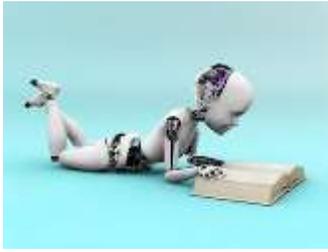**Machine Learning Taps Power of Curiosity**

- *6/6/2017by [Ariella Brown](#)*

(Image: Sarah Holmlund/Shutterstock)

As data analytics become increasingly driven by artificial intelligence (AI), researchers search for a way to drive machine learning. The key ingredient its future development may be a dash of curiosity.

There are all kinds of AI systems currently used by various businesses with different names like [Alexa and Albert to personalize then](#). Perhaps it's time for an AI system named George after the monkey whose curiosity propels him into various adventures.

That would be an apt choice for the Intrinsic Curiosity Module (ICM) developed by a group of four researchers at University of California, Berkeley. The attempt to inject curiosity to achieve self-motivated advances in machine learning was the subject of their paper, [Curiosity-driven Exploration by Self-supervised Prediction](#), that was just submitted to the 34th International Conference on Machine Learning ([ICML 2017](#)).

Their premise is that external rewards for learning are of necessity limited and actually rather rare in real life. That doesn't mean that people stop exploring or seeking out answers even when there are no prizes for doing. They are motivated by their own human curiosity. Infusing that kind of motivation in a virtual agent gets it to test things out for itself even when not directed to do so. The test of the effect was done in monitoring how far it would proceed in two video games, VizDoom and Super Mario Bros. as you see in the demo video here:

As the [paper](#) explains, the utility of curiosity is twofold: one is that it motivates exploration of one's "environment in the quest for new knowledge (a desirable characteristic of exploratory behavior is that it should improve as the agent gains more knowledge)." Another is that it serves as "a mechanism for an agent to learn skills that might be helpful in future scenarios."

Capitalizing on the latter allowed the virtual agent to build on "knowledge gained from earlier experience" to successfully navigate through "new places much faster than" it would have done if "starting from scratch." The researchers demonstrated that effective transfer of knowledge in the agents' extrapolation of what they have learned from one level of the video game to another one even within "a completely new map with new textures. "

Though the agent trained in Super Mario World did not receive any incentive "for killing or dodging enemies or avoiding fatal events," it still came to learn to do these things and arrive at "over 30% of Level-1." The researchers surmise that the motivation for the agent is to do whatever it takes to stay alive in the game in order to be exposed to more of the game environment.

"In order to remain curious, it is in the agent's interest to learn how to kill and dodge enemies so that it can reach new parts of the game space," they say. In that way, "curiosity provides indirect supervision for learning interesting behaviors." There could be far-reaching ramification for such built-in motivation.

The researchers suggest that the course-plotting behaviors learned by the agent like that of learning to avoid "bumping into walls" in the VizDoom game could be adapted to "a navigation system." That would be a design that responds to the stimulus in the environment rather than "reward signals." Given that those occur very infrequently in the "real world," they believe that their "approach excels in this setting and converts unexpected interactions that affect the agent into intrinsic rewards."

In the interest of collaboration, which itself is based on appreciation for shared curiosity, the researchers said they will share their algorithm's code and what was involved in the tests online at no charge. So you can exercise your own human curiosity in learning more about the development of artificial curiosity. Perhaps in future, AI will also entail AC. The advantage of calling it George is that it prevents anyone from confusing it with an air conditioner.