Silver Tree
Consulting & Services

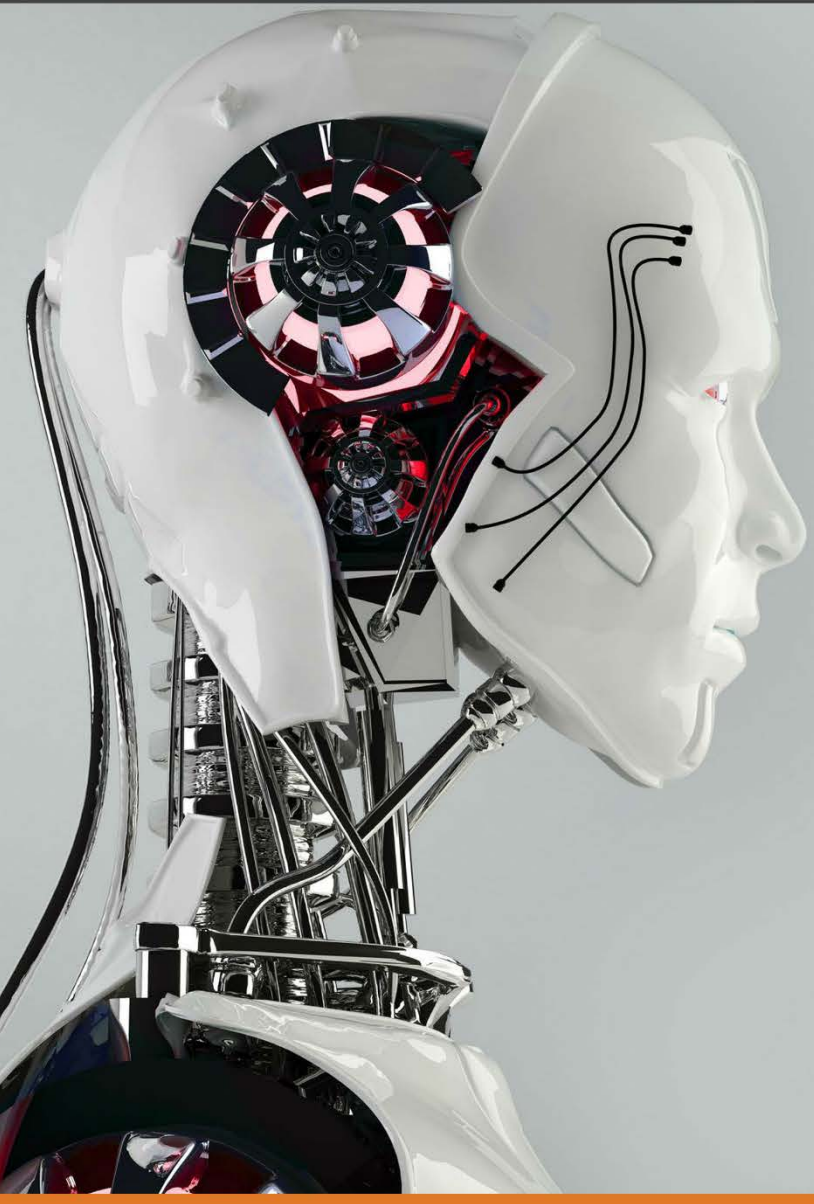

What is AI?

How do I get some of that?

John Petronzi
Advisory Board & Practice Lead,
Intelligent Automation

What is AI? How do I get some of that?

Consider this observation from Warren Bennis, a management consultant:
"The factory of the future will have only two employees, a man and a dog. The man will be there to feed the dog. The dog will be there to keep the man from touching the equipment."

I'm frequently asked by clients "What about AI?", as if it were a product they can buy and plug into their enterprise. Artificial Intelligence (AI) is not a product. You can't buy it off the shelf and configure it. Why would you want it anyway? What would you do with it if you had it?

Ahh, now that is the question! What would you do with it if… ? Without the answer to that question, forget about AI – you don't need it.

The Four Components of AI

AI is something you design to solve very complex challenges that require forms of reasoning approaching (or exceeding) human capabilities. AI becomes relevant and applicable where:

- Decisions must be made at light speed
- Ingesting and analyzing vast amounts of data is important
- Forever learning from the outcomes of previous decisions made by others, or from the AI solution itself, differentiates the solution from something more static.

Those problems, more often than not, address ways to leapfrog the competition by fundamentally transforming a company's business model and anticipating its customer expectations.

Creating Artificial Intelligence solutions typically relies on four primary components:

- Data, data and more data from traditional and non-traditional sources (e.g., social media, the Internet of Things, public research)
- Algorithms that interpret that data in ways that reasoned decisions can be made, problems can be solved, opportunities can be pursued, actions can be predicted or taken, intents can be perceived, learning can happen…
- Technologies that enable actions, understand the environment, interpret audio and text, categorize data, source data, automate processes, integrate systems …
- And, most importantly, an objective with sufficient complexity and strategic importance to justify the cost and effort with clarity of what constitutes success.

Figure 1 - Our Point of View on AI Solution Components

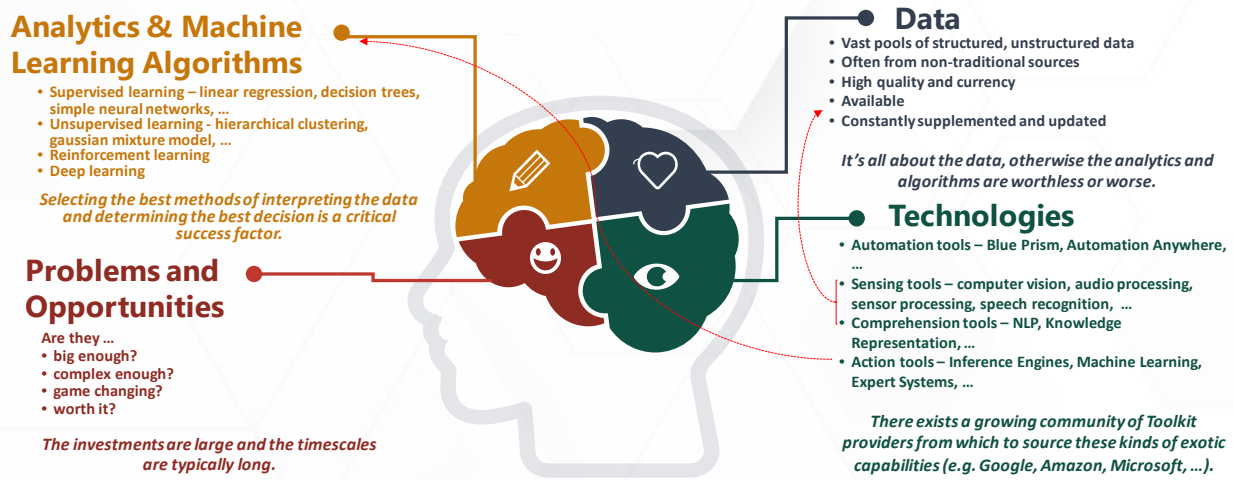

Problems and Opportunities – Where Are You Most Vulnerable?

Identifying a challenge to start with is really important. It's all about innovation, business innovation, and using the latest technologies to enable that business innovation. Conceiving of a challenge that's significantly complex to warrant the cost of developing an AI solution is hard. Fortunately, there are plenty of trailblazers who have developed business Use Cases that demonstrate how such Use Cases can be solved using AI enabled technologies.

And don't forget that those who are trying to disrupt your business model are aiming at the most profitable areas of your business, not where you're losing money. So, a focus on where you are vulnerable to disruption, not necessarily where you can reduce costs, may be the best place to start. Where you start with an innovative Use Case is THE most important challenge you will face when trying to deploy AI solutions in your enterprise.

Let me repeat this – don't think about AI until you have a problem to solve, an opportunity to pursue, or a really cool and disruptive end state as your goal.

Data - There is Plenty, but…

Once you have determined the challenge or opportunity on which to focus, it is time to figure out what kind of intelligence you have and don't have, within or external to your enterprise, that can be used to make your solution compelling. Don't assume that all the data within your enterprise is by definition "structured". There will be plenty of opportunities to turn unstructured data sitting on your email servers or in spreadsheets and PDFs on your file servers into usable data for purposes of supporting an AI solution.

Questions relating to your data, whether internally or externally sourced, include:

- Are there multiple sources of data from multiple systems that require reconciliation?

- What is the quality of that data?
- Can it be trusted?
- Is the data sufficiently current?
- Can it be made available to an AI solution in real time?
- Do we have a plan to maintain the currency and accuracy of the data?

Externally sourced data may be mined or scraped from a variety of sources, such as:

- Internet of Things (IOT)
- Social Media (e.g., Facebook, Twitter)
- Subscription based research libraries
- Publicly available government databases

The fact is that data is ubiquitous. Many enterprises may think that they need to collect their own data, but that is simply not true. There are thousands of free data sets available and ready to be used. Forbes published a list of over 30 examples of data sets that can be used for Big Data analyses. Some of the examples are:

1. US Census Bureau http://www.census.gov/data.html A wealth of information on the lives of US citizens covering population data, geographic data and education.
2. The CIA World Factbook https://www.cia.gov/library/publications/the-world-factbook/ Information on history, population, economy, government, infrastructure and military of 267 countries.
3. Amazon Web Services public datasets http://aws.amazon.com/datasets Huge resource of public data, including the 1000 Genome Project, an attempt to build the most comprehensive database of human genetic information and NASA 's database of satellite imagery of Earth.
4. Facebook FB +4.43% Graph https://developers.facebook.com/docs/graph-api Although much of the information on users' Facebook profile is private, a lot isn't – Facebook provide the Graph API as a way of querying the huge amount of information that its users are happy to share with the world (or can't hide because they haven't worked out how the privacy settings work).
5. Face.com: A fascinating tool for facial recognition data.
6. Google Public Data Explorer includes data from world development indicators, OECD, and human development indicators, mostly related to economics data and the world.
7. Junar is a data scraping service that also includes data feeds.

Yes, we are awash in data. Knowing where to look for it is the real challenge. Most enterprises have a Chief Data Officer or data scientists on staff to address getting the data, ensuring its

integrity and availability, and using it effectively. These are critical success factors associated with launching AI based solutions.

Analytics and Machine Learning Algorithms

I have combined analytics tools and algorithms under this banner, but analytics could have also been placed in the next section under Technologies. The spectrum for analytics is descriptive, diagnostic, predictive, and prescriptive analytics. The last one is the holy grail – optimization decisions.

Analytics typically rely on algorithms and statistical modeling techniques to leverage historical data in order to identify anomalies and outliers and to predict future outcomes. Examples of analytical methods are regression, correlation, segmentation, clustering, and association analysis.

The line between analytical tools and algorithms is a bit blurred. Analytical tools often rely on or are based on algorithms. Many of these algorithms are available from Open Source libraries. Many analytical tools examine past data sets to draw conclusions and insights from that data. They also usually facilitate the process of inspecting, cleaning, transforming, and modeling data with the goal of discovering useful information for insights, suggesting conclusions, and supporting decision making.

Predictive analytical tools do more. They use past data and sophisticated algorithms to predict what could happen next. By changing some of the variables in the algorithms enterprises can process what if scenarios and perform sensitivity analysis as the variables are adjusted. And many of these algorithms have embedded machine learning capabilities that compare what was predicted in the recent past to the actual outcomes in order to refine the algorithm's accuracy in the next round. A recent Scientific American article (March 2018) titled "*Self-Taught Robots*" illustrates the Cascade Prediction process. This process is actually how a child learns to interact with the world around it.

Prescriptive analytical tools go to the final step. They answer what the best choice is or action to take. Optimization can involve arguably the most powerful analytical tool – linear programming embraced by the operations research community. Companies like River Logic offer prescriptive analytic solutions for many common challenges or goals facing companies in a variety of business verticals.

In McKinsey's "*An Executive's Guide to AI*" (March 2018) they provide a handful of examples and the types of algorithms used to produce superior outcomes from the potential options. For example, to predict power usage in an electrical distribution grid there are algorithms known as Random Forest that use regression models to improve on the accuracy of simple decision trees and taking a majority vote to predict the output. Another example cited uses a Hierarchical Clustering algorithm to form a classification system in order to cluster loyalty card customers into progressively more micro-segmented groups for target marketing campaigns. These two

examples, use algorithms that learn, sometimes from their human masters and sometimes from the patterns of data they observe.

Into.AI is a "Wiki-like" website that provides a comprehensive and independent home for Artificial Intelligence and Machine Learning information. They have sections that address algorithms, software tools and data sets as well as other resources for those who are interested in digging deeper into the world of AI or contributing learnings to the website.

For those of you who might want a definition of Machine Learning, I found the following to be an excellent definition – "*Machine learning is a technology used to enable computers to analyze a set of data and learn from the insights gathered. By using complex algorithms an artificial neural network is simulated that enables machines to classify, interpret and understand data and then use the insights for solving problems or making predictions. Once programmed a machine learning algorithm improves and enriches itself based on the data fed to it.*" This definition is from a New Gen Apps blog (https://www.newgenapps.com/blog/machine-learning-vs-predictive-analytics), a mobile application developer.

When someone tells you that their product is a Machine Learning tool or has Machine Learning embedded, ask them to explain more to see if it sounds like it fits the definition.

There are many Machine Learning frameworks that you can leverage once you understand what you are trying to accomplish. Google (Tensorflow), Amazon Machine Learning, Microsoft Distributed Machine Learning Toolkit (DMTK) are three of the more visible ones. But there are many others, with new ones sprouting up every week – Caffe2, H2O, Deep Learning 4j among them. Infoworld produced a list of 13 Machine Learning framework companies in an August 2017 posting titled "*13 frameworks for mastering machine learning* - *Venturing into machine learning? These open source tools do the heavy lifting for you.*"

Technologies – Sensing, Comprehending, Acting and now Automation

Finally, there are many technologies that enable designing, building and launching an AI solution. They are generally divided into groups – sensing, comprehending and acting technologies - and are deployed in integrated or "mashed up" configurations. I have also added a fourth technology to this group – automation technologies, that allow for these other technologies to move from one step in the process to the next at light speed. These technologies, once configured to solve a problem or enable a capability, facilitate a variety of applications, many of which we experience today.

For example, users of Amazon or Netflix, are the subject of what is known as a "recommendation system" that provides "people like you" suggestions for products to buy or entertainment to watch.

Every time you access an electronic device with your fingerprint you are using one of these AI technologies to recognize (sense) your fingerprint and open (act) on that successful recognition by opening the device. And who hasn't used Siri or Alexa to ask a question (comprehend) or direct and action (act)?

I normally break these technologies down into four categories as illustrated in Figure 1:

- **Sensing tools** – technologies that facilitate computer vision, audio processing and sensor processing
- **Comprehension tools** – technologies such as Natural Language Processing (NLP), knowledge representation tools, sentiment analysis tools
- **Action tools** – includes tools like inference engines, machine learning engines and expert systems
- **Automation tools** – we have seen a rapid maturing of this segment with companies like Blue Prism, Automation Anywhere, Work Fusion and others

Although automation tools could be considered a subset of action tools, I prefer to classify it separately, because these tools are capable of ingesting information from sensing, comprehension and action tools and completing a series of actions based on rules or advanced cognitive capabilities built into the tools.

Consider this example of retailer. In a call center the telephony system recognizes the inbound phone number of the customer. It routes the call to a call center representative who views on their computer a deal or offer optimally priced based on algorithms initially designed by the data scientist and applied by AI.

Interaction in an AI Ecosystem

Another area we should discuss is the relationship between these technologies, the Data Sets used to enable AI solutions, and the Analytics and Algorithms discussed earlier. Action tools are basically algorithms packaged to meet a decision / act challenge. Sensing and Comprehension tools contribute to, refine and correct the data sets used by an AI solution. They constantly improve the quality and currency of your data with fresh data or better interpretations of the data you have. And as discussed under the Analytics and Machine Learning Algorithms section above, a rapidly growing community of open source and closed source tool kits are available to select the best tool to meet your AI solution needs.

Debugging AI Solutions. Let's not forget that all these tools are computer software programs, ultimately conceived and written by humans who sometimes make errors or don't anticipate a set of conditions that could occur, even if exceedingly unlikely. So testing these complex mash-ups is difficult - to say the least. And once launched, an AI solution could go horribly wrong. For example, when a self-driving car kills a pedestrian or driver. There is a growing recognition of the importance of "deep" debugging of "deep" learning systems.

In Scientific American's March 2018 issue, there was a short article titled "*I'm Sorry Dave*" (referring to HAL in 2001 Space Odyssey) that discusses this challenge and some emerging work being done at Columbia University by Dr. Junfeng Yang and his colleagues to develop DeepXplore, a system meant to debug AI systems.

And don't forget about "Dave". You will still need people to handle the really complex issues or to audit the veracity of the results, or to interact with customers on a personal basis, something HAL wasn't so good at. The "people" part of your business model will change substantially – roles and responsibilities, skills, staff size – all must be considered as you plow forward into an increasingly AI-enabled business model.

Summary

So, what is AI? My best definition is "*AI is a combination or mash-up of technologies, operating on data sets, to develop capabilities used to solve complex problems that require human-like reasoning and decision-making at light speeds".*

Or as McKinsey has defined it in their "*An executive's guide to AI*" - "*AI is typically defined as the ability of a machine to perform cognitive functions we associate with human minds, such as perceiving, reasoning, learning, and problem solving. Examples of technologies that enable AI to solve business problems are robotics and autonomous vehicles, computer vision, language, virtual agents, and machine learning."*

My advice with AI, as with less complex automation deployments is to work backwards with the end in mind – start with the opportunity, not with the tools or the data. Until you can define your objectives you can't know the right tools to use, or what data you need to support meeting your objectives. And AI challenges or opportunities are generally complex enough, and likely to affect so many areas of your business, that stakeholder involvement and consensus is an absolute must. These solutions, often designed to disrupt competitors, also have the possibility of disrupting your own business model.

Conclusions

My suggestions for those enterprises that are serious about starting on this journey are:

- Start with the end state in mind – a business innovation goal that is capable of leap frogging your competition. This becomes your "north star".
- Gain maturity by deploying automations that drive value. Don't start with AI until you have learned about how it is likely to play in your culture, your industry, within your current business model. Starting with AI is likely to be a bridge too far for most enterprises. Even the moon shot started with several planet orbiting excursions.
- Establish a "sandbox" for R&D and business innovation, but don't staff it with just techies. Ensure that the R&D group has clear goals but can withstand failures and learn from them.

- Beef up your organizational and technology change management capabilities because AI deployments are likely to change everything.
- And finally, don't operate in a vacuum. To the extent it is practical, create a stakeholder community with whom to communicate and solicit advice and counsel. The best ideas sometimes come from the least likely sources.

Silver Tree
Consulting & Services

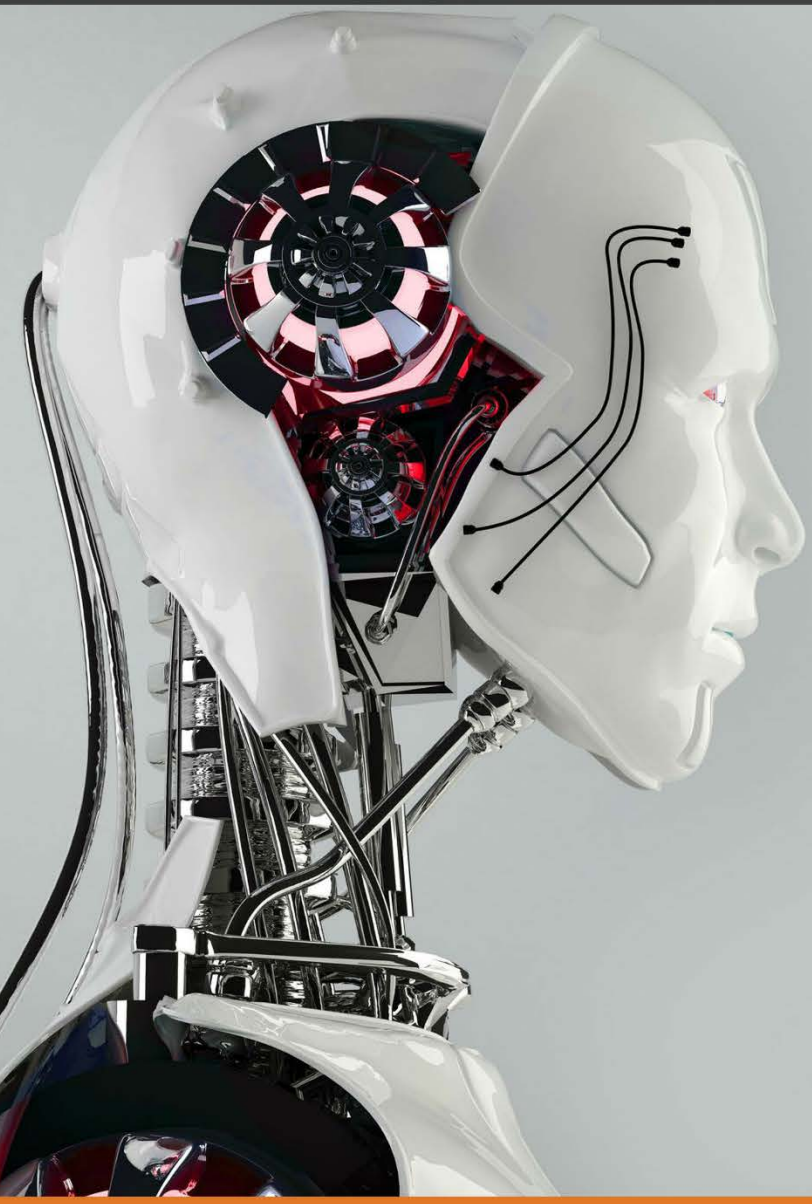

John Petronzi

Advisory Board & Practice Lead,
Intelligent Automation

John.petronzi@silvertreeservices.com
Mobile +1 (917) 353-7470

silvertreeservices.com