Artistic Intelligence

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

Abstract

Machine Learning (ML) has been applied to financial, medical, and educational realms to make, for example, smart stock predictors, hospital robots, and virtual assistants. However their use in the most human of endeavors, that of creative expression, has been relatively unexplored. Current applications of ML to artistic endeavors mostly employ artificial agents to extend human capabilities to realms where extensive data access provide opportunities for associations previously unexploited by human artists. These examples take the human point of view first and merely expand their abilities, include 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, and joining poetic text based on phonetic similarities. While these applications rely on ML as a data mining agent over unexplored domains, they fail to exceed the limit of human expectations of what they 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 meds. I propose to use ML to give novel behaviors to objects we interact with, allowing these behaviors to vary using predefined parameters for training that are unknown to the audience. Applying ML to unexpected forms of interactions subverts what we think machines are capable of, creating situations where AI truly goes beyond human expectation of what machine intelligence should mean to us, making them oddly, Artistically Intelligent.

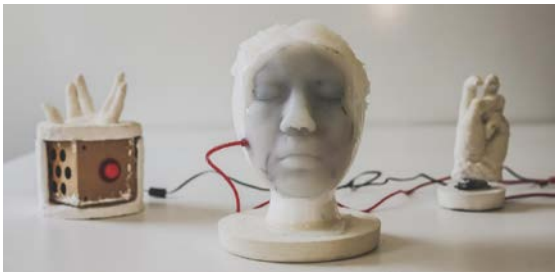

Fig 0. Artistic Intelligence by Ray LC: sculptures imbued with machine learning for creative expression. Source: Ray LC.

Introduction

Technology is taking over much of our daily lives. Instead of memorizing epic poems passed down through generations like Homer, we invented books to record them. Now instead of using physical paper as media, we record information digitally, no longer needing books. We went from talking, singing, and memorizing, to recording, archiving, and searching when we need something. The tools we have took over human capabilities and made them more powerful, making the experiences and findings of all previous generations available on our finger tips. 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 evolved to be overtaken by our tools?

The most unique thing about humans is their ability to express themselves 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 even more limited in the way they create works of imagination. Studies have found cells in monkey cortex that react to use of tools [3], but non-human primates have limitations on what they can do in open ended cognitive tasks, such as an inability to compose in a picture

making task [10]. Humans, on the other hand, can make entire worlds dreamt up in their minds, invent hypothetical scenarios and stories and evaluate them, and think of futures that may not correspond to realities. We mentally use ideas imaginatively much as we use tools, talking about hypothetical futures based on "what if" questions [2]. Can this fundamentally human ability be one day offloaded to tools that we invent? Will we make Artificial Intelligence (AI) that create with us, or even more capably, create for us? Can we make an AI for Artistic Intelligence?

The uniquely human creative potential comes not from particular domains like painting or theatre, for many cultures exhibit creativity without having the same venues to express them. Instead, creativity can be defined in terms of the ability to shape and improve ideas adaptively in changing environments [5], a task suitable for Machine Learning (ML) once the goal state of adaptation has been established. Tasks with simple goal states like "winning chess game" has comparatively simple ML solutions, because algorithms may simply find more and more effective ways of searching in state space for a solution to "win the game." In creative endeavors, the goal state is less obvious to humans, and hence we are not able 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 hence-so-far not been taken over by ML algorithms, because it's not clear what those algorithms should be aiming 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. Another approach is to make ML agents part of a human ecosystem of creative works, exploiting our assumptions about what machines that have humanoid behaviors can or should do, creating a new 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 is programed 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, begging the question of what would happen after Cohen's death. Would AARON stop learning, and if so, was it ever really creative, but rather simply following patterns? Cohen's contention is that art didn't require constant creativity, but rather devising rules to follow and allowing the pattern of rules to take over [4]. If that's the case, AARON is only a translator from patterns to artifacts, albeit 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. Pikazao 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 pikazo's website shows clear filter-like manipulation of images using styles of various artists. Project Magenta dispenses with idea of machine creativity and instead focus 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 always assumes that ML is used to heighten what humans can do by creating novel interfaces and creative combinations of basic palettes enabled by artists, but not by having the algorithm generate ideas. Similar efforts in the textual domain were undertaken as 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 uses 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 undermine how people as creatives can interact with the world. ML systems like Deep Mask and Tensor Flow enable online systems to categorize people into stereotypical forms and use their private data to form conclusions about their lives [7]. With a future of machine surveillance pending, 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 that included a jawbone. Then Bolink made a plaster copy of the jawbone, and allowed Goggles to interpret it, in this case as a hand. The complete 20 object series are placed together as a representative as what machines interpret human art works to be, 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 [6]. 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 the raw material for powerful creative acts [9]. A counter argument is that more data is not necessarily useful, for great artists have often been given constraints to their point of view which make their work particularly expressive of their limited scope, and can evolve powerful emotion in those who had similar experience. Perhaps artistic genius comes from a combination of ML-like exploration and human-like constraints.

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

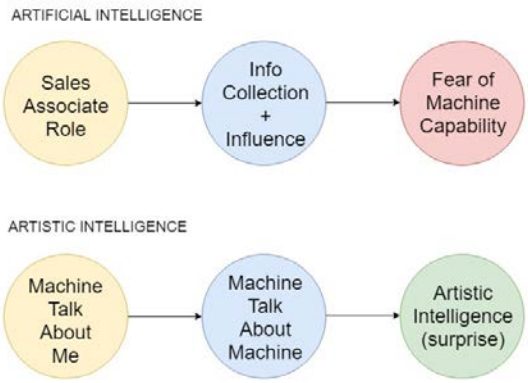

Fig 4. Example of AI in general domains. The machine is intended to be programmed for one area (sales specialist) but shocks the audience with human level knowledge of another (influence, collecting our data). In artistic domain, AI does something unexpected based on preconception.

All the works discussed so far have been applying 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 really allow machines to go beyond the 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 artifacts. 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 their boundaries and 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 chat with a "sales specialist." After asking her a few questions, you get the feeling that she is not American, that perhaps she is contracted from a foreign country, because her replies are accurate but has unusual use of phrases. As the order proceeds you realize that she has uncanny ability to know exactly what you have been search for and your online identity and history from the past several months. Is she a person or a machine? Does it matter? Predictable AI does not make for creative ones, and the truly creative AIs will possess an aura of mystery, that neither the programmer nor the machine alone can explain. Using ML to subvert what we think about ML puts us in a world where machines and humans are equals in their abilities to influence: one is better at data, another is better at language, one is better with analysis, the other at emotional response. The AI is unexpected, and hence feels creative.

Process

To unleash the power of ML for creating smart objects that exhibit unexpected interactions with the audience, 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 they usually sit inside a museum. What's more, they are usually considered serious and high-brow due to their association with classical works of art and intellectualism. I chose sculpture as the domain of experimentation because I want to buck these two stereotypes about sculpture by creating pieces that interact instead of being sedentary, and that exhibit quirky and unexpected behavior instead of being profound and unexciting.

To begin, I observed that ML algorithms start from the premise of using observable states coupled to desired outcomes to predict future observations, using a learning algorithm to update the network in order to make the predictions more accurate [8]. I asked that if ML agents really are 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 a sculpture, and prompts the audience to make the same gesture in response. When the audience comes close to the hand sculpture, it detects that a presence is close by using an ultrasonic sensor, and turns to face right or left randomly. However, the distance between the left and right sides with respect is the sensor is different, so the ML can use whether the audience actually is left or right of it as data 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 itself, and will move there in anticipation. The predictions become more and more accurate over time as data is accumulated to drive the ML (see https://recfreq.wordpress.com/portfolio/ai-artistic-intelligence/).

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

In user test, I found that it was difficult to have people stay in interaction with the sculpture to see the effect of the training. The ultrasonic distance sensor is also occasionally finicky, making the data filtering necessary to maintain accuracy of sensor data for prediction. Moreover, the sculpture direction can be randomly correct earlier on, because there are only two possible states, and the error rate is only 50% even without learning, it can mask the progress that the sculpture has made over time. However, if committed to seeing the development over time, one can surmise the learning undertaken by the ML agent. Audiences also find the statue engaging, because plaster hands don't usually move, and statues that have interactive components are considered "cute" by some of those in observation. Many were also surprised by its ability to move, and those who had the patience to observe found the adaptability of the statue to be evocative. The canonical view of an immobile sculpture is replaced by an interactive element, which I 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 response contrary to its stuffy classical stereotype.

Fig 6. A "Star Trek" signaling hand sculpture fortified by a raspberry pi running the Google Speech API. It interprets audience voice input and repeats back a reply that distorts the initial intent. The audience is prompted to press the red button and say anything with the word "sculpture" in it. The sculpture turns the phrasing to its own need. Sources: Ray LC.

Next, I wanted to take the unsuspected sculptural agency idea one step further by making a talking sculpture that appear 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 mold of a hand doing *Star Trek* Vulcan "peace and prosper sign." The *Star Trek* reference here is intentional, for it evokes future and technology thinking 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 that does ML to recognize words, I get the statue to answer back not only repetitions of what the user says, but to say it as if it has agency (see video).

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

For example, whenever the user says "sculpture," it talks back with a different noun that first appears it is referencing the user. However, as interaction proceeds, pronounces and verbs are also changed, and the audience is seen to notice that the sculpture is using the previous noun to refer to itself, not to the user. The statue is seen to have made a creative transformation in the audience's view, not by the way it has changed its interaction style, but in the way that the audience discovers what is algorithmically already there. In user tests, I only tell folks to say anything they want referencing "sculpture," yet what occurs is that users learn more and more about the rules of engagement undertaken by the statue. One user said that she thinks the statue was subservient and complimentary at first, but then over the course of the interaction, it became more "sassy." The rules didn't change, only the potential for the ML agent to surprise (and annoy) the audience. 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 actually had a contentious exchange where both user and statue claimed to be the superior agent. Thus I have created speech-producing statue capable of surprising and emotionally evoking audiences.

As a final exercise, I wanted to extend the idea of creative production further than simply surprising interactions. I decided to focus on visual representations after having explored the physical and language arenas previously. 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 appear 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 audience is. To evoke perception of creativity, I decided to let the machine take on the persona of a human face. Humans are distinguished by their ability to manipulate and communicate using language and in their ability to creatively express themselves. I put both of these agencies in a traditionally inanimate sculpture by putting a LED matrix behind the silicone-based sculpture. The wood-grain-embedded silicone retains the form of a classical statue yet 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 is detected by an attached camera. The animations are dependent on where the human face is. I wanted to make a connection between human speech and machine data processing. Whereas we express our creativity by make speeches, writing novels, creating worlds by language, the machine analog 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 also don't know 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.

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 color of the matrix evokes a type of bloodiness associated with the mouth. They found the pattern of the matrix display mesmerizing, for it tends to change form when they put their fingers on different parts of the silicone. The computer vision interaction allows audiences a feeling of agency for the light only comes on when they are near, and appears to track their face, a type of digital productivity. Unlike traditional sculptures, my piece evokes a creative potential that contrasts with the classical form. One user said that it reminds 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 that 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 consequences of business investments, from enabling communication over long distances 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 centered 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 disregard for how non-human processes work [1]. 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 model. In this view, the context and situation of the use of ML is just as important as algorithms, enabling a world where creative machines appear to permeate. The more we know about our tools, the more we know ourselves, and our Artistic Intelligence.

References

1. Bratton, B. "Outing AI: Beyond the Turing Test." New York Times Opinionator: February 2015. Online: https://opinionator.blogs.nytimes.com/2015/02/23/outing-a-i-beyond-the-turing-test/
2. Dunne, A., and Raby, F. *Speculative Everything*. Boston: MIT Press, 2013.
3. Ferrari, P. F., Rozzi, S., and Fogassi, L. "Mirror neurons responding to observation of actions made with tools in monkey ventral premotor cortex." *J. of Cognitive Neuroscience*. 2006: 17(2).
4. Gayford, M. "Robot art raises questions about human creativity." MIT Technology Review, 2016. Online: https://www.technologyreview.com/s/600762/robot-art-raises-questions-about-human-creativity/

5. Hondzel, C. D., and Hansen, R. “Associating creativity, context, and experiential learning.” *Journal of Education Inquiry*. 2015: 6(2).
6. Kurzweil, R. *The Age of Intelligent Machines*. Boston: MIT Press, 1990.
7. Paglen, T. “Invisible images (your pictures are looking at you).” LA Time: April 2014. Online: https://thenewinquiry.com/invisible-images-your-pictures-are-looking-at-you/
8. Rumelhart, D. E., Hinton, G. E., and Williams, R. J. “Learning representations by back-propagating errors.” *Nature*. 1986: 323.
9. Thompson, C. *Smarter Than You Think: How Technology is Changing Our Minds for the Better*. London: Penguin Press, 2013.
10. Vancatova, M. “Creativity and innovative behavior in primates on the example of picture-making activity of apes.” *NFU Psychology*. 2008: 2(2).