Proceedings of the European Society for Aesthetics

Volume 13, 2021

Edited by Vítor Moura and Connell Vaughan

Published by

European Society for Aesthetics

Proceedings of the European Society for Aesthetics
Founded in 2009 by Fabian Dorsch
Internet: http://proceedings.eurosa.org
Email: proceedings@eurosa.org
ISSN: 1664 – 5278

Publisher
The European Society for Aesthetics

esa European Society for Aesthetics

Department of Philosophy
University of Fribourg
Avenue de l'Europe 20
1700 Fribourg
Switzerland

Internet: http://www.eurosa.org
Email: secretary@eurosa.org

Proceedings of the European Society for Aesthetics

Volume 13, 2021

Edited by Vítor Moura and Connell Vaughan

Table of Contents

European Society for Aesthetics

Extended Aesthetics: Art and Artificial Intelligence

Emanuele Arielli[1]

IUAV University of Venice

ABSTRACT. In this paper[2], I will argue that developments in machine learning and artificial intelligence (AI) applied to aesthetics have relevant implications for philosophical aesthetics, in particular concerning the discussions about the nature of creativity and authorship. The automatic generation of aesthetic artifacts, as well as the development of software increasingly supporting the work of artists and designers, call into question the uniqueness of individual creativity and artistic imagination in an unprecedented way. Moreover, in a scenario in which formal properties of artifacts seem to be easily replicable by machines, the debate on the relationship between aesthetics and the nature of art seems also revitalized. Overall, diverging positions on this issue oscillate between the view of the machine as an Other competing with human capabilities, and, on the contrary, an interpretation of technology as an extension of human potentialities through the externalization of mental processes. AI and machine learning would be in this sense a direct practical manifestation of an *extended aesthetic mind*, in which traditional cognitive limits of the biological mind can be overcome also in areas related to aesthetic creation.

1. Introduction

Since the beginning of the 21st century, computation, data analysis, and artificial intelligence have gradually entered the aesthetic realm. We see this first in what we could call consumer aesthetics, where algorithms are increasingly able to predict what we like and recommend in accordance with our taste, like in music streaming services such as Spotify, or video platforms

[1] Email: arielli@iuav.it.

[2] Parts of this contribution anticipate an in-depth investigation of the relationship between computation and aesthetics with the provisional title *Artificial Aesthetics* (forthcoming), by Lev Manovich and Emanuele Arielli.

like YouTube and Netflix, that keep track of user behaviour, preferences and automatically recommend content we may like (Arielli, 2018). Netflix also uses its data to determine what kind of content to produce in order to maximize the success chance of its movies. Automatic photo improvement is a standard feature in all photo editing in mobile and desktop apps, and also computational analysis of image preferences by users and by professionals has been used to train Networks to evaluate, predict and, eventually, generate images according to their aesthetic value. A further issue is how those systems are pervasively used in the selection and diffusion of images we are exposed to, especially in the digital world, and how this can have a feedback effect on our development of taste, sensitivity, and preferences. An important aspect of these developments is the fact that machine learning systems, neural networks, and what is commonly called Artificial Intelligence (AI) seem to show a self-learning ability and almost autonomous behaviour, reproducing or even surpassing humans in many areas of human competence.

The encounter between AI and aesthetics is crucial because aesthetics is considered a quintessentially human domain and its intractability and complexity have long appeared as not susceptible to algorithmic reduction. Aesthetic phenomena involve a complex relationship between all human faculties, from low-level perceptual mechanisms to higher-level affective, cognitive, and cultural processes. Specifically, art is seen as the pinnacle of human creativity and therefore as one of the last domains in which to test human-machine differences. This is not the place for trying to predict how far AI could develop in the domain of art and aesthetics. Rather, this contribution aims to argue how the contact between those technologies with aesthetics highlights in a new way some traditional issues in our understanding of artistic and aesthetic phenomena.

2. A Map of the Relationships between AI and Aesthetics

First, it is useful to map out the landscape in which machine learning, AI, and aesthetics are connected. In fact, computational technologies could be used both in *describing* artworks, images, and texts (as so-called digital humanities do) and also for *generating* them. An example of this latter case is the deep-learning algorithm that was trained in 2016 to learn Rembrandt's style by analyzing his 346 known paintings and then was asked to generate a brand-new

portrait. The result looked uncannily like a real Rembrandt painting.[3] Another notable example is the painting *Edmond de Belamy* a generative adversarial network image constructed in 2018 by the French arts-collective Obvious. Printed on canvas, the work gained widespread notoriety after Christie's sold it in an auction for almost half a million dollars. Computational approaches can focus on *objects* (such as images, music, texts, artworks, and so on) but also on *subjects*, namely in the analysis of human preferences, people's aesthetic choices, and behaviour.

By crossing the two pairs - description vs generation and object vs subject - we obtain four different domains of application of AI to aesthetics:

	Pattern recognition (analysis and description)	**Pattern generation** (production and prediction)
Objects	**"Studying objects"**	**"Generating objects"**
Subjects	**"Studying subjects"**	**"Generating subjects"**

1.) "Studying objects" means describing them, through the analysis and extraction of objects' formal features, as in the case of the deep-learning algorithm analyzing Rembrandt's paintings and learning his style.

2.) "Generating objects": the system, having been trained with the dataset of Rembrandt's painting, is then used to produce new Rembrandt-looking images.

3.) "Studying people" concerns analyzing people's choices, preferences, and tastes, building their psychological profile of aesthetic experience. An example is the domain of consumer aesthetics, where users of streaming services like Netflix and Spotify have data collected and analyzed by algorithms to determine their habits – resulting in recommendations that

[3] https://www.nextrembrandt.com/

appeal to their tastes.

4.) "Generating Subjects". Recommendation systems in online platforms use models that predict what a user would appreciate. However, by modelling a person's aesthetic judgment, not just his aesthetic behaviour can be predicted, but it is also possible to generate his behaviour and judgment. Modelling listeners' preferences and aesthetic responses allow in principle to simulate how people would behave and react in front of specific objects. If a composer (or the AI itself) creates a new musical variant in J.S. Bach's musical style, an artificial system trained according to people's aesthetic model can formulate its own evaluations and judge if this variant will be appreciated by the people, without surveying real subjects anymore. In this regard, we could imagine a growing use in the future of "artificial judgment" systems that autonomously evaluate if a design artifact, a fashion item or a song have higher or lower aesthetic value. An artificial judge does not simply tell us "what we may also like" (as in traditional recommendation systems), but it would tell us "how much people would appreciate" a specific aesthetic artifact we submit to the system, how people would judge it, or even predict what they would tell us about it.

A typical example of artificial judgment are automated systems for predicting image aesthetic score that make use both of objective metrics (like the quality of the image, sharpness, optimal contrast and colors, etc.) and subjective evaluations. To create such a system, large numbers of people rate lots of images, and this data is used to train a neural network, which can then automatically rate new images.[4] Automatic photo improvement is a standard feature in all photo-editing programs, and also computational analysis of image preferences by users and by professionals has been used to train networks to evaluate, predict and, eventually, generate images according to their aesthetic value. By analyzing people's aesthetic choices, and by extracting the formal features they seem to prefer, we can even speculate how these algorithms could be able to identify aesthetic properties (on the side of objects) and individual preferences (on the side of subjects) of which people are

[4] See https://www.diyphotography.net/rise-machines-googles-ai-will-decide-photos-aesthetically-pleasing/ and https://towardsdatascience.com/deep-image-quality-assessment-30ad71641fac.

not even aware of, but that are manifested in their appreciative behaviour.

3. Between Production and Reproduction

AI is widely used to generate new synthetic artifacts including artworks, music, design items, and texts. Besides the already mentioned generation of a painting in Rembrandt's style, in the same year, researchers at the Sony Computer Science Laboratories in Paris developed a neural network, called DeepBach, that produces choral cantatas in the style of J.S. Bach.[5] Since then, other music-generating algorithms have been developed. In 2019, Deutsche Telekom put together a team of international experts in music and AI to complete Beethoven's unfinished 10th symphony, celebrating the 250th anniversary of his birth. The completed symphony, "Beethoven X - The AI Project", premiered on October 9, 2021, in Bonn. Also, in 2019, an AI used the computing power of a new smartphone model to finish Schubert's "Unfinished Symphony" (n. 8, 1822), although this was accomplished with the help of a composer who cherry-picked the best-generated melodies.[6] A prevailing opinion holds that developments like those mentioned above simply mimic existing styles and are not creative at all. They are a sophisticated kind of reproduction, "computational mannerism" so to speak, not of production, and could be considered just more advanced innovations in the long tradition of generative computer art. In those examples, computers analyze pre-existing works and generate variants conforming to their patterns, while trying to introduce some level of variation. These algorithms do not generate styles of music or painting that are completely new: however, it could be just a matter of time until even the experts are deceived and an AI produces artworks that are judged as aesthetically superior to their human variants.

Artificial generation of artworks like those just mentioned could be considered cases of technical reproduction in Walter Benjamin's sense, in which not a particular artifact, but a whole style is reproduced. According to Benjamin, the aura of artworks' unicity had been diminished by photographic reproduction: today, the unicity of an art style or an artist's

[5] https://arxiv.org/abs/1612.01010
[6] https://www.classicfm.com/composers/schubert/unfinished-symphony-completed-by-ai/

oeuvre has been further reduced by the infinite possibilities to generate similar variations.[7] One should bear in mind that the examples mentioned above involve sets of artworks with a good amount of repetition and similarity: low variability enables neural networks to easily extract general features and generate new examples. In other words, it seems particularly straightforward to produce traditional or classical artworks, as they tend to display a clear, recognizable style and follow the specific patterns of an artist, school, or tradition. Machine learning systems are ideally suited to analyze numerous occurrences of an object type with small variations and extract the relevant features. On the contrary, it would be very difficult to reproduce something like the Duchamp-style body of work, since the AI would have to start with the very heterogeneous dataset of this artist's oeuvre, encompassing Fountain, Bottle Rack, the Large Glass, the late *Étant donnés*, and so on.

Typically, conservative views on art consider technical mastery as a criterion for "real art", and many people still do not consider to be art something that does not require higher technical ability. Technical ability, however, means procedural knowledge, and AIs are designed to deal with precisely this kind of knowledge. Clearly, recognizable styles are well-defined problems that can be reduced to computational tasks, while the generation of variants that don't follow compositional rules (like Duchamp's works) results in ill-defined tasks that have no easy procedural solution. "My kid could have done that!", the popular cliché directed at contemporary art, seems now, in an ironic reversal, to turn against the great and stylistically complex - but computationally tractable - art of cultural tradition: even an AI could do that. It is the Duchamp that remains outside AI's creative abilities, at least for now.

This point highlights the nature of our modern understanding of the relationship between aesthetics and art. Reproducing formal features may be relevant for any domain where aesthetics in its sensorial and perceptual form plays a significant role – like in landscaping, design, decoration, or clothing – but art deals primarily with meaning, symbolic and cultural value, and conceptual depth, like Duchamp's intuition against "retinal art" was primarily about. A computer composing Bach sonatas impresses us, but people would probably be less impressed by the generation of the repetitive melodic

[7] This also allows overcoming copyright issues by means of algorithmically generated content: we cannot use a Beatle song without asking for permission and paying a fee, but we could freely make use of an automatically generated Beatle-like song. Online services offering the generation of "musical replicas" are already available.

patterns of Steve Reich's *Piano Phase*. Or a "Frank Stella"-AI generating black paintings would not make any headlines, and neither would do a "Fontana"-robot equipped with an industrial mechanical arm that cuts canvases. Arthur Danto, theorizing about the ultimate detachment of art from aesthetics, did not deny that artworks have some kind of appearance, but he stressed that art differs from non-art in being about some idea or meaning conveyed by an "atmosphere of artistic theory" and a proper institutional context. AI and machine learning applications in art generation, by being completely off base about this point, exacerbate the contemporary rift between a formalist aesthetic view of artifact production and an "aesthetics of meanings" that need deep cultural references and not simply the mimicry of formal and stylistic patterns.

4. Aesthetic Turing Tests

When a machine paints a Rembrandt or composes a Bach sonata, many are tempted to say that this is neither original nor art. This raises questions about the nature of art that are already being asked in philosophical aesthetics: what is the nature of creativity? Which kind of recombination of ideas, unusual analogies, and conceptual connection are considered the mark of originality? If an artifact or an image is the product of devices, algorithms, and technological extensions that generate and reinterpret an artist's or designer's intention, to whom should we attribute authoriality? If the chain of production is mediated by increasingly complex intervention from third-party software (as in post-production image filters or retouching algorithms), how can we determine *where* and by *whom* creative innovation has taken place?

A testing ground in judging the advancement of artificially generated artworks concerns the comparison with human artifacts. In 2020, for example, an undergraduate student at Princeton University used a so-called Generative Adversarial Network (GAN) to produce traditional Chinese landscape paintings that were capable of fooling humans in a visual Turing test.[8]

A Turing test is primarily concerned with the possibility to reproduce human-like artifacts or content that a human judge cannot tell apart from human ones. The relevance

[8] https://arxiv.org/pdf/2011.05552.pdf, "Among the 242 participants, paintings from our model where mistaken as human-produced over half the time."

of the Test has been widely debated in past debates in AI and, similarly, one can question whether it offers a relevant criterion in aesthetics. Producing something that looks similar to the human hand is, as such, not particularly surprising: for example, there are already digital filters or applications being able to turn a photographic image into an oil painting or a pencil drawing. Music can be produced by simple combinatorial pattern generation (see Pachet, Roy, 1999). Moreover, being human-like is not a specific aesthetic value: What about innovative, beautiful, compelling design, or artforms that appear *non-human*?

If we want to judge the aesthetic value of an artifact, we may focus on qualities like "being pleasant", or even "beautiful", "meaningful", "striking", "overwhelming", "funny", but also "shocking", "uncanny", "moving", and so on. Aesthetic judgment is a complex and multi-faceted variable; it is context and domain-dependent and based on individual taste: evaluating tapestry patterns is different from reacting to artwork or poetry. In contrast to well-defined problems like chess, where programming advancement can be rigorously evaluated in terms of playing strength, in aesthetics, there is no single and well-defined utility function that an algorithm could maximize.

Therefore, we are led to revise the aim of a Turing test beyond the simple "imitation game" they are originally based on, and define its purposes differently. For example, we could say that a machine passes such a test if:

I) achieves *superior* human performance (that is, produced something that is ranked higher in beauty, pleasantness, "amazingness" etc.), without regard to similarity to human behaviour. Or, alternatively

II) it manifests the ability to be *creative*, that is, to generate novelty. Or

III) it shows *autonomous* behaviour, in which the machine seems able to produce something unexpected, distant from the programmers' initial parameters and inputs.

A notorious example of superior performance (I) in AI outside aesthetics are programs beating humans in games like chess or go. But even in art, the ability to produce something that is judged to be superior to humans is not new: as early as 1966, an algorithm was programmed to generate Mondrian-like paintings that were judged by the public to be aesthetically more pleasing than the real Mondrian canvases (Noll 1966).

A *creativity* (II) Turing test (actually a so-called *Ada Lovelace*-test, according to some remarks that the mathematician Lovelace already made in the 19th century) would test something generated by the machine by a human public of judges. However, we are not yet inclined to attribute true creativity to machines, no matter how surprising their artifacts are, since we consider them essentially deterministic devices. Someone would even consider purposefulness or consciousness a presupposition for attributing creativity, something machines don't seem to possess (yet). An interesting exception is the case of *AlphaGo*, a program developed by Google DeepMind. The program beat Go world champion Lee Sedol in March 2016. Sedol claimed that the program was able to make incredibly creative moves and that it revealed how certain moves or game strategies that humans' thought were creative, were actually not. In particular, during the second game of the challenge, AlphaGo made a move (n. 37th) that many commentators described as unusually creative, by catching the player off-guard and then allowing the computer to win.[9]

Similar to the issue of creativity is the idea of *autonomy* (III), that is, how much a machine seems to be independent in its behaviour from the programming and the input settings, and how much its products are unexpected and unpredictable to the human observer. On one side, there is no clear-cut criterion for autonomy: is a mono-cellular organism autonomous? What about an insect? In attributing autonomy, the line we draw is not based on pre-determined criteria, but is based on a great deal of subjectivity: there are cultural and subjective reasons behind the fact that we attribute autonomy (or "will") to a fruit fly (or even a bacteria), but not to machines like AlphaGo, no matter how complex these are.

More generally, these examples show how in this domain any technological progress remains dependent on what the human being can evaluate and understand. This is a crucial point, because our ability to understand what a machine generates and to determine how innovative and extraordinary its products are, remains the ultimate filter (leaving the dystopic scenario aside in which only machines will evaluate what other machines do). Our capacity to make sense and interpret what the machine is offering us is limited and

[9] According to one professional player: "All but the very best Go players craft their style by imitating top players. AlphaGo seems to have totally original moves it creates itself", https://www.latimes.com/world/asia/la-fg-korea-alphago-20160312-story.html

inherently "human": a machine generating inhuman or *alien* aesthetics, transcending our understanding, would go unnoticed or simply deemed by us as meaningless.

5. Externalized and Extended Aesthetics

Autonomy, free will, and unconstrained creativity reveal how independence from human intervention is considered a crucial benchmark in the evolution of artificial systems. From this perspective, decisional autonomy is one aspect of the technical evolution of devices: from a thermostat "deciding" when to stop heating according to a target temperature to a self-driving car regulating its speed according to traffic conditions. However, if we solely focus on the question of autonomy against or independently from a human decision, we would miss the main perspective from which we have always observed technological development: namely expansion, integration, and enhancement of human action and experience. Under this premise, machine learning and AI should be considered a further step in our *tool-making* ability in expanding humans' skills through devices. Therefore, their impact should be measured not in their theoretical degree of autonomy from human choice, but in their influence and contribution to human potentialities. Technical devices account for the possibility to externalize cognitive and manual tasks and to extend human capability beyond its natural limits. We are used to the notion of the external (and extended) mind in philosophy and cognitive science when we refer to writing and memory devices. Technological development allows for externalization and automation of memory and reasoning, but we could also speculate that aesthetic production happens more and more "outside", employing tools that take up part of the creative work. The so-called "extended mind" paradigm could be here coupled with an understanding of technology not as a separate entity from human nature, but as a process of integration and augmentation between mind and technology (Simondon, 1958).

In the specific case of art, tools like paper, canvas, pencils, chisels, music instruments, and photo and video cameras, all made it possible to produce artworks that otherwise would have been not possible to craft. AI used to produce artworks would represent in this sense a further evolution. Artist Mario Klingemann, a leading exponent of AI-art, concisely said: "If you heard someone playing the piano, would you ask: 'Is the piano the artist?' No. So, same thing here. Just because it is a complicated mechanism, it

doesn't change the roles […] The typewriter enables someone to write a book. For me, the keyboard enables me to write code, […] there are neural networks involved that maybe you could say that they are my brushes that I learn to use."[10] Klingemann does concede the possibility to consider the machine like a child one puts "out in the world alone", trusting that he will keep on doing what the artist/parent hoped he would do., Unlike a piano, there is indeed a moment in which the artist may decide to take his hands "off the keyboard" and let the machine follow its course without human intervention.

Externalization and automation also involve our imaginative potentialities. Imagination is the capacity of internal generation of non-existent things, through novel combinations and transformation of elements, and is a crucial faculty for creativity. Technologies like those mentioned above contribute to visualizing images that we were before able to see only with the "mind's eye" of our imagination. Now, we can see a photo transformed into a painting in the style of Van Gogh, or into a new Monet landscape, or read a new text generated in the style of Shakespeare[11], and so on.

"Artificial imagination" would in this sense be a specific extension of our creative potential. We have biologically limited visual acuity, but microscopes and telescopes allowed us to amplify the realm of the visible. Similarly, our cognitive calculation skills and memory have upper limits, but calculators and computers expanded those skills.[12] Along this line of argumentation, one could speculate that also aesthetic skills (imagination, perceptual sensibility, capacity of formal innovation) have human limits. We could even suppose the existence of natural "peak creativity" levels determining "bounded" aesthetic skills in humans. Would then an artificial extended aesthetics allow augmenting our aesthetics skills, both in deepening our sensibility and our creative process? We could envisage a future in which composers or writers stuck in their creative process could make use of systems giving them suggestions, evaluating alternative ways to go on with their work, and testing if their idea could meet the positive response from the public. Advanced computational systems would then be a further evolution of devices that are already used in creative disciplines, such as graphic programs, computer-aided design technology, music

[10] https://www.sothebys.com/ens/articles/artificial-intelligence-and-the-art-of-mario-klingemann, interview conducted for the installation *Memories of Passerby I*, 2019.
[11] https://spectrum.ieee.org/this-ai-poet-mastered-rhythm-rhyme-and-natural-language-to-write-like-shakespeare.
[12] This is the classical view of human "bounded rationality" defined by Herbert Simon and of technology as a means of overcoming the natural limitations of our brains.

software, and so on. Our engagement with technology, by expanding and modifying the way we create, would eventually influence how cultural evolution is shaped.

We must also keep in mind that these extensions affect not only the creative process (see above: "Generating objects"), but also the way we describe and analyze cultural and aesthetic products ("Studying objects"). Technologies based on the analysis of big data, supported by the evolution of deep learning algorithms and AI, potentially make it possible to extract new patterns in complex cultural phenomena by analyzing huge amounts of information and detecting previously unthinkable connections. For example, researchers from MIT's and Microsoft created in 2020 an algorithm able to discover unnoticed similarities and hidden formal connections between works of art on display at the New York's Metropolitan Museum of Art and Amsterdam's Rijksmuseum.[13]

Here we can conclude by making an analogy to Walter Benjamin's notion of the "optical unconscious", according to which visual reproductive technologies allowed us to see things differently, magnifying them in greater details, slowing them down, or presenting them in new configurations through montage and so on. Similarly, AI and machine learning make undetected connections visible, like a microscope or a telescope allowing the perception of things that were not possible to see with the naked eye. In a word, they allow bringing *data unconscious* to light. *Unconscious* is however an ambiguous notion and could be both interpreted as bringing buried and hidden structures to the surface and as an ex-nihilo construction of patterns and Gestalts that almost creatively make use of data, like star constellations. These would constitute two sides of the same coin: seeing connections and building new connections, where the tools that are used for the analysis of huge cultural data are the same employed for the generation of artificial artworks. *Extended* aesthetics, therefore, involves processes that are relevant for both sides of reception and creation, and focusing on technological advancements in this domain appears crucial for contemporary debates in (philosophical) aesthetics as well.

References

Arielli, E. (2018), 'Taste and the Algorithm', *Studi di Estetica*, 46,3, pp. 77-97.

Benjamin, W. (2009), 'The Work of Art in the Age of Mechanical Reproduction', In *One-way*

[13] https://news.mit.edu/2020/algorithm-finds-hidden-connections-between-paintings-met-museum-0729.

street and other writings, London, New York: Penguin, pp. 228-258.

Danto, Arthur. (1998), *Beyond the Brillo Box: The Visual Arts in Post-Historical Perspective*, Berkeley: University of California Press.

Noll A.M. (1966), 'Human or machine: a subjective comparison of Piet Mondrian's 'Composition with lines' (1917) and a computer-generated picture', *The Psychological Record*, 16, pp. 1-10.

Pachet F., Roy P. (1999), 'Automatic Generation of Music Programs', In Jaffar J. (eds) *Principles and Practice of Constraint Programming. CP 1999. Lecture Notes in Computer Science*, vol 1713, Berlin, Heidelberg: Springer.

Simondon,G. (1958), *Du mode d'existence des objets techniques*, Paris: Aubier.

Turing, Alan (1950), "Computing Machinery and Intelligence", *Mind*, 49 (236): pp. 433-460.