Embrace the unexpected: To teach AI how to handle new situations, change the rules of the game

April 8 2021, by Mayank Kejriwal

The Monopoly version of the author's AI novelty environment can trip up AI's that play the game by introducing a wealth tax, rent control and other unexpected factors. Credit: Mayank Kejriwal, <u>CC BY-ND</u>

My colleagues and I changed a digital version of Monopoly so that instead of getting US\$200 each time a player passes Go, the player is

charged a wealth tax. We didn't do this to gain an advantage or trick anyone. The purpose is to throw a curveball at artificial intelligence agents that play the game.

Our aim is to <u>help the agents learn</u> to handle unexpected events, something AIs to date have been decidedly bad at. Giving AIs this kind of adaptability is important for futuristic systems like <u>surgical robots</u>, but also algorithms in the here and now that decide <u>who should get bail</u>, who should <u>get approved for a credit card</u> and whose resume <u>gets</u> <u>through</u> to a hiring manager. Not dealing well with the unexpected in any of those situations can have disastrous consequences.

AI agents need the ability to detect, characterize and adapt to novelty in human-like ways. A situation is novel if it challenges, directly or indirectly, an agent's model of the external world, which includes other agents, the environment and their interactions.

While most people do not deal with novelty in the most perfect way possible, they are able to to learn from their mistakes and adapt. Faced with a wealth tax in Monopoly, a human player might realize that she should have cash handy for the IRS as she is approaching Go. An AI player, bent on aggressively acquiring properties and monopolies, may fail to realize the appropriate balance between cash and nonliquid assets until it's too late.

Adapting to novelty in open worlds

Reinforcement learning is the field that is largely responsible for "superhuman" game-playing AI agents and applications like self-driving cars. Reinforcement learning uses rewards and punishment to allow AI agents to learn by trial and error. It is part of the larger AI field of machine learning. The learning in machine learning implies that such systems are already capable of dealing with limited types of novelty. Machine learning systems tend to do well on input data that are statistically similar, although not identical, to those on which they were originally trained. In practice, it is OK to violate this condition as long as nothing too unexpected is likely to happen.

Such systems can run into trouble in an <u>open world</u>. As the name suggests, open worlds cannot be completely and explicitly defined. The unexpected can, and does, happen. Most importantly, the real world is an open world.

However, the "superhuman" AIs are not designed to handle highly unexpected situations in an open world. One reason may be the use of modern reinforcement learning itself, which eventually leads the AI to be optimized for the specific environment in which it was trained. In real life, there are no such guarantees. An AI that is built for real life must be able to adapt to novelty in an open world.

Novelty as a first-class citizen

Returning to Monopoly, imagine that certain properties are subject to <u>rent protection</u>. A good player, human or AI, would recognize the properties as bad investments compared to properties that can earn higher rents and not purchase them. However, an AI that has never before seen this situation, or anything like it, will likely need to play many games before it can adapt.

Before computer scientists can even start theorizing about how to build such "novelty-adaptive" agents, they need a rigorous method for evaluating them. Traditionally, most AI systems are tested by the same people who build them. <u>Competitions</u> are more impartial, but to date, no competition has evaluated AI systems in situations so unexpected that not even the system designers could have foreseen them. Such an evaluation is the gold standard for testing AI on novelty, similar to <u>randomized controlled trials</u> for evaluating drugs.

In 2019, the U.S. Defense Advanced Research Projects Agency launched a program called <u>Science of Artificial Intelligence and</u> <u>Learning for Open-world Novelty</u>, called SAIL-ON for short. It is currently funding many groups, <u>including my own at the University of</u> <u>Southern California</u>, for researching novelty adaptation in open worlds.

One of the many ways in which the program is innovative is that a team can either develop an AI agent that handles novelty, or design an openworld environment for evaluating such agents, but not both. Teams that build an open-world environment must also theorize about novelty in that environment. They test their theories and evaluate the agents built by another group by developing a novelty generator. These generators can be used to inject unexpected elements into the environment.

Under SAIL-ON, my colleagues and I recently developed a simulator called <u>Generating Novelty in Open-world Multi-agent Environments</u>, or GNOME. GNOME is designed to test AI novelty adaptation in strategic board games that capture elements of the real world.

Our first version of GNOME uses the classic board game Monopoly. We recently <u>demonstrated the Monopoly-based GNOME</u> at a <u>top machine</u> <u>learning conference</u>. We allowed participants to inject novelties and see for themselves how preprogrammed AI agents performed. For example, GNOME can introduce the wealth tax or rent protection "novelties" mentioned earlier, and evaluate the AI following the change.

By comparing how the AI performed before and after the rule change, GNOME can quantify just how far off its game the novelty knocked the AI. If GNOME finds that the AI was winning 80% of the games before the novelty was introduced, and is now winning only 25% of the games, it will flag the AI as one that has lots of room to improve.

The future: A science of novelty?

GNOME has already been used to evaluate novelty-adaptive AI agents built by three independent organizations also funded under this DARPA program. We have also built GNOMEs based on poker, and "war games" that are similar to <u>Battleship</u>. In the next year, we will also be exploring GNOMEs for other strategic board games like <u>Risk</u> and <u>Catan</u>. This research is expected to lead to AI agents that are capable of handling novelty in different settings.

Making novelty a central focus of modern AI research and evaluation has had the byproduct of producing an <u>initial body of work</u> in support of a science of novelty. Not only are researchers like ourselves exploring definitions and theories of novelty, but we are exploring questions that could have fundamental implications. For example, our team is exploring the question of when a novelty is expected to be impossibly difficult for an AI. In the real world, if such a situation arises, the AI would recognize it and call a human operator.

In seeking answers to these and other questions, computer scientists are now trying to enable AIs that can react properly to the unexpected, including <u>black-swan events</u> like COVID-19. Perhaps the day is not far off when an AI will be able to not only beat humans at their existing games, but adapt quickly to any version of those games that humans can imagine. It may even be capable of adapting to situations that we cannot conceive of today.

This article is republished from <u>The Conversation</u> under a Creative Commons license. Read the <u>original article</u>. *This story is part of* <u>Science</u> <u>X Dialog</u>, where researchers can report findings from their published</u> research articles. <u>Visit this page</u> for information about ScienceX Dialog and how to participate.

Provided by The Conversation

Citation: Embrace the unexpected: To teach AI how to handle new situations, change the rules of the game (2021, April 8) retrieved 29 April 2024 from <u>https://sciencex.com/news/2021-04-embrace-unexpected-ai-situations-game.html</u>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.