An Introduction to Artificial Intelligence (AI)

By: Alex Feild, Chief Scientist – Applied Analytics Strategy, Compass Point Federal, an Akima Company

Introduction

Artificial Intelligence (AI) has had several common definitions throughout its history. However, a commonly accepted definition for AI is something we have not yet achieved, much to the chagrin of AI researchers developing useful expert control systems in the 1980s, voice recognition and signature matching in the 1990s, and improvements to search in the early 2000s.

A recently used definition has been anything with apparently intelligent, complex, or complicated behavior. This lack of a commonly accepted definition has allowed any company who wants to deliver AI to label their "sufficiently fancy" solution as such. This is akin to saying that all buildings that use electricity are electrical buildings.

That said, AI does have a formal definition. AI consists of an agent that observes its environment (takes input) and performs an action on the environment (produces an output), while a reward function (or error function, fitness function, etc.) affects the behavior of the agent by mapping the environment and agent's output to some value that will be minimized or maximized. In reinforcement learning, this is analogous to how a dog learns what their owner wants it to do. The input is the dog's environment including commands from its owner. The output is whatever the dog does. The reward function is the affection and treats of the owner.

The power of this model is that the reward function does not need to know how the goal is accomplished. In the example of the dog, this is how the dog thinks and controls its motor functions. The flexibility of this model comes from the fact that the agent can be a general problem solver without specific knowledge on the desired outcome built in. In this same example, the dog learns through experience what it should do, without having to comprehend the thoughts and desires of a human being.

AI and Optimization

AI is similar to optimization; however, one major difference is that optimization takes a fitness function (cost/reward plus constraints) as its input, while an AI agent takes the environment as its input. Both optimization and the efforts of an AI agent are still hill climbing. The goal is always to maximize or minimize some function of some number of dimensions. In optimization, the algorithm takes in a gradient or other function (a swarm technique) that encodes the slope of the environment (input). This allows the optimization algorithm to calculate the fitness of its current state and proposed states (for some methods). Conversely, an AI agent will have to know or figure out for itself what a slope is.

Another major difference between AI and optimization is the goal. In optimization, the result is the goal (e.g. designing a car).

In AI, the goal is the agent (e.g. a car recognition agent). Note that in machine learning, the process of training the agent is an optimization problem. We can achieve swarm control with both AI and optimization; however, if we can state the swarm's fitness function directly, it is easier and more robust to use optimization. It is possible that some form of deep reinforcement learning would perform better in this situation, but only after significant development and training, particularly if we replace swarm agents with a single control agent.

Another difference between AI and optimization is how you formulate the problem. If you can formulate the problem into a fitness function of the environment and a design/control variable, optimization is likely a better choice. Why? Because compared to AI, the approach is more easily validated and estimates for the solution space are easier to generate. Example applications for optimization include swarm control, engineering design, photogrammetry, and reverse rendering. If formulating your problem completely as a fitness function would be difficult, or you desire a more general approach, AI may be better suited. Examples for AI include classification and function approximation in graphics and language, probabilistic generative modeling, unsupervised learning, reinforcement learning, symbolic reasoning, and expert systems.

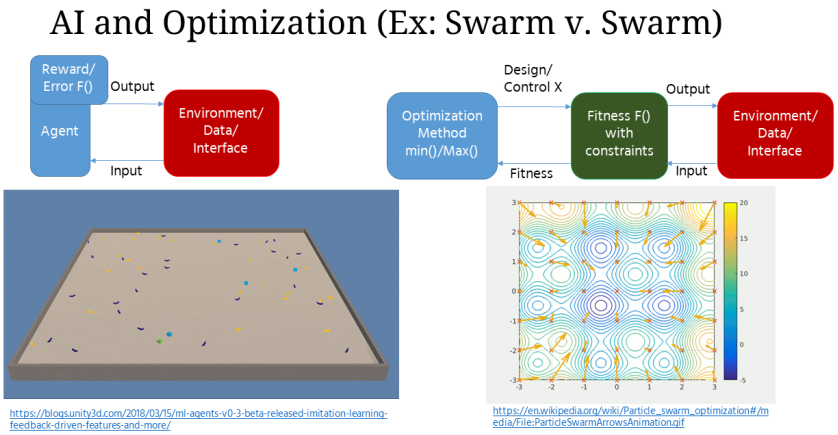

Figure 1: The agents on the left take direct observations of their environment and are a form of AI. The agents on the right take the gradient of surface they are on from the fitness function.

The History of AI

The beginnings of AI date back to the 1930's and the work of Kurt Gödel, Alonzo Church, and Alan Turing. The Church-Turing Thesis states that a function is effectively computable if a machine can evaluate it. Another way to look at this is, some machines exist that can evaluate all functions.

The first work into neural networks began in the 1940s and 1950s. Symbolic reason, logic, search methods, semantic networks, and game AI all began in the 1950s. Government funding of AI became significant in the 1960s, and the first "AI winter" hit in the late 1970s/early 1980s. All AI winters, including future AI winters, are a result of overinflated expectations and a lack of acknowledgement of the progress that has actually occurred. Expert systems were the rage in the 1980s followed by another AI winter late in that same period. The 1990s saw intelligent agents, including agent based modeling, Deep Blue and enhanced search, and the introduction of rigorous statistical and formal methods such as Bayesian Networks and Markov Models. There were also many deployed successes of AI in the 1990s including signature recognition, speech recognition, data mining, and robotics, among others. Back propagation, otherwise known as the training method of deep neural networks, gained recognition in 1986, and by 2007 deep learning became accessible to everyday researchers with the deployment of CUDA by Nvidia.unsupervised learning, reinforcement learning, symbolic reasoning, and expert systems. From 2009 to 2012, vast amounts of data and successful papers made deep learning significantly rise in popularity. Fast forward to 2019 and major new developments in AI are in interpretability, predictability, and reliability.

Types of AI Work

The Venn diagram on the right is from the book "Deep Learning." It depicts the sets of AI, Machine Learning, Representation Learning, and Deep Learning.

Machine Learning

Arthur Samuel, a pioneer in the field of AI, defines machine learning as a "field of study that gives computers the capability to learn without being explicitly programmed." A mathematical model is trained to perform a task from data. Machine learning is used to solve problems that are easier to describe through data and training than explicitly. Methods used in machine learning include Bayesian networks, support vector machines, neural networks, and more. Machine learning is used for supervised and unsupervised learning, representation learning, reinforcement learning, generative models, predictive modeling, and more. Classification and clustering are classic machine learning use cases.

Care should be taken to avoid leaking testing data into the training dataset in order to mitigate memorization. In general, the more training data and the simpler the model, the less over fitting will occur. Validating certain machine learning models such as neural networks can be difficult and perhaps intractable depending on the method used. The real world performance of non-rigorous methods can be unpredictable beyond the testing dataset.

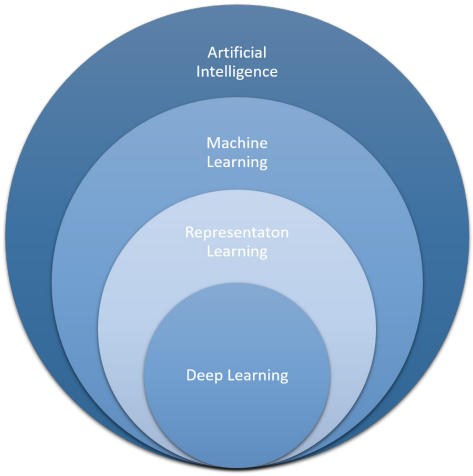

Figure 2: From Deep Learning by Ian Goodfellow, Yoshua Bengio, and Aaron Courville

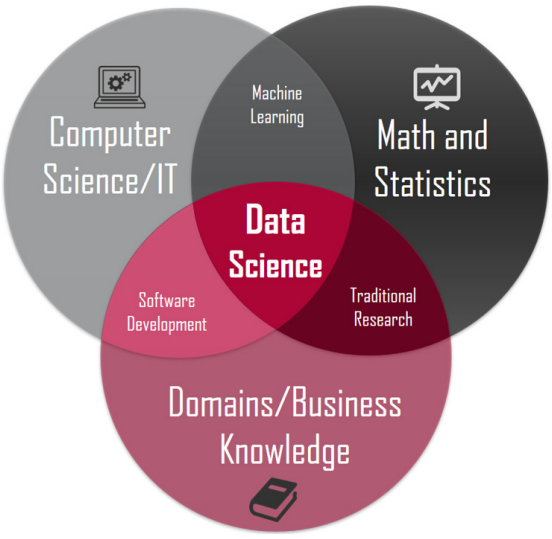

Figure 3: Venn Diagram showing the relationship between disciplines that research and use machine learning. From towardsdatascience.org.

Representation Learning

Representation learning is a set of techniques that allow a model to automatically learn the representations required for feature detection in raw data. Methods such as Bayesian networks, support vector machines, and random forest require that feature extraction be performed prior to classification or function fitting. For example, images would need to be turned into a collection of edges, hulls, or other shapes. Sound might have to go through a fast Fourier transform and a time series run through a prior model. Representation learning, such as autoencoders, takes in raw data and performs the task, including feature representation, using that raw data. In short, you replace resolution with sematic meaning. Modern computer vision with deep neural networks rely on the network to learn feature detection and shape representation on their own. Representation learning is useful for when it is more difficult to describe the feature representation appropriately than it is for the model to learn it.

Neural Networks and Deep Learning

The fundamental computational unit of a neural network is the inner product of learned weights with the input, plus a bias, followed by a nonlinear function. This is premised on the universal approximation theorem that a sufficiently wide neural network can approximate any function. Neural networks are trained via back propagation of error with respect to the weights.

A deep neural network is a neural network with more than three layers and more than one hidden layer. The reason for previous neural networks being shallow is due to loss of the error gradient mention above after repeated derivation through the layer. Modern hardware, application specific architectures, and training advancements have made deep neural networks possible. Modern deep neural networks have millions upon millions of weights. The name was adopted to avoid a stigma that had developed around neural networks. Most machine learning challenges (imagery, video, speech, text) are dominated by deep neural networks. They are used for image and video classification, annotation, segmentation, target tracking, speech recognition, and text generation among many other applications. New applications and discoveries are found constantly. Personally, I have worked in violent behavior detection, target pose and vulnerability, target behavior prediction, synthetic data and transfer learning, and generative models.

For deep learning to work, the output needs to be solely a function of the input and thus require no additional information. Deep neural networks can memorize training data and therefore care should be taken to ensure overfitting or a false sense of confidence has not occurred. Overly deep neural networks can learn unhelpful features due to a vanishing gradient and the compensation power of the later features. Importantly, deep neural networks can suffer from data attacks and that failure is brittle (unexpected).

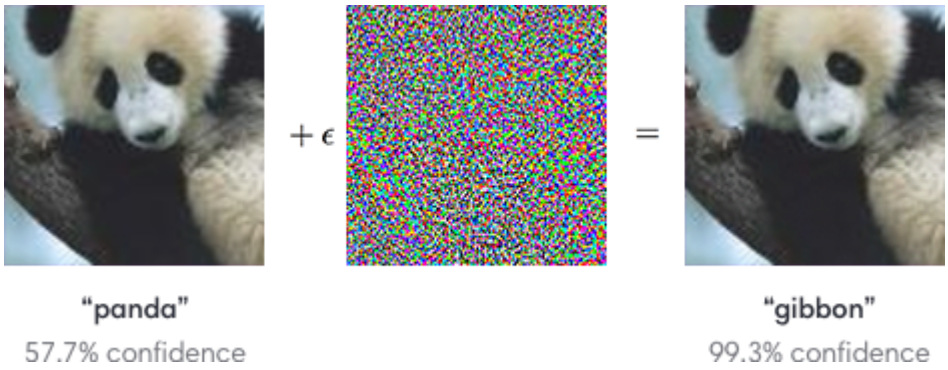

Figure 4: Data hack from Explaining and Harnessing Adversarial Examples by Ian Goodfellow, Jonathon Shlens, and Christian Szegedy.

What's Next for AI?

The best performing methods of AI suffer from a lack of explainability and interpretability. We have seen a focus on mending this problem in order to enable validation and ensure trust. Human-machine teaming and confidence has been a recent effort of the Defense Advanced Projects Agency. For example, it would be extremely useful for a machine to tell us "It is raining so I am reduced to 95% accurate." In general, there is an increasing awareness of machine learning safety, covering reliability, explainability, failure prediction, formal verification, data bias, input hacking, and privacy.

Like any significant technology, there are benefits and dangers in AI development and deployment. It is incumbent upon us to move responsibly and make informed decisions. In this way, we can make the best use of massive compute, data, and human creativity and capacity for rational thought.

Looking to deploy AI solutions at your agency? Akima subsidiaries Cloud Lake Technology and Compass Point Federal deliver the expertise agencies need to turn data into insight and insight into action. As 8(a), Small Disadvantaged Businesses (SDB), and Alaska Native Corporations (ANC), we can partner with you to deliver the agile, flexible solutions you need to deliver on your missions. Contact us today to learn more.