

Defining AI Arts: Three Proposals

Lev Manovich

June 2019

1.

On first sight, coming with a definition for “AI arts” does not sound hard. AI (an abbreviation for the term Artificial Intelligence) refers to computers being able to perform many human-like cognitive tasks, such as playing games of chess and Go, recognizing content in images, translating between languages, selecting best candidates in a job search based on their CVs, and so on. This is how AI has been traditionally understood, and we can extend this concept to the arts. Following this logic, “AI arts” would refer to humans programing computers to create with a significant degree of autonomy new artifacts or experiences that **professional members of the art world recognize as belonging to “contemporary art.”** Or, we can teach computers skills of artists from some earlier historical period and expect that professional art historians recognize new artifacts the computer creates as possible art from this period. (In one study, computer scientists asked art historians to evaluate images generated by a neural network to simulate styles of particular artists. (See Artsiom Sanakoyeu, Dmytro Kotovenko, Sabine Lang, Björn Ommer, "A Style-Aware Content Loss for Real-time HD Style Transfer", <https://arxiv.org/abs/1807.10201>.)

In fact, we **can extend the famous Turing test to AI arts** - if art historians mistake objects a computer creates after training for the original artifacts from some period, and if these objects are not simply slightly modified copies of existing artifacts, such computer passed “Turing AI arts” test. This sounds simple and logical. Let’s refer to this idea as our first proposal for the definition of “AI arts.” In this definition, art created by an AI is something that professionals recognize as

valid historical art or contemporary art.

Unfortunately, this logical approach is not sufficient. In fact, on closer inspection, its clarity dissolves. For example, there is no commonly accepted definition of "art" today among the professionals such as art critics, art theorists, philosophers of art, or sociologist of culture. So how can we program a computer to independently create something which we can't even define?

The development of modern art during the 20th century involved systematic questioning of the boundaries of what counts as art, and then going outside these boundaries - from Marcel Duchamp's readymades to happenings, performances, land works, and installations of the 1960s, to Internet art of the 1990s. But to understand what things can expand the boundaries of what counts as "art" at a given moment in a meaningful way requires knowledge of art history and development of the arts until the present - and **this is something nobody so far tried to program into a computer.**

Instead, most of the attempts to use AI techniques in the arts relied on (usually implicit) understanding of "art" that was popular before second modernist revolution of the 1950s-1960s (if we count 1880s-1920s as the first revolution). In other words, the artists, writers, composers and computer scientists taught computers to create objects in the formats that were accepted as art in the modern people up to late 1950s - single images, poems, music compositions. (By a strange coincidence, these experiments in AI arts begun at the same time as modern art enters its second revolutionary period, i.e. late 1950s. So, while some artists move beyond art as it existed up until that time, other artists start programming computers to create "traditional art," i.e. objects rather than processes, situations, and performances.)

This tendency is still with us. If we look at what recently (2015-) has been celebrated as achievements of AI in visual arts, these are often single images that look like modernist paintings. They may deliberately simulate visual appearance of some well-known modern artist, or simply look like some variations of expressionism,

cubism, post-impressionism, etc.

If we follow this conservative tendency, we have to accept that “AI arts” only simulate the historical art. It is not capable of executing the main strategy of modern art - **constantly expanding what counts as art**. (The genre of “interactive computer installation” that developed in the 1990s is one important exception). Of course, it is also possible to argue that in the early 21st century this strategy of expanding art lost its energy, we entered the period of pluralism, and creation of the “new” is no longer relevant. Still, this does not invalidate my main point - what has entered history as the achievements of “AI arts” during last six decades are **simulations of historical art** created before this AI arts work starts.

2.

Let’s try another approach. Instead of thinking about the outputs of an “art computer,” let’s consider the process of creation. Given that computers have been used in the arts in lots of ways for six decades, is there something unique about “AI arts”? Is it possible to make a clear distinction between “computer arts” (or “digital arts”) and “AI arts”?

One of the most popular method for using computers in the arts and design is writing computer programs that generate objects in various media (text, image, video, 3D shapes, graphic designs, logos, urban plans, music, etc.) Such programs can take variety of forms - simple instructions to draw a sequence of shapes, algorithms that generate fractals, cellular automata algorithms, genetic algorithms (Karl Sims), and so on. For example, the pioneering computer artists of the 1960s - **Vera Molnár, Desmond Paul Henry, Frieder Nake, Georg Nees, Michael Noll, Sonia Sheridan**, and others - wrote programs that generated geometric black and white patterns using precise instructions, while also sometimes incorporating random parameters. In design and architecture worlds, use of algorithms is often called procedural, generative, or parametric" design. This approach to design is widely used today in all design fields and it is responsible for some of the most

famous cultural creations of our times such as works by Zaha Hadid Architects.

Is there some fundamental distinction between such methods of computer arts that have been used for decades, and another paradigm that became very popular in 2010s - use of **machine learning and deep neural networks**? Note that AI field includes many approaches developed since 1950s. Machine learning and neural networks are only two among them. They became dominant in the industry in 2010s.

Neural networks paradigm includes a number of methods and some of them were adopted for generation of cultural artifacts. In one approach, the single network is trained using a large set of examples such as images in one style. Following the training, the network can generate more images in the same style.

In another approach called GAN (Generative Adversarial Network), generation of new artifacts involves two networks. One trained on a set of example creates new artifacts. These artifacts are evaluated by a second network and it selects the ones that are similar to the training examples.

In yet another approach called "style transfer," the network learns how to transfer a style from a single or a series of images to new images (or video) - for example, transferring a "style" of one van Gogh's painting to a photograph. (For examples of this work, see <https://arxiv.org/pdf/1807.10201.pdf>.) I think that this approach has a conceptual problem, because **an artist such as van Gogh does not have a "style"** - i.e. a form that exist independently from works' content. The particular transformations of visible world we see in van Gogh paintings are content specific - sky is transformed in one way, trees in a different way, etc. Therefore, van Gogh like images generated via style transfer method do not capture the real logic of his art, and the same holds for other examples generated with this method.

On the one hand, neural networks approach indeed departs from the methods of computer art and design developed earlier. With this technology we don't program a computer explicitly to generate new objects using a sequence of steps, a system or rules, or in some other way that we have to specify in all details. Instead, a

network itself extracts deep structure from a set of cultural artifacts and then generates new artifacts. Does this mean that we finally have real “artistic AI,” the true “art intelligence”?

Maybe not yet. There are at **least three points in this process where a human author makes explicit choices and controls** what computer would do. First, a human designs network architecture and also an algorithm used to train a network (or selects from the existing ones). Second, the human creates the training set. Third, the human selects what in her/his views are most successful artifacts from many more the network generates.

Given all this human curation and control, we can't claim that generation of cultural artifacts via machine learning / neural networks is more “intelligent”, i.e. shows the higher level of autonomy than any other computer art method. Each of these methods also includes human decisions & choices and execution of algorithms. Thus, machine learning is not more advanced form of artistic AI than geometric drawings of first computer artists, cellular automata artworks, or many interactive computer-driven installations. In fact, I think that machine learning approach is more restrictive than the earlier approaches, since a human makes decisions in so many points in the process. (And if we recall our earlier discussions of expanding boundaries of art, interactive installations are more interesting than a computer that generates van Gogh like images.)

How do we translate these arguments into another possible definition of “AI arts”? We can now say that **all methods developed in computer art since 1950s are equally valid instances of “AI arts”** - from a program in Processing generating geometric simple patterns or d3 code generating interactive data visualization to a deep neural network trained on very big data. What defines whether something is “AI” is not a method but **the amount and type of control** we exercise over algorithmic process.

3.

For our third attempt at “AI arts” definition, let’s focus now on the **key idea of supervised machine learning that uses neural networks**- the computer automatically extracting common patterns from a group of artifacts. This aspect of machine learning is indeed a new thing in long computer art history. A computer that by itself can learn the structure of the world is an impressive proposition, even such a computer still (or maybe be always) quite different from a human child doing this - because the world of training set objects carefully curated by a human engineer is an artificial world, far removed from the heterogeneity, diversity and noiseless of the real world that a child is exposed to (and also because we have to construct the network layers for extracting patterns ourselves, as opposed to the network evolving and constructing itself.)

So, shall we get excited if a computer that learned patterns from a training set can generate new artifacts with the same patterns? It is a satisfying proposition at first because here we see a computer that appears to replicate human cultural behavior and capturing its essence. What is it? Over many thousands of years of human culture making, the diverse cultural expressions that developed in different geographic areas, in different materials, by anonymous groups or later by named authors all have one thing in common: cultural expressions created in one area, in one period or by one group share some common patterns. Their ornaments, clothing, decorations, designs, music, performances, rituals and so do not vary arbitrary - they have a “style,” i.e. a system of rules, constraints and affordances. They define what is possible within a given style, what is less likely, and what is impossible. The coherence of a style in traditional cultures is very strong, and this is why styles of artifacts are used by archeologists to date periods of human civilizations and understand their development.

A particular style system is visible not only across the artifacts remaining from a particular civilization (which is always on a closer look is a meeting point of cultural vectors from different places) but also within a single artifact. Consider a pattern covering the clothing or a vessel from some historical civilization. If they are covered with some ornament, the style of this ornament does not change dramatically across the surface it covers. In fact, if we select a smaller area of this

ornament, we can write a computer program that can predict pretty well the rest of the ornament.

The phenomenon of a systematic style was present in all historical civilizations and periods that I am aware of. Surprisingly, it did not disappear in the modern art and design, despite modernists' revolt against traditional aesthetics (e.g., refusal of symmetry, adoption of dynamic composition, text without capital letters, valuing shock over harmony, etc.) Whether it is a painting by **Sonia Delaney, Lyubov Popova, or Jackson Pollock**, the style system does not change across one painting. In the same way as in traditional ornaments and decorations in ancient and folk arts, here the patterns operating in one large part of an image are the same we find in other parts. Jackson Pollock's mature abstract expressionist paintings are indeed almost like traditional ornaments, with one part containing all DNAs of the whole painting. (Note that there are also artists such as **Delaney or Malevich** who structure their works differently: some patterns remain the same across the full painting, while on the scale of the overall composition, some elements may not be predicted from examining only small parts.)

Why did humans throughout their history kept creating artifacts that have this **single meta-pattern**, i.e. a systematic rigid style within one group of artifacts, and also within a single artifact? Why we are not interested to create images that have one aesthetic system in one image corner and completely different systems in another corner? As I already mentioned, this deep structure of human culture was not challenged by modernist inventions, including **collage** and **montage** tactics develop in the 1920s. Later **remix** practices (1980-) made possible by electronics and later yet digital computers also did not challenge it. Yes, a remix can move between samples drawn from very different aesthetic systems - but once you listen to a part of a remix song, the system established there typically does not change in the rest of the remix song. The same is true of music videos.

Given this, when we teach computers to extract patterns from large sets of artifacts in a single aesthetic system, and then generate new artifacts that belong to the same system, is this really radical? **We force computers to create like us** - like we did for tens of thousands of years. In my opinion, it would be more radical to use

computers to break away from this meta pattern of human culture. Let's teach computers to do something we humans can't do - to move between different systems and aesthetics within a single work, or from work to a work in a series. Modernist revolution that had its high moment a hundred years ago already started questioning some of the basic assumptions of human aesthetics, so maybe computers can help us to continue this process.

One relevant example of AI research is **MuseNet**, "a deep neural network that can generate 4-minute musical compositions with 10 different instruments and can combine styles from country to Mozart to the Beatles." (Payne, Christine. "MuseNet." OpenAI, 25 Apr. 2019, openai.com/blog/musenet.) The system can generate new music in the style of a particular composer and also combine these styles. In one instance, "the model is given the first 6 notes of a Chopin Nocturne, but is asked to generate a piece in a pop style with piano, drums, bass, and guitar."

Exploring such direction is only one of many possible ways to push computers to do something that both appeals to us aesthetically and semantically and at the same time was never yet done in human civilization. It's a common thing to say that if computers can be programmed to create really novel art, we will not recognize it as art, or not will understand it. But maybe this is not that interesting or ground-breaking. Instead, we may want to focus on what lies between such "art for computers", non-comprehensible to humans, and the universe of all aesthetic possibilities already realized in human civilizations (including our own modernist and contemporary periods). Certainly, so many possibilities can be explored in this vast "in between."

This is, then, my third definition of "AI arts." **AI art is type of art that we humans are not able to create because of the limitations of our bodies, brains, and other constraints.** One such possibility I sketched above is computer generated objects, media, situations and experiences that do not have the usual systematicity and predictability of human arts - but they are not random either, they don't mechanical juxtapose elements just to shock, and they are not simply instances of remix aesthetics. I imagine that they will have another systematicity that we have not seen yet even in the most radical modern music, sculpture, architecture,

photography, etc. Something that we would deeply love once we see it. Something that, as all great art before, will expand who we are as humans.