Article from:

Forecasting & Futurism

July 2012 – Issue 5

Artificial Intelligence: What Is It and How Can I Use It?

By Brian Grossmiller

Recently I had the opportunity to participate in an online class which provided a very thorough introduction to the field of artificial intelligence (see *https://www.ai-class.com/* for the course materials). There are several exciting potential applications to an actuary's practice. In this article I will share some of the key definitions along with some possible applications.

A key definition in the field of artificial intelligence is an intelligent agent. This is what we are trying to build, an agent or system that (ideally) behaves optimally in its environment. Intelligent agents can vary considerably in their complexity, from simple agents that respond with a reflex reaction to agents that actually learn from their environment and can adjust their actions for unexpected impediments. The components involved in an intelligent agent will of course depend on the application; these can include software, robotics, cameras, keystroke inputs and computer files.

An intelligent agent consists of three parts: a sensor, a control policy and an actuator. Sensors can be cameras, optical character readers, or an input section of a computer program and are the means by which the agent perceives its environment. The control policy is the element that decides what action to take based on the agent's perception. As a simple example, a search engine would perceive a keyword entered into it and the agent would decide what list of URLs to display based on a set of rules. Actuators are the means by which an agent responds to its environment; these can be anything from robotic arms to simple text outputs on a monitor.

One of the more basic intelligent agents is a problem-solving agent, which attempts to reach a goal while maximizing its performance according to a metric. This type of agent can be constructed through searching, where it is typically designed to find the optimal path from a starting point to its defined goal state. A familiar example is GPS navigation, which finds the shortest path from your current location to your destination. In this case the agent finds a path designed to minimize distance or travel time.

An interesting type of search strategy includes a heuristic function, which provides an estimate of the cost to reach the goal state from any point along the way. This function is combined with the known cost to reach each point in an effort to find the cheapest solution in the shortest amount of time. A key criterion for a heuristic to work in this fashion is that it has to be "optimistic," that is, it never overestimates the actual cost.

Generating heuristics can prove to be an interesting problem in and of itself. Some strategies include solving a "relaxed" problem, where some restrictions of the actual problem are ignored, or looking at a subset of the problem. There are also techniques for learning heuristics from examples of solutions.

A very powerful and exciting application of artificial intelligence is constructing agents that learn from examples. This approach can build programs to solve problems that are excessively difficult or tedious for a programmer to design directly. The main categories of machine learning are reinforcement, unsupervised and supervised learning.

Reinforcement learning depends on some metric which determines whether an outcome is favorable or not. In a game such as chess, this would be provided in the form of a win or a loss. Over a large number of games the agent can determine which of its actions tend to lead to a favorable outcome. This approach produces agents designed to take actions to optimize their expected result.

In unsupervised learning the data provided are not labeled. The agent attempts to learn the structure and features of the data. Notably, agents are not given feedback as in the other methods; the result is usually a summary of the data.

Supervised learning involves collecting pairs of inputs and outputs. The outputs can be generated from humans, or perhaps be a set of related measurements. The dataset is used to train the agent to infer the output for inputs not in the training set. Actual applications may blur the distinction between these categories. For instance, semi-supervised

learning problems typically have a labeled subset of data and a much larger unlabeled dataset. Results can have greater accuracy than a fully unsupervised agent without the potentially considerable expense of labeling a large dataset.

One method that has shown promise in actuarial applications is Genetic Algorithms (GAs), which are search heuristics modeled on natural evolution. GAs utilize a fitness function and develop optimal solutions over a series of generations by combining and mutating the top performers. Typical problems involve a large number of possible solutions with a readily calculable fitness function where computing each solution directly would be prohibitively time consuming.

A first generation is developed randomly and scores under the fitness function are computed. The top performers are randomly combined together to develop a new generation, usually with a random mutation of a small number of genes. The new generation is scored under the fitness function and the process is repeated recursively. When the initial population contains a diverse set of solutions, the child solutions tend to be radically different from the parents', which often leads to significant improvements in performance in the first few generations. The solutions become increasing similar, since they are drawn from the same gene pool, so later generations typically show more marginal improvements in performance.

When the generations become too similar, the GA version of inbreeding occurs, which limits further gains. Part of the "art" in developing GAs involves striking good balances between retention of the top performers and introduction of mutations. Keeping the best of a generation and affording them "breeding rights" helps to prevent their children from regressing. However, mutations can sometimes bring about innovative advances.

This technique is being applied to selecting an optimal provider network for a health plan. The fitness function in this instance scores each provider according to their relative cost efficiency and produces solutions that maintain an adequate panel of providers in each area of medical practice. In this situation, there are a large number of providers and many possible solutions entailing different combinations of groups of providers. This provides an excellent starting point for developing a narrow panel, as selecting a panel by hand can be very tedious and may overlook better performing alternatives.

In addition, a GA approach has been successfully implemented in Life Insurance Asset/Liability Management. In this application, the fitness function measures volatility driven by shocks to the interest rate curve; the GA minimizes this volatility by generating optimized asset allocation strategies. This has produced strategies which provided superior minimization of interest rate risk over traditional methods (see Ben Wadsley's article, "Are Genetic Algorithms Even Applicable to Actuaries?" in the July 2011 issue of the *Forecasting & Futurism Newsletter*, page 6).

There are many interesting applications of artificial intelligence being deployed today; I invite you to review the course materials online. In addition, a free series of excellent computer science courses can be found on Udacity's website at *http://www.udacity.com/*. ▼

REFERENCES

Russell, Stuart Jonathan., and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Upper Saddle River: Pearson, 2010. Print.

Introduction to Artificial Intelligence, taught by Professor Sebastian Thrun and Professor Peter Norvig. https://www.ai-class.com/

Brian Grossmiller

Brian Grossmiller, FSA, FCA, MAAA, works at ODS Health Plans in Portland, Ore. He can be reached at *brian.grossmiller@yahoo.com*.