AI in Supply Chain

Some Definitions

Plenty of obtuse definitions exist on the web,
but none of them tell you what AI does for Supply Chain.
These definitions definitely will.

AI in SUPPLY CHAIN—SOME DEFINITIONS

Introduction

We hear and read over and over that AI is a technology designed to work like the human brain. *No, No, No!* That is not exactly right. The smell of fresh rain. The feel of those 400 thread count sheets. The taste of fresh hot bread and butter. The sense of accomplishment. These are all processed by the *human* brain. Of course, in the broadest sense, creating a way for a machine to *mimic these responses* is a goal of AI, but we know the *actual organic experience* and emotions stored about these experiences is very different from digital information.[1]

Part of the problem with the public definitions is they don't really work for Supply Chain. And so much of the press is about consumer apps, call centers, HP systems, etc. For our purposes, as supply chainers who are actually trying to decipher what AI[2] is—what it's good for and why we need it—we require practical definitions.

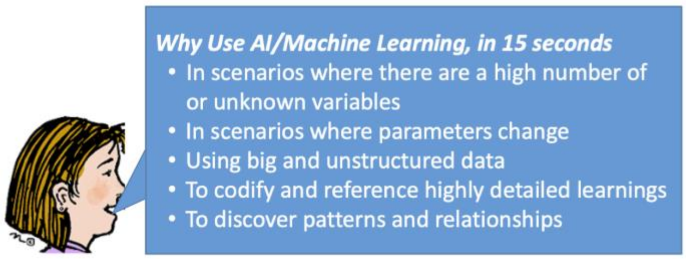

In this post, therefore, we attempt to define AI in practical terms for supply chain professionals. We will stay focused on applicability in supply chain,[3] since AI and Machine Learning is a huge domain with many types of technologies and applications. To cover broader topics requires a series of textbooks.

Today, it is important to learn a little about AI and where to apply it. In the future, AI technologies will already be within the supply-chain applications and analytics you use every day such as planning, demand management, logistics, supplier management, and so on, which will just keep getting better. So, let's get started.

Artificial Intelligence for Supply Chain, an Overview

AI is an umbrella set of technologies, from robotics to analytical systems. Within AI we have various subgroups such as machine learning, deep learning, natural language processing, robotics, and so on.

There are a lot of myths about AI due to Hollywood. But AI systems are not unconstrained or spontaneous. There is no HAL from *2001*, or Number 5 from *Short Circuit* or AMEE from *Red Planet* going rogue and taking on human characteristics of feelings and actions.[4] Humans have defined the programs that operate AI systems and, therefore, set the boundaries;[5] the meaning and relevancy of inputs; and then, based on those, how it can operate. We humans, on the other hand, can draw knowledge from experiences beyond some pre-set boundary.

---

[1] A human will *enjoy* the silky feel of the sheets and may decide, probably after one encounter, they prefer 400 count. The computer will need a repetitive test—or training routine—to statistically record that given a choice, people consistently select the higher thread count.

[2] There are also huge economic and social implications of AI which are relevant to supply chain professionals, but that is a topic for another time. For more on AI and the implications read here: *"Thinking Machines, or People?"* and *"The Supply Chain Planning Department of the Future"* and *"Will AI Bring Another Gilded Age?"*

[3] In future posts we will provide use cases.

[4] This article "*How Far Are We From Achieving Artificial General Intelligence?*" says most experts believe AGI is at least 40 years away.

[5] Today's AI systems provide "narrow AI" (aka applied AI), i.e., designed to solve a specific problem. We do not yet have Artificial General Intelligence, i.e., a machine that can learn any intellectual task that a human being can.

The story about data is essentially intertwined with AI algorithms. Within the AI umbrella we have an adjacent field of *data science*, the ability to understand various types of data—analog and digital, structured and unstructured—taking us beyond the "tables and rows" of our digital systems. IoT sensors, movement, voice, pictures, video, social, and so much more can be incorporated into the data sources that AI systems often use. And that is one of the great benefits of AI: being able to make use of the plethora of data sources, which keep accelerating in growth,[6] the sheer size of which traditional systems are not equipped to handle. Through these new sources, we have innovative new processes that make us more aware.

Finally, it is important to make some distinction between AI methods—those that are algorithmic driven (picking a "math" routine to find an answer) and those that are more data driven (making sense out of data to find patterns and so on). This will become more clear as you read through the definitions.

> **To note:** even the AI/Machine Learning geeks often have trouble clarifying these terms. So, we did seek out peer reviewers[7] for this "dictionary."

AI for the Supply Chain—Definitions

Artificial Intelligence

(AI) is the design and evolution of algorithms and "intelligent agents" designed to derive patterns, learn, and create optimal insights, choices, make predictions and take action. Traditional programming techniques, fundamentally, are systems that use fixed logic-/scenario-driven code—if, then, and specific rules that predetermine the kind of outcome expected. Given that the data and code are correct, this routine is run once. There will be no variation in outcomes.

AI analytic information systems are employed in scenarios with unknown or highly dynamic variables, changing parameters, or areas of inquiry (new questions) to discover patterns and relationships (associations, causation), and derive insights through the mega-iterative processing of data.

What is *learning*?

We often use the word *learning* as the key attribute in AI and ML systems. Learning in a computer sense is statistically, consistently achieving the desired outcome after processing large and changing data sets.

AI does not think like a human. Think of AI as a powerful *observational engine* taking advantage of mega-computing power.

Like humans, a neuron or network of neurons (neural network) can be built from the learning and referred to again and again.

[6] And which are expected to accelerate even more in the coming decade as we learn to understand and use these data sources in both traditional and inventive applications.

[7] See reviewers in last section, Peer Review

In robotics in fixed environments,[8] AI is employed *during* a training phase, but, later, the robot will not encounter too much variation. However, today, we are letting loose a new generation of robots—so-called autonomous—that do absolutely encounter highly variable environments. In reality, any autonomous entity is highly reliant on the store/database of learning—either embedded or remote (in the cloud)—in order to assess dynamic situations.

Within AI environments, data continues to be processed with refinements, over time, discovering the rule or logic that should fit the scenario. So, the rule is also an outcome we might be seeking, not just an information bit.

Those pesky Google ads!

AI adapts through progressive learning. This is how Google learns what to advertise to you or which products to recommend. If you wonder why Google may get it wrong with you, the algorithm needs training and data.

You provide that through your searches and purchases (or by directly giving feedback on the ad). If you don't provide data, Google will still get it wrong—a lot!

AI systems, therefore, maintain *successes and failures*, data associations and identifiers, to improve and self-correct future recommendations and actions. Often, libraries of intelligent agents (code) are available from the technology provider or can be found on open source libraries.[9]

Machine Learning

(ML) is an application of AI that provides systems with the ability to automatically learn and improve from experience. ML focuses on the development of computer programs that can access data and use it to learn for themselves.[10] ML is not only seeking, through refinement, better and better answers, but better and better questions. In other words, is this the right algorithm to achieve the result? What rule or case actually applies in this circumstance? Whereas in traditional systems, we have preselected the type of routine or algorithm, and that never varies.[11]

There are supervised, unsupervised, and reinforcement learning modes within AI/ML.

Data Structures

Structured data: tables and rows, tabular, that fit nicely into various database schemas, with strict definitions of content, field lengths, and so on. These are text and numbers which may or may not be labeled.

Semi-structured: Does not have a schema and fixed length. Content is searchable with some kind of tag and other elements which can be indexed for search (for example, XML on web pages). These are text and numbers.

Unstructured: There is no schema. The package/file may have a tag (often hypertext to locate it) for filing and searching, but the content is not searchable. These are often pictures, and streaming voice and video.

[8] for example, a manufacturing line

[9] In these instances, a best fit algorithm may be recommended.

[10] For an excellent explanation of the Machine Learning algorithms please reference "Machine Learning in Layman's Terms," by Audrey Lorberfeld, in *Towards Data Science*.

[11] To note here: in AI/ML we doing training by providing appropriate sample data and ensuring that the best fit algorithms are applied and that the results make sense.

Supervised Learning is based on methods that train the system with human direction.[12] Data is categorized (labeled) and the system pulls the data into a specific model. Traditional business systems use data models that are rigid. In a departure from traditional fixed models, with ML, the relevancy of a specific attribute may change with a different query or parameter change; thus, the data might be assigned a new category.

The scope of data can widen with the use of additional data sources, applying broader data categorizing and labeling techniques. So, a necessary step is to ensure you have the data and sources defined.

Adding more data and the corresponding labels can be hugely labor intensive in certain applications, so users may employ data annotation services.[13] Knowing up front, to the extent one can, what might be relevant data is important in setting up AI systems.[14] This data expansion is an important concept to understand in AI, in general. (See Machine Learning Data-Management definition, page 5.)

Unsupervised Learning is extremely useful to process big and unstructured data where the system can explore and organize data, finding underlying patterns and associated data. In unsupervised learning, we are often dealing with unlabeled data we may get from partner documents, or find on web text, social dialogue, images, voice, and so on. (See Deep Learning definition, page 6). Here, users can explore much broader data sources—both structured and unstructured—to learn new things. Thus, unsupervised learning is a growing area of use.

Reinforced or Reinforcement Learning is where the system learns from repetition. In the past, this was often the domain of robotics, with the goal of performing a given set of tasks under varying conditions. AI can structure and classify through repetition, developing a "skill." If the robot keeps bumping into a wall, very soon it will record a wall and develop precision on how close to get before it turns.

Data in AI Speak

In ML this is *data science*, not just data mining. It includes the meaning and value of data in a specific case. ML systems focused on data management are an essential step in processing and preparing the data for use.

Learn the data lingo:

Labeling: assigning a specific tag/name to the type of data (plant, man, women, vehicle and so on). Labeling should be as elemental as possible to be discoverable by many types of processing.

Categorizing: assigning another layer—for example, man, women, and child are human.

Attributes: are an observable characteristic used to assign meaning and association of data. For example, products have attributes such as color or size. People have attributes like age.

Clustering: grouping like data and/or content to find patterns, and building relationships among data and hierarchies. Really smart clustering can assess how close or distant the relationship may be between the attributes.
For example: Attributes like age and work status are closely related, whereas a preference for ice cream is not closely age related.

Training data: An ML system is only as good as the data it is given. ML systems will learn and be validated during the training phase. Feature engineering—selecting the correct, relevant data and constructing the training set—is one of the biggest tasks in building a successful machine learning system. For training, support from those who have *domain-specific knowledge is critical* in getting the right information and the quickest path to success.

[12] The path to AI is an easy step up from traditional processing in supervised learning, as most organizations already have defined data and sources.

[13] When using unstructured data on the web—such as visual/video/voice, crowdsourcing, deciphering text and categorizing it, refining the meaning and translation of text in social networks and blogs, and so on, these types of services can be useful.

[14] Learning from others is critical here, since one can get lost in the world of data. Getting to relevant data quickly is essential for initial project success. Later, with experience, users can expand their repertoire. We recommend users seek out data curating services that can access, clean, and store data relevant to the application.

Machine Learning Data Management

Understanding data and how to manage it is essential in the AI/ML world. Data management is an advanced science in technology and a service category in the ML world.

From semantics to understanding the meaning, value, relationship among data, and even its emotional import, there are specific algorithms and methods to improve the quality and *remove subjectivity* from the data. These can be *expert systems* and can also rely on *real people* to help understand the data and its use cases, to categorize and label it.

We know data accuracy is a big issue even within traditional systems. But once we access broader and fuzzier sources like web data, accuracy issues become even more acute. If we are to depend on external data sources, even in traditional methods such as collaborative forecasting and demand/supply automation, and expect our systems to learn and operate somewhat autonomously, we have to be assured that the *data in* is not *garbage in.*

Machine Learning Data Management has more advanced capabilities beyond tagging/labeling data, such as:

- *Reading and assessing* data—discovering data relevancy; building it into an actionable framework (system, workflow, report, etc.)
- *Semantics*—data meaning and value
- *Quality and completeness*—ML has methods to resolve missing or inaccurate data and make predictions about what the source is likely to input
- *Building a knowledge base based on the category of data*—many types of documents could be assembled by multiple categories (Google is an example of this.)
- *Dimension reduction*—Many people think now with all the web, IoT, and temporal event data (which can be a huge continuous stream of data) that we have to boil the ocean—churning carillons of data points to find our one needle in the haystack. In dimension reduction, ML's function is to reduce the number of variables and filter for only applicable data. The value of computing power is indispensable in iterative processing until this filtering succeeds at getting important and relevant data.

Why would adding ML be better than our traditional optimization module?

The *real* world is multi-dimensional, with a lot of unstructured data.

With ML, we can use interesting and important, yet challenging, data such subjective or fuzzy data in which we may derive *trends*; or "non data"—blind spots—that traditional systems don't account for. These can be real impediments to performance in the physical world.

This has obvious applicability in spatial applications like warehouse and transportation.

In areas where there are multiple and dynamic constraints, ML may do a better job of problem solving.

And interestingly, supply chain users are learning to layer these various insights in new ways to understand consumer demand and the logistics-chain dynamics.

Consider this: If the pattern is already self-evident, then we don't really need AI/ML.

As organizations move beyond the boundaries of traditional systems, understanding the broader context of data[15] and how AI plays such a pivotal role, is fundamental.

[15] More on supply chain data: [It's All About The Data]

Neural Networks (NN)

NN is a category of AI that focuses on learning to recognize patterns in the data. Neural networks are an area of AI that is conceptually more understandable[16] and more "like humans." We humans learn something—an association between two points, events, people and so on—and the neuron is created. And there can ultimately be many associations between one point and another, creating a network in our brain.

In essence, artificial neural networks[17] work similarly. Neural networks can segment data and assess, weighing the value, correctness, and relevancy of input data based on parameters that are fed to the algorithm. These analytic processes then correlate and build that neuron. As with humans, it may take a couple of iterations to understand the relationship. Think of these as intermediate steps, sub-processes creating nuggets of input for use in the next step. Neurons can then be layered, going through many iterations, creating interrelationships and dependencies among the data, ultimately forming a whole picture. In practice, neural networks are applied to complex data sets where there are a lot of variables.

So, a question here might be: which variable or combination of variables is most significant for my discovery goal? Finding anomalies, errors and useful patterns all fall under the category of neural networks. Once the neural network has been trained, it can make predictions by detecting similar patterns in future data.

Deep Learning

Deep learning is a recent addition within AI/ML.[18] Deep learning should be considered a superset of machine learning and neural networks whereby multiple layers of processing cooperate to solve a problem. Deep learning is applied to problems of scale and complexity, requiring ultra-computing power.

Most machine learning routines use labeled data to complete their tasks. Deep learning is almost always applied to problems that use unstructured data. In deep learning, the system may leverage multiple neurons, with neurons being those that use the same algorithm or those that use different, unique algorithms. (These intelligent layers of neurons can be part of a mature NN environment.)[19]

---

[16] But in practice, as more and more nodes are created though iterative processing, it does become difficult to peer through it all. To note here: Defining neural networks can be a bit tricky if you rely on the web, due to the diverse existing definitions—either hopelessly useless for understanding what the computational essence is, or extremely technical—for advanced computer scientists and mathematicians.

[17] You may see either term used: artificial neural networks or just neural networks.

[18] The main difference between deep learning and machine learning is, machine learning models become better progressively, but the model still needs some guidance. If a machine learning model returns an inaccurate prediction, then the programmer needs to fix that problem explicitly. But in the case of deep learning, the model does it by itself. An autonomous driving system is a good example of deep learning. *Source: GeeksforGeeks*

[19] Again, many terms get used in marketing or the media that have the same or overlapping meaning, which you may encounter.

Let's take the example of ML teaching itself to play chess. In a preprogrammed game, the rules of chess would be coded. A knight can only make these types of moves, a rook these moves, and so on. This game proliferated decades ago for the PC. But we often could beat the system because it did not have the *ability to learn*. Now, applying deep learning, the system observed thousands of games and millions of moves. Backed up by mega computing power and memory to process thousands to millions of strategies and *what ifs*, it can recommend decisions that produce a superior outcome.

Deep learning has additional properties of associating and discovering attributes that may not be obvious without deep searching. For example, in deep learning, a program, from processing web data over and over again can associate data in many different scenarios or context, for example, the word "red" with fire, traffic lights, hair color, fabric dye, planets, the sky, boiling temperature, to emotions (he was seeing red), and so on.

Deep learning is also being used as an underpinning in such applications as speech or facial recognition.

Computer Vision

Computer vision is a field of artificial intelligence that trains computers to interpret and understand the visual world. Using digital images from cameras, videos, and sensors, AI algorithms and deep learning models can accurately identify and classify objects—and then react to what they "see." In today's Augmented Reality systems, security systems, and autonomous vehicles, this application of AI is the foundation. In supply chain, computer vision is applicable in manufacturing, quality/defect detection, transportation, and warehouse management.

Cognitive Computing, Insight Engines, and Sentiment

These are the areas within AI that are most subjective today. Cognitive Computing, for example, is not optimizing or building "correct" answers, but seeking out patterns of human behavior and trends. An insight engine is search plus AI. Sometimes, you may hear the term *inference engine* or deep learning used here. Sentiment analyzes a string of text to assign a positive or negative value.

Insight Engine platforms and applications such as sentiment are a bit fuzzy in some of the conclusions they draw. But isn't it true that humans are a bit fuzzy too? Our language is peppered with *well, maybe, if, might, but,* and so on. And we are notorious for *consistently being inconsistent.* Therefore, seeking a definitive, absolute answer should not be an expected outcome.

Users may interact with insight engines using natural language queries, and then apply the powerful search capabilities. There is a lot of data categorizing, filtering, and so on to make use of this data. It requires a long-term investment in the data and the application to fine tune these applications and then apply them in a practical manner. Thus, an insight engine must be a platform of capabilities to be successful.

Supply Chain Scientist

Here is a term we hope will enter the lexicon of supply chain definitions. Already, a professional supply chainer must learn a tremendous amount of analytics, graphical representations, and algorithms—mathematics. With the addition of AI tools, we are upping our game and our profession to even a more scientific level, including the three-dimensional and sensor world—physical, geospatial intelligence, or psychology—adding to the body of knowledge and to our daily work. Thus, we become Supply Chain Scientists.

Peer Review

This document was designed as a general-purpose guide for the supply chain professional. It has been reviewed by outstanding professionals in the field of Supply Chain and AI/Machine Learning.

With much gratitude we recognize these contributors:

Harve Light	Churchill Systems
AI R&D team	Halo Business Intelligence
Justin Siefert	Logility
AI R&D team	Descartes Systems Group
Mavi Silveira	Descartes Systems Group
Kristen Franks	First Insight

…and a few anonymous reviewers….

ChainLink AI Collection highlights:
http://www.chainlinkresearch.com/AI/index.cfm

About ChainLink Research

ChainLink Research, Inc. is a Supply Chain research organization dedicated to helping executives improve business performance and competitiveness through an understanding of real-world implications, obstacles and results for supply-chain policies, practices, processes, and technologies. The ChainLink 3Pe Model is the basis for our research; a unique, multidimensional framework for managing and improving the links between supply chain partners.