

## Chapter 1

---

# Artificial Intelligence

**W**e consider artificial intelligence (AI) to be a “big wave” phenomenon due to its potential universal application in today’s organizations. The rapid evolution of AI in terms of machine learning, natural language, and neural networks combined with massive processing power introduce capabilities to potentially drive all aspects of commerce. Whether automating internal processes, picking stocks, predicting (perhaps manipulating) voting behavior, or routing traffic, the application of AI is limited only by our ability to imagine the future. As the analogy intends, just like surfers, organizations must position themselves on this wave of opportunity or risk missing an epic ride enjoyed by others.

Despite its long scientific history, AI has only recently become a hot topic for most organizations (Austin 2017; Austin et al. 2016; Andrews et al. 2016). Although AI has been

*theoretically* possible for some years, it is due to the convergence of three conditions that it is now making it *practically* possible for its commercial use. First, powerful hardware is now accessible through cloud computing and super computers such as Watson (Austin 2017). Second, the huge amounts of data that drive AI are now available through apps, sensors, and other devices (Rometty 2016). Third, new software such as powerful algorithms, natural language, and machine learning have matured to a point where they can do certain jobs as well as or better than humans (Austin 2017; Rometty 2016).

Excitement about AI is growing in the business world (Andrews et al. 2016) and it is also beginning to creep into our daily lives. Whether it is autonomous driving, medical diagnoses, or stock picking, it is clear that AI is inexorably moving from the realm of science fiction to reality. But, as with any new technology, AI raises some important questions for how it is to be used in today's organizations. "We're really at the very early stages of this technology," said a manager. "But we know it's going to be incredibly transformational. We're just not sure exactly how."

It is widely agreed that AI will be disruptive in a number of ways. First, it is a platform shift—away from the traditional structured ways of doing business towards a more "natural" interaction that adapts to context and conditions (Austin et al. 2016). Second, it requires new ways of managing technology. It "learns" and is "trained"; it isn't programmed. It finds things we didn't know were there and arrives at outcomes we don't anticipate. It will need to be controlled, managed, and

monitored in ways we aren't ready for at present (Andrews et al. 2016). Third, it replaces a considerable component of many knowledge workers' jobs and this will have repercussions for both organizations and society (Klotz 2016). And finally, when organizations cede some or all decision-making in an area to AI, they are going to have to address a number of legal and ethical issues related to privacy, accountability, governance, responsibility, and decision-making transparency, as well as a host of social policy considerations (Bernstein 2016).

In this chapter we explore some of these challenges and IT's role in selecting, designing, implementing, and managing AI. We start by reviewing our current understanding of AI, its strengths and weaknesses, and where it fits with other types of technology in use in organizations. Next we look at how organizations are using AI and where they see it will be valuable to them. Following this, we examine a number of dimensions of AI management that will be important during its introductory phase in organizations. Finally, we outline several recommendations for managers who seek to “get AI right.”

## What Is AI?

---

AI is the broad term that refers to many types of “smart” machines that enable organizations to tackle more complex problems than traditional, structured systems currently allow. It encompasses a number of subtypes, including:

- **Robotics**—where machines do physical work, which can range from manufacturing, to supply delivery in a building, to complex laboratory processing such as blood tests.
- **Machine learning**—where algorithms that improve automatically through experience are used for processing and decision-making (Quora 2017).
- **Neural networks**—inspired by biological neural networks and which aim to mimic brain connections. Machine-based neural networks have developed strong pattern identification skills and can also be a type of machine learning—discovering rules, developing new rules, and tolerating noise and variability in data.
- **Natural language processing**—which includes a variety of technologies that facilitate conversational interfaces between humans and a machine. The most well known is IBM’s Watson that used these processes to win *Jeopardy* in 2010. Natural language processing is being introduced to the public through digital “assistants” like Siri and Cortana, and a variety of chatbots.

AI works differently from other forms of technology primarily because it doesn’t function in the binary fashion (“if-then-else”) used by modern computer programs and doesn’t require fixed commands (Andrews et al. 2016; Austin et al. 2016). As a result, it can be used to *augment* human intelligence to make better decisions, analyze massive

amounts of data, identify anomalies, and proactively predict events (Moore 2016; Rometty 2016; Hoffman 2016).

Data is the fuel of AI and one of AI's key values is helping us deal with the massive amounts of data currently being created—both structured and unstructured. “In many ways we live in an era of cognitive overload characterized by an exponential increase in the complexity of decision-making” writes one expert. “It’s impossible to create protocols, algorithms, or software code to successfully anticipate all potential permutations, trajectories, and interactions” (Rometty 2016). The focus group agreed. “The driver in IT for AI is our data lakes and how we use them to provide value,” said a manager. “As we improve with our big data and analytics initiatives, it is exposing more and more opportunities for how to use AI.”

Although AI makes machines “smarter,” the focus group stressed that today’s AI applications are narrowly focused on a single function, such as image recognition, pattern identification, or a particular task. While they can be spectacularly good at these types of tasks, they still lack general executive functions and, because they are trained by humans, can incorporate human biases into their actions (Austin 2017). For example, Facebook’s nudity recognition engine ran into problems when it banned pictures of breastfeeding moms and Michelangelo’s statue of David.

As one expert points out,

Things that are so hard for people, like playing championship-level Go and poker, have turned out to be

relatively easy for the machines. . . . Yet at the same time, the things that are easiest for a person—like making sense of what they see in front of them, speaking in their mother tongue, the machines really struggle with. . . . General intelligence is what people do . . . we don't have a computer that can function with the capabilities of a six year old or even a three year old, and so we're very far from general intelligence. (Higginbotham 2015)

However, AI is evolving rapidly, the focus group noted. "It's all about the human-machine interface really," said one manager. "This line is moving. Twenty years ago, cheque recognition was cutting-edge AI. Today, we just take it for granted." The group also noted that the line between AI and other forms of technology is unclear. "At what point do we say it's AI?," one asked. He noted that self-driving cars use a variety of "AI-like" technology, such as image recognition to identify specific patterns (e.g., a stop sign), a rules engine for things that rarely change (e.g., what to do at a four-way stop sign), and machine learning that uses probabilities and judgment to determine if you should stop at the sign, and for how long, in order to avoid an accident. "Right now, if it's stuff we can't do, we call it AI," another concluded.

## What Are Organizations Doing about AI?

---

Although of high interest to many organizations, AI has not yet made the leap into organizational practice, except in very experimental ways. Many members of the focus group were tinkering or testing AI applications, such as chatbots, but as one noted, “we haven’t put them in charge of anything yet.” “Right now, it’s all fuzzy and experimental,” said another. “We’re testing right now so that we can learn and make sure it works,” said a third.

Focus group members see AI as being part of a continuum that begins with big data, predictive analytics, and business intelligence and expands from there. “As we improve with these, AI will be part of a natural progression,” said a manager. “But we believe existing data technologies can still improve our business before we need AI.” Another noted, “We’re still struggling with how to connect AI to our business.” Nevertheless, exploration and experimentation are important to help both business and IT find the right role for AI and the right approach to investing in it. “One of the challenges is that AI is both a technology and a solution,” said a manager. “And it flips between the two. We have to learn how to manage it effectively.”

Another challenge for organizations is that “most valuable AI platforms are built on narrow, proprietary platforms, while most broad, general-purpose AI platforms lack ready-made

valuable AI applications and require buyers in every enterprise to fund the redevelopment of new applications” (Austin et al. 2016). For this reason, most enterprises are not yet exploiting AI. “[They] want and need lower cost, lower risk, faster to deploy and easier to manage solutions that are built on a common technical infrastructure” (Austin et al. 2016). As a result, “The zeal of the possible is tempered by the practical,” said a member. The focus group agreed that AI is still so expensive most organizations can’t tackle it on their own. Therefore, they are relying on partners and cloud applications to acquire the solutions they are experimenting with.

Members also pointed out that, in these early days of AI, it is difficult to see the full shape of how it might be applied in the future. “Our implementations will be very different ten years from now,” said a manager. “At this stage we’re just using it to help us build faster processes,” said another, “much like the car was initially envisioned to be a faster horse.” As a result, many of the common applications of AI at present are in areas where massive amounts of data are required, such as with legal or medical information. In the focus group, two companies are using AI in this way to monitor their security logs for abnormalities. However, “if it recognizes a pattern that doesn’t fit, it calls a grown-up (i.e., a human),” said a manager.

Another company is experimenting with using AI for basic underwriting using simple rule sets. “We’re just taking baby steps right now,” said a manager. “We need to get our data in shape, develop use cases, train the system up, and then



reinforce it with more data.” And one firm is exploring natural language voice response in its call center but again: “We need more data before it’s successful.” The focus group stressed that AI still works on the garbage-in, garbage-out rule so improving data and monitoring outcomes are paramount in their AI work. “We need to make sure it works because how would you know it’s broken?,” explained a manager.

Another organization is taking a different approach, working with IBM’s Watson computer to see if it can identify data that might be relevant to its products and services and how it would apply. “We hope to use it to predict new products,” the manager said. “We want to use it to differentiate insights from noise.” Companies are also learning how *not* to use AI. “We have been looking at robo-investing,” said a manager. “But we’ve learned so far that people don’t really want robots making their decisions. We’re therefore re-vectoring this technology to stress that it provides *advice* which someone can take or not.”

Although companies are experimenting widely with AI for a variety of reasons, members believe that the primary driver is business cost reduction—involving people and processes. “In the long run, we can expect to see a decay of jobs over time as AI is used to replace not only factory workers but also knowledge workers,” said a manager. And the cadence of change is speeding up, they said. As they look into the not-so-distant future, they see a world where low-level knowledge workers will be largely replaced by AI, where decisions of all

types are made and/or supported by AI, and where data and data scientists will be kings.

## Dimensions of AI Management

---

The focus group identified a number of areas related to AI management that their companies are beginning to explore. Members stressed that AI is not a project but represents a fundamental change in how work is done. As a result, the issues outlined below are merely some of the initial challenges that IT and business management should consider when preparing for AI. As these are addressed and as AI evolves, new dimensions of AI will undoubtedly arise.

- ***Digital transformation.*** Most of the focus group companies are well down this path, which involves embracing new and different technologies in ways that challenge operational and value assumptions and integrate them with existing technologies to deliver new products, services, business models, revenue streams and/or customer/stakeholder experiences. What many may not have grasped yet is that “digital is not the destination. Rather it is laying the foundations for a much more profound transformation to come. Within five years . . . all major business decisions will be enhanced by cognitive technologies” (Rometty 2016).

- **Data management.** There is no question that AI requires more and better data to make it effective.

Data is the lifeblood of AI. To train computers to learn . . . you have to feed them tens of thousands of examples of something. The computers try to understand what elements of those examples define what makes a cat a cat in an image or what gives meaning to a certain word. The algorithm then gives a statistical weight to each guess that helps the computer “learn” what the right answer is. The computer scientist helps train the algorithm by giving feedback and more examples along the way. (Higginbotham 2015)

“AI is forcing better data management,” said a manager. One writer suggests: “Data [is] the world’s great new resource. What steam power, electricity, and fossil fuels did for earlier eras, data promises to do for the 21st century—if we can mine, refine and apply it” (Rometty 2016). However to do this, companies need to build strong data functions and address the perennial problems of data ownership, privacy, security, and data classification, as well as the newer areas of big data and its management.

- **Business value.** As with other new types of technology, organizations want to understand how to use AI to deliver value. This value could come from replacing or augmenting human labor or from new products and services that have yet to be conceived. At present, organizations are

mostly exploring how other firms are using AI in applications or in specific industries and identifying areas in their own firms that could benefit from AI use (Andrews et al. 2017). Supporting human decision-making is a primary area of interest (Andrews et al. 2016). The focus group added that working with big data and business rules exposes potential opportunities for AI. “We should attack these opportunities as they arise but also allow for synergies to develop and serendipitous sources of value to surface,” said a manager. Others pointed out that to truly deliver business value, business practices will likely have to evolve as well. Although it will be IT that will make AI happen, the members agreed, AI implementation must be an enterprise-wide initiative with CEO sponsorship and funding. “AI is going to happen quickly and everyone must be together on this,” said a member.

- ***Skills development.*** There is general agreement that we really don’t know much about the specific skills that will be required to work with AI (Austin et al. 2016; Bernstein 2016). What we do know is that they are scarce. The focus group believes that data, analytics, data mapping, quality control, and quantitative measurement will be key emergent skills. “We will also need algorithm and modeling skills,” added a manager. Methods for working with AI do not yet exist and this inhibits organizations’ understanding of what skills to look for (Austin et al. 2016). At present, the best advice is to seek broad

problem-solving skills and the ability to work on fluid teams (Bernstein 2016). In the shorter term, Gartner predicts that by 2019 more than 10 percent of IT hires in customer service will be writing bot scripts (Andrews et al. 2016). The focus group was optimistic about their organizations' abilities to acquire AI skills. "We have some staff already who have AI training and who are eager to use it," said a member. "And if we ask, our people will want to learn new skills." Members stressed the importance of having both business and data expertise in the future, as well as technical skills. "Our skill sets will change and entry level skills will be replaced with automation, but we will train at a more senior level so humans will still have 'skin in the game,'" one concluded.

- ***IT's role.*** IT's role in the organization will change because it is likely that responsibility for AI will be split between different parts of the enterprise (Andrews et al. 2016). The focus group agreed that IT's role will be primarily that of creating the right conditions for enabling AI and integrating it with existing systems. That said, IT will play a part in many dimensions of AI work including: helping to identify opportunities, clarifying the purpose of an AI implementation, cleansing and identifying data, managing and maintaining the AI environment, helping to select and implement appropriate algorithms and maintaining them, ensuring quality outcomes, building bridges between applications, and coordinating with

privacy, legal, and security groups. In addition, it is clear that current IT practices will need to evolve with the advent of AI, particularly enterprise architecture, vendor selection, software development, and business intelligence (Andrews et al. 2017). AI implementation is best suited to iterative development, noted the focus group, so it is important to develop a competence in this area. Finally, it will be essential to create a governance structure that will enable decisions about AI to be made effectively and to ensure accountability for these decisions. The focus group anticipates that, at minimum, an AI implementation will need several sets of approvals from business, IT, legal, HR, and advisory groups.

- ***Testing and audit.*** Testing is especially critical when working with AI because AI can detect unique and unanticipated patterns in data (Austin 2017). Therefore, traditional methods of testing all code paths aren't sufficient. AI outcomes must be monitored over much iteration with multiple data sets. Sometimes, with the best of intentions, AI simply yields a wrong result. It is therefore the responsibility of testing to ensure that no wayward results occur. Experts caution against the assumption that, once deployed, smart machines will need no further attention (Austin 2017). This is a fallacy that needs to be corrected. AI, for all its strengths, will need to be retrained and retested as new data are collected. Failing to consider “the challenges of continuously maintaining

and monitoring an implementation will lead to failure in many enterprises” (Austin 2017). The focus group also noted that many industries require a clear audit trail that will justify decisions made and that can be used to root out discrimination. “Whether or not we participate in decision-making, we are still responsible for those decisions,” said a member. “Therefore, where a decision is important, it’s essential to have an audit trail built into all AI algorithms.”

- ***Cognitive ergonomics.*** The focus group noted that any AI implementation must fit into an organization’s social fabric as well as address specific business opportunities. One of the biggest emerging issues in AI is therefore how this technology will interact with humans. In a world where machines and humans will collaborate on problem solving, decision-making, and customer services, organizations need to find the right way to blend humans and machines successfully (Andrews et al. 2016). Cognitive ergonomics is a new field that considers the why and how of AI implementation, taking human and social factors and design principles into account.

## Driving Opportunity with AI

---

One of the most challenging aspects of AI is that it is leading managers and society at large to consider some of the broader

impacts of technology adoption. “There are real social impacts to this technology,” said a manager, “and we need to adopt it in a way that is mindful of them.” The focus group identified five big issues that need to be tackled when adopting AI:

1. *Ask the bigger questions.* “AI raises many questions that still need to be resolved,” said a manager. Within organizations, the focus group easily listed a number of issues that are being discussed by IT leaders, including:
  - Should a bot self-identify as a robot?
  - Should AI be transparent about how it makes decisions? Is informed consent needed to use AI?
  - Who accepts decision responsibility and accountability? The organization? The algorithm supplier? The data scientist?
  - Should waivers be required in some cases? Are they ethical?
  - What are some of the “back door” implications of AI (e.g., smart TVs that leak data)?

“We don’t understand what norms for using AI will be acceptable,” said a manager, “and these will likely vary around the world. They are likely to evolve more slowly than the technology itself.”



There are even bigger questions that must also be addressed. AI and robotics are beginning to affect labor markets (Bernstein 2016), and most predictions point to increasing levels of job losses over the next ten years (Hoffman 2016). One writer notes, “The disconnect with past work models is happening a lot faster than in the past. . . . We’ll soon see enormous waves of workers put out of work and ill prepared to take on very different jobs” (Bernstein 2016). The focus group was very aware of the potential societal dangers involved in such massive economic displacement and the fact that our institutions are ill-equipped to deal with these changes. “Our welfare, unemployment, retirement systems and our universities all need to adapt,” said a manager. Another added, “We need to work to maximize ‘friendly’ AI to extend human intelligence and open new fields of employment.” “There are real social impacts to AI and we all need to work together to identify the questions that need to be asked, establish norms for its use, and reform our social and educational institutions,” a third manager concluded.

2. ***Beware of anthropomorphism.*** “Anthropomorphism is the attribution of human traits, emotions, and intentions to non-human entities and is considered to be an innate tendency of human psychology” (Wikipedia 2017). As computers become more and more human-like in their ability to interact conversationally, it is

natural to ascribe human characteristics to them. AI developers try to leverage anthropomorphism to make computers easier to use. While not necessarily unwise, experts warn that the inappropriate use of anthropomorphic metaphors creates false beliefs about the behavior of computers such as overestimating their “flexibility” (Wikipedia 2017). For example, a customer service call center with “chatbots” could lead to disaster if it is unable to address complex human needs. This is a real danger, said the focus group, when companies are under constant pressure to reduce costs.

3. ***Work to develop trust.*** It is critical that people and organizations be able to trust what technology is able to do (Austin et al. 2016). At these initial stages of AI, this trust must be constantly tested. Organizations can expect vendors to oversell their capabilities as well, leaving people skeptical of what AI can really do (Andrews et al. 2017). Furthermore, trust will vary by context. One manager noted: “The level of trust required depends on the types of decisions AI is making—less is needed when determining the best route to work—much more when there are safety implications.” In addition, trust can be misplaced or abused and this should never be forgotten. “We had total confidence in our automated airplane tracking system until Malaysia Airlines Flight 370 completely disappeared. Such occurrences reveal inappropriate assumptions that are

temporarily threatening and require a complete reassessment of how we are using technology,” stated a manager.

4. ***Build multiple work models.*** As the world changes in response to new work practices resulting from adoption of AI, organizations will have to integrate their legacy systems and practices into this fast-paced, rapidly changing environment (Bernstein 2016). No one knows what models will be effective in this new world, so the best advice is to experiment with multiple ones (e.g., crowdsourcing, distant manufacturing or transaction processing, and contract work). Experience with different work models will help develop flexibility and agility and start to modify organizational cultures (Bernstein 2016). Having an adaptive culture will give organizations much more than a one-time advantage. It could be key to their very survival. One manager quoted Charles Darwin: “It is not the strongest of the species that survives, nor the most intelligent, but the one most responsive to change.”
5. ***Consider open AI.*** Open AI is a non-profit artificial intelligence research company, supported by Tesla’s Elon Musk, which aims to carefully promote and develop friendly AI in order to benefit, rather than harm, humanity as a whole. It is also an open-source

project aimed at creating specifications for AI and associated programs and tools. Its short-term goals are to build tools and algorithms that will be shared publicly and longer term, to develop better hardware that can perform more like a human. An open-source model is a cheaper way to address AI problems, and if it works it could help advance AI for everyone (Wikipedia 2017; Higginbotham 2015).

## Conclusion

---

AI has now moved beyond the realm of science fiction and is just about ready for prime time. It is appropriate to ask ourselves important questions about how it could be used wisely or unwisely in organizations, and what needs to be done to mitigate the larger social impacts it is likely to cause. Preparing for AI is a daunting task. Thoughtful business and IT leaders must not only consider its potential value, but also its broader costs. Anticipating that economic pressure will eventually force AI adoption, organizations should seize the opportunity now to educate themselves about the different ways they can deploy it, and develop principles for its use that will take the larger social context into account. In addition, it is incumbent upon all organizations to work collaboratively with governments and researchers to ensure that the negative impacts of AI are addressed and remediated.

## References

---

- Andrews, W., F. Karamouzis, K. Brant, M. Revang, M. Reynolds, J. Hare, and D. Berman. “Predicts 2017: Artificial Intelligence.” Gartner Research Report, ID: G00317025, November 23, 2016.
- Andrews, W., D. Berman, A. Linden, and T. Austin. “Artificial Intelligence Primer for 2017.” Gartner Research Report, ID: G00318582, February 3, 2017.
- Austin, T. “Smart Machines see Major Breakthroughs after Decades of Failure.” Gartner Research Report, ID: G00291251, January 4, 2017.
- Austin, T., M. Hung, and M. Revang. “Conversational AI to Shake up your Technical and Business Worlds.” Gartner Research Report, ID: G00315689, September 30, 2016.
- Bernstein, A. “Globalization, Robots, and the Future of Work: An Interview with Jaffrey Joerres.” *Harvard Business Review*, October 2016, 74–79.
- Higginbotham, S. “Here’s Why Elon Musk and Everyone Else is Betting on AI.” *Fortune*, December 16, 2015.
- Hoffman, R. “Using Artificial Intelligence to Set Information Free.” *MIT Sloan Management Review* 58, no. 1 (Fall 2016).
- Klotz, F. “Are You Ready for Robot Colleagues?” *MIT Sloan Management Review Digital*, July 6, 2016. <http://sloanreview.mit.edu/article/are-you-ready-for-robot-colleagues/>.

Moore, A. “Predicting a Future Where the Future is Routinely Predicted.” *MIT Sloan Management Review* 58, no. 1 (Fall 2016).

Quora 2017. “Machine Learning.” <https://www.quora.com/topic/Machine-Learning> (accessed March 29, 2017).

Rometty, G. “Digital Today, Cognitive Tomorrow.” *MIT Sloan Management Review* 58, no. 1 (Fall 2016).

Smith, H. A., and J. D. McKeen. “Developing a Digital Strategy.” Smith School of Business, Queens University, IT Management Forum, 2015. <https://smith.queensu.ca/it-forum/index.php>.

Wikipedia, 2017. “Anthropomorphism.” <https://en.wikipedia.org/wiki/Anthropomorphism>.