Deep Dreams are Made of These

Over the summer, the Internet went into a collective frenzy for what looked like psychedelic images more fitting for a late 1960s dorm room, not for the most popular tech news sites. Images of dogs sprouting other dogs in fractal delight, and other hallucinatory apparitions took over the Internet. I'm describing Google's "DeepDream" image software.¹ It seemed like everyone and their uncle were posting crazy, deepdreamed images on my Facebook feed for a few weeks. Of course these computers weren't dreaming, or at least not the way we do. A complex artificial neural network

UPCOMING ISSUES

Winter 2015

[December issue]

Internet of Things

Spring 2016

[March issue]

Digital Fabrication

Summer 2016

[June issue]

Cultures in Computing

trained on images, not REM sleep, is responsible for creating these bizarre creations.

What began as a submission for the 2014 ImageNet Large-Scale Visual Recognition Challenge, became software created to identify objects and patterns. The DeepDream software takes that code and optimizes input images so that its system can find features in images for easier classification.

Deep learning has been advancing the field of artificial intelligence at Google and beyond. These levels of artificial neural networks (ANN), which

are generated from unsupervised learning, have existed for a long time (ANNs were hot in the 1980s), but now we have more data. In many cases the features the system finds make little or no sense at one level, but contribute meaningfully many levels up. But now deep learning is starting to make big progress in speech and image recognition. Graphics and HCI researchers are starting to use these tools to produce exciting results; in a recent meeting, a colleague suggested using deep learning to classify some of my data. Once the HCI people start using it, you know it must be a big deal.

All of this comes out of an AI approach that starts to look a lot more like how our brain works than tradi-

¹ <http://googleresearch.blogspot.co.uk/2015/07/deepdream-code-example-for-visualizing.html>

Looking toward the horizon, these evermore intelligent machines mean new, evermore intelligent applications in almost every domain.

tional machine learning approaches. Of course, we still have little understanding of how the brain works on a fundamental level—many of the interesting questions in neuroscience are

often answered by a “we don’t know.” But more and more researchers are seeing the importance of lots of data, or exposure to events, as central to our learning and perception. Philosophers like Alvin Noë have described the importance of embodiment in learning and cognition. AI researchers have long aspired to embody these principles in their research, but for the past few decades support-vector networks (SVNs) and machine learning have prevailed. Now, big data and robotics might change this.

Looking toward the horizon, these evermore intelligent machines mean new, evermore intelligent applications in almost every domain. But there may be darker applications: ImageNet can be useful for labeling YouTube videos, but it can also be used to identify enemy combatants. That is why a number of AI researchers and technologists have recently been pushing a ban on AI weapons that have the ability to decide who or what to target.² Some say an AI arms race may push us in an undesired direction, and that “autono-

mous weapons will become the Kalashnikovs of tomorrow,” easy to produce and hard to control. Signatories include Elon Musk, Stephen Hawking, and AI experts like Stuart Russell, Eric Horvitz, and Peter Norvig.

Of course, we are not there today. As tech journalist John Markoff pointed out during a recent interview at the Computer History Museum in Silicon Valley, robots can barely open doors. The recent DARPA robotics challenge, although clearly displaying some impressive strides in the field, had a human operator in the loop and still resulted in viral videos of robots falling over.

Further down the line we will need to think about not only the ethics of robot use, but also the rights of robots as well. In a clear effort to beef up my nerd cred, I’ve recently become addicted to “Star Trek: The Next Generation.” One of the most interesting story lines on the show, or at least for an arm-chair computer science ethicist like myself, is that of Lt. Commander Data—one of two androids in the galaxy—and his quest to understand what it means to be human. In the episode “Measure of a Man,” Starfleet wants to make copies of Data for military means but to do so may involve potentially disassembling him. Data must convince a Starfleet court that he is alive and sentient. Just as the current fight for animal legal “personhood” for captive apes or circus animals has made us question the rights of animals, would my Roomba or a future version of it need such freedoms?

Now is the time to think about these issues. Can we hard code ethics and morals into these robots? How can we teach ethics to robots if we can’t yet teach it to the engineers who are programming the robots? With large-scale engineering scandals like the Volkswagen Diesel exhaust regulation fraud—the automaker was caught with systems that falsified pollutant numbers when tested—it seems clear that when there are motivations for financial gain ethical lines become blurred. Before we can have ethical robots, we have to have more ethical engineers.

—Sean Follmer

With a focus on network security, the MSCS positions students with advanced technical skills essential to today’s programmers.

- One-year or part-time formats
- Evening classes
- Internship and networking opportunities
- Competitive tuition
- Scholarships and financial aid
- Optional test scores

www.merrimack.edu/gradcs
graduate@merrimack.edu • 978-837-3563

Merrimack
COLLEGE
NORTH ANDOVER, MA