Artistic Intelligence

Ray LC [Luo]
Parsons School of Design
Brain and Mind Research Institute, Cornell Medical College
rayLC@newschool.edu [rayLC.org]

Abstract

Machine Learning (ML) has been applied in the financial, medical and educational fields to make, for example, smart stock predictors, hospital robots, and virtual assistants. But their use in one of the most human of endeavours, creative expression, has been relatively unexplored. Current applications of ML in artistic endeavours employ mostly artificial agents to extend human capabilities to realms where access to extensive data provides opportunities for associations previously unexploited by human artists. These examples take the human point of view first and merely expand human ability, by generating novel musical combinations based on a simple palette of tones, analyzing image content to pick out styles that serve as training for further image transformations, or joining poetic text based on phonetic similarities, for example. While these applications rely on ML as a data-mining agent in unexplored domains, they fail to exceed the limits of human expectations of what they can do. There's another arena in which ML enables artistic expression: using Artificial Intelligence (AI) in *unexpected* ways in everything we interact with. Imagine, for example, talking to a human whose responses are generated by Google Assistant, or interacting with a robot who secretly wants to make you take medication. I propose using ML to give novel behaviours to objects we interact with, allowing these behaviours to vary using predefined parameters for training, which are unknown to users. Applying ML to unexpected forms of interactions changes what we think machines are capable of, creating situations where AI goes beyond human expectations of what machine intelligence means to us, making objects oddly, Artistically Intelligent.

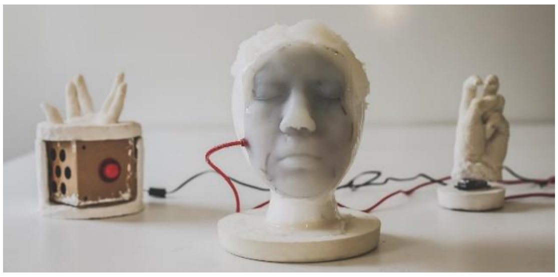

Fig 0. Artistic Intelligence by Ray LC: sculptures imbued with machine learning for creative expression. Source: Ray LC.

Introduction

Technology is taking over much of our daily lives. Instead of memorizing epic poems passed down through generations, as in Homer's era, humans invented books to record them. Now instead of using physical paper as media, we record information digitally and no longer need books. We went from talking, singing and memorizing, to recording, archiving and searching when we need something. These new tools have become integrated with human capabilities and made us more powerful, with the experience and findings of all previous generations available at our fingertips. If previously human capabilities, like way-finding, calculating and memorizing, can be overtaken by GPS, computer programs and the internet, what other fundamentally human abilities will be overtaken by the tools humans create?

The most unique thing about humans is our ability to express ourselves by creating. Animals and plants can transform their environments the way we do, but they have limited means of making tools to do their work, and they are even more limited in the way in which they create works of imagination. Studies have found cells in the monkey cortex that react to the use of tools, [1] but non-human primates are limited in what they can do in open-ended cognitive tasks, such as the inability to compose a picture. [2] Humans, on the other hand, can create entire worlds in their minds, invent hypothetical scenarios and stories, and evaluate them, and think of futures that may not correspond to reality. We use ideas imaginatively much as we use tools, talking about the hypothetical future based on "what if" questions. [3] Can this fundamentally human ability one day be transferred to tools we invent? Will we make Artificial Intelligence (AI) that creates with us, or even more capably, creates for us? Can we make an AI for Artistic Intelligence?

Our uniquely human creative potential comes not from particular domains, like painting or theatre, for many cultures exhibit creativity without having venues in which to express them. Instead, creativity can be defined in terms of the ability to shape and improve ideas adaptively in changing environments, [4] a task suitable for Machine Learning (ML) once the goal state of adaptation has been established. Tasks with a simple goal state, like winning a chess game have comparatively simple ML solutions, because algorithms can simply search for more and more effective ways of searching for a solution to winning the game. In creative endeavours, the goal state is less obvious to humans, so we are unable to create machines that do the task for us, just by virtue of the ambiguity of what that task is actually trying to do. A sculptor may create a sculpture as much for its likeness to someone in her life (a well-defined goal) as for a need to expose societal prejudices (a goal much harder to define digitally). Hence, creative expression has so far not been taken over by ML algorithms, because it's not clear what the algorithms should aim to achieve.

One approach is to use ML to achieve what human artists achieve by learning (copying) the process of artifact creation. In this scheme, any future "invention" by machines is coded for by the creator, and ML is only a tool for template-based creation. In contrast, another approach is to make ML agents part of a human ecosystem of creative works, exploiting our assumptions about what machines that have humanoid behaviours can or should do, giving voice to the machine's own Artistic Intelligence.

Background

The first approach of using ML to mimic human creativity started with computer programs used to make "novel" images. Harold Cohen's AARON robot was programmed by its creator to make abstract drawings based on predefined styles. Over the years AARON's output looked a lot like Cohen's own evolving style, leading to the question of what would happen after Cohen's death. Would AARON stop learning, and if so, was it ever really creative, or simply following patterns? Cohen's contention is that art does not require constant creativity, but rather devising rules to follow and allowing the pattern of rules to take over. [5] If this is the case, AARON is only a translator from patterns to artefacts, with some randomness added.

Fig 1. AARON: a robot used by artist Harold Cohen to make abstract images autonomously using a routine programmed to mimic Cohen's own style. Source: technologyreview.com.

Other examples of ML art based on emulating human styles and customizations include ventures in digital image processing, like the Pikazo app, which combines an image and a style embodied by a painter in the history of art or an uploaded texture to make a novel image combination. The role of ML in the app is to perform the combination process in a seamless manner, using image recognition algorithms. However, there's no creativity for the AI in this approach. Images from the Pikazo website show clear filter-like manipulation of images using the styles of various artists. Project Magenta dispenses with idea of machine creativity and instead focuses on algorithms that augment what human creators can do. For example, in Beat Blender, beat rhythms for music can be generated by drawing a path through a spatial-temporal state space of beats, allowing the musician to make creative content using an intuitive feel for beats in time and patterns in space. Project Magenta assumes that ML is used to heighten what humans can do by creating novel interfaces and creative combinations of basic palettes enabled by artists, not by having the algorithm generate ideas. Similar efforts in the textual domain have been undertaken to create machine-generated novels, such as Allison Parrish's *Our Arrival*.

Fig 2. Pikazo: an app that creates new images based on a preselected style and an image to be modified. Source: pikazoapp.com

While the majority of ML art projects use ML to drive creativity, another segment of artists have focused on what AI will do to the creative process by focusing on understanding the machine. In particular, they aim to understand what is it about machine data mining that undermines how people, as creatives, can interact with the world. For example, ML systems like Deep Mask and Tensor Flow enable online systems to categorize people into stereotypical forms and to use their private data to make inferences about their lives. [6] With machine surveillance fast becoming part of our future, artists like Merijin Bolink are wondering how best to understand machines in order to coexist with them. In his "Google's Eyes" project, he used Google's Goggles app to iteratively identify a sculptural object. First he created a ceramic tire, which when interpreted by Goggles, returned a list of items, which included a jawbone. Then Bolink made a plaster copy of the jawbone and had Goggles interpret it, which it identified as a hand. The complete 20-object series are placed together as a representation of how machines interpret human art, showing how the human creative potential may be subverted by machine recognition.

While some artists like Bolink fear the rise of ML in the creative process, others herald it as the next phase of our evolution. In an early treatise on machine creativity, Roger Schank suggests that creativity can be defined as innovative problem solving, and that looking for "near misses" allows machines to hone in on these miss patterns and come up with creative modifications. [7] In a similar vein, arguments have been made that human creative power can be supplemented by machine interfaces, which have access to a larger scope of data, which can serve as raw material for powerful creative acts. [8] A counter argument is that more data is not necessarily useful, for great artists have often been constrained in expressing their point of view, which makes their work particularly expressive given their limited scope. This can create powerful emotion in those who have had a similar experience. Perhaps artistic genius comes from a combination of ML-like exploration and human-like constraints, much like the way trans-humanism puts machines and humans together.

Fig 3. Google's Eye project: Each object is iteratively shown to the Google app Goggles, which gives suggestions related to what it sees using image recognition. Each suggestion then becomes the next object. Source: fastcodesign.com

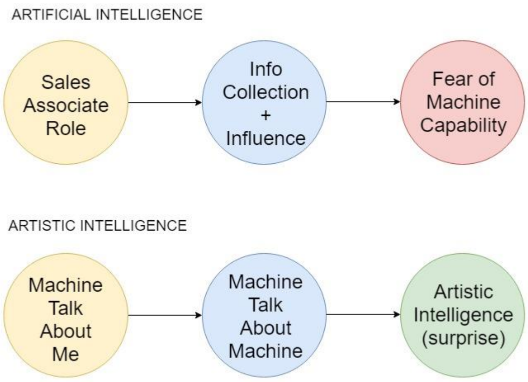

Fig 4. Example of AI in general domains. The machine is intended to be programmed for one area (sales specialist) but shocks the audience with human-level knowledge in another field (influence, data collection). In the artistic domain, AI does something unexpected based on preconceptions.

All the works discussed so far have applied ML to enable or enrich creative processes. A different approach to human-machine creativity interaction is to realize that our reaction to machines and what they are supposed to be capable of in human terms can be used to imbue them with intelligence and perceived emotionality and creativity. To allow machines to go beyond human creative potential, we have to go beyond just what machines are capable of, and instead, think about what is it in humans that makes us think that this is what machines can do. Creativity is about remaking processes, not artefacts. What makes this process unique is that by using what humans believe about machines to subvert our preconceived notions, we are making both humans and machines more creative. We are more creative because we can make tools that transcend boundaries to allow them to work closely with us. Machines are more creative because to the audience, they are doing more than what stereotypical machines do.

There's a natural consequence to the approach of using ML to transform what we think machines should do, which is that our fears about machines posing as humans or knowing our every move will manifest itself as uncertainty as to which part of the machine's response is from the machine and which is from its programmer. This point is akin to going to a website that offers interactive chats with a "sales specialist." After asking her a few questions, you get the feeling that she is not from your country and that perhaps she is contracted from a foreign country, because her replies are accurate but she uses unusual phrases. As the order proceeds you realize that she has the uncanny ability to know exactly what you have been searching for and knows your online identity and buying history from the past several months. Is she a person or a machine? Does it matter? Predictable AI does not produce a creative machine. Truly creative AI machines will possess an aura of mystery, which neither the programmer nor the machine can explain. Using ML to subvert what we think about ML puts us in a world where machines and humans are equals in their ability to influence: one is better at data; the other is better at language; one is better at analysis; the other is better at emotional response. If the AI machine is unexpected, it seems creative.

Process

To demonstrate the power of ML for creating smart objects capable of unexpected interactions with people, I created a set of sculpture pieces that incorporate digital technology using ML to predict and control, and occasionally, to surprise. Sculpture has the connotation of being inactive, because it usually remains inside a museum or in a fixed public space. What's more, it is usually considered serious and high-brow due to its association with classical works of art and intellectualism. I chose sculpture as the domain of experimentation because I wanted to challenge these two stereotypes about sculpture by creating pieces that interact instead of being sedentary, and that exhibit quirky and unexpected behaviour instead of being profound and unexciting.

To begin, I observed that ML algorithms start from the premise of using observable states coupled with desired outcomes to predict future observations, using a learning algorithm to update the network to make the predictions more accurate. [9] I asked if ML agents are really making predictions based on observations, how would a humanoid version that behaves similarly be interpreted by humans. I made a hand sculpture that rotates either left or right using an embedded servo motor. The gesture is meant to convey the act of "looking" by the sculpture and prompts the audience to respond with the same gesture. When a person moves close to the hand sculpture, the sculpture uses an ultrasonic sensor to detect the person's presence, and turns to face right or left randomly. However, the distance between the left and right sides with respect to the sensor is different, so ML can train it to determine whether the person is on its left or right and to train itself to adapt to the sequence of human hand movements. Using this data, the sculpture learns to predict whether the next hand motion from the human will be to the left or right of it, and will move there in anticipation. The predictions become more and more accurate over time as data is accumulated to drive the ML. The algorithm takes the average of recently detected locations and forms a maximum likelihood estimate of where the hand will be next, which is a form of time series prediction (see https://recfreq.wordpress.com/portfolio/ai-artistic-intelligence/).

Fig 5. Hand sculpture that predicts where the interaction with it will stem from. The servo for rotation is controlled using a microcontroller that detects user distance using an ultrasonic sensor. The learning algorithm predicts future user positions by keeping track of the averaged time series of previous responses. Source: Ray LC.

In user tests, I found that it was difficult to have people continue to interact with the sculpture to see the effect of the training. The ultrasonic distance sensor is also occasionally finicky, making data filtering necessary to maintain the accuracy of the sensor data for prediction. Moreover, the sculpture direction can be randomly correct early on, because there are only two possible states, so the error rate is only 50% even without learning. This can mask the progress that the sculpture makes over time. However, if observers are committed to watching the development over time, they can see the learning undertaken by the ML agent. Audiences also find the statue engaging, because plaster hands don't usually move. The statues with interactive components were considered "cute" by some observers. Many were also surprised by its ability to move, and those who had the patience to observe its learning found the adaptability of the statue to be evocative. The canonical view of an immobile sculpture was replaced by an interactive element, which I will continue to explore in other modalities. Thus, I have shown that a motorized sculptural piece capable of learning about its audience can use ML to enrich its interaction and evoke positive unexpected responses, contrary to its stuffy classical stereotype.

Fig 6. A “Star Trek” signalling hand sculpture, fortified by a raspberry pi running the Google Speech API. It interprets user voice input and replies, distorting the original meaning. The user is prompted to press the red button and say anything with the word “sculpture” in it. The sculpture adapts the words accordingly. Source: Ray LC.

Next, I wanted to take the unsuspected sculptural agency idea one step further by making a talking sculpture that appeared to have some capabilities of creative speech production. I used the ML in the Google Cloud Speech API executed on a raspberry pi as a starting point to create my own style of machine speech interface. The audience is prompted to press a button and say something involving or about “sculpture.” A computerized voice reply comes back from the sculpture, which is a plaster mould of a hand doing the *Star Trek* Vulcan “peace and prosper sign.” The *Star Trek* reference here is intentional, for it evokes future technology and thought in a traditional sculptural form. The peace and prosperity metaphor also subtly prompts the audience to talk to the sculpture as if it is a character in a movie with agency, and evokes the sensibilities of smart devices that serve human needs and work cooperatively without conflict, much as the vulcans in *Star Trek* operate. Using speech recognition and custom routines based on the Google Speech API, which uses ML to recognize words, I trained the statue to answer not only repeating what the user says, but saying it as if it has agency (see video).

Fig 7. A head sculpture that uses computer vision to see where the user is, and replies using digital code embodied as an LED matrix that sweeps across the mouth of the sculpture, representing machine communication.

For example, whenever the user says “sculpture,” the sculpture replies with a different noun, which first appears to be referencing the user. But as the interaction proceeds, it also changes the pronunciation and verbs, and the user notices that the sculpture is using the previous noun to refer to itself, not the user. The statue is seen to have made a creative transformation in the user’s view, not by the way it has changed its interaction style, but in the way in which the audience discovers what is algorithmically already there. In user tests, the only instruction I gave was to tell the users to say anything they wanted referencing “sculpture,” but what occurred is that the users learned more and more about the rules of engagement undertaken by the statue. One user said that she thought the statue was subservient and complimentary at first, but then over the course of the interaction, it became “sassier.” What changed was not the rules, but the potential for the ML agent to surprise (and annoy) users. The form of the hand gesture as a *Star Trek* symbol was key as well, for users say that they expected the statue to be “high-minded and calm,” but they actually had a contentious exchange, in which both user and statue claimed to be the superior agent. Interestingly, it’s not the ML part (voice recognition and understanding) that made the surprising results possible, but rather the human intervention that involved swapping the text. Thus, I created a speech-producing statue capable of surprising and evoking an emotional reaction from users.

As a final exercise, I wanted to extend the idea of creative production further than simply unexpected interactions. I decided to focus on visual representations after having previously explored the physical and language arenas. Although inspired by the ML algorithms for image association used by Google and Pikazo, I wanted to situate the piece so that the sculpture is the agent behind the “deep dreaming” undertaken by ML agents. Unlike previous efforts, I wanted to create a physical interface that appears to be producing the creative output, so that it’s not a computer using user input to create modified dreams, but the sculpture itself which makes content based on who and where the user is. To evoke the perception of creativity, I gave the machine a human face. Humans are distinguished by their ability to manipulate and communicate using language and by their ability to creatively express themselves. I put both of these agencies in a traditionally inanimate sculpture by putting an LED matrix behind the silicone-based sculpture. The wood-grain-embedded silicone retains the form of a classical statue, but forms a mesh that has hidden within it the ability to express itself. The LED matrix appears to respond to human touch due to its proximity to the silicone layer. Using Arduino to control the matrix, I created custom animations that evoked visual creation from the mouth of the statue when the user's face was detected by an attached camera. The animations depend on where the human face is. I wanted to make a connection between human speech and machine data processing. Whereas we can express our creativity by make speeches, writing novels, or creating worlds using language, for example, the machine analogue is not human language as we know it, but a machine code that we can only visualize across a layer that blurs communication. Just as we as 3D beings cannot contemplate life in 4D, we don't understand machine creative processing and the ways it can express itself as a form different from human conception. As humans, we can only hope to visualize the data machine produce across a layer of uncertainty. Again, it's not the ML aspects (computer vision, image recognition, etc.) that made the sculpture surprising, but the human intervention that appeared to reveal machine "thought" using the LED matrix.

Fig 8. The head sculpture lights up when a face is detected, but also moves its pixels based on where the face is in space. In this example, the face of the person whose face was cast for the sculpture is detected by the statue.

Users found the silicone face and LED matrix frightening at first. The red matrix evokes a type of bloodiness associated with the mouth. They found the pattern of the matrix display mesmerizing, because it tends to change form when they put their fingers on different parts of the silicone. The computer vision interaction provides users with a feeling of agency, because the light comes on only when they are close to the sculpture, and appears to track their face, a type of digital productivity. Unlike traditional sculptures, my piece evokes creative potential that contrasts with the classical form. One user said that it reminded him of the way machines would speak to each other if they were to communicate, because it "doesn't say the same thing twice." The silicone layer masks the lit up digital LEDs, so the effect is a filtered view of what machines would do creatively if they were creative. In summary, I created a digital machine metaphor for human creativity, which can be experienced through a filter established by classical forms.

Directions

The tools we create are taking over our lives. From recording our memories onto physical pages to analyzing the consequences of business investments; from enabling communication over long distance to interpreting our speech and predicting our desires, digital machines enabled by ML are going from helping us to enabling us to *thinking* for us. Will the most unique characteristic of humans, that of creative expression, be the next bastion to fall? Experiments with machine creativity have centred on using ML to help or imitate the human creative process. This strategy, however, is based on an anthropomorphic view that the way humans express themselves is the basis for all types of creative works, including those of machines, much as the Turing Test inherently situates machines within the human space with no regard for how non-human processes work. [10] I proposed that machine artistic expression can emerge instead from exploiting what humans think of objects and devices, allowing ML to subvert traditional forms, coalescing into a system of creative expression beyond simply generating data from modifying previous models. In this view, the context and situation of the use of ML is just as important as algorithms, enabling a world permeated by creative machines. Indeed, we may be making machine creative expression possible not by simply coding it into their algorithms, but rather by changing the way we think about machines and how they operate. In short, the more we know about our tools, the more we learn about ourselves and our own Artistic Intelligence.

References

1. P. F. Ferrari, S. Rozzi, and L. Fogassi, "Mirror neurons responding to observation of actions made with tools in monkey ventral premotor cortex," *Journal of Cognitive Neuroscience* 17, no. 2 (2006).

2. M. Vancatova, "Creativity and innovative behavior in primates on the example of picture-making activity of apes," *NFU Psychology* 2, no. 2 (2008).

3. Anthony Dunne, and Fiona Raby, *Speculative Everything* (Boston: MIT Press, 2013).

4. C. D. Hondzel, and R. Hansen, "Associating creativity, context, and experiential learning," *Journal of Education Inquiry* 6, no 2 (2015).

5. Martin Gayford, "Robot art raises questions about human creativity," *MIT Technology Review* (2016). https://www.technologyreview.com/s/600762/robot-art-raises-questions-about-human-creativity/.

6. Trevor Paglen, "Invisible images (your pictures are looking at you)," L.A. Times, April 2014. https://thenewinquiry.com/invisible-images-your-pictures-are-looking-at-you/.

7. Ray Kurzweil, *The Age of Intelligent Machines* (Boston: MIT Press, 1990).

8. Clive Thompson, *Smarter Than You Think: How Technology is Changing Our Minds for the Better* (London: Penguin Press, 2013).

9. D. E. Rumelhart., G. E. Hinton, and R. J. Williams. "Learning representations by back-propagating errors." *Nature*. 1986: 323.

10. Benjamin Bratton, "Outing AI: Beyond the Turing Test," *New York Times* Opinionator, February 2015. https://opinionator.blogs.nytimes.com/2015/02/23/outing-a-i-beyond-the-turing-test/