



xCoAx 2019

Conference on Computation,
Communication, Aesthetics & X

Milan, Italy

Tsila Hassine

Tsila.Hassine@gmail.com

Université Paris 1 Pantheon-
Sorbonne, Paris, France and
Shenkar College of Engineering,
design Art, Tel Aviv, Israel.

Neeman Ziv

zneeman@gmail.com

independent scholar

The Spectre of Zombies is Haunting AI Art: How AI Resurrects Dead Masters, and Alternative Suggestions for AI Art and Art History

Keywords

AI
art
deep learning neural networks
generative art
Rembrandt
deep art

In the past few years deep-learning AI neural networks have achieved major milestones in artistic image analysis and generation, producing what some refer to as ‘art.’ We reflect critically on some of the artistic shortcomings of a few projects that occupied the spotlight in recent years. We introduce the term ‘Zombie Art’ to describe the generation of new images of dead masters, as well as what we term ‘The AI Reproducibility Test.’ In conclusion, we propose new directions for both AI-generated art and art history, in the light of these new powerful AI technologies of artistic image analysis and generation.

1 INTRODUCTION

AI has been in the public eye and imagination for many years already, with endless scenarios describing the disappearance of different jobs and human skills, which will be taken over by intelligent machines. Artistic creation is no exception to this, as the question of mechanized artistic creation has been tantalizing human imagination for some decades. Recent breakthroughs in machine learning—especially in popular accounts—herald the achievement of this goal. While we applaud the progress in machine learning, neural nets, image recognition and manipulation, we question whether they constitute a major artistic breakthrough, at least in their current form. We suggest that by rethinking their conceptual goals and uses, more interesting AI generated art may be created. We further foresee a new frontier of AI based art history. For the purpose of our discussion we rely on three AI art projects which have attracted a substantial amount of media attention recently.

The Dutch *Next Rembrandt* project created by a multidisciplinary group of researchers, uses custom created AI to analyze the style and content of a large number of Rembrandt's paintings, then used them to produce a 'new Rembrandt portrait' (ING et al, 2016). The *DEEPART* project created by a German group of computer scientists. This project provides proof of concept that AI can successfully separate the content of an image from its style, and combine the style of one image with the content of another. Their well known example is an image which reproduces a picture of the contemporary city of Tübingen in the style of Van Gogh's *Starry Night*. The same software also generated images in the style of other artists based on that same Tübingen photograph (Gatys Ecker and Bethge, 2016a, 2016b).¹ Slightly different, the Parisian collective Obvious generated painterly portraiture images based on a large dataset of 14th-20th century portraits analyzed by deep learning neural net. It received a lot of publicity when one of its generated images, *Portrait of Edmond de Belamy*, was auctioned by Christie's in 2018 for 432,000\$ (*Obvious*, 2018; *Obvious, Explained*, 2018; Schneider and Rea, 2018).²

In the following discussion we will refer to "artistic interest," "significance" and "value." Yet we consider art and its appraisal as a cultural phenomenon whose function and meaning together with its evaluations change through time and across cultures. Hence attempting to define such qualities is extremely difficult. However, we can delineate a bare minimum of connected traits that can stand as a correlate of artistic interest, significance and value in the context of contemporary culture. These include: creativity, innovation and a sense of surprise.³

2 THE SPECTRE OF ZOMBIES IS HAUNTING AI ART

2.1 AI as Forgery?

As much as we applaud these advances in AI Neural Networks, one might consider these projects as generating what we may call 'Zombie Art.' 'Zom-

1. Currently it has opened a website, where people can select and send two images, one for content and the other for style, and the software generates a third image which is the amalgamation of the two images; see: <https://deepart.io/#>

2. There has been controversy around *Obvious'* use of its GAN algorithm, since it is virtually a copy of an algorithm created by Robbie Barrat and uploaded by him to *GitHub*. (Flynn, 2018).

3. The selection of these traits follows to a large degree, the article by Bringsjord, Bello and Ferrucci (2001), 5, 8.

4. The term ‘deepfakes’ is currently a catch-all term for forgeries of images, videos, audio, etc., that are exceedingly difficult to detect due to their high quality achieved by advanced AI deep learning neural networks to generate or manipulate them.

bie’ because these machine algorithms generate paintings by masters that have been dead for centuries. Hence, a new painting in the generic style and content of Rembrandt, images in the unmistakable style of van Gogh, or any given image in the style of Munch’s *Scream*. The possibilities are literally endless. We consider such images as technical assemblages existing in the space between the past and the present, life and death. These images are like the living dead or specters: they’re zombie images, at once ‘dead’ and ‘alive.’ While it may certainly be considered an achievement to create a new artistic category such as AI generated zombie art, we question its actual artistic significance and interest. Aren’t these images simply “deepfakes” (Hao, 2018).⁴ Or simply put: machine made forgeries?

Zombie art is not limited to machines alone. Human artistic forgeries of dead masters, such as the case of van Meegeren’s fake Vermeers from the 1930s can also be considered as zombie art. The only difference is that the human forger injects a least a modicum of creativity to the forgery (though the goal of creating a passable fake will tend to limit such creativity). In the van Meegeren case, human forgery still relied on the artistic prowess and creativity of the forger, making the forgery unique. Therefore, human generated forgeries might be more accurately, and less provocatively, termed simply as ‘forgeries.’

The creators of the Dutch *Next Rembrandt* project themselves describe their project along the lines of creating a forgery:

“because a significant percentage of Rembrandt’s paintings were portraits, we analyzed the morphology of the faces in these paintings, looking at factors such as gender, age and face direction. The data led us to the conclusion that the subject should be of a Caucasian male with facial hair, between 30-40 years old, in dark clothing with a collar wearing a hat, and facing to the right. (ING et al., 2016)

Such a description of Rembrandt’s (or any other artist’s) characteristic subject matter that neatly corresponds to popular perceptions about him, constitutes the *sine qua non* of forgery; the real challenge of course is to fool the experts.

2.2 Averaging the Grand Masters

Our central criticism of these projects is that their forgery-like memetic aim constitutes precisely the reason why they are artistically underwhelming. Though we realize, that the emergence of such AI generated works already questions the current meaning of artistic value. These projects may be experimentally interesting from a technical perspective, but reading into the algorithmic process itself, we come to a conclusion that this process actually undermines the value of the original artworks themselves, before they were transformed into datasets. The generation of “new” Rembrandt paintings based on the datafication of his original oeuvres emphasizes the repetitive dimensions of his creativity, in a way that has so far eluded the

human viewer of his work, thus diminishing the singular interest and value of Rembrandt's actual paintings.

However, we find a deeper problem with all three projects. Because the *aim* of these projects is to *emulate* the style and/or content of a specific artist's oeuvre, such image generation will inexorably zero in only on the most clearcut, characteristic and recognizable perimeters of an artist's style and/or content. By definition, this leads such projects to focus on the most obvious and redundant subjects and/or style traits of an artist, in order to generate a signature style and/or content. However, examining the Dutch Rembrandt project illuminates well this inescapable drift towards an artist's most distinct and popularly known traits—alas, also the most trite characteristics of his oeuvre.

Yet Rembrandt did not acquire his reputation through a mere repetition of subject and style. As art Historian Christopher Wright writes in his book on 17th century Dutch painting: “one of the secrets of Rembrandt's subsequent reputation is [the] variety” of his oeuvre (1978, 172). Such variety is often divided into: history paintings of Biblical and classical subjects; landscapes; animals; self-portraits; portraits of family members; genre scenes of Dutch life; and portraits (Rembrandt Painting Net). Against this rich variety, *Next Rembrandt's* highly circumscribed focus on portraits of a Caucasian male in dark clothing etc., appear as limited.

Moreover, as the Dutch Rembrandt project video explains, once it was decided that the ‘new’ Rembrandt would be a portrait, they used various algorithms to extract average shapes of facial features such as eyes, noses and mouths from Rembrandts' portraits, and their facial proportions.

Next Rembrandt's project's drastic limitation of the image content and the idealized averaging of facial features as input data is the major reason why its ‘new’ Rembrandt portrait is underwhelming. While the ‘new’ portrait achieves a high level of painterly technique, this attainment is undermined by the very statistic averageness of its subject and style. Due to this averageness, for us, *Next Rembrandt's* ‘new’ generated portrait is ultimately dull since it does not contain any artistic surprises or novelty.

Similar dynamics are operative in the image generation of both the *DEEPART* and *Obvious* groups, although their algorithmic method is different. *DEEPART* attempts a balancing or averaging between content and style, in order to generate “visually appealing images” (Gatys, Ecker and Bethge, 2016a, 2419). But, creating pretty images by balancing the content/style parameters does not necessarily make for significant artistic images in our opinion. We find that actually many of the images created during the process, displaying unbalanced weightings of the content/style parameters are of greater artistic interest than the featured balanced ones, since they contain more surprises than the end result (Gatys, Ecker and Bethge, 2016a, Fig. 3).

In regards to *Obvious*, this dynamic of averaging or limitation is constituted in a different manner. *Obvious'* generation of portraits was achieved by inputting its deep learning neural net with “training data set of more than

5. In portraiture painting, the focal point is always the face of the person (otherwise it isn't deemed a portrait); the face almost always looks at viewer or is slightly turned; there are only three central formats: full figure, 'half-shot' (only the head and torso are pictured), or it is a 'head-shot (showing only the face and shoulders); the figure is nearly always either standing or sitting, generally in an interior. Caselles-Dupré, of the *Obvious* collective is quoted saying that: "We did some work with nudes and landscapes, and we also tried feeding the algorithm sets of works by famous painters. But we found that portraits provided the best way to illustrate our point, which is that algorithms are able to emulate creativity." (Im, 2018) We suggest this is an implicit confirmation of our argument about their choice of genre.

15,000 portraits created between the 14th and 20th centuries" (Schneider and Rea, 2018). Yet the supposed variety of its input data is not that broad for two reasons. The first is artistic: historically, the genre of portraiture is the most durable and least changing genre in art history, due to its highly circumscribed conventions.⁵ The second factor is the selection of input images; in *Obvious*' website they write: "[w]e carefully select a large number of input images with common visual features. The goal is to create a new sample that shares these features" Together these two factors emphasize commonality, rather than the variety of the input images.

Hence, all three projects are confined conceptually and operationally in a variety of ways, including limited inputs, the search for common features, averageness or an emphasis on an artist's most redundant traits. In our opinion, all these conspire to limit and restrain artistic creativity, novelty and surprise.

3 OPEN-ENDED EXPERIMENTATION VERSUS PRESPECIFIED GOALS

The projects mentioned above can be seen as an evolution of computer based generative art, which started with the early computer age. Those early artistic experiments with computers and current day AI art share many common features: creating the algorithms or neural nets; tweaking them; selecting the best images from a large output of generated images. As such it is instructive to place these works in the genealogy of generative art pieces. Yet there is a significant difference between these two approaches to using the computer creatively. The significant difference is the use of deep learning networks rather than non-learning algorithms. In other words, earlier generative art did not set forth to reproduce old masters, and therefore did not have to "learn" anything.

This highlights a significant difference between early computer-generated art, from the 1960s-1970s, and this new type of generative art. Early computer art was undertaken in the spirit of open-ended experimentation, without a specific goal in mind. As, Max Bense and Reinhard Döhl proclaimed, "The artist today realizes accomplishments on the basis of conscious theory and deliberate experiment[ation]" (1964, 9)⁶

In contrast, the projects of *Next Rembrandt*, *DEEPART* and *Obvious* are all directed towards their predetermined and specific goals, thus determining the modus operandi of these projects. Of course, these projects included substantial experimentation, yet this type of experimentation was most likely motivated by engineering rather than artistic purposes. Experimentation was not open-ended, but was rather of an instrumental kind, in order to achieve their pre-determined goals of imitative forgery-like artistic representations. Indeed, these projects' well defined teleology, constitutes one overarching reason that their results have only limited artistic value.

6. See also Nake, 2005, 60, 93; Nake, 2012, 77.

4 ART IN THE TIME OF AI: SUGGESTIONS FOR ALTERNATIVE USES OF DEEP LEARNING AI ART NETS

As we stated previously, we consider art and artistic value to be cultural and historical phenomena. Hence the emergence of such powerful new AI image analysis and generation technologies changes the current artistic ecosystem. Indeed, throughout history technology has always influenced and impacted art, for example: the ancient production of pigments, the invention of oil-based colors or the invention of photography.

We suggest that the use of deep learning AI visual networks could be utilized as an analytical tool, as well as, we believe, offering a potentially better way for generating artistic images. We begin with the question: what if, for example, Marcel Duchamp's entire oeuvre with all its different styles, mediums and genres was inputted into a deep learning neural network that has been trained to extract or distill a single artists' style, content, or both. What are the chances that such an AI neural network might succeed in this task and reproduce a new, yet recognizable Duchamp? We believe it would not be able to do so in any satisfactory manner. Most artists have a single 'mature style.' Yet there are artists that among their prominent signature is their simultaneous (or rapidly changing) creations in many artistic styles, genres and mediums. The names of Francis Picabia, David Smithson, Gerhardt Richter and Sigmar Polke, among others come to mind as being such multitudinous artists.⁷ Inputting their entire oeuvres into such a deep neural network, as described above, might not yield impressive mimetic results. Therefore, we suggest the possibility of what we call '*The AI Reproducibility Test.*' The test would consist in seeing whether a deep learning AI net, inputted with the entire oeuvre of a single artist, will be able to generate novel images commensurate with that artist's oeuvre, or not. It is perhaps possible, that such a hypothetical test would yield a clear demarcation between artist's whose oeuvre allows for such generation of images, against those that don't. However, more realistically, we think the outcome of such a test would be a spectrum of results, ranging from high ratings for artists whose oeuvre will easily abet the generation of new images in their signature style and/or content to those artists with which the neural networks will only achieve limited or unsatisfactory results and receive a lower rating. We suggest that this would perhaps constitute the beginning of new forms of engagement of AI with art history, that might well lead to interesting new insights regarding artistic practices. We can even imagine a possible future scenario in which the relative ranking of contemporary, working artists in '*The AI Reproducibility Test.*' would become significant; thereby, likely influencing artists to attempt creating oeuvres that would produce lower 'Reproducibility Test' ratings.

In the context of the projects we discussed above, which ultimately generate new images based on the input of already existing ones, we propose that in regards to individual artists' oeuvres, the more interesting results will come from those artists whose '*Reproducibility Test*' ratings are at the lowest part of the spectrum; i.e., cases where neural nets will not be able

7. Other contemporary multitudinous artists include the likes of: Bruce Nauman, Albert Oehlen and Martin Kippenberger.

to satisfactorily distill their style and/or content. Such supposed ‘failures’ will serve as kind of constraint on the neural net’s tendency to focus on the most reoccurring features of an artist’s work. Just as importantly, it will likely diminish, severely skew or avert the seemingly insistent drift toward the averaging by the nets’ operations. Thus, generating novel images based on hard to reproduce artists will generate images that would be ‘off-kilter’ and this will likely generate more surprising, unexpected and potentially more creative images in our view.

To conclude, we find the strength and significance of these AI based projects, not in the production of new, out of context paintings by dead masters, but rather in the creation of a new approach to art history through the eyes of 21st century intelligence. If our machines can now paint a ‘new Rembrandt,’ and separate between style and content, can we use them to learn new things about the processes, significance and meanings of artistic creation throughout human history?

REFERENCES

Bense and Reinhard Döhl.

1964. “Zur Lage,” quoted in, Suter, Beat. “Lineages of German-Language Electronic Literature: The Döhl Line” 105-18, 109, accessed from 2013 ELO Conference, <http://conference.eliterature.org/critical-writing/lineages-german-language-electronic-literature-d%C3%B6hl-line>.

Lovelace Test.

“*Minds and Machines* 11: 3–27.

Gatys, Leon A., Alexander S. Ecker and Matthias Bethge.

2016a. “Image Style Transfer Using Convolutional Neural Networks.” *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2414-2423. Accessed 28.1.19 from: https://www.cv-foundation.org/openaccess/content_cvpr_2016/html/Gatys_Image_Style_Transfer_CVPR_2016_paper.html.

Hao, Karen.

2018. “Inside the World of AI that Forges Beautiful Art and Terrifying Deepfakes.” *MIT Technology review, The Algorithm*, accessed 31.1.19 from: <https://www.technologyreview.com/s/612501/inside-the-world-of-ai-that-forges-beautiful-art-and-terrifying-deepfakes/>

Bringsjord, Selmer,

Paul Bello And David Ferrucci. 2001. “Creativity, the Turing Test, and the (Better)

Flynn, Meagan.

2018. “A 19-Year-Old Developed the Code for the AI Portrait that Sold for \$432,000 at Christie’s.” *Washington Post*, accessed 7.2.19 from: https://www.washingtonpost.com/nation/2018/10/26/year-old-developed-code-ai-portrait-that-sold-christies/?utm_term=.373360d4adbb.

Gatys, Leon A., Alexander S. Ecker and Matthias Bethge.

2016b. “A Neural Algorithm of Artistic Style.” arXiv:1508.06576v2.

Im, Jimmy.

2018. “This Portrait Made by A.I. Just Sold for \$432,000—That’s 40 Times the Original Estimate.” *CNBC*, accessed 2.7.19 from: <https://www.cnbc.com/2018/10/25/portrait-made-by-artificial-intelligence-sold-for-432k-at-christies.html>.

ING, Microsoft, TU Delft, Mauritshuis, Museum Rembrandthuis.

2016. *The Next Rembrandt*. Video. Accessed 28.1.19 from: <https://www.nextrembrandt.com/>.

Nake, Frieder.

2012. "Construction and Intuition: Creativity in Early Computer Art," eds. McCormack, Jon and Mark, d'Inverno, *Computers and Creativity*. 61-94. Berlin: Springer https://doi.org/10.1007/978-3-642-31727-9_3.

Obvious, Explained.

2018. From *Obvious* Website, accessed 31.1.19 from: <https://medium.com/@hello.obvious/ai-the-rise-of-a-new-art-movement-f6efe0a51f2e>.

Schneider, Tim and Naomi Rea.

2018. "Has Artificial Intelligence Given Us the Next Great Art Movement? Experts Say Slow Down, the 'Field Is in Its Infancy'." Artnet News, accessed 2.2.19 from <https://news.artnet.com/art-world/ai-art-comes-to-market-is-it-worth-the-hype-1352011>.

Nake, Frieder.

2005. "Computer Art: A Personal Recollection." Conference paper, DOI: 10.1145/1056224.1056234

Obvious.

2018. Website, accessed 31.1.19 from: <https://obvious-art.com/>.

Rembrandt Painting Net.

2018. "Complete Rembrandt Catalogue," accessed 2.2.19 from http://www.rembrandtpainting.net/complete-catalogue/start_history.htm.

Wright, Christopher.

1978. *The Dutch Painters: 100 Seventeenth Century Masters*. Woodbury, N.Y.: Barron's.