CHAPTER

2

Reimagining Competitive Advantage in the AI World

AI as an area of research and fascination has been around for a long time. It's a field that once disappointed its proponents in terms of progress and is now at the forefront of solving real-world problems. It's rapidly entering into our daily lives as it expands into activities commonly performed by humans, often raising concerns about our very existence.

As AI is beginning to realize its potential in achieving human-like capabilities, businesses are asking how they can harness AI to take advantage of the specific strengths of human and machine.

Sensing, thinking, and acting have been key to human evolution, but then, when you say machines can do these things as well, it raises quite a lot of eyebrows. Humans tend to exhibit heightened capabilities in general intelligence (we call it *fast parallel processing and pattern recognition*), whereas they tend to slow down when it comes to logical reasoning (we call this *sequential processing*), owing

S. Mohanty and S. Vyas, *How to Compete in the Age of Artificial Intelligence*,
https://doi.org/10.1007/978-1-4842-3808-0_2

to careful evaluation of implications/rewards. On the other hand, computers/machines are becoming good at pattern recognition in narrow fields (called *narrow intelligence*), yet they are superfast at logical reasoning provided there are clearly defined inputs and expected outputs.

Based on the characteristics it exhibits and the contours of interaction boundaries as well as its limitations, AI can be placed into four categories, discussed in the following sections.

The Sense-and-Respond Machines

The most basic AI systems are purely reactive; they are designed to sense an event and programmed to provide a pre-determined response. These basic AI systems neither have the ability to understand what led to the event, nor do they have the ability to form memories or use past experiences to make informed decisions. Deep Blue (IBM's chess-playing supercomputer), which beat international grandmaster Garry Kasparov in the late 1990s, is the perfect example of this type of machine.

Deep Blue was designed to sense the pieces on a chess board and determine what move its opponent is playing. Based on the opponent's move, it was programmed to respond with the most optimal move from among the various possibilities. However, apart from having access to a huge library of chess moves and rules, it didn't have any concept of the past, nor any memory of what moves were made before. It was designed and programmed to sense the present moment and respond with an appropriate move.

Similarly, Google's AlphaGo (based on neural network techniques), which has beaten top human Go experts, uses a much more sophisticated analysis method than Deep Blue's, but it still can't evaluate all potential future moves. By applying neural networks and sophisticated algorithms, these machines have progressively become better at playing specific games and beating human experts at their own games, but they aren't adaptive to other situations easily, whereas humans are. In other words, these machines can't function beyond the specific tasks they're trained to work on and can be easily fooled.

They can't adapt to changing scenarios (if you change the rules of the game, they will stop working and need to be reprogrammed). Instead, these machines will behave exactly the same way every time they encounter the same situation. Actually, this behavior of being consistent is their strength in a different context. It brings about a heightened trust quotient—consider that you would want your autonomous car to be a reliable driver all the time, every time.

Limited Memory Machines

This type of machine has limited memory and limited reasoning capabilities. Self-driving cars do some of this already. For example, they observe other cars' speed, proximity, and traffic signals. These observations are added to the self-driving cars' reference library to simulate real-world representations and act based on what is pre-programmed. But this library of information about the real-world representations are only transient, it lacks the ability to learn from the environment in a continuous manner, the way human drivers compile experience over years behind the wheel. For self-driving cars to achieve human-level cognition, they need continuous updates to their pre-programmed representation of the world.

So how can we build AI systems that are capable of full representations of the world, remember the experiences, improvise to adapt to the changing environmental conditions, and learn how to handle new unseen situations?

Theory of Mind

This is the point where machines start to truly exhibit thinking and reasoning capabilities on their own. Machines in this more advanced class not only have the capability to form representations about the world, but also have the capability to interact and learn from other machines. In psychology, this is called "theory of mind"—thoughts are continuously evolving and reasoning capabilities are deciding trade-offs.

The theory of mind is central to human evolution, because it allowed us to learn from social interactions. Assuming someday machines will achieve the same thinking and reasoning capabilities as humans, they need also to be aware of human motives and intentions; otherwise, the human-machine coexistence will become at best difficult, and at worst impossible.

Self-Aware Machines

The pinnacle of AI development is to ultimately build systems that can not only think and reason but also exhibit consciousness.

This is, in a sense, an extension of the "theory of mind" possessed by machines that are self-aware. For example, understanding and interpreting "I want that item" is a very different statement from "I wish I had that item". It's all about being conscious and being able to predict other people's feelings. For example we may assume that aggressive drivers are either in a hurry or are impatient, so we let them go. Our self-driving cars need to exhibit the same understanding and give way. Without a theory of mind, we can't make these kind of inferences.

While we are probably far from creating machines that are self-aware, it is obvious that we will need quantum leaps in processing power for machines to exhibit a combination of vastly different types of sensing, thinking, and problem-solving capabilities—the hallmark of human intelligence. For example, today's much hyped self-driving car doesn't exhibit what we would consider common sense, such as pulling aside to assist a child who has fallen off her bicycle.

It begs the question: How can businesses leverage AI to gain a competitive advantage?

AI in the Boardroom

The very idea of artificial intelligence (AI) taking a seat in the boardroom may seem preposterous and far-fetched. After all, experience, lateral thinking, judgment, shrewdness, business acumen, and an uncanny ability to spot trends ahead of time are the critical skills required for the kinds of complicated matters that boards often deal with. However, at the same time, it is also true that AI is addressing some extremely nuanced, complicated, and important decision making processes. So, AI taking a seat in the boardroom may not be a bad idea at all!

Businesses are going through heightened sense of volatility and there are strong headwinds in the form of technology initiated disruptions. We all know that the cost of bad decisions is high, and if these bad decisions are made at the board level then the implications are further magnified. While there are numerous examples of successful new product launches, mergers and acquisitions, and digital transformations, at the same time there is an even larger number of instances where many of these decisions have turned into failures, dragging the company into the ground. A report from Innosight (a growth strategy consulting firm) shows an alarming trend— "The 33-year average tenure of companies on the S&P 500 in 1964 has narrowed to 24 years by 2016 and is forecasted to shrink to just 12 years by 2027".

Over the past five years alone, companies that have been displaced from the S&P 500 list include many iconic corporations such as Yahoo, DuPont, Staples, Dell, EMC, and Safeway, to name a few. While the reasons are many and some were beyond the control of corporate leaders, it is true that many of these disruptions were driven by a complex combination of technology shifts and economic shocks.

The signs were there, but these companies missed the opportunities to adapt or change. For example, they continued to apply old business models to new markets (similar to the tiller effect discussed in Chapter 1), were slow to respond to disruptive competitors in low-profit segments, or failed to adequately envision and invest in new growth areas that often takes a decade or longer to pay off. In the mean time, the explosion of "decacorn"

companies across industry sectors like transportation (Uber and Lyft), financial services (ANT Financial and SoFi), aerospace (Space-X), real estate (We Work), healthcare (Outcome Health), and energy (Bloom Energy), as well as everything in the technology space has been accelerating the disruption by introducing new products, business models, and services to new markets and customers.

In short, businesses have become too complex and there are blind spots—most leaders are under the false notion that future competition is coming from existing players, whereas the reality is that new competitors are entering, serving your customers in an entirely new way. You need to respond appropriately to these new set of rivals and for that your board and CEO need to make good decisions with the support of intelligent systems.

Leaders across every industry are failing to understand these disruptive forces and need serious help in developing some kind of methodical approach to the madness. The Innosight paper recommends five essentials:

1. *Spend time at the periphery*: A good place to start is to continuously scan the horizon and understand what the emerging unicorns are doing. Your corporate strategy and research teams need to spend time online or in physical spaces experiencing new products and services and reflect back on the impacts to your core business model. By doing so, you can spot early warning signs and opportunities that could cause massive shifts in value.
2. *Focus on changing customer behaviors*: You can no longer be under stable assumptions; you need to become customer obsessed, observe your customer habits minutely, and seriously evaluate their changing behavior patterns, no matter how tiny these patterns appear to be. While you may think you are providing a large bouquet of services, in reality you might find there are new entrants solving essential jobs at cheaper rates and thereby eating away at your business.
3. *Adopt a future-back strategy*: The conventional approach to strategy starts by analyzing what has happened historically, assessing the implications of future trends, and arriving at a SWOT kind of analysis. Then you extrapolate from today to formulate the strategy for tomorrow. While this approach has worked reasonably well in the past, given current business conditions where VUCA (volatility, uncertainty, complexity, and ambiguity) is the new normal, this approach constrains strategic choices and

can prevent a company from objectively assessing their current model. In today's dynamic business conditions, leaders are overwhelmed and worried about not living up to their responsibilities.

To stay future ready, organizations need to adopt a different strategy. Instead of starting from the present and looking forward, they have to envision a future environment and business portfolios based around changing customer needs. The vision of the future only serves as the starting point. You then need to move backward in time to develop a set of strategic initiatives and innovation interventions for the present. The underlying assumption is that tomorrow may not resemble today.

4. *Embrace dual approach*: Transforming the core business is only going to give companies incremental linear growth. Hence, companies need to adopt a dual approach like the digital platform companies have shown. Discover new growth opportunities outside the core, and then invest and govern them separately.
5. *Assess the cost of inaction*: All said and done, it requires a strong conviction to prioritize and marshal your resources including capital and talent to embrace the dual approach. When it comes to decision making, leaders tend to take the safe way out—they push new growth ideas to the back burner. While there may be valid and perfectly acceptable reasons to do so, you need to continuously validate your strategy given the pace of technology evolution and multiple forces of disruption at play. There is no longer the notion of annual corporate strategy and blueprints. Nowadays you need to have a quarterly corporate strategy and blueprints and estimate a price tag for inaction by measuring the impact of lost opportunities.

The question is, how do you prepare for such dynamic and fast changing scenarios?

One plausible solution is to incorporate AI supplements into the practice of corporate governance and strategy. This is not about automating leadership and governance, but rather about augmenting intelligence using AI, for both strategic decision-making and operations decision-making. How do you do this?

AI could be used to improve strategic decision-making by analyzing capital allocations and highlighting concerns—maybe your company is cutting spending on research and development while your competitors are increasing investment in the R&D and M&A areas, thereby actively funding startup ecosystems and becoming better at market intelligence to identify potential new entrants moving into your market space. Similarly, AI can be used to improve operational decision-making by analyzing internal processes and systems and determining employee productivity, predicting churn, highlighting inefficiencies, etc.

Being a new technology, AI has its fair share of naysayers—executives are uncertain of its business case and believe that human capabilities are critically important to capture the returns from AI in the enterprise.

Having lots of data and producing insights is not the key. How you consume the analytical results is the real game-changer. In the digital era, no doubt organizations have matured in the data and analytics areas. By investing in infrastructure, tools, processes, and talent, they were able to improve their analytical insights production capabilities more quickly than they were able to improve the consumption abilities. As a result, despite the fact that their analytics production capabilities were improving, the analytics consumption gap is widening—the opposite of what you would hope and expect.

Managing AI technology demands new leadership skills, including those required to implement, govern, and make right use of the analytical insights.

Using AI to Create a Strategic Advantage

"Datafication" of the world and rapid evolution in prediction technology have gotten us to solve well-defined but complex problems like recommending movies, diagnosing cancer, and creating autonomous vehicles, to name a few. However, just using the technology, no matter how advanced, does not deliver a competitive advantage. For technology to advance the business strategy, it must be embedded into business processes, customer interaction touch points, partner ecosystems, and into the day-to-day life of the enterprise workforce.

No doubt, technology has gotten better and smarter. As a result, AI has forayed into tasks that require creativity and intelligence, such as creating movie trailers, composing music, performing facial recognition, and detecting emotions. They have also gotten better at conversing with humans, including interpreting and extracting insights from images, machine generated data, text, and unstructured data. Given such rapid progress, all in the last few years, there are good reasons to get excited that technology might deliver a "AI strategist" app that can directly assist CEOs in their strategic decisions.

How do we build a "AI strategist"? It can't be developed by itself, at least for now. It requires significant human involvement to begin with. We human beings are unique in many ways; we can think outside the immediate scope of a task or problem. AI lacks this capability. After a lot of training, it can flawlessly execute a well-defined task, but if you change the problem definition, it will falter.

In other words, for AI to exhibit human-like thinking, it needs to act in concert with humans and learn how to seamlessly execute conceptual and analytical operations—including problem definition, signal processing, pattern recognition, abstraction and conceptualization, analysis, and prediction. Of course, this is not to say that AI is incapable of learning these higher-order skills—but that state is far away, too far away.

For AI to demonstrate effective business strategy development, it must demonstrate *reframing skills and continuous learning capabilities*—the process of redefining and reanalyzing the problem with feedback.

Amazon provides an excellent example of how AI and humans, in a continuous learning mode, execute the business strategy. The company has several AI systems (supply chain optimization, inventory forecasting system, sales forecasting system, profit optimization system, recommendation engine, and many others). These systems are intertwined with one another and with human strategists to create an integrated, well-oiled AI ecosystem. For example, if the sales forecasting system detects that the popularity of an item is increasing, it triggers a series of communications: update the inventory forecast, which triggers an optimization process in the supply chain system to update inventory across warehouses. This causes the recommendation engine to push the item more, the profit optimization system adjusts pricing, and the marketing system and launches real-time campaigns and discounts. The resulting impact in turn feeds back to the sales forecast system. These are only some of the first-order effects; there are numerous further interactions between the systems and partner ecosystems that occur downstream. While many of the operations happen automatically, human beings also play an important role: they design experiments and review data traces to make the AI systems continuously learn and improve. Humans also pay specific attention to extract higher-order insights from anomalies and patterns captured by AI, that essentially serves as food for thought and as the impetus for Amazon's next strategic moves. For example, if the item gaining popularity is not usual, then it requires a certain amount of investigation. Is it a long tail item but gaining popularity because of certain uniqueness that was not observed earlier? Is it because the item suddenly found a fan following due to another related but costly item falling out of favor? How seasonal or short-lived is this moment?

How can businesses create an effective AI strategist app? At a broad level, there are three requirements, described next.

A Future-Back Goal

To start with, you must define the desired outcome. Since strategy can encompass anything and everything, it is important to have a well-defined area to begin with; for example, you could focus on competitive intelligence with a special focus on unicorns. Human strategists must provide the initial set of questions and help resolve ambiguity. The AI app can then learn to improvise based on continuous feedback and eventually become better at what it is supposed to do.

The opposite of not having a specific goal is to go on a wild goose chase. There is always this danger with powerful technology—it leads us to become preoccupied in thinking about what it *can* do rather than what it *should* do. We then invariably get caught up and get driven by AI's capabilities, rather than addressing the problems that we need to solve. In other words, if all you have is a hammer, then everything looks like a nail.

A Human-in-the-Loop Approach

The problem statement, data, algorithms, conversational interfaces, feedback mechanisms, automation capabilities and most importantly the human strategist (either to intervene or to refine) must form a tightly integrated continuum. This integration is critical, because the human and the technology should have only one objective—to optimize for the global outcome rather than for individual tasks. Why? Let's take the problem statement as, "evaluate a new business opportunity". This will need deeper research around competitive threats, strategic fit to the organization's vision, evaluation of whitespaces, build vs. buy decision prudence, etc. All these areas are related but are also a sub-domain on their own. If they are analyzed separately, the AI strategist will be no closer to an answer unless there is a mechanism to integrate the findings and the human is involved as an integral component to resolve trade-offs, thereby generating new insights.

As the AI strategist takes on increasingly complex questions, human beings, instead of running around to collect data (which is where they spend a lot of time today), with their unique ability to understand the broad context and connect insights from disparate spheres, will spend more time making sense of the data and feeding those lessons back to the AI strategist, which eventually will help make the AI strategist perform better.

At this rate, it is not absurd to think that one day the AI strategist will become smart enough to surpass the human strategist. It will be an interesting day for sure.

A Well-Designed No-UI

The AI strategist must be able to communicate its findings and recommendations to people; conversely, people should be able to understand, examine, and validate those recommendations and provide feedback to the AI strategist.

Humans are able to think outside the box and carry out higher-order tasks of reasoning, validating, and reframing because we communicate effectively. If the AI strategist operates as a black box and creates outputs that are incomprehensible, it loses its value. People cannot interpret the outputs and therefore cannot build deeper and richer insights through successive reframing. To avoid this situation right at the onset, the AI strategist must be designed to be transparent and conversational.

How Important Are Cost, Quality, and Time?

AI as a discipline has progressed rapidly. It has emerged out of the labs and is already making dramatic changes to our day to day life. Some AI applications have even reached human-level performance in many aspects of vision, conversational speech, and problem-solving. As a result, every industry is actively considering AI initiatives within its enterprise. However, since we are in the midst of such unprecedented disruptions, there is no methodology or prior use cases one can follow. Important questions include: How much does it cost? How do you define ROI? How do you measure the quality? How much time should it take? There are many more related questions, including what will happen to our jobs!

No doubt, with today's robust AI technologies, we are transforming the way we have been working. In many sectors, we have automated activities that are repetitive and mostly back-office operational tasks. Interestingly, AI has also started to lend its intelligence to fields that require less repetitive manual labor and once seemed immune to automation, such as law, education, journalism, etc. In short, AI has started augmenting human minds, not just muscles.

In the midst of all these happenings, it is important to acknowledge that there's no shortage of work that can be done only by humans (refer to Chapter 1—human judgment becomes invaluable). The challenge that organizations need to address is not a "world without human involvement" but a "world with rapidly changing human labor-related skills".

Certainly costs and quality parameters are important for any business to stay competitive and meet customer expectations, but the AI development lifecycle needs to be looked at differently compared to other business application development lifecycles. The initial AI development cost may appear to be expensive, but once it's developed and the AI system starts running in the environment, it can scale to accommodate enormous workloads without a fuss.

The other important advantage is that AI systems continuously learn from data as well as from the human-interactions, thus making the ongoing run and maintenance costs negligible, if not zero. For example, once you have developed an AI system to analyze the lab reports and report findings, you can scale it to do the same analysis for millions of lab reports at minimal incremental cost. This aspect of AI paves the way for organizations to replace or increase productivity of expensive knowledge workers as AI automates many tasks. The result? An opportunity to reexamine currently expensive bottlenecks and see how AI can overcome them.

With machines, we can expect execution discipline, precision, and consistent results. Your AI radiologist will deliver the same output every time, all the time. With this consistency in place, organizations can incrementally refine and improve quality over time.

While reducing costs and improving quality are important, AI also addresses a fundamental business goal—faster time to market. For example, the current loan approval processes typically take 10-30 days. Getting a loan, even a pre-approval, doesn't happen overnight. With data increasingly available to support all the information that goes into a loan processing application, AI can dramatically reduce the time required to process the loan. If you find a house to your liking in the morning, there is no reason why your loan application will need a month to be processed. Ideally you should be getting the keys to that house by the afternoon. Loan approving officers (if they continue to exist) will simply need to be able to swipe your credit card to close the transaction.

Lawsuits are at the other end of the extreme. They can take considerable amount of time to resolve. Some of the time is in gathering data and preparing the case, some is in delays due to court congestion, yet more is due to deliberation and settlement. This entire process, with multiple bottlenecks and dependencies, could be shortened by applying an AI solution that can research and prepare the case along with references to past judgments and then provide an early indication of how much time it would take to arrive at the outcome. Certainly, it will help people from just going round and round in circles, wasting time and money, if they are provided with enough factual information about where their case stands!

Medical diagnoses are another example of cumbersome processes. Many times the delay happens because of non-availability of past medical history. Secondly, even if you have the past medical history, you are dependent on availability of the right physician in your insurance covered healthcare provider network. What if you need a simple diagnosis to understand how bad your condition is? AI can shorten the lag time considerably by analyzing your past medical history and matching your current symptoms with a vast library of diagnosis results. The potential for faster preliminary diagnosis can certainly help hospitals and clinics and at the same time improve health outcomes through earlier detection of conditions that require quick medical attention.

As AI's usefulness begins to outweigh its uncertainties, what do organizations need to watch out for?

Too Fast to Manage

We are fortunate that the slow pace of our current processes allows us time for monitoring and management. If (when!) something doesn't work as expected, we usually have time to perform root cause analysis, find any bottlenecks, instigate fixes, or perform course corrections. However, when machines run our processes, we may not have time to step in and manage the situation, because everything will be running at a machine pace. We may not get the opportunity to intervene before the money is gone from a risky loan, the guilty walk away, or the patient suffers from an incorrect treatment. In the race to optimal speed, the machine pace (extreme automation) may lead to risky outcomes.

Too Fast to React

Besides many other reasons like improved productivity and improved responsiveness, one of the key reasons for expediency is to stay ahead of your competitors. With AI you can create stiff entry barriers for your competition. However, it won't likely work out that way. Just like you are seriously thinking about AI, your competitors are as well. The result? When everybody wants to stay one-up, it creates a huge amount of instability in the ecosystem. You no longer have the luxury of market leadership or competitive advantage; it suddenly has become a level playing field.

Too Fast to Learn

AI solutions are data hungry. The more data there is and the more varied it is, the better the AI outcomes. In a bid to outweigh competition, if every enterprise seriously acquires as much data as they can, suddenly you end up in a situation where there is no longer any differentiation due to data. The other issue is that the value of data may decay quickly. The rate of decay may exceed the rate of availability. The Netflix prize is a classic example: the company offered a prize for a substantial improvement in its algorithm that recommends movies based on historical customer viewing behavior (DVD-by-mail rentals). After awarding the prize, Netflix found that the algorithm was not useful on the new data (video streaming behavior). When AI depends on data, the rapid decay of the data's value may actually render the algorithm less useful.

Are AI Learning Scenarios Unpredictable Enough?

What if a collision happens between an AI operated vehicle and a human operated vehicle? Quite naturally, the sensational headlines would play to our biased thinking that the machine is to blame. This scenario is not fiction. It actually happened in Las Vegas, where a minor collision happened between a self-driving shuttle and a human operated truck when the delivery truck backed into the front bumper of the shuttle. The AI operated vehicle was trained to stop if there is an obstruction on its path, and it behaved exactly as it was supposed to behave.

This incident highlights a crucial gap in AI and human interactions. AI systems are typically trained in context where there are no external factors with adversary objectives. The focus is to learn from observations and, while designing such systems, our first intent is to get something working.

Consider how the incident between the AI operated vehicle and the human operated vehicle unfolded. The AI operated shuttle accurately recognized the obstruction, predicted a potential collision scenario, and stopped. This is the strength of AI—observing many inputs and processing those quickly to arrive at a reasonably accurate decision.

Since the collision happened anyway, the question is: what else could the AI operated vehicle have done to avoid the collision? This is where it becomes difficult. It could have honked loudly, which might seem like common sense and fairly a risk-free option to humans. It could have swerved away from the approaching truck, but this seems more difficult and riskier than a honk. For imperfect AI (still learning and improving) faced with uncertainty, a reasonable decision is to stop and do nothing.

The core of the problem is information asymmetry—perfect information versus imperfect information. When we are thinking about human-machine interactions from a game-theory perspective, access to information (or lack of information) changes the outcomes radically. If the AI system and the human don't know what the other will do—both have imperfect information—then the situation is extremely volatile and unpredictable. If AI has perfect information, that is, it knows what the human will do, then the uncertainty no longer exists (at least for AI with perfect information). The situation plays to the AI's strengths. Similarly, if humans know what the AI will do, but the AI systems have imperfect information, we are creating a scenario that plays to the AI's weaknesses.

Let's take the example of job applicants. Once the job applicants figured out that AI systems were looking at an absence of certain keywords to filter out job applications, they got creative and included every possible keyword that would show up as an important and critical skill in their resume. Awareness of how the algorithm worked meant applicants could manipulate the system to their benefit.

There are heightened concerns that the preoccupation with autonomous driving and creating human-appearance like robots is distracting us from the more enterprise relevant, and potentially more transformational, applications of artificial intelligence in business. While there is nothing wrong with solving problems that we had never imagined we would be able to solve, prudence must be applied by business leaders. We must not stay stupefied by these fascinating applications of AI, but learn from these examples and apply novel ideas to solve business problems.

How can managers can bring sanity into the AI operated world? Let's take some inspiration from a scenario that played out during the period of the Cold War.

Lieutenant Colonel Stanislav Petrov noticed that the Soviet information systems were sounding alerts about incoming nuclear missiles from the United States. He was responsible to make a well-informed decision and he had to act fast: Should he authorize a retaliatory attack or hold on and make further investigations? Fortunately for all of us, Petrov chose to investigate. He realized that a real attack was unlikely because of several other factors—one of which was the small number of "missiles" reported by the system. After his death in May 2017, a report credited him with "quietly saving the world" by not authorizing a retaliatory attack.

The real reason behind the false alarm was due to system's inability to accurately distinguish the light signatures of sun's reflection off clouds to light signatures from missiles once they become airborne.

Businesses face similar (although hopefully less consequential) questions about whether and when to remove humans from their decision-making processes. There are no simple answers. As the false alarm incident demonstrates, a human can add value by carefully evaluating a system's recommendations and its implications before taking action. We can surely question the efficacy and robustness of the AI algorithms. Humans develop the algorithms, hence humans could add more value by helping the classification system prevent misclassification.

If our world was full of autonomous machines, when and why would they want our assistance? In the nuclear attack scenario, the machines, while correctly predicting the imminent attack should also highlight the consequences of the desired action—worldwide destruction. Perhaps we should build enough checks and balances to weigh in the consequences of a certain action, even if it is the right thing to do.

With immature AI, the machines should assess and recognize their own state of inaccuracy and initiate human intervention automatically. For example, if there are insufficient observations, humans are more likely to utilize their breadth of experience, confer with other humans, and learn from other

instances in ways that machines cannot (yet). Classification does not have to be a binary state (missile or no missile); AI systems, when lacking confidence, should request human help.

Fortunately, business decision-making often differs considerably from the fully automated machine world. The pace of digital business may have accelerated, yet many business decisions can still afford time for a second opinion where humans can play the role to confirm or deny. In business situations where time does not allow additional steps of human intervention, the AI systems should keep asking for additional training to correct errors.

The key question is: Since AI is still maturing, should we take a wait and watch stance or should we take few pragmatic steps, clearly knowing that there are implications?

Managing Immature AI

Despite the advancements in AI technology and even as organizations continue to push the boundaries of what's possible with AI, there is a wide gap between the promise and reality of AI.

While at one end, we have many successful implementations of AI solutions, at the other end, we do have many unsuccessful attempts. This is largely because AI is still a young, immature field.

Algorithms are competing with, and often winning against, expert humans in complex games such as chess and Go. But at the same time, simple customer service chatbots are failing to address the real pain points with customers and at times can be more annoying than helpful.

AI has the potential to apply data and algorithms to offer quick decision making; however in many scenarios it seems the algorithms are themselves biased, thus not giving an impartial view on the outcomes. Facing this contradiction, it can become difficult for executives to get a clear view about the usefulness of AI, and hence there is the temptation to adopt a wait and watch policy.

So, how are executives pragmatically incorporating AI into current business processes?

Clearly, AI can take the information you have and generate information you don't have. From this perspective, you can expect AI's primary job to enhance a knowledge worker's ability to make use of the predictions (there is a high probability of someone defaulting on a payment), whereas the primary job of the knowledge worker is to apply judgment (seems like an anomaly, as this person has never defaulted) to improve AI's prediction performance.

For quite some time, the set of recognized prediction problems were statistical and operations research related, such as inventory management and demand forecasting. However, over the last few years, image recognition, driving, asset maintenance, talent acquisition, health care, and many other tasks that were solely the purview of humans have also been framed as prediction problems.

As more and more tasks that were solely under the purview of human-prediction work get reframed as machine-prediction problems, organizations need to shift from training their employees in prediction-related skills to judgment-related skills. In short, organizations need to focus on creating an ecosystem and processes to enable teams of judgment-focused humans and prediction-focused AI agents working in concert to deliver business outcomes.

The effectiveness of AI technologies will be only as good as the data they have access to, and the most valuable data may exist beyond the boundaries of one's own organization. Organizations will need to look at their current inward looking data strategy and open up to develop strategic alliances that depend on data sharing.

For example, BMW, Daimler, and Volkswagen, even if they compete with each other, formed an alliance to buy a Berlin-based digital mapping company (HERE) to create a real-time platform that would track and monitor driving conditions (traffic congestion, estimated commute times, accidents, or vehicle break-downs on routes, emergency services, public utilities, and weather patterns). These are based on data collected from cars and trucks from each brand. Alone, it would have been prohibitively costly for one company to build such a platform, but together they created a sufficiently robust data platform that provides improved customer service and creates new revenue models, such as subscription based fees for municipalities, private emergency service providers, citizen services companies, insurance companies, and many more.

AI will affect different organizations in different ways, and many of the changes will have direct implications to managers, such as how they plan work allocations, how they review the work in progress, how they provide interventions, how they manage a combination of AI and employees, and so on. Additionally, these changes bring considerable risks as well.

- *Replacement risks*—AI promises to provide assistive intelligence to human performance in a number of ways, which is good as it improves human productivity, but at the same time AI also redefines certain activities, thereby completely eliminating human involvement. The reality is that human roles and positions are not immune any more. Insurance companies have already started using AI (instead of human agents) to offer customers insurance plans. As the push for smart machines, smart factories,

and smart plants continues to progress at a rapid pace, the manual labor force, whose workday is largely intertwined with operating machines, will find their jobs are taken away by AI agents.

- *Dependence risks*—In the long run, as AI becomes completely embedded into business processes and in a sense runs the enterprise functions, a strong dependency will emerge. Unless enough care and strong governance models are in place, this dependency will create new vulnerabilities, creating deadlock situations between AI agents and human-in-the-loop and leading to potentially ineffective operations. Errors and biases may creep undetected into algorithms over time.
- *Security risks*—Today, we manage security through surveillance, monitoring, and access controls. Our current security management practices are designed to detect unauthorized activities by humans. We are beginning to realize the implications of security threats from software. Imagine when enterprises are fully running on AI agents and use sophisticated algorithms that deal not only with mission-critical applications but also with sensitive data. The risk of AI stealing information from other AIs will be very real.
- *Privacy risks*—As AI takes on the role of information worker in an enterprise, the AI itself will become a source of information over time—collecting, analyzing, and governing fairly sensitive corporate and customer data. Widespread use of AI will certainly raise many data privacy and ownership issues (we are already seeing stricter data management and protection regulations in the form of GDPR for the EU). The black box nature of AI systems will pose additional challenges for human-machine collaboration scenarios. Unless you know why the AI system did what it did, you really won't be able to trust the AI system. This is where organizations need to come out with clear ethical standards and draw the line between AI autonomy and human ownership.

How will AI impact and disrupt the way enterprises do business? And how should executives plan for the upcoming decade of disruption?

Innovative organizations reinvent themselves; they demonstrate that they are able to fail fast and learn faster. If you are graduating with an AI technology degree from a prestigious school or you have invested your time wisely to get trained on the latest AI technologies, would you pursue a startup that is focused on a niche AI-based product offering, or join a company that wants to build innovative AI applications, or use your academic skills to collaborate with scientists in other fields to conduct advanced research?

The opportunities presented by the first two options are glamorous (thanks to the press) and highly rewarding, consequently driving the salaries of AI skilled resources through the roof. Why is there such a rush to hire AI talent? The answer lies in the fact that companies are sold on the idea that AI will magically create new business segments or will help widen the gap between them and their competitors.

In our view, the third option of advanced research is equally important for long-term sustainable and ethical use of AI in all spheres of life and society. We keep hearing optimistic views that AI (and other digital trends) will create new industries and new job categories that will outshine whatever job losses it causes. There are two open questions regarding these views: How will these new industries be created? How soon will they come?

New industries are not created overnight; it takes decades to transform pathbreaking ideas into viable commercial value propositions, which eventually leads to setting up new industries and creating new jobs. Given this long journey, the real question is, will these new industries and new job categories emerge fast enough to maintain sufficiently high employment levels in the economy and offset the jobless impact due to automation and AI?

Companies lapping up AI talent are primarily focused on creating AI solutions to automate and optimize, not to create new industries. There are numerous examples already proving this mentioned point: Uber, Lyft, Airbnb, Amazon, and many other algorithm-powered companies are either focused on improving efficiency out of current business models or creating new business models by removing process bottlenecks (intermediaries). These two methods, while introducing step-change efficiencies to businesses as well as end-users like you and me, are not necessarily creating new industries or new jobs.

Hence, it is equally important for companies and governments to find a way to use the already scarce AI talent to contribute to the discovery and pursuit of scientific and technological advances that could lead to creation of new industries and new types of jobs.

Skills to Succeed: EQ Skills for Evolving AI Economy

Baron-Cohen popularized the empathizing-systemizing theory. People bring different skills and strengths to do their jobs effectively. Empathizing skills come to aid at identifying, understanding, and resolving conflicts, motivating teams, and responding to the mental states of others, whereas systemizing skills come to aid at analyzing, understanding, reasoning, troubleshooting, planning, and predicting outcomes. According to Baron-Cohen, women score higher on empathizing and men score higher on systemizing. This doesn't mean that women can't excel at systemizing or men can't excel at empathizing.

Continuing our conversation on the impact of AI on job roles, the roles that are likely to disappear over the next decade are not limited to a particular industry or cadre in the corporate hierarchy—the impact is far and wide and across all levels. AI agents will replace not only truck drivers and assembly line workers, but also radiologists, financial planners, insurance agents, and security guards, to name a few—all traditionally male-dominated roles. It is also abundantly clear that as more and more work becomes automated, there will be a significant shift in demand for the skills like judgment, empathy, compassion, influence, and engagement (skills in which women more often excel). For simplicity, let's call these as emotional quotient (EQ) skills.

We are heading toward a world where AI would act as a great leveler. How? Take the example of radiology. The AI agent may determine that your radiology scans indicate cancer, and you will be at ease when a human sits down with you and explains the best course of action. Imagine, instead of a human, the machine tells you in a boring and monotonic voice what your treatment plan and next steps should be! Similarly in case of business, the AI agent may suggest what operational improvement actions are required in the company to deliver better margins. It is still much more effective for a human to lead the charge to persuade people to execute the recommendations. Imagine the machine sending out prescriptive actions to your employees!

In the AI economy, where machines will increasingly do the "systemizing" type of work (where men excel over women), there will be a greater need for "empathizing" type of skills (where women excel over men). Usually we put a premium on our technology skills and often we tend to downplay our EQ skills. However, given the AI led future knocking at the immediate horizon, all of us—men, women, and organizations—need to start paying attention to the importance of the EQ skills.

How can you get started developing your EQ skills? Just like any technical skill, a person's EQ skills can be improved through training and by providing them with the right support systems. Here are three steps to get started:

- *Create your own EQ baseline*: Humans are social beings. We work in teams; we communicate with each other and together we solve problems. However, when it comes to providing feedback, most of us are very hesitant to criticize someone's interpersonal skills because such feedback pushes us into uncomfortable situations and we become defensive. Many sharp, effective people have no idea that they need to improve on their EQ skills because they simply haven't paid attention to the subtle indicators from their peers and teams. Hence, the first step is to become self-aware of your own EQ quotient and pay attention to feedback you've been given, especially comments along the lines of, "You are difficult to work with," "You are too argumentative," "You need to do a better job of reading others body language."
- *Acknowledge that EQ skills are important too*: We have always paid way too much attention to systemizing skills, whether it is systems, technology, engineering, mathematics, or related fields. Although every role has an EQ angle to it, historically we have pushed it to the back burner. For example, doctors are well trained to identify and treat disease. Even though they are expected to sit down with patients and personalize the treatment plan to suit the patient's preference and lifestyle, they tend to ignore or delegate this specific aspect of relating to someone else. You must first be clear about the outcome you are delivering, not the activity—is getting the diagnosis right the most important measure of success? Or is it actually improving the patient's health? If the latter is true, then it is equally important for you to start focusing on your EQ skills.
- *Reframe learning and knowledge management as EQ management*: Learning and knowledge management functions have always taken a myopic view of business needs. The focus is to train you for tasks to be done not for roles of the future where more than your technical skills your interpersonal skills will become much more important. You need to honestly assess your EQ quotient. You need to find a coach who will give you honest feedback and mentor you, and you need to consciously work on the improvement areas. The difficult part of EQ is that we are hesitant to agree that we have EQ gaps. None of us want to admit our EQ needs work, and we have this notion firmly printed on our minds that our EQ is inborn and unchangeable. We are wrong on both accounts.

As AI becomes increasingly embedded into everything we do in our professional and personal spheres, the "softer" aspects of our skills need to become much more important.

Use Case: AI and Amazon Flywheel

More than any data-driven company, Amazon has delivered stupendous services and products with its audacious vision for an AI-powered enterprise. At the center is Amazon's "flywheel," which ensures various parts of its enterprise functions work in tandem, by feeding into each other to deliver business outcomes. Another interesting concept that propels Amazon to keep delivering amazing innovative products and offerings is the "future-back-to-the-present" approach. At Amazon, any initiative starts with the end in mind—a "six pager" pitch starting with a speculative press release describing the finished product or service.

AI is at the core of this flywheel, where algorithm powered innovations in one part of the company are extensively leveraged to accelerate the efforts of other teams. Amazon meticulously examines every opportunity to monetize what they build. For example, during the early days of Amazon, they built robust data management and AI platforms to efficiently run their own business. Later they realized if they can offer these internal capabilities as a paid service to outsiders, they would not only make money, but they will also get an opportunity to understand how people are using the AI services and what they are doing with it. That way they could churn out more meaningful offerings in the future.

Amazon had its humble beginning as a bookseller. First it established a marketplace and launched a powerful e-commerce platform. Perhaps Amazon was the first company to truly realize the importance of a platform. It opened up its marketplace to other retailers, thus establishing a sharing economy that took advantage of its e-commerce platform. It built warehouses to fulfill orders for customers, and then offered fulfillment services by Amazon to its network of marketplace business partners, logistics, and distribution companies and other channels. Amazon found a way to monetize the excess computing capacity (AWS) they had built to support the business during the busiest shopping seasons.

"Amazon Go" leverages the computer vision and AI algorithms to offer a cashier-less shopping experience to supermarket shoppers. The intent is not to open thousands of Go stores across the country, but to offer this AI-powered retail infrastructure to shopkeepers as a subscription. Another monetization idea in the making.

Right in the formative years Amazon realized the potential of AI and started investing in it. In the early days, Amazon's AI talent was scattered across divisions and cross-leverage of the innovations and ideas were sporadic. The top-down push to institutionalize AI across the company brought in significant changes. The company put "serving its customers" as the key driver for anything they would do, which resulted in islands of AI innovations to collaborate across projects and share their solutions with other groups, thus putting the flywheel in motion.

The Conversational Interface Effect

Amazon's Echo line of products powered by the voice platform Alexa also sprang from a six-pager, delivered to Bezos in 2011 during an annual planning process called Operational Plan One. The vision was to come out with a low-cost computer with all its brain in the cloud that you could interact with over voice—you speak to it and ask questions, and it engages in a conversation with you and provides answers.

Building Echo required a level of AI prowess that the company did not have at that point. However, thanks to its approach of working backward from an imagined final product, the high-level blueprints included features that were critical to making Echo. The voice recognition feature in particular demanded a level of conversational AI that was yet to be invented. For example, activating the AI system by using a "wake word" ("Hey Alexa!"), hearing and interpreting natural language commands, filtering out noise and delivering non-absurd answers, etc. Were not conceptualized anywhere.

In addition, for Echo to be commercially successful, it has to be cheap, extremely effective in voice recognition capabilities, and be able to deliver never before experienced end-user satisfaction. For all this to happen, Amazon had to build an AI system that could understand and respond to natural language queries in noisy conditions. The biggest bottleneck was where Amazon would get massive amounts of conversational data to train the AI.

The answer to this question was found in the Amazon flywheel. It started leveraging AWS (its cloud infrastructure offering) to create a speech recognition service (Alexa) that became a valuable asset beyond its original scope of fulfilling the Echo's mission. Once they developed Echo as a far-field speech recognition device, they saw the opportunity to do something bigger. First, they integrated Alexa into their other products and offerings: you converse with Alexa to access Amazon Music, Prime Video, your personalized recommendations from the Amazon shopping website, and other services. Second, they expanded the scope of Alexa to become a voice service by allowing developers to create their own voice-enabled applications (known as "skills") to run on Echo.

In the beginning Amazon struggled to find conversational data to train Alexa, but once customers began using Echo, Amazon steadily gathered real conversational data. This conversational data became a powerful learning and experimental test bed—once you have a device in the market, you get access to real data that is so fundamental to improving everything, especially the underlying platform.

Democratizing AI

Amazon realized that if they could externalize what they were doing internally using their cloud platform and their own data scientists, they could generate tremendous value. They started focusing on how to simplify AI. Another epic six-pager was in the making.

While the tactical aspect was to add machine-learning services to AWS, the broader aspect was a grand vision. How could AWS become the destination choice to build AI applications for all and sundry? You don't have to invest on platforms, you don't have to have data scientist teams of your own—you just need to define the problem you are trying to solve. The rest is all on AWS.

In 2016, AWS released new machine-learning services—Polly (a text-to-speech component), Lex (a natural language processing engine), Rekognition (a deep learning offering for face recognition)—and launched a more comprehensive AI platform called SageMaker. These offerings allow AWS customers (spanning from giants like Pinterest and Netflix to tiny startups) to build their own AI-powered applications.

AI Way-of-Working

Amazon democratized AI within the company and various units in a decentralized way. They created a specialist central group to define standards, architectures, and best practices. This group's primary objective is to evangelize and promote machine learning across the enterprise in a highly collaborative fashion. Several examples highlight how different units within the company leveraged lessons from other units to solve problems.

- The fulfillment team wanted to better predict which of the eight possible box sizes it should use to package a customer order. They leveraged the algorithms developed by other groups and applied them to solve their own problem.
- Amazon Fresh, the company's grocery delivery service, needed a better way to assess the quality of fruits and vegetables. They leveraged new algorithms developed by their Berlin-based team using IoT sensors in the delivery vans to determine the freshness of fruits and vegetables in transit.

- Amazon Go leveraged a new AWS service called Kinesis to process streaming data from hundreds of cameras to track the shopping activities of customers in the store.
- Amazon's Prime Air drone-delivery service, still in the prototype phase, built a new AI solution to enable its autonomous drones, drawing on knowledge from the rest of the company and figuring out what tools to use.

A New Business Model

Shopping using Amazon has become a part of our life. We visit the website, shop for items, place them in the "basket," and pay for them, and then Amazon ships them to us. Right now, Amazon's business model is shopping-then-shipping.

We all have noticed Amazon's recommendations when we shop; the basic idea is to increase the basket size. The recommendation engine offers suggestions of items that you may find interesting to buy. It is nothing but a prediction system using "product-user similarity" algorithm. Considering that there is a match making happening at scale (millions of items to match with millions of customers' buying preferences), the prediction system does a reasonably good job by accurately predicting what we want to buy about 5% of the time. In other words, we actually purchase about one out of every 20 items it recommends. Not bad!

Amazon continuously collects information about us. In addition to our searching and purchasing behavior on their website, it knows not only what we buy, but also what is our preferred time to do online shopping, where we are located, how we pay, how many times we browse through before we finally make a decision to pay, how many times we have returned products, and many more. The prediction system uses all these data points to continuously refine its algorithms. The goal is to keep improving the metric from browsing to actually buying.

What happens when the prediction system's accuracy level comes significantly closer to 100%?

Will it become more profitable for Amazon to ship you the products it knows you will want to buy rather than wait for you to order them? A change in the business model from shopping-to-shipping to shipping-to-shopping!

Because the prediction system has now reached the threshold of being 100% accurate, Amazon can ship you products it knows you will want. You don't have to order anymore. You decide in the comfort and convenience of your own home which items you want to keep and which ones you want to return. This new approach offers two benefits to Amazon. First, since Amazon proactively delivers products to your door step, it indirectly persuades you

to not to go elsewhere to do your shopping. Second, by delivering products to your doorstep, it influences you to swipe your card, which otherwise you were considering to do for some time but might not have gotten around to. In both cases, Amazon benefits by gaining a higher share-of-wallet. Turning the prediction accuracy far enough creates a new and profitable business model for Amazon.

What are some of the challenges of this new business model?

In the current shopping-to-shipping model, Amazon picks up the products we want to return. However, in the new business model of shipping-to-shopping, there will be many more returns than observed today. So, Amazon will have to invest in infrastructure (perhaps scale up their fleet of trucks or expand partnerships with local logistics companies) to manage product delivery and returns process efficiently.

You can ask, what if they were to launch this new business model at the current prediction accuracy (that is, 5%)? By launching sooner, Amazon will gain a lot in terms of access to newer data and will improve its prediction systems accuracy much faster. In addition, it will create a new entry barrier for competitors. The new business model will attract more shoppers, more shoppers will generate more data to train the AI, more data will lead to better predictions, and the cycle continues.

What do we learn from the Amazon flywheel example? First, when AI is embedded into enterprise functions and processes, it improves the prediction capabilities and in turn has a significant impact on the business strategy. Second, there are increasing returns to investments on AI; however the timing of adopting AI as a business imperative matters a lot. Adopting too early could be highly experimental and costly, but adopting too late could be fatal.

Companies face two questions in light of all of this. First, they must develop a better understanding of how fast and how far they can embed prediction capabilities in their businesses and applications. Second, they must focus on developing a strategy to continuously assess and gain business benefit of applying prediction capabilities to their business.

Putting AI to Work

AI has emerged out of the labs and entered the world of mainstream business. AI as a service offered on the cloud is simple enough for all to adopt; however, there is a lot of hard work that goes into managing the interplay of data, processes, and technologies, and solving problems that are relevant for your industry. If you are under the impression that you can simply subscribe to AI platforms on the cloud, throw all your data at the pre-built algorithms, and you will magically produce "intelligence," then you are completely wrong.

The following sections discuss AI examples across industries and corporate functions, to give you a sense of how AI is used to solve problems.

Marketing and Sales

Gone are the days of mass marketing campaigns. In this increasingly digital world the best way to draw attention is to personalize. AI enables us to offer personalized services, advertising, and interactions. We have already seen the impact of successful personalization in the retail industry. Using not only the data shared by customers but also a large variety of data that is available in the web, retailers, for example, have developed loyalty apps that can reside on our smart phones. Based on our location, time of day, and proximity to supermalls or retail outlets, the loyalty app can push hyper-personalized offers to use in real time.

In sales and marketing, AI can provide augmented intelligence capabilities correlating millions of individual data points with information on general consumer trends and then build a real-time marketing system that will potentially deliver thousands of customized offers in real time. This is a complex task that marketers today would probably take a week to deliver.

"Next best offers" are positioned at the intersection of a customer's needs and the company's catalog of offerings. The key is to understand what the customer really wants (explicitly said and unsaid), do a match with the company's offerings, and close the deal to the satisfaction of the customer as well as the company. In case of the insurance industry, this would mean building a model that reflects the needs of customers as they pass through various life stages. The model will leverage complex algorithms that can crunch more than 1,000 static and dynamic variables consisting of demographics, prior policies, prior agent interactions, economic conditions, segmentation of other customers, thus matching this customer and what products they have taken. The result? A potential to significantly increase cross-selling. The insurer can also use the model to reflect on how its agents are performing, thereby designing training programs to upskill agents who are not up to the mark.

These examples demonstrate the importance of a rich supply of contextual and specific customer data. If you have a well-defined problem statement and data, you can generally validate a proof of concept within four to six weeks and then put together a detailed plan for the data infrastructure and the resources required for a full rollout.

Research and Development

R&D is a fascinating field and the problems in this field are generally complex. They require deep domain and technical expertise and are often driven by many experimentation cycles. AI solutions in the R&D area can expedite the

experimentation cycles, often predicting the outcomes by simulating millions of data points.

For example, in industrial manufacturing, AI can correlate real-time sensor data from machines with past maintenance data and OEM benchmark data to predict when the machine is going to fail.

The problem in the R&D space is lack of availability of data. Companies wanting to better their product designs need to digitize machine operations data. Consumer electronics sector will need to focus a lot more on 3D design data and user experience data. The pharmaceutical and healthcare sector will need to focus on clinical trials data and drugs efficacy data. Given the knowledge and expertise required to develop AI solutions purpose-built for R&D, companies must engage domain experts to systematically curate data and help improve the prediction capabilities.

Operations

Operational practices and processes have standard operating procedures, generate a wealth of data, and have measurable outputs against each task. Quite naturally, operational tasks are well suited for AI supplements. Rapidly evolving technology is transforming every field, and manufacturing is no exception. Cost pressures are rising as industries become increasingly competitive. Heavy industries like chemical, power, and thermal, with longer gestation times and long-term ROI, are facing the brunt, not only from within but also due to the fast-moving external factors. For example, labor that was once readily and cheaply available is now hard to get.

Manufacturing trends in itself are changing; made-to-order is the new way. Just-in-time, lean manufacturing calls for less inventory and a more real-time operations. How, then do manufacturers adopt and embrace this shifting scenario? The answer lies in becoming smarter and better through artificial intelligence.

Artificial intelligence can be applied to the entire lifecycle of manufacturing, right from problem identification to problem communication and then resolution. Automation is necessary to streamline repetitive tasks such as scheduling and rescheduling, planning and data tracking.

Prediction to preempt anomalies in operations goes a long way in bringing to notice aberrations and thereby avert critical catastrophes! Such AI-enabled systems must be instilled for material, machine, and equipment updates and also extended to the systems that they interact with, such as customer orders and supply-procurement.

Procurement and Supply Chain Management

The rise of AI in optimizing procurement function is astounding. Organizations today don't just measure what they spend, they are pushing the traditional approach of spend analysis to measure total value contribution to the business, taking advantage of both conventional and newly accessible data sources. How have they done this?

Artificial intelligence is adept at automating even complex tasks. From identifying new markets and tracking exchange rate volatility to managing risks and assessing the best suppliers, procurement as a corporate function is leveraging AI to streamline processes and improve decision-making. Spend analytics along with contract analytics have already gone through significant automation activities (processes of collecting, cleaning, classifying and analyzing an organization's expenditure data). As a result, the low hanging fruits of identifying areas where savings can be made and point to paths of greater efficiency have already been achieved, but this is only scratching the surface of the transformative powers of AI in procurement.

AI can bring in advanced strategic reasoning and strategic sourcing in to the mix. AI supplements such as cognitive procurement advisors (CPAs) and virtual personal assistants (VPAs) that use natural-language processing (NLP) and natural-language generation can further increase automation and efficiency in procurement. By facilitating creation of requisitions from free-form data like that contained in an email or PDFs (not just data entered into a requisition interface), AI can help bring more purchases under direct control. Speech recognition (such as Siri and Alexa) can change the requisition and approval process for busy and off-site workers. Also, using image recognition technology, employees may opt to send a picture of an item they wish to order or reorder, again simplifying the requisition process and making it likely that more spend is being actively managed. AI can play a significant role in reviewing and even generating and monitoring supplier contracts, automating a typically labor-intensive set of tasks. In addition, AI technologies can look for cost/price discrepancies, unusual order quantities or frequencies, compare contract data to orders and invoices, and a whole host of other patterns that will help companies detect potential fraud or errors, detect and predict purchasing patterns, and identify the top-performing trading partners.

Lastly, artificial intelligence is capable of "learning" your sourcing patterns and behaviors and making best supplier recommendations for a specific project or spend. It can even be used to predict market prices, identify and analyze new potential vendors, and help you evaluate the success of your relationships with existing suppliers.

If AI is to transform supply chain, it should be driven by issues that affect operations today. So, the first question to ask is "What needs to be improved in supply chains right now?" For example, typical retail/CPG supply chains carry 60-75 days of inventory, the average service level in the store is about 96%, with promoted item service levels much lower at the 80% range. The restaurant and casual dining industry, on the other hand, carries around 12-15 days of inventory with relatively high waste and high cost of goods sold.

The following are few challenges in SCM that AI can effectively address:

- *Hard to plan for demand*: Requires multiple iterations to arrive at a plan.
- *Excessive safety stocks and bullwhip effect*: Complex integrations and calculations at each step in the process and at each node in the supply chain network.
- *Supplier unreliability and transport network unpredictability*: Huge opportunities hidden in the network because they are locally sub-optimized.

In different shapes and forms, AI can help with all these problems. However for AI to offer optimal value in supply chain, it is important to ensure the following:

- *Access to real-time data*: The most important business imperative in SCM is to improve legacy batch oriented planning systems, but for that to happen we must eliminate the stale data problem. Most supply chains today attempt to execute plans using data that is days old, which results in poor decision-making that sub-optimizes the supply chain planning and requires frequent manual user intervention to stay current. Without real-time information, an AI supplement is just going to make bad decisions faster.
- *Access to varied data*: Unless the AI supplement can see all relevant constraints and bottlenecks due to external influence in the supply chain, the results will be no better than a traditional planning system. Hence, having access to all touch points related data, including external data, will help improve visibility into supply chain processes and inter-dependencies.

- *Last mile visibility and user engagement*: SCM goes through a number of process areas that are not integrated. The end consumer is the only consumer of your finished goods products. During the supply chain process, each function has their own goals and objectives, which one way or other stay aloof of the last mile impact. Hence, the primary goal of the AI engine must be to improve consumer service level at lowest possible cost but taking into account the entire supply chain touch points.

Data in a multi-party real-time network is always on a fast moving lane. In addition data variability and high latency are recurring problems. Hence, the AI engine can't afford to stay on a batch mode; it must be looking at the problems continuously and should learn as it goes on how to best set its own performance parameters to fine tune its abilities.

Significant value can only be achieved if the algorithm can make intelligent decisions and execute them. Furthermore, the algorithms need to execute not just within the enterprise but where appropriate, across the ecosystem of partners. This requires your AI system and the underlying execution systems to support multi-party execution workflows and, to a large extent, those workflows need to be automated.

Lastly, AI should not operate as a "black box." Users must get visibility to decision criteria, propagation impact, and they should also get a view of the issues that the AI system cannot solve. The users, regardless of type, should play an active role in monitoring and providing additional input to override AI decisions when necessary.

Shared Services Functions

Shared services functions in the corporate world are seen as cost centers, hence there is a natural tendency to partially outsource the support activities to drive cost optimization. Beyond the cost optimization, the service organizations need to shift focus to embed intelligence and automation in processes in order to offer higher-value services and improve service levels.

Many service organizations are starting to recognize the benefits of combining AI with robotic processing automation (RPA). They are using rules-based software bots to replace repetitive manual human activities adding flexibility, intelligence, and learning via AI. This approach combines the rapid payback of RPA and the more advanced potential of AI.

Conclusion

What are the keys to a successful AI strategy?

For ages, we have indulged in statistics, operations research, and automation to solve business problems and gain competitive advantage. For example, pricing is core to any business and the only goal is to increase the company's sales and margin. In simpler times, companies adopted different strategies for pricing that more or less stayed static for an extended period of time. To their pricing strategy, they had built-in scenarios, rules, heuristics, econometrics, and competitive intelligence. The more holistic and flexible the pricing strategy was, the more it served as a source of competitive advantage. However, today we are not in simpler times.

Ubiquity of digital in consumer lives, connected devices, the always online generation, and in general an increased level of awareness means consumers are no longer captive to one product or one service or one company. Today businesses are forced to focus on dynamic pricing (the right price to the right customer, at the right time, for the right product, through the right channel) to solve the same old problem statement—to increase sales and margin. This is where AI comes in.

Companies looking to achieve a competitive edge through AI need to reimagine their business models. They need to identify what machines can do better than humans and vice versa, develop complementary roles and responsibilities for each, and redesign processes accordingly. Executives need to identify where AI can create the most significant and sustainable advantage.

You need to ask four questions:

1. Is there a customer need?
2. How do we leverage technological advances?
3. Do we have access to data sources?
4. Can we decompose our processes?

First, understand the real needs of your customers. AI is an overhyped topic today, hence it always makes sense to return to the fundamental business questions. Do you have a view of your current or potential customers explicit or implicit unmet needs? Even the most disruptive recent business ideas, such as Uber and Airbnb, address people's fundamental requirements.

Second, incorporate technological advances. To do AI, you need to acquire and process new sources of data and push aggressive automation targets across your enterprise functions. The general availability of AI platforms and services can come to your aid, you can also take a different approach, of building your own AI infrastructure in house. The important thing is you need to put AI strategy on your board's agenda.

Third, create a holistic architecture that combines existing data with new or novel sources, even if they come from outside. The stack of AI services has become reasonably standardized and is increasingly accessible through intuitive tools. Even non-experts can use large data sets.

Finally, break down processes and offerings into relatively smaller and isolated elements that can be automated, instead of looking at a huge process flow and getting lost in the complexity. Once you have figured out a way to infuse AI supplements into these smaller processes, you can then do an end-to-end orchestration to optimize and automate the entire process flow.

For many organizations, these steps can be challenging. To apply the four questions systematically, companies need to have a clear picture of current and emerging capabilities of AI. A center for excellence can serve as a place to incubate technical and business acumen and disseminate AI expertise throughout the organization.

If you are the CDO of a company and are responsible for ushering in the AI transformative capabilities to your organization, what would you do? Besides the four questions that you must ask, you also need to focus on additional three areas, discussed next.

Develop a Clear Line of Sight to the Business Value

Start by assessing the relevance of AI from a business value in relation to specific operations and IT challenges.

Business value is an imperative. Many organizations become enamored with AI capabilities, but in the process they fail to determine the most strategic value drivers. You need to be certain about where to apply critical resources, such as data scientists and new solutions, that would benefit from AI and then firm up your plan to build capabilities where longer-term business outcomes are desired.

Expand your strategy with frameworks that will help you determine AI's applicability to business processes. Business process assessment frameworks establish a common language for describing your organization's existing business model. It also aids in assessing and proposing changes to individual components—improving cost structures, enabling data-driven revenue streams, or identifying new key partnerships where data and analytics can play a prime role. It also can you help identify changes to interrelated components that support potential extensive business model changes.

Harness Disruptive Potential in Customer Experiences

AI presents several opportunities for gaining insight, creating personalization, and enhancing the customer experience, which is one of the best opportunities for the use of AI and machine learning. Assessing its disruptive potential gives you the opportunity to engage customers in new ways, deepening your understanding of customer behavior and shaping the future of customer experience.

There are many opportunities to improve customer experience with AI, including developing customer insights and customizing their journeys, as well as predictive analytics for marketing. You'll need to leverage approaches such as journey mapping and outcome-driven innovation to identify unmet customer needs and opportunities.

Address Organizational, Governance, and Technological Impact

Prepare for the organizational, governance, and technological challenges imposed by AI. Lack of the necessary skills is often seen as a primary hurdle to AI adoption, so developing the necessary competencies will be critical. The obvious impact is with the development of data science skills and refactoring the CDO's organization to foster the creation and use of intelligence.

Many of the benefits of AI will come from the predictions rendered by machine learning. Yet, organizations are woefully ill-prepared to use these insights rather and evaluate and use probabilistic assessments of outcomes in decision making. This underscores the equal, if not greater, importance of developing a data-driven culture and the ability to "speak data" from a business perspective.

Using AI to gain insight into areas that humans can't underlies advancements in predictive analytics, natural-language processing, computer vision, image recognition, and many other displays of seeming intelligence. Numerous business scenarios certainly benefit from AI-generated insights and capabilities, but governing them may be a challenge due to a lack of transparency in how some of these approaches attain their results, a lack of processes to ensure quality results and appropriate use.

For example, it's possible that the same data with the same analytics may be governed differently based on use—one is ethically okay and the other is potentially not. The same may also be true for security, privacy, compliance, and retention.

To address these challenges, you need to develop a data-driven culture; be mindful of regulatory and ethical considerations; and steer clear of dangerous myths, all while fostering a learning laboratory for AI capabilities.

In this chapter, we discussed a variety of topics associated with AI, including some high-level pointers for developing your own AI strategy. In the next chapter, we discuss a very interesting question whereby the board asks the CEO to define the AI strategy.

References

1. https://www.bcg.com/publications/2017/competing-in-age-artificial-intelligence.aspx
2. https://sloanreview.mit.edu/article/ai-in-the-boardroom-the-next-realm-of-corporate-governance/
3. https://www.bcg.com/publications/2016/strategy-technology-digital-integrated-strategy-machine-using-ai-create-advantage.aspx
4. https://www.wired.com/story/amazon-artificial-intelligence-flywheel/
5. https://hbr.org/2017/10/how-ai-will-change-strategy-a-thought-experiment
6. https://psmag.com/news/artifical-intellgience-is-the-key-to-understanding
7. http://www.scmr.com/article/8_fundamentals_for_achieving_ai_success_in_the_supply_chain
8. https://www.gartner.com/smarterwithgartner/prepare-for-the-impact-of-ai-on-procurement/