

Climbing the AI Ladder

“The first characteristic of interest is the fraction of the computational load, which is associated with data management housekeeping.”

—Gene Amdahl

“Approach to Achieving Large Scale Computing Capabilities”

To remain competitive, enterprises in every industry need to use advanced analytics to draw insights from their data. The urgency of this need is an accelerating imperative. Even public-sector and nonprofit organizations, which traditionally are less motivated by competition, believe that the rewards derived from the use of artificial intelligence (AI) are too attractive to ignore. Diagnostic analytics, predictive analytics, prescriptive analytics, machine learning, deep learning, and AI complement the use of traditional descriptive analytics and business intelligence (BI) to identify opportunities or to increase effectiveness.

Traditionally an organization used analytics to explain the past. Today analytics are harnessed to help explain the immediate now (the present) and the future for the opportunities and threats that await or are impending. These insights can enable the organization to become more proficient, efficient, and resilient.

However, successfully integrating advanced analytics is not turnkey, nor is it a binary state, where a company either does or doesn't possess AI readiness. Rather, it's a journey. As part of its own recent transformation, IBM developed a visual metaphor to explain a journey toward readiness that can be adopted and applied by any company: the AI Ladder.

As a ladder, the journey to AI can be thought of as a series of rungs to climb. Any attempt to zoom up the ladder in one hop will lead to failure. Only when each rung is firmly in hand can your organization move on to the next rung. The climb is not hapless or random, and climbers can reach the top only by

approaching each rung with purpose and a clear-eyed understanding of what each rung represents for their business.

You don't need a crystal ball to know that your organization needs data science, but you do need some means of insight to know your organization's efforts can be effective and are moving toward the goal of AI-centricity. This chapter touches on the major concepts behind each rung of the metaphorical ladder for AI, why data must be addressed as a peer discipline to AI, and why you'll need to be creative as well as a polymath—showcasing your proficiency to incorporate multiple specializations that you'll be able to read about within this book.

Readying Data for AI

The limitations can be technological, but much of the journey to AI is made up of organizational change. The adoption of AI may require the creation of a new workforce category: the *new-collar* worker. New-collar jobs can include roles in cybersecurity, cloud computing, digital design, and cognitive business. New-collar work for the cognitive business has been invoked to describe the radically different ways AI-empowered employees will approach their duties. This worker must progress methodically from observing the results of a previous action to justifying a new course of action to suggesting and ultimately prescribing a course of action.

When an organization targets a future state for itself, the future state simply becomes the current state once it's attained. The continual need to define another future state is a cycle that propels the organization forward. Ideally, the organization can, over time, reduce the time and expense required to move from one state to the next, and these costs will be viewed not as expenses but as derived value, and money will cease to inhibit the cycle's progression.

Worldwide, most organizations now agree that AI will help them stay competitive, but many organizations can often still struggle with less advanced forms of analytics. For organizations that experience failure or less than optimal outcomes with AI, the natural recourse seems to be to remove rigor and not increase it. From the perspective of the AI Ladder, rungs are hurried or simply skipped altogether. When an organization begins to recognize and acknowledge this paradigm, they must revisit the fundamentals of analytics in order to prepare themselves for their desired future state and the ability to benefit from AI. They don't necessarily need to start from scratch, but they need to evaluate their capabilities to determine from which rung they can begin. Many of the technological pieces they need may already be in place.

Organizations will struggle to realize value from AI without first making data simple, accessible, and available across the enterprise, but this democratization

of data must be tempered by methods to ensure security and privacy because, within the organization, not all data can be considered equal.

Technology Focus Areas

Illustrated in Figure 1-1, the level of analytics sophistication accessible to the organization increases with each rung. This sophistication can lead to a thriving data management practice that benefits from machine learning and the momentum of AI.

Organizations that possess large amounts of data will, at some point, need to explore a multicloud deployment. They'll need to consider three technology-based areas as they move up the ladder.

- **Hybrid data management** for the core of their machine learning
- **Governance and integration** to provide security and seamless user access within a secured user profile
- **Data science and AI** to support self-service and full-service user environments for both advanced and traditional analytics

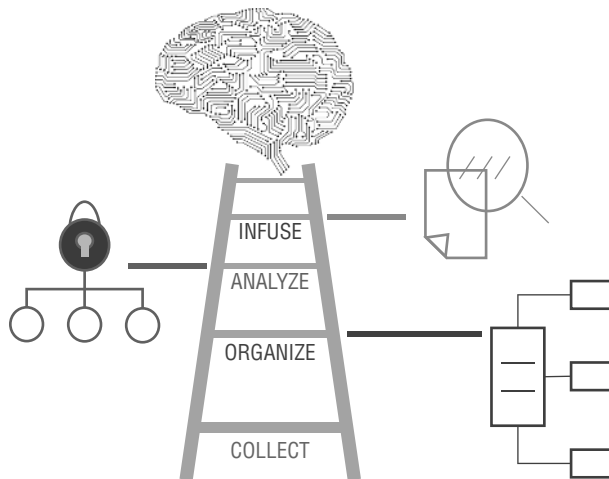


Figure 1-1: The AI Ladder to achieve a full complement of data and analytics

These foundational technologies must embrace modern cloud and microservice infrastructures to create pathways for the organization to move forward and upward with agility and speed. These technologies must be implemented at various rungs, enabling the movement of data and delivering predictive power from machine learning models in various types of deployment, from a single environment to a multicloud environment.

Taking the Ladder Rung by Rung

As shown in Figure 1-1, the rungs of the ladder are labeled Collect, Organize, Analyze, and Infuse. Each rung provides insight into elements that are required for an information architecture.

Collect, the first rung, represents a series of disciplines used to establish foundational data skills. Ideally, access to the data should be simplified and made available regardless of the form of the data and where it resides. Since the data used with advanced analytics and AI can be dynamic and fluid, not all data can be managed in a physical central location. With the ever-expanding number of data sources, virtualizing how data is collected is one of the critical activities that must be considered in an information architecture.

These are key themes included in the Collect rung:

- Collecting data with a common SQL engine, the use of APIs for NoSQL access, and support for data virtualization across a broad ecosystem of data that can be referred to as a *data estate*
- Deploying data warehouses, data lakes, and other analytical-based repositories with always-on resiliency and scalability
- Scaling with real-time data ingestion and advanced analytics simultaneously
- Storing or extracting all types of business data whether structured, semi-structured, or unstructured
- Optimizing collections with AI that may include graph databases, Python, machine learning SQL, and confidence-based queries
- Tapping into open source data stores that may include technologies such as MongoDB, Cloudera, PostgreSQL, Cloudant, or Parquet

The *Organize* rung infers that there is a need to create a trusted data foundation. The trusted data foundation must, at a minimum, catalog what is knowable to your organization. All forms of analytics are highly dependent upon digital assets. What assets are digitized forms the basis for what an organization can reasonably know: the corpus of the business is the basis for the organizational universe of discourse—the totality of what is knowable through digitized assets.

Having data that is business-ready for analytics is foundational to the data being business-ready for AI, but simply having access to data does not infer that the data is prepared for AI use cases. Bad data can paralyze AI and misguide any process that consumes output from an AI model. To organize, organizations must develop the disciplines to integrate, cleanse, curate, secure, catalog, and govern the full lifecycle of their data.

These are key themes included in the Organize rung:

- Cleansing, integrating, and cataloging all types of data, regardless of where the data originates
- Automating virtual data pipelines that can support and provide for self-service analytics
- Ensuring data governance and data lineage for the data, even across multiple clouds
- Deploying self-service data lakes with persona-based experiences that provide for personalization
- Gaining a 360-degree view by combing business-ready views from multicloud repositories of data
- Streamlining data privacy, data policy, and compliance controls

The *Analyze* rung incorporates essential business and planning analytics capabilities that are key for achieving sustained success with AI. The Analyze rung further encapsulates the capabilities needed to build, deploy, and manage AI models within an integrated organizational technology portfolio.

These are key themes included in the Analyze rung:

- Preparing data for use with AI models; building, running, and managing AI models within a unified experience
- Lowering the required skill levels to build an AI model with automated AI generation
- Applying predictive, prescriptive, and statistical analysis
- Allowing users to choose their own open source frameworks to develop AI models
- Continuously evolving models based upon accuracy analytics and quality controls
- Detecting bias and ensuring linear decision explanations and adhering to compliance

Infuse is a discipline involving the integration of AI into a meaningful business function. While many organizations are able to create useful AI models, they are rapidly forced to address operational challenges to achieve sustained and viable business value. The Infuse rung of the ladder highlights the disciplines that must be mastered to achieve trust and transparency in model-recommended decisions, explain decisions, detect untoward bias or ensure fairness, and provide a sufficient data trail for auditing. The Infuse rung seeks to operationalize AI use cases by addressing a time-to-value continuum.

These are key themes included in the Infuse rung:

- Improving the time to value with prebuilt AI applications for common use cases such as customer service and financial planning or bespoke AI applications for specialized use cases such as transportation logistics
- Optimizing knowledge work and business processes
- Employing AI-assisted business intelligence and data visualization
- Automating planning, budgeting, and forecasting analytics
- Customizing with industry-aligned AI-driven frameworks
- Innovating with new business models that are intelligently powered through the use of AI

Once each rung is mastered to the degree that new efforts are repeating prior patterns and that the efforts are not considered bespoke or deemed to require heroic efforts, the organization can earnestly act on its efforts toward a future state. The pinnacle of the ladder, the journey to AI, is to constantly modernize: to essentially reinvent oneself at will. The Modernize rung is simply an attained future state of being. But once reached, this state becomes the organizational current state. Upon reaching the pinnacle, dynamic organizations will begin the ladder’s journey anew. This cycle is depicted in Figures 1-2 and 1-3.

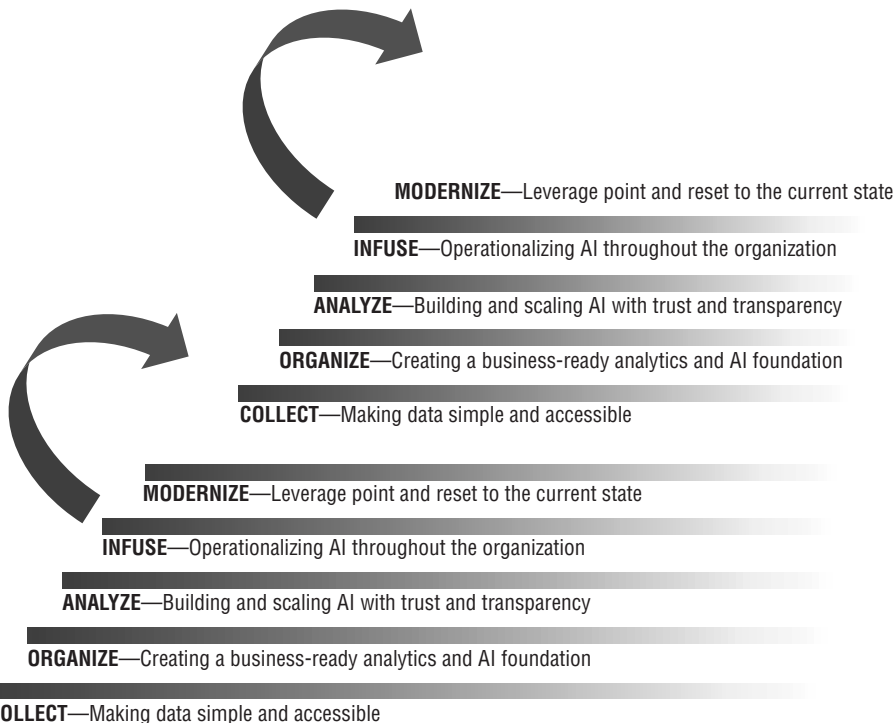


Figure 1-2: The ladder is part of a repetitive climb to continual improvement and adaptation.




Figure 1-3: Current state ⇌ future state ⇌ current state

These are key themes included in the Modernize rung:

- Deploying a multicloud information architecture for AI
- Leveraging a uniform platform of choice across any private or public cloud
- Virtualizing data as a means of collecting data regardless of where the data is sourced
- Using DataOps and MLOps to establish trusted virtual data pipelines for self-service
- Using unified data and AI cloud services that are open and easily extensible
- Scaling dynamically and in real time to accommodate changing needs

Modernize refers to an ability to upgrade or update or, more specifically, to include net-new business capabilities or offerings resulting from transformational ideas or innovation that harness reimaged business models. The infrastructural underpinnings for organizations that are modernizing are likely to include elastic environments that embrace a multicloud topology. Given the dynamic nature of AI, modernizing an organization means building a flexible information architecture to constantly demonstrate relevance.

THE BIG PICTURE

In agile development, an *epic* is used to describe a user story that is considered far too big to be addressed in a single iteration or a single sprint. Therefore, an epic is used to provide the “big picture.” The big picture provides an end-to-end perspective for what needs to be accomplished. The epic can then be decomposed into a series of workable stories that can be worked on. The epic serves to ensure the stories are threaded appropriately.

In the AI Ladder, the ladder represents the “big picture.” The decomposition is represented by rungs. The ladder is used to ensure that the concepts for each rung—*collect, organize, analyze, infuse*—are threaded appropriately to ensure the best chance to succeed and deliver value.

Constantly Adapt to Retain Organizational Relevance

As the famous adage goes, “If you dislike change, you’re going to dislike irrelevance even more” (www.network54.com/Forum/220604/thread/1073929871). In a world of rapid change, the endgame is to remain relevant. *Relevance* is an *ends*. Accordingly, transformation and disruption can be viewed as a *means*. Knowledge of the means provides the how (a course of action and the directives). Knowing the ends gives you a target (the desired result) for which to aim, as shown in Figure 1-4.




Figure 1-4: Ends and means model

The ends and means model can be iterated upon to adjust to continuous change and continuous course correction to drive toward sustained and improved relevance.

If an organization can combine relevance with uniqueness, then the path forward may offer greater opportunities, as the organization is unlikely to be viewed as another commodity player in a given space.

Formulating the ends is potentially more difficult than formulating the means. Words that are substitutes for *means* include *transformation* (including digital transformation) and *disruption*. An organization wanting to transform might simply do so by focusing on select technologies or capabilities that are newer; for example, cloud computing, agile development, big data, consumer-oriented platforms, blockchain, analytics, and even AI. Regardless of which technologies and capabilities are picked as a means, the question would remain: but, to what ends?

Relevance can be more difficult to articulate in terms of an *ends*, especially in the light of newer technologies and capabilities. The struggle with articulation might result from subject-matter experts and end users having minimal insight or experience with how a new solution can actually be positioned and leveraged.

Disruption can be viewed through its two adjectives: disruptive and disrupted. A disruptive organization is more likely to garner a leadership position in handling change, while a disrupted organization is likely to have a lagging position. For example, traditional publishing companies and journalists were found flat-footed in the wake of the digital publishing era. For more than a

decade, profits and viability have dwindled in traditional publishing due to social media, blogs, and the proliferation of web-based content.

Another example of a business type in the midst of being disrupted is the hospital. In the United States, the population size has grown by approximately 100 million people in the past 35 years and by nearly 200 million people in the past 70 years. It would not be unreasonable to assume that the demand for hospital beds has steadily risen alongside population growth, especially with an aging population. However, the demand for hospital beds is approximately the same now as it was in the 1940s. In the United States, the need for hospital beds peaked in the early 1980s.

Not all that long ago, a hospital was not exactly the safest of places to go and get well. In one account from the 1800s, it was noted that “hospitals are the sinks of human life” (www.archive.org/stream/proceedingsconn08socioog). Through the use of anesthesia and the adoption of sterilization techniques, along with the advent of X-rays in 1895, hospitals turned a corner from being a highly risky place to get treated. But now, in a throwback to the 18th and 19th centuries, hospitals are once again seen as a less-than-desirable place to receive therapeutic medical treatment, as cited in a CDC report on hospital-acquired infections (www.documentcloud.org/documents/701516-cdc-hai-infections-deaths.html).

Facilities currently challenging the traditional hospital include walk-in urgent care centers, imaging facilities, surgical centers, physician offices, and so on. Hospitals are being forced to consider mergers and acquisitions as well as downsizing. Hospitals are being disrupted and need to seek nontraditional ways to remain relevant. Could the use of advanced analytics be part of the approach?

Ultimately, if AI is going to augment human intelligence, AI will be part of the *means* to transform or disrupt. While AI can potentially hypothesize about what can be relevant, AI is likely going to be challenged to convey a de facto direction as what needs to be done to remain relevant. For AI and for humans, collaboration is an augmented opportunity to address a defining issue of our times: relevance.

Therefore, the sole purpose of an information architecture for AI can be postulated, as an aid in the transformation and disruption of an organization that is on a ladder to achieve sustained or regained relevance, whereby each point of leverage is based on data and the organization is willing and capable of harnessing the insights that can be derived from that data.

ECONOMICALLY VIABLE

Without data there's no AI. Period. AI works well because organizations now have the means to economically collect and hoard immense quantities of digital information. Augmenting our work with machines and AI is the norm.

For example, when was the last time you went a day without *Googling*? Chances are, you don't give machine learning a second thought when you're interacting with

AI-infused capabilities. We are augmenting our work with AI, and in large part, it feels totally natural.

Organizations still have a long way to go to fully realize how they can augment all of their processes with AI. For companies that are successful, the result should feel natural.

Data-Based Reasoning Is Part and Parcel in the Modern Business

Advanced analytics, including AI, can provide a basis for establishing reasoning by using inductive and deductive techniques. Being able to interpret user interactions as a series of signals can allow a system to offer content that is appropriate for the user's context in real time.

To maximize the usefulness of the content, the data should be of an appropriate level of quality, appropriately structured or tagged, and, as appropriate, correlated with information from disparate systems and processes. Ascertaining a user's context is also an analytical task and involves the system trying to understand the relationship between the user and the user's specific work task.

For an industrial-based business application, a user might have a need to uncover parts and tools that are required to complete maintenance on a hydraulic system. By using adaptive pattern-recognition software to help mine a reference manual about hydraulic systems and their repair, a system could derive a list of requisite tools and related parts. An advanced analytic search on hydraulic repair could present content that is dynamically generated and based on product relationships and correlated with any relevant company offerings.

Pulling content and understanding context is not arbitrary or random. Aligning and harmonizing data across an enterprise or ecosystem from various front-end, mid-end, and back-end systems takes planning, and one of the results of that planning is an information architecture.

Advances in computer processing power and the willingness for organizations to scale up their environments has significantly contributed to capabilities such as AI to be seen as both essential and viable. The ability to harness improved horsepower (e.g., faster computer chips) has made autonomous vehicles technologically feasible even with the required volume of real-time data. Speech recognition has become reliable and is able to differentiate between speakers, all without extensive speaker-dependent training sessions.

There is no hiding that AI can be a complex subject. However, much of the complexity associated with AI can be hidden from a user. While AI itself is not a black art, AI benefits when traditional IT activities such as data quality and data governance are retained and mastered. In fact, clean, well-organized, and managed data—whether the data is structured, semistructured, or unstructured—

is a basic necessity for being able to use data for input into machine learning algorithms.

There will be many situations when an AI system needs to process or analyze a corpus of data with far less structure than the type of organized data typically found in a financial or transactional system. Fortunately, learning algorithms can be used to extract meaning from ambiguous queries and seek to make sense of unstructured data inputs.

Learning and reasoning go hand in hand, and the number of learning techniques can become quite extensive. The following is a list of some learning techniques that may be leveraged when using machine learning and data science:

- Active learning
- Deductive inference
- Ensemble learning
- Inductive learning
- Multi-instance learning
- Multitask learning
- Online learning
- Reinforcement learning
- Self-supervised learning
- Semi-supervised learning
- Supervised learning
- Transduction
- Transfer learning
- Unsupervised learning

Some learning types are more complex than others. Supervised learning, for example, is comprised of many different types of algorithms, and transfer learning can be leveraged to accelerate solving other problems. All model learning for data science necessitates that your information architecture can cater to the needs of training models. Additionally, the information architecture must provide you with a means to reason through a series of hypotheses to determine an appropriate model or ensemble for use either standalone or infused into an application.

Models are frequently divided along the lines of supervised (passive learning) and unsupervised (active learning). The division can become less clear with the inclusion of hybrid learning techniques such as semisupervised, self-supervised, and multi-instance learning models. In addition to supervised learning and unsupervised learning, reinforcement learning models represent a third primary learning method that you can explore.

Supervised learning algorithms are referred to as such because the algorithms learn by making predictions that are based on your input training data against an expected target output that was included in your training dataset. Examples of supervised machine learning models include decision trees and vector machines.

Two specific techniques used with supervised learning include classification and regression.

- *Classification* is used for predicting a class label that is computed from attribute values.
- *Regression* is used to predict a numerical label, and the model is trained to predict a label for a new observation.

An unsupervised learning model operates on input data without any specified output or target variables. As such, unsupervised learning does not use a teacher to help correct the model. Two problems often encountered with unsupervised learning include clustering and density estimation. Clustering attempts to find groups in the data, and density estimation helps to summarize the distribution of data.

K-means is one type of clustering algorithm, where data is associated to a cluster based on a means. Kernel density estimation is a density estimation algorithm that uses small groups of closely related data to estimate a distribution.

In the book *Artificial Intelligence: A Modern Approach, 3rd edition* (Pearson Education India, 2015), Stuart Russell and Peter Norvig described an ability for an unsupervised model to learn patterns by using the input without any explicit feedback.

The most common unsupervised learning task is clustering: detecting potentially useful clusters of input examples. For example, a taxi agent might gradually develop a concept of “good traffic days” and “bad traffic days” without ever being given labeled examples of each by a teacher.

Reinforcement learning uses feedback as an aid in determining what to do next. In the example of the taxi ride, receiving or not receiving a tip along with the fare at the completion of a ride serves to imply goodness or badness.

The main statistical inference techniques for model learning are inductive learning, deductive inference, and transduction. Inductive learning is a common machine learning model that uses evidence to help determine an outcome. Deductive inference reasons top-down and requires that each premise is met before determining the conclusion. In contrast, induction is a bottom-up type of reasoning and uses data as evidence for an outcome. Transduction is used to refer to predicting specific examples given specific examples from a domain.

Other learning techniques include multitask learning, active learning, online learning, transfer learning, and ensemble learning. Multitask learning

aims “to leverage useful information contained in multiple related tasks to help improve the generalization performance of all the tasks” (arxiv.org/pdf/1707.08114.pdf). With active learning, the learning process aims “to ease the data collection process by automatically deciding which instances an annotator should label to train an algorithm as quickly and effectively as possible” (papers.nips.cc/paper/7010-learning-active-learning-from-data.pdf). Online learning “is helpful when the data may be changing rapidly over time. It is also useful for applications that involve a large collection of data that is constantly growing, even if changes are gradual” (Stuart Russell and Peter Norvig, *Artificial Intelligence: A Modern Approach*, 3rd edition, Pearson Education India, 2015).

LEARNING

The variety of opportunities to apply machine learning is extensive. The sheer variety gives credence as to why so many different modes of learning are necessary:

- Advertisement serving
- Business analytics
- Call centers
- Computer vision
- Companionship
- Creating prose
- Cybersecurity
- Ecommerce
- Education
- Finance, algorithmic trading
- Finance, asset allocation
- First responder rescue operations
- Fraud detection
- Law
- Housekeeping
- Elderly care
- Manufacturing
- Mathematical theorems
- Medicine/surgery
- Military
- Music composition
- National security

- Natural language understanding
- Personalization
- Policing
- Political
- Recommendation engines
- Robotics, consumer
- Robotics, industry
- Robotics, military
- Robotics, outer space
- Route planning
- Scientific discovery
- Search
- Smart homes
- Speech recognition
- Translation
- Unmanned aerial vehicles (drones, cars, ambulance, trains, ships, submarines, planes, etc.)
- Virtual assistants

Evaluating how well a model learned can follow a five-point rubric.

- **Phenomenal:** It's not possible to do any better.
- **Crazy good:** Outcomes are better than what any individual could achieve.
- **Super-human:** Outcomes are better than what most people could achieve.
- **Par-human:** Outcomes are comparable to what most people could achieve.
- **Sub-human:** Outcomes are less than what most people could achieve.

Toward the AI-Centric Organization

As with the industrial age and then the information age, the age of AI is an advancement in tooling to help solve or address business problems. Driven by necessity, organizations are going to use AI to aid with automation and optimization. To support data-driven cultures, AI must also be used to predict and to diagnose. AI-centric organizations must revisit all aspects of their being, from strategy to structure and from technology to egos.

Before becoming AI-centric, organizations must first identify their problems, examine their priorities, and decide where to begin. While AI is best for detecting outcomes against a pattern, traditional business rules are not going to disappear. To be AI-centric is to understand what aspects of the business can best be addressed through patterns. Knowing how much tax to pay is never going to be a pattern; a tax calculation is always going to be rule-based.

There are always going to be situations where a decision or action requires a combination of pattern-based and rule-based outcomes. In much the same way, a person may leverage AI algorithms in conjunction with other analytical techniques.

Organizations that avoid or delay AI adoption will, in a worst-case scenario, become obsolete. The changing needs of an organization coupled with the use of AI are going to necessitate an evolution in jobs and skillsets needed. As previously stated, every single job is likely to be impacted in one way or another. Structural changes across industries will lead to new-collar workers spending more of their time on activities regarded as driving higher value.

Employees are likely to demand continuous skill development to remain competitive and relevant. As with any technological shift, AI may, for many years, be subject to scrutiny and debate. Concerns about widening economic divides, personal privacy, and ethical use are not always unfounded, but the potential for consistently providing a positive experience cannot be dismissed. Using a suitable information architecture for AI is likely to be regarded as a high-order imperative for consistently producing superior outcomes.

SCALE

On occasion, we are likely to have experienced a *gut feeling* about a situation. We have this sensation in the pit of our stomach that we know what we must do next or that something is right or that something is about to go awry. Inevitably, this feeling is not backed by data.

Gene Kranz was the flight director in NASA's Mission Control room during the Apollo 13 mission in 1970. As flight director, he made a number of *gut feel* decisions that allowed the lunar module to return safely to Earth after a significant malfunction. This is why we regard AI as augmenting the knowledge worker and not an outright replacement for the knowledge worker. Some decisions require a broader context for decision-making; even if that decision is a gut feel, the decision is still likely to manifest from years of practical experience.

For many businesses, the sheer scale of their operations already means that each decision can't be debated between man and machine to reach a final outcome. Scale, and not the need to find a replacement for repetitive tasks, is the primary driving factor toward needing to build the AI-centric organization.

Summary

Through climbing the ladder, organizations will develop practices for data science and be able to harness machine learning and deep learning as part of their enhanced analytical toolkit.

Data science is a *discipline*, in that the data scientist must be able to leverage and coordinate multiple skills to achieve an outcome, such as domain expertise, a deep understanding of data management, math skills, and programming. Machine learning and deep learning, on the other hand, are *techniques* that can be applied via the discipline. They are techniques insofar as they are optional tools within the data science toolkit.

AI puts machine learning and deep learning into practice, and the resulting models can help organizations reason about AI's hypotheses and apply AI's findings. To embed AI in an organization, a formal data and analytics foundation must be recognized as a prerequisite.

By climbing the ladder (from one rung to the next: collect, organize, analyze, and infuse), organizations are afforded with the ability to address questions that were either previously unknown (When will a repeat buyer buy again?) or previously unanswerable (What were the influencing factors as to why a given product was purchased?).

When users can ask new questions, users can benefit from new insights. Insights are therefore a direct means to empowerment. Empowered users are likely to let specific queries execute for multiple minutes, and, in some cases, even hours, when immediate near-zero-second response is not fully required. The allure of the ladder and to achieve AI through a programmatic stepwise progression is the ability to ask more profound and higher-value questions.

The reward for the climb is to firmly establish a formal organizational discipline in the use of AI that is serving to help the modern organization remain relevant and competitive.

In the next chapter, we will build on the AI Ladder by examining considerations that impact the organization as a whole.