clouds (with many providers), private clouds, and traditional on-premises deployments. Hybrid, multicloud platforms address this data proliferation challenge that many organizations are facing—information distributed across multiple silos, databases, and clouds. If an organization is deploying and running AI projects in these environments, they are forced to use the AI tools from *those specific cloud providers*. This represents modern day “vendor lock-in,” which can stifle innovation and prevent businesses from scaling AI efforts. Hybrid platforms provide the foundation necessary to support all the capabilities a business would need to build, deploy, and manage AI models at scale.

Finally, given the dynamic nature of AI, organizations need to automate AI life cycles across an array of contributors through collaborative workflows. By adopting an agile, cloud-native platform, organizations can successfully make their data ready for AI, put open source to work, and infuse AI across teams.

After you’ve modernized your data architecture, you’re ready to determine where you are on your AI journey. As we’ve pointed out earlier, there are four rungs of the AI Ladder: collect, organize, analyze, and infuse. We cover each of these in the next sections.

Collect: Make Data Simple and Accessible

After an organization has modernized its architecture, it must make its data simple and accessible. For too long, data has been held captive within systems of record, and isolated by rigid platforms, segregated business functions and data types. This results in siloed data, which is difficult to access, making it impossible to gain true analytical insight. The challenges become worse as businesses evolve. The possibilities of AI to transform an organization are made moot because insights are only as good as the data.

Here are some examples of how companies can make their data simple and accessible:

Use all data types

If a retail company currently has access to only its historical, transactional data, that means its AI models are being trained on only one data source. But, what if that retail company could take its structured transactional data and combine it with unstructured data from sources like social media or live-time clickstreams, or both? Now, the retail company can build AI models that will tell it what its customers bought last week, how they feel about it, and what they're shopping for this very second—giving it a more three-dimensional view of its business.

Invest in AI-infused data management tools

The database used to be just a system of record and a data-storage repository, but that time has passed. Now, we expect intelligence on every level of the modern technology stack, from frontend user interfaces with natural language and speech recognition, throughout our selection of applications, and all the way back to our data-management layer. An AI-powered data management platform offers organizations a broad suite of capabilities to translate queries across vendors, languages, locations, and structures. This means that organizations can build the right foundation to access all of their data and ensure that the platform will grow as their data does.

A hybrid data management strategy not only can address some of the current challenges with data quality, but also can provide access to more data to fuel smarter AI.

For example, a manufacturing company is undergoing a digital transformation where it is not only looking to restructure the cost to its customers, but also reduce its total cost of ownership (TCO) and IT. As a company that generates petabytes of data—varying from reporting performance metrics, to scorekeeping how the company is doing in terms of revenue and profitability—the cost of its legacy database management software was not sustainable. So, it makes the decision to migrate its database systems to a more flexible and cost-effective platform. By using a hybrid data management system to collect and provide better access to the data, the manufacturing company has seen lower storage usage, performance improvements, and faster transactional response times. Whether it's an analytics

platform or a transactional platform, the company not only has a reliable and available database, but more important, has benefited from the TCO reductions it expected (which includes database, platform, and managed services cost).

Now, with this firm data foundation, the company's IT department can give the business reports it can actually use, as opposed to somebody pulling the data out of the system and then trying to figure out how they want to manipulate the data. This foundation also sets the groundwork for the company to look ahead at AI and machine learning solutions to continue on its path of digital transformation.

Organize: Create a Business-Ready Analytics Foundation

The cloud and mobile revolutions have accelerated the pace of data creation—structured and unstructured. However, many organizations don't know what data they actually have, where it resides, what processes use it, and how to stay compliant with current data-related laws and regulations.

At this step on the AI Ladder, there are three key issues an organization must consider.

The first is to *address data quality* and determine whether your data is “business-ready,” meaning it's been cleansed, it's not incomplete, and it's compliant, so it is ready to use to build AI models. When data is not business-ready, finding, understanding, and putting data to productive use is a constant challenge for all data consumers including data scientists, analysts, and line-of-business users.

In [Forrester's Predictions 2019: Artificial Intelligence report](#), 60% of the decision makers at firms adopting AI cite data quality as the number one challenge when trying to deliver AI capabilities. Today, [organizations spend 80% of their time preparing data](#) for productive use, creating a bottleneck for business agility, competitiveness, and profitability. A properly designed and governed [data lake](#) eliminates data inconsistencies, resolves duplicates, and creates a single version of the truth for users to access. By managing and mastering data in the data lake, organizations are able to find and trust the quality of the data they're working with.

Next, an organization must *organize and catalog data*. To understand why this step is important, an analogy of a library is useful. If

a library was simply a room filled with thousands of books, it would have little value to the average reader. Libraries are useful because they are organized and provide a catalog to help find information in various ways—you can find all the books by a specific author, in a specific genre, or on a specific topic. This catalog must be maintained every time a new book is added (or removed) from the library. Data is the same. Organizations need to have a catalog of their data to provide information about its source, who owns it, metadata mapped to its business context, and so on. When new data becomes available for users by integrated or replicated data sources (just as new books are published and available in a library), without proper catalog processes and up-to-date inventory, the data becomes difficult or impossible to know, trust, and use.

Finally, an organization must *govern the data* to ensure that only permissioned users have access. Many businesses underestimate the pitfalls that poor data governance can create. For example, a major bank learned this the hard way when it discovered that a sales manager and her team adjusted numbers that ended up costing the bank millions. Because the company didn't have its data properly organized, it fell victim to the implications of broad access control. Just as organizations must govern access to data, they must also govern the data itself.

Leaders can't be confident in their AI without ensuring that the data fed to it is trustworthy, complete, and consistent. Data is the bread and butter of AI, and the quality of that data directly affects the results of your AI. That's because your AI is only as good as your data. If you have bad data, you're not going to have trusted, transparent AI. Organizations need their data to be cleansed, organized, cataloged, and governed to ensure that only the people who should be able to access it, can access it.

Here are some examples detailing how companies can work to organize their data into a trusted, business-ready analytics foundation:

Focusing on data preparation and quality (DataOps and AI)

DataOps, a moniker for Data Operations, is a methodology that orchestrates people, process, and technology to deliver a continuous business-ready data pipeline to data consumers at the speed of the business. This enables collaboration among data consumers and data providers, creates a self-service data cul-

ture, and removes bottlenecks in the data pipeline to drive agility and new initiatives at scale.

Governing the data lake

A properly designed and governed data lake eliminates data inconsistencies, resolves duplicates, and creates a single version of the truth for users to access. By managing and mastering data in the data lake, organizations are able to find and trust the data they're working with, thus making it easy for companies to build AI models and extract actionable insights from the data lake they can trust.

Modernizing applications

Companies want their data to be secure and meet compliance regulations to protect customers and users. Companies also want to be able to deploy higher-quality applications in less time, and at a lower cost, by provisioning and refreshing test data environments on-premises or in the cloud.

Ensuring data privacy and regulatory compliance

Organizations are under two competing business goals: running a profitable business and conforming to ever-evolving laws and regulations. It's essential to meet privacy obligations and protecting personal data, which requires the discovery and classification of different types of data across the business.

Providing 360-degree information-driven insights

With multiple sources of competing customer and product data across an enterprise, organizations need a data management solution to establish a single, trusted view of information, and to use real-time data as a critical asset.

Analyze: Build and Scale AI with Trust and Transparency

After an organization has been able to collect its data and organize it in a trusted, unified view, it can now tap into that data to build and scale AI models across the business. This allows companies to glean insights from all of their data, no matter where it resides, and engage with AI to transform their business—putting themselves at a clear competitive advantage.

This step in the ladder presents its own unique challenges. Hasty investments can create some serious issues in tools, people, and process. Point solutions create complexity in integration, maintenance,

and support, leading to increased technical challenges and costs. The surging demand for data scientists makes hiring and retention a huge challenge. In addition, a lack of a unified analytics platform often results in poor return on investment (ROI) and frustration on the part of leaders looking for more positive impact. Businesses need to trust their AI before they fully scale and automate across their enterprise organization.

To build AI models from the ground up and scale them across the business, organizations need capabilities covering the full AI life cycle. This includes the following:

Build

This is where companies build their AI models. Consider this like an artist's studio, or a carpenter's workbench. It's where companies create and train models they can then use for predictions. At this phase in the AI life cycle, it is critical to ensure that companies use the right algorithms to build their models for making predictions.

Run

After a model has been built, it needs to be put into production within an application or a business process. When a model is deployed, it's now running in the organization, so it's making claims decisions, pricing decisions, and so on and can be retrained as needed.

Manage

After a model is built and running, the question becomes: how can it be scaled with trust and transparency? To address complex or diverse build-and-run environments, enterprises need a tool that not only can manage that environment, but explain how their models arrived at their predictions and scale it across their organization. By having the management in place, organizations can track who changed the model, when the model was deployed, and the lineage on the model. By tracking all of these items, organizations can ensure that their models are not biased and that they're explainable and transparent.

In today's world of regulations, GDPR, and data privacy laws, the way organizations engage with AI is under intense scrutiny. Organizations need to manage their AI across the entire AI life cycle in order to explain either to a consumer or another business how their systems came to a decision and why. For example, a bank needs to

be able to tell a consumer what the factors were behind their loan being denied, and what the consumer would need to do to change that decision.

An organization will never truly be able to scale its AI across its workflows without first establishing trust and transparency in the AI. Here are some examples showcasing why companies need visibility into what their AI is doing:

Insurance

Insurance underwriters can use machine learning to make more consistent and accurate risk assessments around claims, ensure fair outcomes for customers, and explain AI recommendations for regulatory and business intelligence purposes.

Telecommunications

Data scientists can build machine learning models and work with their IT operations teams to confidently recommend proactive asset maintenance for telecommunications providers.

Health care

Organizations in the health care industry must maintain regulatory compliance by tracing and explaining AI decisions across workflows, and intelligently detect and correct bias to improve outcomes.

Fraud

Fraud activities have continued to increase at a rapid pace over the years. Fraud is difficult to detect. That's because the data used in fraud detection tasks is complex and often unstructured. And, if data is siloed, it can be difficult for businesses to get a 360-degree view of it, leading to false positives or missed alerts, costing companies hundreds of millions of dollars.

Organizations can engage with AI to help them handle the aforementioned issues with predictive insights, real-time analysis, more sophisticated modeling techniques and automation technologies. By investing in a hybrid, multicloud platform, organizations can accomplish this all in a governed and secure environment.

Infuse: Operationalize AI Throughout the Business

The infuse rung of the ladder is about being able to apply AI across your enterprise. Meaning, organizations need to be able to put AI to

work in multiple departments and within various processes—from payroll, to customer care, to marketing, drawing on predictions, automation and optimization—to help advance their business agenda. In many cases, this will require organizations to develop entirely new workflows and business processes across every department. Imagine the vast dispatch network of on-demand transportation companies and automated warehouses. These are systems infused with AI, enabling a complex dance of humans and machines cooperating to complete ever more complex tasks at enormous speed and scale.

Here are some examples of how companies across different industries are infusing AI into their business processes:

Customer care

An AI assistant can access and “learn from” hundreds of customer touchpoints and can share this data with other systems in an integrated way so that the customer enjoys a consistent experience and is made to feel that the company “knows” her. Well-executed AI-powered chatbots improve Net Promoter Score (NPS) through speed-to-resolution and round-the-clock accessibility and through the channels in which users are most comfortable with, be it web chat, SMS (text messages), social media like Facebook, or a phone call to customer service.

Knowledge workers

AI can be used to scan and compare enormous volumes of documents to understand and find relationships, insights, and discrepancies that humans would never notice. With AI, procurement professionals and lawyers, for instance, are able to pore over billions of pages of contracts and legal documents with less effort and more accuracy, saving both time and money.

Marketing

AI excels at finding insights from vast amounts of disparate data, discovering correlations that a human analyst would never notice. In today’s world, creative marketing teams include data scientists because the amount of data marketers have access to, and the ability to apply AI to it, is changing the way marketers work—from predicting shopping trends, and optimizing marketing strategies, to delivering advertisements that are engaging, relevant, and personal.

Supply chain

With AI, companies now have visibility into their operational data so that they can derive unprecedented insights to make better-informed decisions and mitigate disruptions before they affect their ability to deliver business results. Ultimately, this improves the efficiency of the entire supply chain for an increased ROI and better customer satisfaction.

Human resources

A company's success, depends on its employees. Now, with AI, organizations have a clearer picture of predicting, with accuracy, which employees are at risk of attrition, and help identify the right candidate for a job opening. It can also be used to help determine which employees deserve job promotions and raises.

Insurance

With AI, the claims approval process for insurance carriers is more consistent and automated, can more accurately predict risk assessments, and ensure fair outcomes for customers. This allows their employees to focus on higher-value tasks while still allowing their employees to explain how those AI decisions were made.

Enterprise financial planning

By investing in an AI-powered financial analytics solution, organizations can develop next-generation CFOs who can apply predictive analytics and unlock the value previously hidden in their data. This allows organizations to break free of manual, error-prone work.

These examples demonstrate the impact AI can have as companies continue to find ways to gain more value from their data and transform business.

Take, for instance, a large, independent energy company that has been a global leader in oil and gas for more than half a century. What's its secret? It hires and develops some of the world's foremost engineering experts. This formula has helped the company build some of the largest structures on the planet, in some of the most remote parts of the ocean, and safely transport the energy it produces to people around the globe to ensure that the next generation could successfully carry the torch.

This energy company knew that it had to harness the institutional knowledge of its best employees, and by harnessing the power of AI, the company was able to extract meaningful insights from 30 years of dense and complex engineering data and put decades of knowledge at the fingertips of employees across the company. This helped the organization answer tough questions faster to enable fact-driven decision making on complex projects.

Now, this is just an example of how a company can infuse AI into its core business. However, this same approach can be replicated across any industry because it begins with putting AI to work. It's about automating and optimizing processes in whole new ways and applying AI to unlock new value for its business, regardless of size or industry.

Conclusion

It's clear that in today's hybrid multicloud world, for organizations to succeed with AI, it's imperative that they modernize the architecture in which information is ingested, stored, organized, accessed, and analyzed.

As organizations seek to adopt AI, it's important to remember three things. First, to *start with your problem*. Whether you're just starting on your AI journey or you're well underway, always come back to the core business problem that you're trying to solve. In a lot of cases, it begins with your needs and pain points. Think about how AI can help build exceptional, personalized customer experiences.

Second, remember that *you can't have AI without IA*. Organizations need a modern information architecture to connect data from all necessary sources, to make it accessible to users across teams, to dynamically build and deploy AI models, and to simplify and unify data and AI services across cloud environments. The AI Ladder was developed to help organizations build an information architecture and ultimately reach their AI goals.

Finally, *AI is not magic*. It's hard work. It requires the proper tools, methodologies, and mindset, to overcome the gaps that companies are facing (data, skills, and trust) to truly embrace an AI practice and put it to work across your organization.

AI is the biggest opportunity of our time, and yet there's still a certain fear in the market that AI is going to replace jobs. However, the

reality is this: AI is not going to replace managers. Rather, the managers who use AI will replace the managers who do not.

About the Author

Robert D. Thomas is the general manager of IBM Data and Artificial Intelligence. He directs IBM's product investment strategy, sales and marketing, expert labs, and global software product development.

Major product brands under Rob's leadership include Watson, Db2, Netezza, Cognos, SPSS, and InfoSphere. Since joining IBM's software unit, Mr. Thomas has held roles of increasing responsibility including business development, product engineering, sales and marketing, and general management. He has overseen four acquisitions by the firm representing more than \$2.5 billion in transaction value.

Mr. Thomas trained in economics at Vanderbilt University, earning his BA in Economics. During his MBA from the University of Florida, Mr. Thomas worked in equity research, learning applied economics, finance, and financial analysis.

Mr. Thomas has published two books: *Big Data Revolution: What Farmers, Doctors, and Insurance Agents Can Teach Us About Patterns in Big Data* (John Wiley & Sons, 2015), and *The End of Tech Companies* (self-published, 2016), educating business leaders on how to navigate digital disruption in every industry. Today, he writes extensively on his blog robdtomas.com.