

Machine Learning in the Creative Arts

Australian National University
Three Minute Presentation Competition

Tyson Jones

Feb 2017

You've probably heard of "machine learning". Computational models which gives computers the power to find sophisticated patterns in huge volumes of data, and make intelligent decisions. With them, we can do amazing science and engineering; classify DNA sequences, predict stock values, interpret images and drive cars.

It's a very deliberate name though: "machine learning", emphatically different to "human learning". Apparently it's hard to make computers think like humans. Although machines have now beaten our best humans at Chess and Go, these are just *silly* and *quantifiable* games, no? Surely a deterministic machine could never possess the qualitative creativity of a human being. Surely the spontaneity and imagination invoked in paintings, poetry and the other arts are all inherently and exclusively human. Surely?!

We challenge this, by teaching computers to compose music; a previously assumed human feat. We take neural networks - models of statistical inference inspired by our understanding of the brain - and expose them to tens of thousands of songs; "human music". Machines can then articulate patterns discovered in the music, into new music. My work specifically explored rendering transcribed music in a meaningful, seemingly 'human' way. A simple three note melody might render quite plainly in a naive MIDI player, but a trained violist would perform it much more convincingly, intuiting the dynamic, playing with vibrato, intonation and with particular timbre. We looked at re-purposing a recent breakthrough by Google text-to-speech to render MIDI files.

Our work is a very small part in something much greater; generating creative art with machines. The results can be profoundly convincing. Robot-written novels have come disturbingly close to winning blind literary competitions. Computers have identified characteristics of famous painter's style that humans missed, and used them to imitate original, stylistic painting.

It turns out, machines are very good at imitating human creativity. But that's *all* this is, right? *Imitation?* This turns the spotlight on us. How do we humans compose music? We don't create in vacuums; artists listen to thousands of hours of other music. Mozart *influenced* Chopin. My own composing is little more than stolen progressions from Rachmaninov and arpeggios from Matt Bellamy. Who is to say that the human creative process is anything different from how machines statistically digest their training data?

How far does this extend beyond music? Aren't the ways we develop accents, senses of humour, empathy and other assumed-human traits - that is, via lots of experience - disturbingly akin to machine learning? Is there nothing machines cannot learn in this way? Fear? Love? The ability to distinguish between butter and *I Can't Believe It's Not Butter!*

Although we might *feel* like human summer scholars, making amazing friends and new experiences... And although I was sent here to teach machines to compose music...

Maybe *I'm* the machine.
LeArNinG hOW To LovE