Applications of Artificial Intelligence in Learning

By Paul Cummings, Vice President of Innovation & Technology

As a computational scientist, I’ve always been interested in the interaction of human and machine intelligence. Specifically, which areas of human performance that can be enhanced by algorithms, technology, and analysis. There has been a dearth of information about how intelligent computation using artificial intelligence (AI) will spur a new generation of smart computing systems that will help us solve complex problems. I will be presenting a few examples of how AI may be used to enhance our ability to design, structure, and analyze learning content in the future.

So, let’s get started.

Several years ago, I was involved in a project with a large rail firm that was interested in building learning strategies for their new conductors. Many of their seasoned conductors were leaving the firm, taking with them unique knowledge about how to do their jobs. As a conductor, one of the more interesting challenges was to build a train from rail cars located on different tracks in the yard. In order to sort the cars, several procedures would need to be implemented, including the mechanics to support safety stops, secure equipment, foul equipment, shove protection, and switch positions; in addition, a three-step protection would need to be considered. All the while, each move of a car costs money, so minimizing the switching of cars was vital.

In the past this knowledge was handed down much like one would expect. A new conductor worked with an experienced conductor as a mentor who would teach them the ropes over the course of several years until the fledgling conductors were able to do the job on their own.

In the past this knowledge was handed down much like one would expect. A new conductor had a mentor that would teach them the ropes over the course of several years until the fledgling conductors were able to do the job on their own.

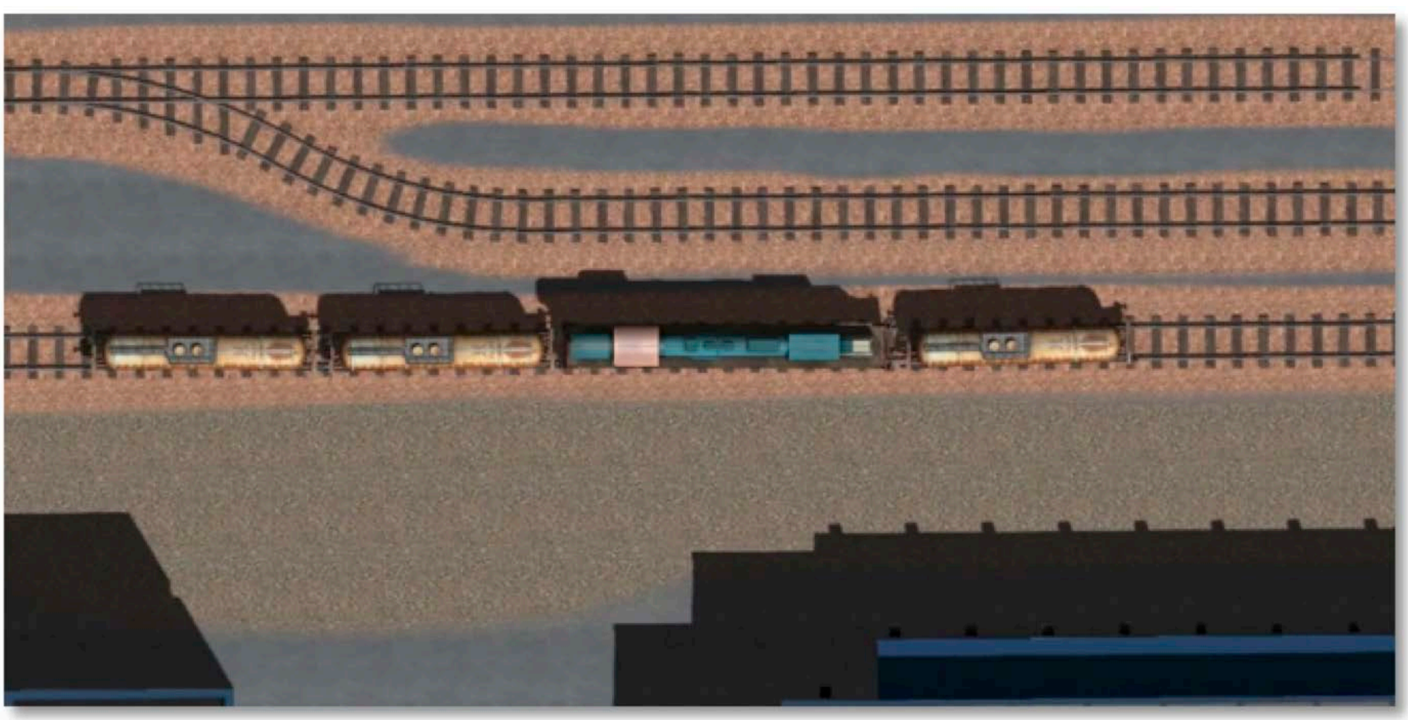

Train Simulator Teaching Conductor Cognitive Strategies

Our initial approach was to decompose the intelligence of the seasoned conductors into learning objectives that became strategies to be learned and tested. The process of deriving learning objectives within the simulation was quite elegant and truly exemplified the importance of how a well-designed learning objective-driven training could aid in complex decision-making. Important learning strategies, such as “Most Logical Action” and “Constant Forward Thinking,” were built as both learning descriptive content (didactic) and as applications within the simulation.

That is, the basic approach was as follows:

1. Build learning objectives based on strategies
2. Discuss how to train these objectives with a subject matter expert
3. Build strategies into training simulator

So, what’s the problem?

Although the training was designed with the utmost rigor in many ways, there were still challenges when developing the training approach.

Expert Knowledge: The simulation still relied heavily on a small subgroup of experts. Learning objectives and strategies were important from a declarative knowledge perspective, but the application in complex tasking created a greater challenge.

Scalability: We could not guarantee that the vignettes built for the simulation could translate into the millions of permutations that may occur in the field.

Complexity: In a domain where the learning objective is how to change a tire, a procedure-based approach to learning strategies is a fairly straight-forward process. Learning objectives easily map to declarative knowledge as well as procedures. In the example of the rail yard sorting, the process gets much more challenging. Understanding and applying strategies had several cases that were not always quantifiable, and many special cases needed to be considered.

Introducing AI (and Machine Learning)

The topics of Artificial Intelligence and Machine Learning are vast and complex. We will begin by asking what is Artificial Intelligence (AI) and what is Machine Learning (ML)? AI and ML are often spoken in the same breath, and with little forethought there is in fact a difference, although there are areas of interconnectedness. We can succinctly say that *Artificial Intelligence i*s a broad topic where machines produce seemingly intelligent behavior. *Machine Learning*, on the other hand, is subset of AI that describes a process for how algorithms learn without giving explicit instructions on how and what to do. Certainly, with a well-crafted methodology and a bit of computational know-how, there is some truth to that, but we need to first examine areas that can use the power of AI in learning.

Exploring Solutions with AI

What we're interested in now is determining how to use AI to help us solve problems, and for this, we use a type of AI called *Reinforcement Learning*. The best way to think about Reinforcement Learning is from the perspective of child development. As children, we learn by interacting with the world around us. A baby walks, touches, puts things in its mouth, but doesn't have a set of pre-programmed rules that it must follow; however, it does have direct sensorimotor association with the world. Through simple interactions, it produces a wealth of information about interactions and effects, and within time learns to hone these actions to achieve goals. Generally speaking, we carry these processes with us throughout our entire lives. We experiment, receive feedback and decide if the behavior is getting us what we want.

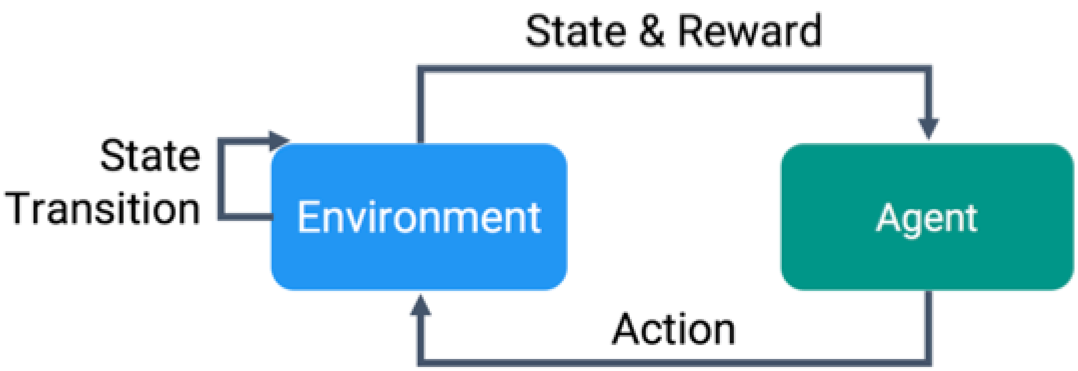

Standard Reinforcement Learning Cycle

The idea behind Reinforcement Learning is that an entity (or *agent*) will learn from the environment by interacting with it and receiving rewards for performing actions. The agents are not told which actions to take, but instead must discover which actions yield the highest rewards by trying out "actions" within the environment. Why is the goal of the agent to maximize the expected cumulative reward? Reinforcement Learning is based on the idea of the reward hypothesis. In our case we'll use a reinforcement learning technique called Proximal Policy Optimization (PPO). PPO uses a neural network to approximate the *ideal approach* that maps an **agent's observations** to the **best action** an agent can take in a given state.

Applying Reinforcement Learning Algorithms

We'll begin by generating a very simple environment where our intention is to model a 3-rail yard with 3 rail cars that we'll call agents. There are 3 agents in the environment. In the case of the railyard concept, what we can now do is build a reward system that looks as simple as this:

1. **Red Agent** (left car): if the center car is to its right, give it a reward of 1.

2. **Green Agent** (right car): if the center car is to its left, give it a reward of 1.

3. **Blue Agent** (center car), if the red agent is to its left and the green car is to its right, give it a reward of 1.

4. **Default:** provide a small penalty for not reaching its objective of -0.05 times the distance between agents.

So, what is going on here?

As the model is trained over time, the agents will converge on solutions that map its observations to rewards. The graph below demonstrates how a machine learning model in Tensorflow increases accuracy in its training model over time. Once the model is built, the agents use this model to determine where to be on the tracks and where to move based on this simple set of rules. It may be now apparent that this automated system with some level of training and detail, can calculate optimal solutions, rather than rely solely on the expertise of a conductor.

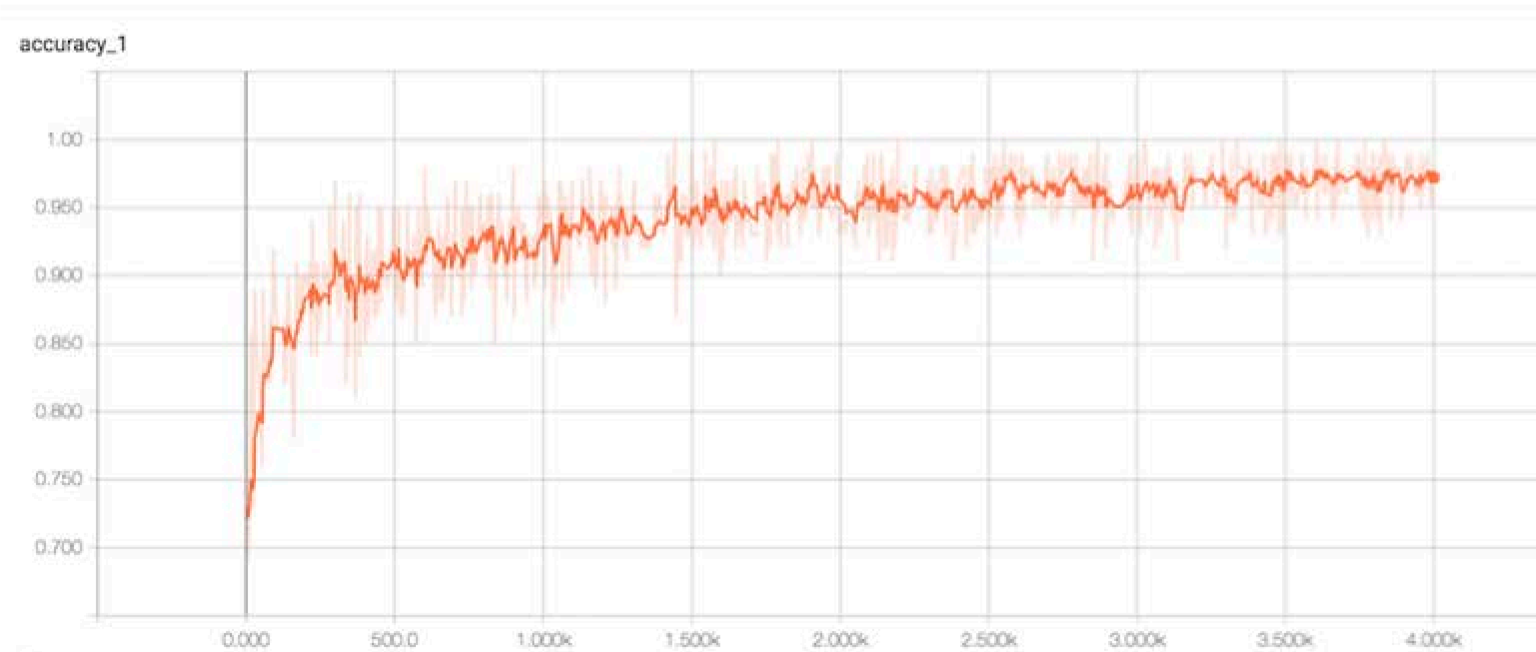

Tensorflow Accuracy Results

Scaling the Problem to Real World Scenarios

Once properly trained, the machine learning model can provide an automated way of solving problems like "What are the least number of moves it will take for me to sort rail cars in a specified order given the initial configuration?". This saves an extraordinary amount of time training conductors to be car sorting experts, thereby freeing up their time to do more important tasks.

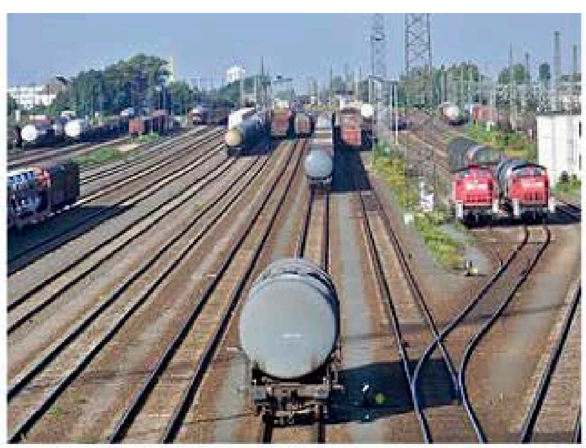

Next Steps

This was a very short introduction to the value of machine learning in training systems, and although I tried to keep the discussion concise there is still quite a bit to unpack. In the future I will be looking at additional AI topics as described below:

- **Expert Modeling/Imitation Learning**: Developing models that can translate expert behaviors into intelligent tutors and information within a training domain. This includes the development of representations of general human profiles and behaviors.
- **Tailored Scenario Creation**: Methods that help to generate dynamic and scalable scenarios for training. Here we consider the importance of how AI can produce procedural, adaptive and automatic content.
- **Data Analysis and Measurement**: Broad discipline of using data to measure performance and provide detailed descriptive, statistical and predictive behavior analysis.
- **Real-Time Dynamic Scenarios**: Methods to develop scenario changes on the fly to support tailored, scalable, and incrementally challenging complexity to the learner.
- **Artificial Human Behavior:** The design of AI agents that mimic realistic, repeatable, and measurable human behavior in synthetic environments
- **Intelligent Tutoring, Remediation and Feedback**: The use of intelligent data analysis to provide learning-driven feedback to training participants.

Mathematical Side Note

If you're interested in a detailed mathematical discussion I'd be happy to sit down with you to discuss its basis in dynamical systems theory and the Bellman equation as it relates to machine learning, but to be frank, it's not completely relevant to getting the basics down. Oh, and it's a bit dry.

(sorry math, we still love you).