



Guest

Machines are becoming more creative than humans

RISTO MIIKKULAINEN, SENTIENT ([HTTP://VENTUREBEAT.COM/AUTHOR/RISTO-MIIKKULAINEN-SATIENT/](http://venturebeat.com/author/risto-miikkulainen-satient/)) APRIL 3, 2016 12:23 PM

TAGS: ARTIFICIAL INTELLIGENCE (/TAG/ARTIFICIAL-INTELLIGENCE/), MACHINE LEARNING (/TAG/MACHINE-LEARNING/), RISTO MIIKKULAINEN (/TAG/RISTO-MIIKKULAINEN/)




Image Credit: Mopic/Shutterstock

Can machines be creative? Recent successes in AI have shown that machines can now perform at human levels in many tasks that, just a few years ago, were considered to be decades away, like driving cars, understanding

spoken language, and recognizing objects. But these are all tasks where we know what needs to be done, and the machine is just imitating us. What about tasks where the right answers are not known? Can machines be programmed to find solutions on their own, and perhaps even come up with creative solutions that humans would find difficult?

The answer is a definite yes! There are branches of AI focused precisely on this challenge, including evolutionary computation and reinforcement learning. Like the popular deep learning methods, which are responsible for many of the recent AI successes, these branches of AI have benefitted from the million-fold increase in computing power we've seen over the last two decades. There are now antennas in spacecraft (https://en.wikipedia.org/wiki/Evolved_antenna) so complex they could only be designed through computational evolution. There are game playing agents in Othello, Backgammon, and most recently in Go that have learned to play at the level of the best humans, and in the case of AlphaGo (<http://www.sentient.ai/google-on-alphagos-win-over-world-go-champion-we-landed-it-on-the-moon-how-impressive-is-the-win-and-what-does-it-really-mean-for-the-future-of-ai/>), even beyond the ability of the best humans. There are non-player characters in Unreal Tournament that have evolved to be indistinguishable from humans, thereby passing the Turing test— at least for game bots. And in finance, there are computational traders in the stock market evolved to make real money (<http://video.cnbc.com/gallery/?video=3000501862&play=1>).

Passing the Turing test for Video Games



Passing the Turing test for video games: The AI is indistinguishable from human players.

These AI agents are different from those commonly seen in robotics, vision, and speech processing in that they were not taught to perform specific actions. Instead, they learned the best behaviors on their own by exploring possible behaviors and determining which ones lead to the best outcomes. Many such methods are modeled after similar adaptation in biology. For instance, evolutionary computation takes concepts from biological evolution. The idea is to encode candidate solutions (such as videogame players) in such a way that it is possible to recombine and mutate them to get new solutions. Then, given a large population of candidates with enough variation, a parallel search method is run to find a candidate that actually solves the problem. The most promising candidates are selected for mutation and recombination in order to construct even better candidates as offspring. In this manner, only an extremely tiny fraction of the entire group of possible candidates needs to be searched to find one that actually solves the problem, e.g. plays the game really well.

We can apply the same approach to many domains where it is possible to evaluate the quality of candidates computationally. It applies to many design

domains, including the design of the space antenna mentioned above, the design of a control system for a finless rocket (<http://nn.cs.utexas.edu/pages/research/rocket/>), or the design of a multilegged, walking robot (<http://nn.cs.utexas.edu/demos/enso-robots/>). Often evolution comes up with solutions that are truly unexpected but still effective — in other words, creative. For instance, when working on a controller that would navigate a robotic arm (https://www.researchgate.net/figure/2816221_fig1_Figure-1-The-OSCAR-6-robot-arm-OSCAR-is-designed-for-pick-and-place-tasks-and-has-been) around obstacles, we accidentally disabled its main motor. It could no longer reach targets far away, because it could not turn around its vertical axis. What the controller evolved to do instead was slowly turn the arm away from the target, using its remaining motors, and then swing it back really hard, turning the whole robot towards the target through inertia!

The most recent and, in my opinion, the most exciting research in this field focuses on computational design creativity head on. One idea that has emerged, again modeled after biology, is that evolutionary computation should not be set to optimize a particular design objective but instead should be set to simply discover solutions that are novel. Many difficult problems are deceptive — if you try to solve them by making incremental improvements, you will get stuck. Novelty search instead discovers stepping stones, such as candidates that may not perform well but exhibit a highly unique approach. Often a truly creative solution can be found by combining the novel features of several candidates into a single solution that works. For example, it is possible to evolve a fast walking gait for a bipedal robot not by trying to incrementally walk faster and faster but by allowing it to fall on its face as fast and hard as possible and then evolving a way to postpone the fall by taking steps.

Biped Comparison Video



Many new applications have suddenly come within our reach thanks to computational creativity — even though most of us do not realize it yet. If you are facing a design problem where potential solutions can be tested automatically, chances are you could evolve those solutions automatically as well. In areas where computers are already used to draft designs, the natural next step is to harness evolutionary search. This will allow human designers to gain more traction for their ideas, such as machine parts that are easier to manufacture, stock portfolios that minimize risk, or websites that result in more conversions. In other areas, it may take some engineering effort to define the design problem for the computer, but the effort may be rewarded by truly novel designs, such as finless rockets, new video game genres, personalized preventive medicine, and safer and more efficient traffic.

And with all that time saved, we humans will have more time for creative pursuits of our own.

Risto Miikkulainen is a Professor of Computer Science and Neuroscience at the University of Texas at Austin (<http://www.utexas.edu/>) and a neuroevolution pioneer. He is also a fellow at AI startup Sentient Technologies (<http://www.sentient.ai>).