



How Automating Evolution Could Give Machines Imaginations

Written by **JORDAN PEARSON**

August 27, 2015 // 12:49 PM EST

Imagine if, decades from now, a computer spat out a brilliant design for a vehicle so novel that we humans have no choice but to call it an act of imagination.

One day, University of Wyoming AI researcher [Jeff Clune](#) plans to build this kind of artificially intelligent agent, and he already has the blueprints.

Clune calls this kind of agent an “innovation engine,” and while it sounds a little bit like Skynet Jr., it’s really just a clever set-up involving two warring AIs. On one side, there is an [evolutionary algorithm](#) which creates endless mutations of an object, image, or design. This is the “artist.” On the other side, a [deep neural network](#) (a cluster of virtual nodes designed to mimic, in extreme simplicity, the brain) evaluates which of these mutations is both interesting and functional. This is the “judge.”

“We’re trying to recreate what happened on planet Earth,” Clune told me over the phone. “The process of evolution has created all the diversity of amazing creatures we see, from jaguars to hawks and humans.”

["You could challenge the system to produce a huge diversity of](#)

solutions to the problem of [robot] locomotion, or computer chip design, or airplane wing design"

After enough back-and-forth between the artist and judge programs, the innovation engine could eventually spit out something never before imagined by people. The secret sauce in all this is that the engine explores possibilities in all directions, storing interesting solutions along the way, even if they seem like a devolution at first. The result is a system that doesn't get [trapped in local optima](#)—when an algorithm isn't designed to take a perceived step backwards, even if it means eventually leaping ahead.

The innovation engine is further explained [in a paper](#) presented by Clune and his colleagues in Wyoming and at Cornell University at this year's [Genetic and Evolutionary Computation Conference](#) in Madrid, where it won best paper.

The researchers tested a very simple implementation of the innovation engine by applying it to an image generation algorithm. Usually, these kinds of algorithms end up producing patterns that [look like trippy desktop wallpapers](#). After training a deep neural network on real-world data to recognize categories of images, and inserting it into the loop as a "judge," the two AIs worked together to create immediately recognizable images like a skull and a butterfly.

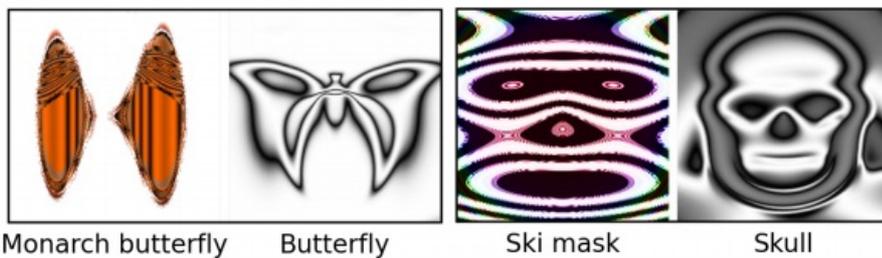


Image: Clune, et. al

Image generation is just an experimental first use case, Clune said. One day, he hopes that this technique could be used to solve real-world engineering problems. "For example, you could challenge the system to produce a huge diversity of solutions to the [problem of \[robot\] locomotion](#), or [computer chip design](#), or airplane wing design, etcetera," Clune said.

Watch more from Motherboard: Rise of the Killer Robots

Huge technical challenges remain before algorithms start pumping out designs for flying cars, however. The team's current approach uses supervised machine learning, which means they told the neural network everything it needed to know beforehand. You might have come across this kind of algorithm if you played around with [Google's Deep Dream tool](#) for creating trippy images. In that case, Google engineers got a neural net to build categories of things it recognizes based on a set

or training images, and it worked, but it wasn't perfect. Its idea of a dumbbell included weird, disembodied arms, so clearly there is more work to be done even on this end.



Image: Google Research

A real breakthrough for an AI innovation engine would come from the use of unsupervised learning techniques, Clune said. Unsupervised learning occurs when a neural network is fed a bunch of data and left to its own devices. Currently, supervised approaches **work much better** than unsupervised ones in most cases, and so unsupervised learning remains a moving target for researchers.

"This is one of those ideas that I have a feeling I will spend my career working at and making improvements on," Clune said. "It would be extremely exciting to get to the long term vision of an automated creativity and innovation engine. We recognize there are many challenges there, but it's not impossible."

--

TOPICS: artificial intelligence, ai, robots, deep learning, deep neural networks, artificial neural network, evolution, algorithm, image generation, google, deep dream, University of Wyoming, Cornell University, GECCO, research, future, machines

SHARE

f FACEBOOK

g+ GOOGLE PLUS

REDDIT

4 COMMENTS

TWITTER

TUMBLR

STUMBLEUPON

RECOMMENDED



The Icelandic Institute for Intelligent Machines Now Has a Unique Ethics Policy



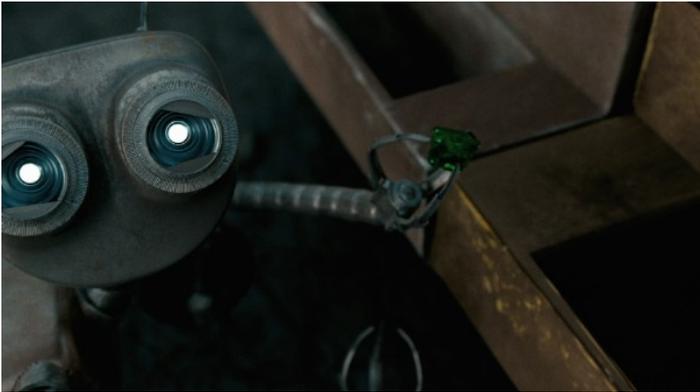
Facebook's Personal Assistant Is Real, and You Can Talk to It Like a Friend



Paging Dr. Watson: IBM Supercomputer Gains Access to 30 Billion Medical Images



Tropical Premises

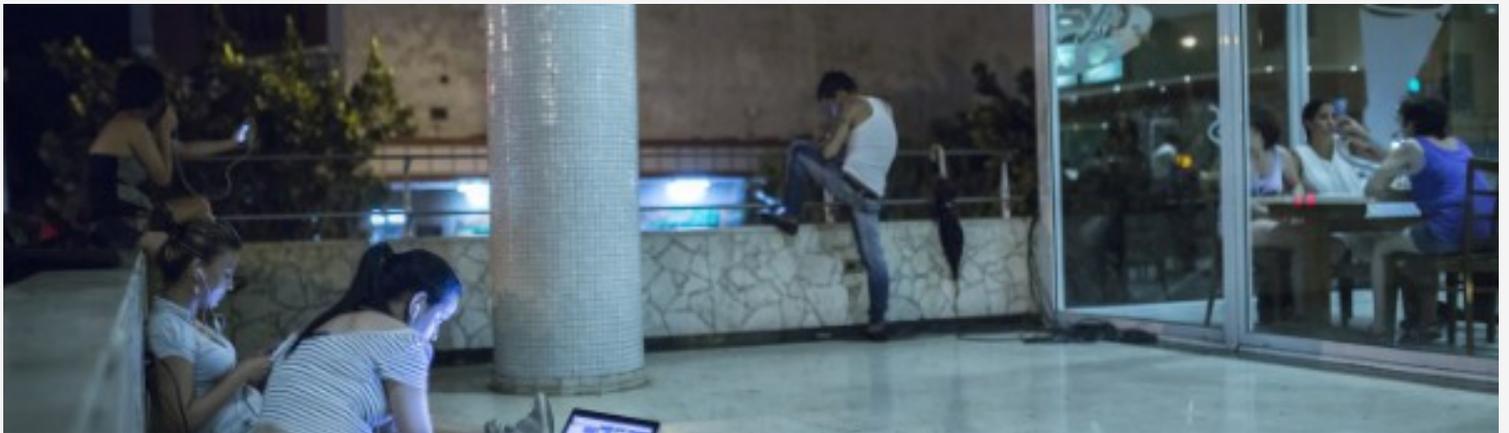


The Most Adorable Film Ever Made About the Perils of Space Capitalism

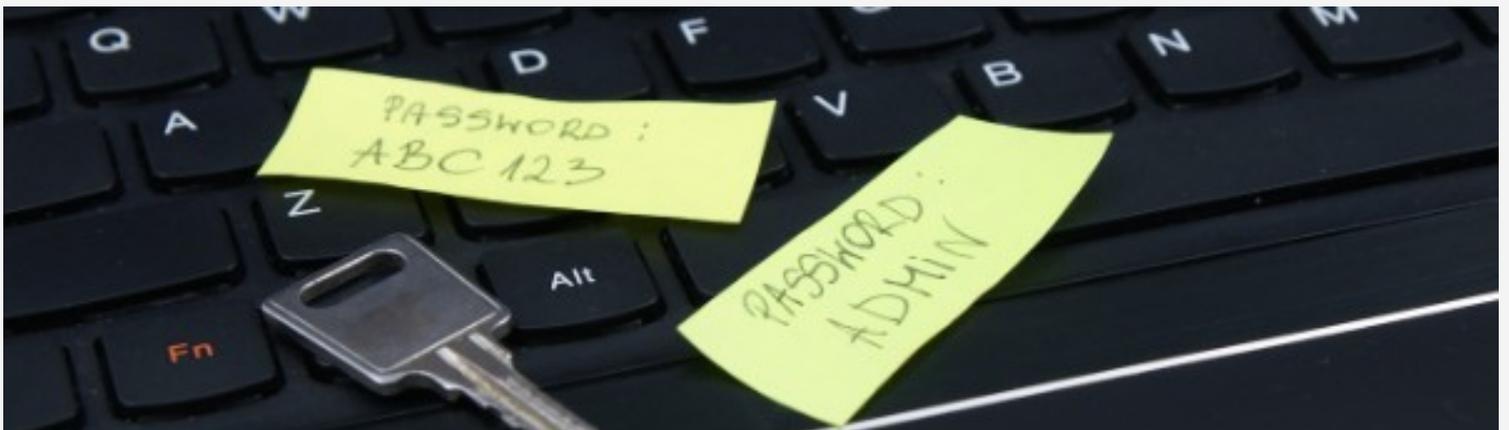


These Animations Were Created with an Ultra-Precise Light-Painting Drone

MOST POPULAR



What It Looks Like to Use the Internet for the First Time



Someone Cracked 4000 Ashley Madison Passwords, And Loads of Them Are Awful



Inside NASA's Giant Flying Laboratory



Ashley Madison Hackers Speak Out: 'Nobody Was Watching'



'I'm Desperate:' Ashley Madison Users Confide in a Security Researcher



The Dark Web's Biggest Marketplace Is Going on Hiatus

COMMENTS

comments powered by [Disqus](#)