ARTIFICIAL INTELLIGENCE IN MAINSTREAM ENTERPRISE APPLICATIONS

Douglas Picirillo[1,*]

[1] Trinity International University, Deerfield, Illinois, USA
*Corresponding author: Douglas Picirillo - doug.picirillo@gmail.com

ABSTRACT

This paper broadly surveys the present adoption of Artificial Intelligence in mainstream enterprise applications and envisions a future in which businesses and organizations of all sizes will derive value from AI-enhanced applications. In this context, Artificial Intelligence refers to a broad spectrum of technologies which have the potential to move the enterprise applications paradigm from one of replacing human effort with machine effort where the machine can simply outperform humans in terms of reliability and operations per unit time, to enhancing human activity by adding artificially intelligent machine activity to model and interact with real world complexity in ways that exceed human ability to synthesize and derive insights from the data tsunami under which many organizations are awash. The present and potential contributions of various primary forms of Artificial Intelligence are considered with respect to applications which drive marketplace innovation, applications which enable marketplace differentiation, and applications which simply enable and improve operations and recordkeeping. A framework is proposed for use by business and technology leaders and enterprise architects for maintaining alignment with organizational goals and values, identifying consumable Artificial Intelligence technologies, and create a value-oriented roadmap for prioritization and implementation.

---

1. What Do We Mean by Artificial Intelligence (AI)?

Ask ten people to define "artificial intelligence" and you will get ten different answers. This is, in some ways, beneficial. Firm definitions can be needlessly constraining to those whose role

it is to innovate. Or according to Stone et al. at Stanford University, "the lack of a precise, universally accepted definition of AI probably has helped the field to grow, blossom, and advance at an ever-accelerating pace." (Stanford, 2016). Nevertheless, to talk about a thing intelligently, it is worth a paragraph here to define it as it is used in the present discussion. The word "artificial" is relatively easy to define. We use it to refer to things that do not occur naturally, things that are created by humans. The word "intelligence" is a bit more difficult. Philosophers, psychologists, anthropologists, and even theologians offer a variety of definitions. For the purposes of this article, we will start with a definition which focuses on problem-solving offered by psychologist Howard Gardner. "Intelligence is the ability to solve problems, or to create products, that are valued within one or more cultural settings" (Gardner, 1983).

Thus, artificial intelligence or AI, as we shall use it here, is an ability, which does not occur naturally, to solve problems or create something of value. This is in the realm of narrow AI, where an intelligent apparatus is created to solve a particular kind of problem or create a particular type of valuable thing. General AI, which theoretically includes non-trivial foresight, independent problem identification, and problem prioritization based on something like a human value system, does not yet exist.

2. What Does AI Offer Humankind?

Information technology in general and AI in particular serve as assistants to humankind in problem-solving and value creation. Humans can articulate problems and propose solutions beyond their capacity to effectively execute the solution. That is, some solutions require computational capacity that exceeds human capacity in scale, complexity, or both.

It is important to distinguish between scale and complexity. Humans are more than capable of adding a column of numbers. Machines are better at scale, making short work of calculating totals and averages of thousands or millions of values. Humans are still better than AI at certain kinds of complex tasks when those tasks are performed on a small scale.

Amazon's Alexa, Apple's Siri, and Google's Assistant use a special kind of AI, Natural Language Processing (NLP), to recognize spoken commands and questions, and to provide synthetic replies. As complex as these tasks might be, they are not as complex as human conversation which considers not only words and grammar, but tone, and non-verbal and paraverbal cues. Human conversations take place in multiple simultaneous contexts. Words and grammar mean different things in different conversations. Thus far, NLP allows machines to process language. It does not yet enable them to converse as humans do, bringing feelings, perceptions, and imagination into the exchange of information. On the other hand, humans can converse with only a limited number of other humans at once at a limited pace. On the other hand, a machine's ability to tirelessly process long streams of natural language, even more than one at a time, is limited only by the technical architecture of the hardware and software on which it is realized.

Figure 1 positions some illustrative problem on axes of scale and complexity. The promise of AI is in the high-complexity, low-scale and high-complexity, high-scale quadrants.
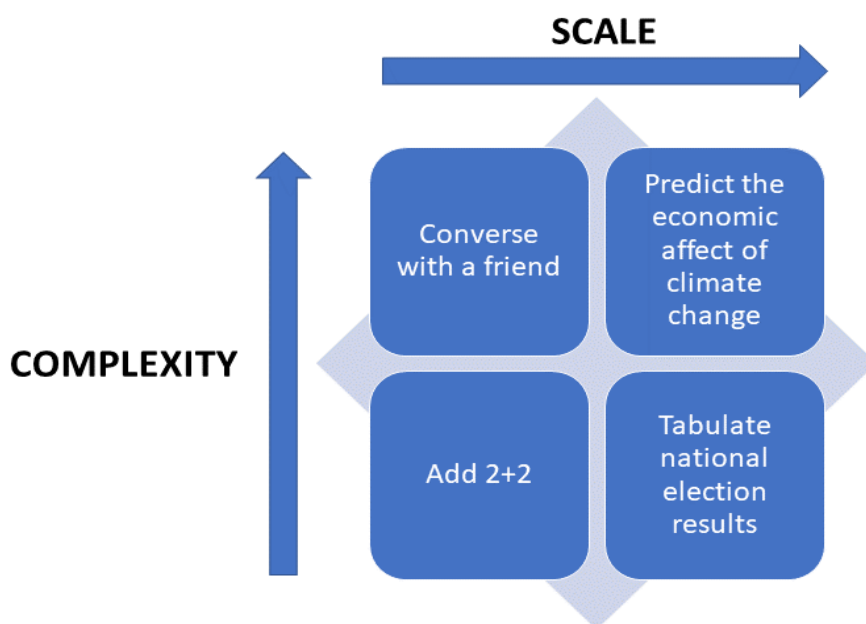

Figure 1. Scale and Complexity

This is what makes widespread applications of AI possible – exponential increases in ubiquitous computing capacity. Computers which filled entire rooms in the 1960s were capable of little more than arithmetic at scale. Today, billions of mobile phones contain the computing capacity required to perform the complex tasks which we refer to as Artificial Intelligence. Except where human "wiring" remains the better tool for some kinds of problem-solving, machines can outperform humans in both scale and complexity. This may seem obvious, but it is exactly the point of all technology. The first sharpened stone was a better cutting instrument than human fingernails. This is what we humans do. We are toolmakers. We create things that allow us to solve problems and create value better, faster, and more reliably than we can with our innate capacities. AI is just a new tool. It may be an exceptionally interesting tool because, more than a sharpened stone, we see a bit of ourselves reflected in it

3. What Does AI Offer Users of Mainstream Enterprise Systems?

Leaders and managers of businesses, non-profit organizations, government agencies and others have something in common. They have organizational goals and organizational stakeholders. Whether those stakeholders are owners, customers, social service clients, or citizen constituents, the organizational goals must align in whole or in part with solving a problem or

creating something of value for stakeholders. Ultimately, as AI becomes widely adopted, everyone stands to benefit. We are all consumers, citizens, and owners. We are all stakeholders.

Here again some definitions are necessary. By "enterprise systems", we mean solutions built of software applications, data, and supporting technologies which enable an organization to achieve its distinct goals and create value for its stakeholders. But we also include in this definition of system, the humans who use the system, the humans who derive some benefit from it, and the policies and operating procedures which govern the use of the technical solution. By "mainstream", we mean those solutions that are widely adopted or that have the potential for wide adoption yielding transformed operational processes across industry sectors.

The field of artificial intelligence (AI) is finally yielding valuable smart devices and applications that do   more than win games against human champions. According to a report from the Frederick S. Pardee   Center   for International Futures at the University of Denver, the products of AI are changing the   competitive landscape in several industry sectors and are poised to upend operations in many business   functions (Kiron, 2017).

It is tempting to omit computer games from this discussion. We might not think that they create real value for the stakeholders. But the consumers who pay for them clearly find value in them as a form of entertainment. Further, the broad genre of computer applications, games and simulations, which incorporate advanced AI technologies, create stakeholder value as training and development platforms. One need to look no further than commercial flight simulators which allow students and pilots of all levels of experience to train and learn without the risk of making inappropriate contact with the ground in a real aircraft with real people aboard.

For over half a century, enterprises of all kinds have adopted successive generations of information technology to create stakeholder value by increasing sales, reducing costs, producing widgets, or delivering services faster and more reliably. Organizations of all sizes today use a computer of some kind to house their Systems of Record – a system for the collection and management of transactional information, for operational needs and management reporting. Some organizations find ways to achieve competitive advantage through clever uses of Information Technology, often doing the same thing as their competitors, just doing it faster or cheaper. We can call this a System of Differentiation. A few companies change the marketplace through Systems of Innovation - innovative applications of Information Technology. A leading Information Technology research firm introduced this three-layer model in 2012, referring to it as "Pace Layered Application Strategy" (Gartner, 2012) because each layer promises to drive organizational change and accelerate value creation at a different pace.

The most well-known example of Systems of Innovation is an entire enterprise – Amazon.com. Beginning as an online bookseller, Amazon innovated so profoundly that they almost single-handedly killed brick and mortar bookstores. It is telling that they are not known as Amazon, Incorporated, or the Amazon Company. They are known by their Internet domain name. Today, their technology innovations have not only transformed how households and business acquire all manner of goods, their technology itself has become a market-disruptive service in the form

of Amazon Web Services (AWS). Further, AWS, along with Microsoft Azure, IBM Watson, and others are now commoditizing elements of their AI infrastructure providing an unprecedented opportunity for enterprises which are not among the "tech giants' to employ these technologies in their pursuit of stakeholder value.

Table 1 lists a sample of the AI services provided by AI that enterprises can integrate into their systems. According to Amazon, "you get instant access to fast, high quality AI tools based on the same technology used to power Amazon's own businesses" (AWS AI Services, 2020).

Table 1. Sample AWS AI Services (source: AWS AI Services, 2020)

Service Name	Service Description
Amazon CodeGuru	Automate code reviews to improve software quality
Amazon Comprehend	Discover insights and relationships in text
Amazon Fraud Detector	Detect online fraud faster
Amazon Lex	Build voice and text chatbots
Amazon Polly	Turn text into lifelike speech
Amazon Rekognition	Analyze image and video
Amazon Transcribe	Automatic speech recognition

Table 2 lists some of Microsoft's AI services provide under a product line called Azure Cognitive Services.

Table 1. Sample Microsoft AI Services

Service Name	Service Description
Content Moderator	Detect potentially offensive or unwanted content
Text Analytics	Detect sentiment, key phrases, and named entities
Translator	Detect and translate more than 60 supported languages
Speaker Recognition	Identify and verify the people speaking based on audio
Computer Vision	Analyze content images
Face	Detect and identify people and emotions in images
Video Indexer	Analyze the visual and audio channels of a video, and index its content

With access to exponentially more powerful computing capabilities and the capacity to cost-effectively generate and store vast amounts of data many organizations are just trying to make sense of the ever-increasing waves of transactional data created by operational activity. Even so, the goal posts keep moving. Organizations must do more than simply keep up with volume. They must find differentiating and innovative ways to extract value from the data, to intelligently steer the enterprise, one operational process at a time, through Systems of Intelligence.

Systems of Intelligence is a fourth layer not explicitly anticipated in the 2012 Gartner report in the sense that it offers the potential of something beyond what we traditionally call innovation. In Gartner's Pace Layers, we think of innovation as something humans do, enabled by technology. The human has a hypothesis about the market or a particular kind of stakeholder.

A System of Innovation, relying on machine learning still depends on a hypothesis formed by a human, but it offers the promise of validating or invalidating many hypotheses, quickly, throughout a complex enterprise.

Data warehouse analytics was once called innovative. Today, the ability to process massive amounts of data and produce meaningful insights are a mature technology – mature enough to be called "traditional". These technologies can create significant stakeholder value, but they suffer from a fundamental limitation. They are all backward facing. Industry research analyst David Floyer said this, "Traditional Data Warehouse Analytics empower the few to steer the ship by analyzing its wake." (Floyer, 2015) Traditional analytics most often deliver value as hindsight. Predictive analytics attempt to add value as foresight. But the accuracy of predictions can be diminished by age of the data. Analyses are often based on data which is days or weeks old. Results are published in reports and dashboard to managers and executives who then decide what to share. In contrast, the vision for Systems of Intelligence is real-time, or near real-time analyses, recommendations, and decisions, in the hands of the information producer and consumer. The idea is, process as much data as possible, as quickly as possible, with the best available algorithms to steer the ship intelligently, collectively, and organically with a forward outlook.

Machine Learning often dominates discussions about AI in which the topic is novel problem solving and modeling and deriving insights from massive and complex data. In the potential for widespread enterprise adoption however, several kinds of AI have the potential to create stakeholder value, separately or in combination.

- Natural Language Processing
- Image Recognition
- Machine Learning based analysis and forecasting
- Robotics

This is not an exhaustive list, but these four kinds of AI are relatively easy to define empirically through well-known examples. Stanford's One Hundred Year Study on Artificial Intelligence — "Artificial Intelligence and Life in 2030" — lists 11 applications of AI each of which are forms of narrow AI. Some of the most compelling applications, or use cases, employ more than one kind of AI. For reasons of space and simplicity, we will go only a little deeper into the found kinds of AI mentioned above.

Natural Language Processing (NLP) is the ability of a machine to process human communication in its natural form, written or audible, and generate human-understandable responses in a natural form. Apple's Siri, Google's Assistant, and Amazon Alexa are well known examples of the tech giants. Many other enterprises are today using some form of NLP. Airline reservation systems are a prominent example. In this case, the organization derives stakeholder value by providing customer service for routine matters at a lower cost than doing so with human agents.

Image Recognition appears in Google's reverse image search at images.google.com. Instead of typing a search term and receiving images as result, a user can provide an image and the technology tells you what it probably is, where to read about it and a selection of visually similar images.

The photo in Figure 2 (taken by the author in February 2020) returned the primary result "Western Wall" even though the most visually prominent object is not the wall but the Dome of the Rock and other structures beyond the wall. A review of the visually similar images suggests an explanation. May of the visually similar images were captured from the same direction but with a wider view. In them, the Western Wall is visible along with the Dome of the Rock. The images appear to have a relatively larger number of metadata tags which refer to the Western Wall rather than the Dome of the Rock. This reveals potential weakness of image recognition technologies. All AI must be "trained" in some way. That is, we give the machine examples and tell it how to recognize important things in the examples. In this instance of image recognition, users contribute images and add keywords or metadata about the image. Because many images in the training data containing the Dome of the Rock are tagged with keywords referring to the Western Wall, Google's algorithm says this is a photo of the wall, even though the wall is not visible. This answer is in the right neighborhood, in more ways than one, but it is not correct. A human observer would say this is a phone of the Dome of the Rock, not of the Western Wall. As impressive as the technology has become, it is not able to overcome poor training or a bias in the training data. Because of errors like this, humans need to supervise the machine's output, especially when that output is used to inform high-stakes decisions. With enough correction provided by human supervisors, the machine will learn more and produce increasingly accurate results.

Figure 1. The Dome of the Rock, Jerusalem (the author source)

4. Machine Learning

Technically, Image Recognition and Natural Language Processing and both forms of Machine Learning. They recognize images based on their training examples – images and metadata, audio waveforms and phonetic models which map waveforms to and from syllables and words. But Machine Learning is commonly used to refer to something different. Rather than data that requires mnemonics to connect data to a human-recognizable image or sound, Machine Learning more commonly focused on the examination of structured or unstructured data for the purposes of analysis and forecasting.

All forms of machine learning, by definition, incorporate algorithms which improve automatically through experience. That is, computations are based on trained models which continue to improve through additional training. When a model's prediction varies from observed reality, the divergence is used to refine the model. The machine learns.

Google describes an interesting development in this field, applying machine learning to the complex and ubiquitous task of weather forecasting. In the application of AI to weather modeling, one line of research involves "physics free" models. That is, forecasts are generated

from computational models which use purely statistical methods to examine the data which describes current conditions in the light of experience (training data). In this approach, the forecasting techniques are indifferent to the physics involved in understanding pressure gradients, heat, and humidity. The promise of Google's research is highly localized, real time precipitation forecasting. They call it "Nowcasting". Writing in Google's AI blog, Jason Hickey, a Senior Software Engineer at Google Research describes the opportunity to use Machine Learn for precipitation nowcasting from radar images.

A significant advantage of machine learning is that inference is computationally cheap given an already-trained model, allowing forecasts that are nearly instantaneous and in the native high resolution of the input data. This precipitation nowcasting, which focuses on 0-6 hour forecasts, can generate forecasts that have a 1km resolution with a total latency of just 5-10 minutes, including data collection delays, outperforming traditional models, even at these early stages of developmen (Hickey, 2020).

4. Robotics

Robotics can be thought of as the field where AI meets the real world. AI is used to intelligently control machines which do real things in the physical world. These machines can be part of a factory's assembly line, in a warehouse or distribution center, or in the field. The machines can be fixed or mobile. The intelligence can be onboard or remote. The robot can be capable of only a narrow range of physical tasks or it can have generalized capabilities. In a narrow example from 2004, the author was a member of a team which created labels for unit dose pharmaceuticals. In an application of image processing, the machine was trained to recognize and reject imperfect labels. This machine had a low degree of autonomy and a single fixed function. Drones and humanoid robots, by contrast have more generalized capabilities, mobility, and higher degrees of autonomy. A ubiquitous household example is floor-cleaning robots. They have limited function but high autonomy. They leave their base, move around the house, detect obstacles, clean the floor, and eventually return to their base to recharge their battery. Arial drones, such as those with military applications, have an elastic and hybrid form of autonomy. They can autonomously complete a mission, but they are typically under continuous remote supervision by a human operator. In civilian utilities, public safety, and emergency response applications, autonomous can be automatically dispatched to the scenes of storm damage, or other emergencies to provide preliminary imagery and data from other onboard sensors.

Robotics can be among the most complex applications of AI, incorporating real-time machine learning, image recognition, natural language processing, and physical devices to accomplish real world tasks. This is particularly true in the real of humanoid robotics.

Japan is fertile ground for innovative and widespread deployment of robotics. This began with rapid reindustrialization following the widespread destruction of World War II. Following its successes in industrial automation and proto-robots, Japan has been a leading exporter of industrial robotics. The Japanese may be more open, as a culture, to humanoid robotics than the West. Western attitudes are often shaped by Hollywood's stories. For many, images of the

Terminator, Robocop, and HAL raise fears of evil robots usurping human control over their own lives. Speaking to the BBC, Martin Rathmann, a Japan scholar at Siegen University in Germany, says Japanese "tend to imagine humanoid robots as intelligent, flexible, and powerful" and not as threats to human order. (Zeeburg, 2020)

Similar to the need to reindustrialize following the war, Japan now faces a new set of practical problems which encourage the use of robotics, humanoid and otherwise. Japan's population is aging and contracting. The median age is 48.4. Population has declined for six consecutive years at an accelerating rate. Japan's population is now declining by more than 1,000 per day. (Worldometer. 2020). This ss creating a growing workforce problem. More elderly need daily care. Fewer young people are available for entry level jobs. A BBC report by … offers several prominent use cases for humanoid robots in everyday Japanese life.

- Social robots that can detect – and predict – healthcare changes in people
- Cute, communicative robots in hotel rooms or restaurants able to assist guests in other languages
- Robots in classrooms across Japan to help teach English.

5. Industry 4.0

"Industry 4.0" is a term used interchangeably with "Fourth Industrial Revolution". According to Deloitte,

The term Industry 4.0 encompasses a promise of a new industrial revolution—one that marries advanced manufacturing techniques with the Internet of Things to create manufacturing systems that are not only interconnected, but communicate, analyze, and use information to drive further intelligent action back in the physical world (Deloitte, 2020).

In Industry 4.0, the Internet of Things (IoT) and other forms of connected robotics enable and enhance manufacturing and other physical processes on an industrial scale. "Industry 4.0" is largely a marketing term. The simple fact that it exists as an idea widely understood by business and technology executives signals the growing interest and adoption of smart technologies in mainstream enterprises.

5. Case Studies

To understand how mainstream enterprises are applying AI to their most complex problems, it is helpful to examine some case studies.

5.1 Case: Entertainment Industry. Machine Learning. Image Processing.

The Disney Company is committed to keeping its archive of nearly a century of drawings, concept artwork, and finished production pieces accessible to writers and animators for reference and inspiration. A team of software engineers and information scientists are building ML and deep learning tools to automatically tag content with descriptive metadata, identifying characters, graphical elements, and scenes. Using this technology, writers and animators can

search for and retrieve characters and their stories from Mickey Mouse to Captain Jack Sparrow. (AWS ML, Disney)

5.2 Case: Financial Services Industry. Machine Learning.

Capital One is one of the largest banks in the United States, and the largest digital bank. Capital One has adopted AI and ML solutions to the full spectrum of its business processes and customer experiences. In one application, Capital One is using ML technology to enhance its fraud detection capabilities. False positives in this area are annoying to customers and Capital One is using ML to reduce them. Their fraud detection models learn from human feedback to improve the detection model and reduce false positives in the future. (AWS, Capital One)

Nitzan Mekel-Bobrov, Ph.D., Managing Vice President of Machine Learning at Capital One spoke of the importance of ML to the company's services.

We've recognized over the past number of years the importance of leveraging machine learning to enhance the user experience, as well as to help us make more informed decisions around engaging with our customers (AWS ML, Capital One, 2020).

5.3 Case: Healthcare Technology Industry. Machine Learning. Image Processing.

One company at the forefront of change is GE Healthcare. In recent years, the company has embraced machine learning as a driver of better patient outcomes, with applications ranging from data mining platforms that draw on patient records to analyze quality of care to algorithms that predict possible post-discharge complications.

As part of its investment in machine learning, the healthcare technology company partnered with clinicians at the University of California, San Francisco to create a library of deep learning algorithms centered around improving traditional x-ray imaging technologies like ultrasounds and CT scans. By incorporating a variety of data sets—patient-reported data, sensor data and numerous other sources—into the scan process, the algorithms will be able to recognize the difference between normal and abnormal results.

Keith Bigelow, General Manager of Analytics, at GE Healthcare says intelligent devices and applications promised improved quality, access, and efficiency, all of which presumably contribute to advances in patient outcomes while also reducing costs.

The more intelligence we can put into medical devices and applications, the more we can increase the quality. It's going to improve access. It's going to improve efficiency and it's going to reduce costs all at the same time« (AWS ML, GE, 2020).

6. COVID-19

No discussion of AI would be complete without some mention of its role in the battle against the COVID-19 pandemic. Every news outlet provides daily reports on case and mortality statistics. An interesting social byproduct may be increased public understanding of some basic

techniques for finding meaning in statistics. Almost everyone has heard about "flattening the curve". People who have never taken a statistics course are also learning about the importance of rolling averages, leading and lagging indicators, and the difference between causation and correlation. Even if they do not understand the underlying mathematics, even if they don't understand how it works, they appreciate the value of predictive analysis.

Writing for Forbes, AI and Big Data Analyst Steve King, describes both dilemmas and opportunities created by the pandemic. According to King, "recent social distancing and isolation practices in place across the globe, mean it has become even harder, if not impossible to keep trying to do so through traditional primary research techniques." At the same time, novel solutions are being developed using the massive data generated by Social Media platforms. Social media offers access to consumer sentiment and related data far beyond the largest focus groups. For example, Sentiment Analysis (based on Natural Language Processing) and other AI-based techniques may offer insights based on real time conversations which may in turn help us understand the impact and spread of coronavirus and some of the societal impact, such as remote work, sense of wellbeing, attitudes about face coverings and more. (King. 2020) This kind of innovation has the potential to influence not only public policymaking processes but market research techniques and voice-of-customer interactions across a wide variety of consumer-facing industries as well.

7. Current Levels of Adoption

Industry analysts agree that AI is becoming increasingly common in mainstream enterprise applications. They differ on the maturity of the solutions, the technologies, and the state of organizational capabilities to derive the most value from AI-based solutions.

In survey results published in March of 2020, O'Reilly Media Inc, an independent technology research and publishing organization, found that the "majority (85%) of respondent organizations are evaluating AI or using it in production. Further, more than half of respondents self-identify as mature adopters of AI. Other findings contradict the maturity claim. Respondents cite numerous factors such as lack of skills as an impediment to adoption but almost 22% cite lack of institutional support as the most significant impediment. This means that organizational leadership at the strategy level does not yet understand or accept the opportunities presented by AI to deliver stakeholder value in excess of implementation costs. Few organizations have implemented formal governance of AI efforts. This means that the unique risks of AI efforts may go unrecognized and uncontrolled. (Magoulas. 2020).

Writing for InformationWeek, James Kobielus takes issue with some of O'Reilly's findings but more importantly identifies additional questions raised by the survey and accompanying analysis.

Artificial intelligence's emergence into the mainstream of enterprise computing raises significant issues -- strategic, cultural, and operational -- for businesses everywhere (Kobelius. 2020).

Kobelius says that a key aspect of immaturity is the lack of established practices (formal governance) for mitigating the risks associated with introducing AI into routine business operations. Among these risks, the lack of a response strategy for unexpected outcomes or predictions is the most worrisome to survey respondents. Potential problems in the statistical models such as bias, degradation, interpretability, reliability, and reproducibility raise the possibility of misguided or suboptimal decision support and automation outcomes. Other governance-related risks include data privacy and information security weaknesses.

8. Drivers and Frameworks for Adoption

From social research to industrial automation, organizations having interest in AI must begin with an organizational in mind. AI implementations are becoming less expensive, but they require significant investments of money, effort, and management attention over time. They must result in stakeholder value. They must solve a problem or create value in some other way for the organization's owners, employees, customers, or citizen constituents.

The realization of the AI-enabled value further depends on the organization's ability to effectively deliver it to stakeholders. The solution must be accessible. The experience of accessing the solution must be acceptable and sustainable. The total cost of delivery must be met or, preferably, exceeded by the value delivered.

The solution must be more than accessible, experientially acceptable, and financially sound. It must also meet quality expectations. Its function must be complete, accurate, and most of all, relevant. An irrelevant solution – a solution for the wrong problem, no matter how well executed, is just a waste of time and money. Unfortunately, the history of commercial activity is replete with examples of products that nobody wanted. Something went wrong in the conversations about stakeholder value. There is no reason to believe that this cannot happen with AI-based solutions. Finally, the solution must perform well. Nobody wants to ask Siri a question and get the answer an hour later. Solutions must be available when they are needed. The must complete within specified periods of time, producing results at the rate of stakeholder consumption. They must do all of this repeatedly and reliably. In particular, through feedback and governance mechanisms, they must avoid the special risks of AI noted earlier. Figure 3 illustrates these relationships.

Figure 2. Drivers and Frameworks for Adoption (author source**)**

9. The Time Is Right

For a time now measured in generations, organizations have been accumulating data about their operations, their stakeholders, and their environments. In most cases, they have not had the tools to do anything truly interesting with it. Now with the growing commoditization of AI, the time is right for organizations to join the data-driven computing revolution.

Organizations can intelligently leverage available data to measurably improve services and products with little technology or implementation risk, and within a reasonable budget and schedule. Organizations often see their data as little more than a resource that needs to be managed. What they have instead is data that is highly relevant, detailed, and an extremely valuable historical record of core organizational processes. This data faithfully and accurately captures the entire service delivery and production lifecycles in thousands or millions of statistically useful observations.

The most powerful use of historical data is to understand the patterns hidden therein and extrapolate those patterns over time to suggest future reality. Data analytics comes in different flavors:

1. **Descriptive Data Analytics** – Aggregation and transformation of data into reports and visualizations (such as dashboards, charts, etc.). This is traditionally called "business intelligence".
2. **Diagnostic Data Analytics** – Built on descriptive analytics, diagnostics go deeper to discover explanations or root causes for known patterns.

3. **Predictive Data Analytics** – Identification of patterns in the data; Leverages machine learning (ML) methods or data science to create models that fit the current data patterns and that can be extrapolated to "predict" new outcomes based on new data.

4. **Prescriptive Data Analytics** – Predictive analytics integrated into a business process to provide users guidance on what to do next based on a given situation. Predictive analytics offer best or likely choices among available options. ("Customers who bought product X also frequently bought product Y" or "Based on your browsing history, you might like product Z" or "Customers like this had a 20% higher likelihood of buying product A when offered service B".)

10. The Power of AI in Business Processes

Unlike predictable transactions which can be modeled as a standard business process and enabled by an electronic workflow, ML can end up uncovering surprises hidden in the data. As more data accumulates and increases in quality, ML methods can deliver even better results. Most data-driven business applications today use Supervised Machine Learning (SML). In this model, there is a human subject matter expert (the "supervisor") who provides guidance on the "right" outputs for a chosen subset of the available data called the "training set". Today's technologies make it very convenient to apply different ML models to the training data set and can then predict the outputs for new inputs that are NOT in the training set.

11. Conlusions: The Journey Ahead

Adopting ML is not an overnight trip. It is a journey during which the critical stakeholders get accustomed to listening to the story being told by the data already in their custody. Ultimately, all ML outputs must be reviewed and validated continuously (or frequently) to ensure that the model is performing well.

In the private sector, two things matter most, sales results and the cost of generating those sales. Over time, a hypothetical ML-based sales assistant application could learn and improve its pre-sales guidance by comparing its guidance to the actual sales results. Likewise, it could learn and improve its recommendations for any activity related to sale, delivery, or post-sale service – any activity which generates data in sufficient volume to create a model. This is a fundamental and prerequisite question for all organizations interested in ML. Do we have the data?

Most enterprises, public sector, and not-for-profit organizations have the data. The tools needed to something innovative, something that creates new stakeholder value, are increasingly affordable and well understood. As a matter of strategic planning and confronting both the known and unknown challenges of the increasingly complex environment in which they operate, all organizations should consider what problems they can solve which were previously unsolvable. What value can they create that was previously out of reach because of complexity and scale?

The definition of AI itself will surely continue to evolve. The bar for what constitutes success in AI adoption is continuously being raised. As the technologies continue to improve and become more widely available, leaders and managers must understand the implications for organizational strategy, AI governance, the impact on the workforce and organizational design, total implementation costs, stakeholder value propositions, ethical questions and data privacy requirements.

References

AWS ML, Disney (2020), "*With Deep Learning, Disney Sorts Through a Universe of Content*", Amazon Web Services (AWS) Machine Learning, available at: https://aws.amazon.com/machine-learning/customers/innovators/disney/(accessed 8 June 2020).

AWS ML, Capital One (2020) "*At Capital One, Enhancing Fraud Protection With Machine Learning*", Amazon Web Services (AWS) Machine Learning, available at: https://aws.amazon.com/machine-learning/customers/innovators/capital_one/ (accessed 8 June 2020).

AWS ML, GE (2020), "*GE Healthcare Drives Better Outcomes With Machine Learning", Amazon Web Services (AWS) Machine Learning*, available at: https://aws.amazon.com/machine-learning/customers/innovators/ge_healthcare/ (accessed 8 June 2020).

Deloitte (2020), "Industry 4.0", *Deloitte Insights*, available at: https://www2.deloitte.com/us/en/insights/focus/industry-4-0.html (accessed 12 June 2020).

Gardner, H.(1983), *Frames of Mind: The Theory of Multiple Intelligences*. Fontana Press, New York.

Floyer D. (2015), "*Managing the Data Tsunami with Systems of Intelligence*", Wikibon, Inc., available at: https://wikibon.com/managing-the-data-tsunami-with-systems-of-intelligence/(accessed 12 June 2020).

Gaughan, D. (2012), "*Introducing the Pace Layered Application Strategy Special Report*", Gartner Inc., available at: https://blogs.gartner.com/dennis-gaughan/2012/02/03/introducing-the-pace-layered-application-strategy-special-report// (accessed 13 June 2020).

Hickey, J. (2020), "*Using Machine Learning to "Nowcast" Precipitation in High Resolution*", Google available at: https://ai.googleblog.com/2020/01/using-machine-learning-to-nowcast.html (accessed 13 June 2020).

King, S. (2020), "Why COVID-19 Is Accelerating The Adoption Of AI And Research Tech", *Forbes*, 11 May, available at: https://www.forbes.com/sites/steveking/2020/05/11/why-covid-19-is-accelerating-the-adoption-of-ai-and-research-tech/#2983b7162140 (accessed 13 June 2020).

Kiron, D. (2017), "What managers need to know about artificial intelligence", *MIT Sloan Management Review*, 25. January, available at: https://sloanreview.mit.edu/article/what-managers-need-to-know-about-artificial-intelligence/(accessed 11 June 2020).

Kobelius, J. (2020), "Enterprise AI Goes Mainstream, but Maturity Must Wait", *Information Week*, 31 March, available at: https://www.informationweek.com/big-data/ai-machine-learning/enterprise-ai-goes-mainstream-but-maturity-must-wait/a/d-id/1337428(accessed 11 June 2020).

Lufkin, B. (2020), "*What the world can learn from Japan's robots*", BBC, 6 February available at: https://www.bbc.com/worklife/article/20200205-what-the-world-can-learn-from-japans-robots (accessed 11 June 2020).

Magoulas, R. and Swoyer, S. (2020), "*AI Adoption in the Enterprise*", O'Reilly Media, Inc, 18 March Available at: https://www.oreilly.com/radar/ai-adoption-in-the-enterprise-2020/(accessed 11 June 2020).

Stone P., Brooks R., Brynjolfsson E., Calo R., Etzioni O., Hager G., Hirschberg J., Kalyanakrishnan S…, and Teller A. (2016), "*Artificial Intelligence and Life in 2030*". One Hundred Year Study on Artificial Intelligence: Report of the 2015-2016 Study Panel, Stanford University, Stanford, CA available at: http://ai100.stanford.edu/2016-report (accessed 17 June 2020).

Worldometer (2020), "*Population of Japan (2020 and historical)*" available at: https://www.worldometers.info/world-population/japan-population/ (accessed 17 June 2020).

Zeeburg, A. (2020), "*What we can learn about robots from Japan*", BBC, available at: https://www.bbc.com/future/article/20191220-what-we-can-learn-about-robots-from-japan (accessed 17 June 2020).