

<https://artnodes.uoc.edu>

ARTICLE

NODE «AI, ARTS & DESIGN: QUESTIONING LEARNING MACHINES»

Generative Art: Between the Nodes of Neuron Networks

Bruno Caldas Vianna

University of the Arts, Helsinki

Date of submission: March 2020

Accepted in: June 2020

Published in: July 2020

Recommended citation

Caldas Vianna, Bruno. 2020. «Generative Art: Between the Nodes of Neuron Networks». In: Andrés Burbano; Ruth West (coord.) «AI, Arts & Design: Questioning Learning Machines». *Artnodes*, no. 26: 1-8. UOC. [Accessed: dd/mm/yy]. <http://doi.org/10.7238/a.v0i26.3350>



The texts published in this journal are – unless otherwise indicated – covered by the Creative Commons Spain Attribution 4.0 International license. The full text of the license can be consulted here: <http://creativecommons.org/licenses/by/4.0/>

Abstract

This article uses the exhibition “Infinite Skulls”, which happened in Paris in the beginning of 2019, as a starting point to discuss art created by artificial intelligence and, by extension, unique pieces of art generated by algorithms. We detail the development of DCGAN, the deep learning neural network used in the show, from its cybernetics origin. The show and its creation process are described, identifying elements of creativity and technique, as well as question of the authorship of works. Then it frames these works in the context of generative art, pointing affinities and differences, and the issues of representing through procedures and abstractions. It describes the major breakthrough of neural network for technical images as the ability to represent categories through an abstraction, rather than images themselves. Finally, it tries to understand neural networks more as a tool for artists than an autonomous art creator.

Keywords

generative art, machine learning, artificial intelligence, representation, algorithms, aura, neural networks

Arte generativo: entre los nodos de las redes neuronales

Resumen

Este artículo utiliza la exposición “Infinite Skulls”, que se inauguró en París a principios de 2019, como punto de partida para hablar sobre el arte creado por la inteligencia artificial y, por extensión, sobre las piezas de arte únicas generadas por algoritmos. Se centra en el desarrollo de DCGAN, la red neuronal de aprendizaje profundo utilizada en el programa, desde su origen cibernético. Se describe el espectáculo y su proceso de creación, identificando elementos de creatividad y técnica, así como la autoría de obras basadas en código abierto, en particular “Edouard de Belamy”, la pintura realizada a partir de inteligencia artificial que se vendió en una subasta de Christie’s por 432 000 dólares estadounidenses. También se enmarcan estos trabajos en el contexto del arte generativo, señalando afinidades y diferencias, así como los problemas de representación mediante procedimientos y abstracciones. Describe el gran avance de la red neuronal para imágenes técnicas como la capacidad de representar categorías a través de una abstracción, en lugar de imágenes en sí mismas. Por último, trata de entender las redes neuronales como una herramienta para los artistas, más que como una obra de arte per se.

Palabras clave

arte generativo, aprendizaje automático, inteligencia artificial, representación, algoritmos, aura, redes neuronales

Introduction

An early 2019, digital artist Robbie Barrat and painter Ronan Barrot collaborated on an exhibition (Bailey 2019) that took place at the L’Avant Gallery Vossen, in Paris. The show consisted on Barrot’s skull paintings, shown side by side with Barrat’s artificial intelligence-based reinterpretation of them. This collaboration is an opportunity to discuss current issues around the production of artworks which rely on machine learning tools. In 2018, an image produced by the Obvious collective with these methods was sold for 432.000 US dollars at the auction house Christie’s in New York (Jones 2018), and several exhibitions in different countries showed machine learning crafted work. How and why did we end up wanting machines to do the work of painters? And are these works mere replicas of great paintings or original works of their own?

Although some of the questions that arise go beyond the possible scope of the article, we would like nonetheless to take the opportunity to understand the processes and the history behind machine learning-based visual arts. We will start by describing the development of the technologies behind the tools used, the development of cybernetics and the quest for autonomy in computing; and then we’ll describe the collaboration and the exhibition itself. Finally we will wrap up mapping a few important questions on how can neural networks be framed within the generative art field, and why does it represent a paradigm change in procedure-based art.

From Cybernetics to DCGANs

The history behind the techniques used in the exhibition may be traced to the beginnings of computer science and the birth of cybernetics. In the founding book of this science, published in 1948, Norbert Wiener states clearly the goals of replacing mankind’s mental and physical workforce by machines. He even acknowledges the inherent risks of this replacement to the point of sharing his concerns about it with labor unions (1965, p. 27-28). He also makes clear references that the key to developing autonomous entities included the study of nature – beginning by the very title “Cybernetics - Or Control and Communication in the *Animal* and the Machine.”

It is not by chance that the first model of a neuron had been developed just a few years earlier, by McCulloch and Pitts in 1943. In the paper that describes it, the authors note how “the activity of the neuron is an ‘all-or-none’ process” (1943). At this time, the standard binary system for computers had not been established; computers themselves were more of an imagined device than an actual tool. The contraction *bit* (for *binary digit*) first appeared in Claude Shannon’s information theory paper in 1948 (1948) and Wiener stated that “in accordance with the policy adopted in some existing apparatus of the Bell Telephone Laboratories, it would probably be more economical in apparatus to adopt the scale of two for addition and multiplication, rather than the scale of ten” (2019, p. 4). In other words: we have, on one side, a biological feature present in most animals and, in the

other side, a mechanical device, and both of them are converging in their way of organizing and processing information.

The next natural step was to think of ways to model – and replicate, if possible – the complex structures of the brain, made up of billions of neurons interconnected in a dense web. A fundamental piece of the evolution of image based artificial intelligence was put in place in 1957 with the development of the perceptron by Frank Rosenblatt (1957). Although the model was first simulated on an IBM 704 computer, Rosenblatt would soon build a special purpose, analog hardware implementation. This apparatus used a grid of 20 by 20 photocells that took the role of a camera or an eye, connected in a network that would store data in potentiometers activated by electric motors. This network would be trained by placing different shapes in front of the camera, and the weights stored could then be used to recognize such shapes (Bishop 2006, 192-196).

This seemingly complex approach was actually criticized, with dire consequences, for being too simple. In a book published in 1969, Marvin Minsky and Seymour Papert (2017) acknowledged the potential of the perceptron, but also demonstrated that Rosenblatt's idea had some concerning limitations. One particularly damaging feature they showed was the fact that the model wouldn't be able to compute one specific logic function, named XOR. As we'll see, neural networks today are also a method to provide a general solution to equations that are non-linear, ie, not solvable by calculus tools. The fact that it wouldn't be able to handle one of the most basic equations in computer science seriously undermined the reputation of the model and helped stall the development of neural networks in general for more than a decade – a period that came to be known as artificial intelligence's winter (Crevier 1993).

What is curious is that, despite the pessimism towards perceptrons brought by Minsky and Papert, the solution to their limitations was pointed in the book itself – the multilayer perceptron. This model, which is still in use by most models in neural nets today (including the one used by Barrat) consisted of connecting the output of a layer of neurons to the input of another layers. The depth (number of layers) of a network can be defined according to the needs of the problem to be tackled, but Minsky and Papert showed that a three-layered net could already implement a XOR function. It should be also mentioned that putting these models to work with the memory and processing power of the devices of that time would be already discouraging.

In some way, this winter also reached the artists that showed interest in cybernetic processes. In the fifties, artist Nicholas Schöffer was heavily influenced by Wiener's ideas. CYSP 1, created in 1956, might be the first piece of art that explicitly embedded the cybernetic concept of feedback into the artwork (Popper 2007). Many others followed suit, like the systems-inspired works of Hans Haacke and Jack Burnham. Yet in seventies these two exponents of cyberart moved away from the field, possibly criticizing the relation between this

science and the industrial military complex or the own “enclosed” aspect of the gallery (Lynch 2018).

The rebound of artificial intelligence would arrive only in the eighties. A movement named *connectionism*, that approached the cybernetics from a cognitive perspective, tried to apply neural networks to cognition models. This movement established two breakthroughs in deep learning that are still in use as of today: *distributed representation*, a strategy that breaks up inference tasks into smaller networks; and the use of back-propagation algorithms to train the net. Back-propagation, in fact, is a nothing more than a sophisticated feedback method – a concept already developed by Wiener in the cited book. Nevertheless, by the mid 90's, AI failed once more to deliver practical applications and the second wave of neural networks ended. It was only in 2006 that another breakthrough would spark its third wave, which we are still riding at of today. Researcher Geoffrey Hinton developed a strategy named greedy layer-wise pre-training to perform what came to be called “deep learning” of the network (Goodfellow and Courville 2016).

Deep learning evolved rapidly after that, with several small but important inventions that made it into one of the richer and most complex fields in contemporary mathematics. It is important to notice that although neurological science will always be considered the original inspiration for neural networks, the brain is no longer an useful reference for AI scientists today. This is due to the fact that beyond simple models of one or a few neurons, it is very hard to analyze the full complexity of the brain and create models to replicate it. At the same time, mathematicians have been developing fantastic methods for different tasks that have nothing related to neurology – at least as an inspiration (Ibid).

In the early 2010's, visual machine learning was still focused mostly in creating tools to recognize shapes and objects, inferring information about the environment in general. Many of such techniques were developed as part of a field named *Computer Vision* – therefore still using the eye and human senses as an inspiration. It was in 2014 that Ian Goodfellow, a PhD candidate at the University of Montreal, developed a strategy that would allow deep learning networks to create images instead– the generative adversarial networks (GANs). It uses two competing networks: one named the *discriminator*, and another one named the *generator*. The generator will be constantly creating images, which the discriminator will try to evaluate as being from the learning set of images or a counterfeit. The method turned out to be extremely successful. While there is some controversy in the sense that the idea might have been developed by others as early as in 2009 (Schmidhuber 2019), it is certain that the method only became widespread after the publication of Good-fellow's paper (2014).

The last puzzle piece of the technique used by Barrat came out in the following year, complementing Goodfellow's GAN with a deep convolutional method. Convolution networks employ the mathematical operation by that name, which happens to be very good in stitching

pieces of images together. It is directly connected to the idea of distributed representation described before, where tasks are broken down in smaller networks (or image areas). Although the paper with the method came out in 2016 (Radford and Chintala 2016), the first development of the code could be found already in 2015 (Chantala 2015).

The development of GANs hasn't stopped there, though. Several new models have been created since then. In the end of 2019, for instance, StyleGAN2 was released and proved itself able to create very realistic high resolution images.

Infinite Skulls

The main primary source of information for this section comes from the videos posted on the Parisian gallery's own site. This trove contains small pieces documenting the creation process, a short lecture by Robbie Barrat on the techniques he uses, an hour-long public discussion on AI and Art with him, the painter, the curator and a lawyer specialized in copyrights. The exhibition lasted only from the 7th until the 11th of February, 2019. It spanned a few rooms, divided in "first and second epochs" (that will be clarified later), a video piece that showed loops of the images being generated and an interactive peephole piece, which showed an unique painting and then immediately erased it forever.

The initiative for the show came from gallerist Catherine Vossen and artist Albertine Meunie, who realized that the obsessive paintings of skulls by painter Ronan Barrot could become a valuable initial set of images for deep learning processes (Gatti 2019, 0:22:00). He has been painting hundreds of those since 2011, and they all share some characteristics such as the orientation and size of the head. Robbie Barrat, on the other hand, was using DCGAN for several of his works, which needed to be fed with thousands of data sets (images in this case) to generate new sets that will resemble the originals. Please notice the irony of needing a huge amount of unique but similar images created by repetitive work to generate endless unrepeatable but similar machine made images.

Barrat used to collect public domains paintings belonging to specific categories such as abstract landscapes and nude portraits, using the technique named *scraping* to massively download them from the web (Barrat 2017). When he accepted the invitation to collaborate on a show with Barrot, they explored possible ways of working with data from the painter, even using the painter's visual references and influences as in input (Gatti 2019, 1:08:00). But they ended up choosing the direct approach of feeding the skulls straight into the network. The consequence of this choice was that the first batch of results came out remarkably similar to the original artwork, to the point of Barrot saying that he wish we would have drawn some of those himself (Avant Gallerie Vossen 2019).

Interestingly enough, at this point, the digital artist decides that the mere replication of the paintings wasn't enough for him to consider that as his own artwork. These images remained on the show in what was called the "first epoch" rooms, borrowing a term used to designate different periods of the learning done by the network. Nevertheless, Robbie decided to manipulate the input set by stretching, rotating the original and produce a second batch in which he could see himself. This selection was exhibited as the "second epoch."

This attitude is repeated throughout Barrat's discourse. When asked how to make the generated works to go beyond the mere repetition of original, his answer is to provoke misinterpretation, to confuse the network with disparaging data: "We must confuse the machine and make it hallucinate a bit" (Gatti 2019, 0:26:48). He also states that his work is mostly comprised of choosing and curating: the choice of the artworks that will feed the database, and curating the output to pick the most interesting results. Barrot, on his side, describes the authorship in these terms: "I chose, Robbie chose, Albertine and Catherine (. . .) In the end, the authors are all four of us." While it is true that there are plenty of discussions in contemporary art regarding the role of the curator as author (or at least co-author) (Lubar 2014), here this duality is exacerbated by the introduction of the machine as a producer of endless choices. We can start to identify the emergence of a particular category of computational artist dedicated to artificial intelligence, with several examples. An inventory of those would, unfortunately, fall outside of the scope of this article.

Barrat, here, is actively playing his role as an artist – not only by making choices on what to feed the algorithm, but also creatively criticizing the output and proposing new iterations. As we will see, selection and curatorship are fundamental skills for creators using GANs.

Generative art and representation

To completely grasp the effect of neural networks in visual expression, we must understand machine learning art as a special case of generative art.

Generative art is closely connected with the idea of using rules and constraints for a creative outcome. We can see it as an essentially algorithmic practice, where the artist's role is to define a method, more than to manually craft a final piece. Philip Galanter (2003) proposed the following definition: "Generative art refers to any art practice where the artist uses a system (...) which is set into motion with some degree of autonomy contributing to or resulting in a completed work of art."

It has existed before and independently from computers. Sol Lewitt, for instance, wrote sets of instructions so that some of his drawings could be executed by anyone. Some even argue that Jackson Pollock's method of action painting could be defined as

algorithmic (Boden and Edmonds 2009). It can also be said that to some extent, every art is algorithmic, in sense that it involves some set of procedures.

Two ideas around generative visual art will be important to our GAN discussion. The first is the fact that the generated art object can be multiple, that is, the same underlying algorithm can produce an infinity of unique works. This might happen for handcrafted works because the instructions can be interpreted in different ways, because the personal skills of each executioner are different and also because of the own nature of handcraft – the same person, following the same instructions, will probably create a different-looking work every time. In computer-based works, the uniqueness must be coded by means of random parameters or heuristic methods.

The other idea is that generative art detaches the creator of the work from its manufacture. In that sense, it becomes very close to conceptual art – and in fact, in this context, an algorithm is a way of representing a concept. This adds to the creative process a seductive notion of democratizing the artwork creation, since anyone could follow the rules of the artist and create his or her unique installment.

These two thoughts are also related to Walter Benjamin's (2008) reflections of about one hundred years ago, when he wrote about the work of art at the age of mechanical reproduction. The advent of mechanographic methods, in special photography and cinema, afforded the creation of numerous copies of an original, and inspired Benjamin to introduce the concept of the aura, the here and now of the work of art, connected to its uniqueness and history. "Reproduction (...) substitutes a mass existence for an unique existence." And now we not only face the mass reproduction of originals, but also the mass creation of unique pieces. Can we call these generated pieces originals? What kind of anti-aura does the mechanical generation of art afford? "(...) the destruction of the aura (...) is the signature of a perception (...) so increased that, by means of reproduction, it extracts sameness even from what is unique." Are we now doing the reverse, that is, obtaining uniqueness from what is the same?

It is clear that generative art has permitted a displacement of the "originalness" from the concept towards the mass produced artwork. But some sort of aura still resides in the mold, in the negative, there where Benjamin found uniqueness. Therefore, the singularity of visual generative art must reside in the system proposed by the artist: the set of procedures, the instructions, the method. If we want to understand GANs and generativity, we must look at its hermeneutics, specially at what is being translated. As McCormack et al (2014) asked, "In what sense is generative art representational and what is it representing?" The field is vast, and I can think of very different examples. CGI works depicting fantasy simulations of real world scenes, graphical syntheses of real world data, abstract visualizations based on mathematical principles, to name a few. But even when it is representing a bedroom, or real time stock market data, the computer (or the craftsperson) must translate something. This something might be lines of code, a

3d file, a recipe, something that in any case implies a *procedure*, the blueprint that *formalizes* an hermeneutic process.

It is not by chance that the aforementioned article (McCormack et al. 2014) also poses the question: "Can human aesthetics be formalized?" The evident conjecture there is related to a specific problem-solving approach in computer science that consisted of logical definitions and rules. "In the early days of artificial intelligence, the field rapidly tackled and solved problems (...) that can be described by a list of formal, mathematical rules." (Goodfellow and Courville 2016). To answer this, McCormack launches a new question: "What kind of aesthetics could be formalized?", implying that some are appropriate to this strategy and some are not. In 1965 Frieder Nake created a computer formalization of Paul Klee's aesthetics, and his "Homage to Paul Klee" today resides in the Victoria and Albert's Museum in London (Smith 2019). In 1966 Michael Noll did a similar project on Mondrian's visuals (Noll 1966).

But the development of neural networks, in special the advent of generative adversarial methods, would change completely the way we can address these questions.

The right tool for the right job

Artificial neural networks have been a method to find approximate solutions for math functions much before their application in visual art. The Cybenko (1989) theorem offered one of the first proofs of this strategy. A clear example of how they can be used to compute *any* function can be found in chapter 4 of Nielsen (2015).

Similarly, generative adversarial nets can be used to model – if not formalize – any visual aesthetics. How does that happen? The process of training a GAN consists of feeding it a number of pictures that belong to one particular coherent category. Ideally this number should be in the order of thousands. The generator network starts with images comprised of pure noise, while the discriminator will try to tell if the generated images belong to the fed category or not. In a process that can take from hours to months, depending on many factors, this feedback process will result in a gigantic statistical model of the input images. This is not, as in the Klee and Mondrian examples, a procedural method of steps to reproduce a given style. It is rather a representation comprised of infinite dormant images that can be brought to surface. Machine learning allowed us to do without procedural strategy, giving us an universal tool that forgoes the development of one algorithm for every style.

The person exploring the latent domain of a GAN can do it much the way a flâneur discovers a city, except that the space visited has much more than 2 or 3 dimensions. If the network was trained with human faces, he or she can stumble onto the neighborhood of an oriental child, that could perhaps be not far from a teenager with native American traces. In any case, no face will be equal to another,

and ideally – if the training wasn't *overfitted* – will also be different from any of those from the input set.

So while both procedural approaches may also be used to create a representation of some aesthetics, machine learning has the advantage of being able to do it for any coherent visual category, much like it enabled the solution of equations that could not be solved by classic methods of mathematical calculus. And while there are underlying procedural methods, as with any computational process, what is being translated in artificial intelligence process is no longer a recipe, but an abstraction, a latent space of endless potential new images - all coherent with a proposed aesthetic style.

As such, generative neural networks could be the greatest shift in the way we fixate images since the photography: if the 19th century invention gave us the ability to represent (and endlessly reproduce) any object, the neural networks has given us the potential to represent (and uniquely instantiate) an abstraction of any coherent category. In the network, the image becomes a sign, an operation analog to language: just like the word book resonates multiple real things that belong in the book category, the net contains endless instances of a given model.

Conclusion: a tool is just a tool

We are about to enter the 2020's, and GANs have given humankind an endless original making machine. The Next Rembrandt project generates paintings that I'd be glad to hang in my house: they could pass perfectly for a piece of the Dutch master and no one else would have a copy. It is art and it is generative, according to the definition we found. But whose art is it? The author-ship of these works resides in a limbo: while it is clearly not a Rembrandt, it could also not be signed exclusively by the engineers who designed the system. In that sense, it is very different from the works of Nike and Noll. Is it original, even if it is unique? And what artistic value does it entail?

These are not questions that can be answered shortly. But we certainly can find some clues. When Robbie Barrat saw the output of network trained with Barrot's skulls, he realizes that this is not his work – and moves on to manipulate them, so to claim it as his own expression. The painting sold by Obvious may be inspired by 19th century portraits, but could never be mistaken for one.

Procedural generation afforded the shift from visual *representation of actual* things to the *abstraction* of visual *sets of things*. These processes, however, were case-dependent, and couldn't be applied to any set. Machine learning is the first method that can be generalized to any coherent style, opening great new territories for artistic exploration. It is clear now the mimicking capabilities of GANs and neural networks in general have great value for scientific and commercial endeavors. But how valuable is it for art itself?

Probably the most intriguing and disturbing images created by machine learning are far from imitations and are only comparable to replicas of nightmares – revealing, perhaps, their neurological features and original inspiration. Some studies in cognitive visual perception proposed statistical models for peripheral vision that are eerily similar to low resolution domains of adversarial networks (Cohen et al. 2016).

Machine learning revealed itself to be a fantastic tool in chartering new frontiers of visual art. But it is most interesting when it becomes a tool for creators, instead of a mechanic, replacement artist. Art is probably what leaks through the gaps between the nodes of the net when the artist makes it “hallucinate a bit.” Its most creative and innovative works delve not into the impressive forgeries made by GANs, but into the mistakes they make, into the imperfect output of poorly trained generators, their glitches and into their nightmarish features.

References

- Avant Galerie Vossen. 2019. Barrat / Barrot, Infinite Skulls #7 // “J'aurais aimé le faire celuilà” Accessed in March 16th, 2020. <https://vimeo.com/309000957>.
- Bailey, J. 2019. “AI Artist Robbie Barrat and Painter Ronan Barrot Collaborate on 'Infinite Skulls.'” Accessed March 16th, 2020. <https://web.archive.org/web/20190206121643/> <https://www.artnome.com/news/2019/1/22/ai-artist-robbie-barrat-and-painter-ronan-barrot-collaborate-on-infinite-skulls>.
- Barrat R. 2017. art-DCGAN. Accessed in March 16th, 2020. <https://web.archive.org/web/20200316134023/https://github.com/robbiebarrat/art-DCGAN/tree/92ffbe92dd9360e2e567e6212c722f8eb63f12e1>.
- Benjamin, W. 2008. The work of art in the age of mechanical reproduction. Penguin UK.
- Bishop, C. M. 2006. Pattern recognition and machine learning. Springer.
- Boden, M.A., & Edmonds, E.A. 2009. “What is generative art?” Digital Creativity, 20(1–2), 21–46. <https://doi.org/10.1080/14626260902867915>.
- Chantala, S. 2015. Dcgan.torch – soooo much win. Accessed in March 16th, 2020. <https://web.archive.org/web/20191022113504/> <https://github.com/soumith/dcgan.torch/commit/45fd6727c36da67ff2fe357aab6e7eaa57ad9209>.
- Cohen, M. A., Dennett, D. C., & Kanwisher, N. 2016. “What is the bandwidth of perceptual experience?” Trends in cognitive sciences, 20(5), 324–335. <https://doi.org/10.1016/j.tics.2016.03.006>.
- Crevier, D. 1993. AI: the tumultuous history of the search for artificial intelligence. Basic Books.
- Cybenko, G. 1989. “Approximations by superpositions of sigmoidal functions”, Mathematics of Control, Signals, and Systems” 2(4), 303–314. <https://doi.org/10.1007/BF02551274>.

- Galanter, P. 2003. "What is generative art? Complexity theory as a context for art theory." In GA2003–6th Generative Art Conference.
- Gatti, E., Barrat R., Barrot, R, Degoulet C., Meunier A. 2019. Discussion about Art and Artificial Intelligence. Video recorded in 2019. Accessed in March 16th, 2020. <https://vimeo.com/325843365>.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., ... & Bengio, Y. 2014. "Generative adversarial nets". In Advances in neural information processing systems (pp. 2672-2680).
- Goodfellow, I., Bengio, Y., & Courville, A. 2016. Deep learning. MIT press.
- Jones, J. 2018. "A portrait created by AI just sold for \$432,000. But is it really art?" The Guardian, October 26, 2018. <https://web.archive.org/web/20190220004707/> <https://www.theguardian.com/artanddesign/shortcuts/2018/oct/26/call-that-art-can-a-computer-be-a-painter>.
- Lubar, S. 2014. "Curator as Auteur". The Public Historian, 36(1), 71-76. <https://doi.org/10.1525/tp.2014.36.1.71>.
- Lynch, G. 2018. "The transformative nature of networks within contemporary art practice." Doctoral dissertation, London South Bank University.
- McCormack, J., Bown, O., Dorin, A., McCabe, J., Monro, G., & Whitelaw, M. 2014. "Ten Questions Concerning Generative Computer Art" Leonardo, 47(2), 135–141. https://doi.org/10.1162/LEON_a_00533.
- McCulloch, W. S., & Pitts, W. 1943. "A logical calculus of the ideas immanent in nervous activity." The bulletin of mathematical biophysics, 5(4), 115-133. <https://doi.org/10.1007/BF02478259>.
- Minsky, M., & Papert, S. A. 2017. Perceptrons: An introduction to computational geometry. MIT press. <https://doi.org/10.7551/mitpress/11301.001.0001>.
- Nielsen, M. A. 2015. Neural networks and deep learning (Vol. 2018). San Francisco, CA, USA:: Determination 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(1), 1-10. <https://doi.org/10.1007/BF03393635>.
- Popper, F. 2007. From technological to virtual art. Mit Press.
- Radford, A., Metz, L., & Chintala, S. 2015. "Unsupervised representation learning with deep convolutional generative adversarial networks." arXiv preprint arXiv:1511.06434.
- Rosenblatt, F. 1957, "The Perceptron-a perceiving and recognizing automaton." Report 85-460-1, Cornell Aeronautical Laboratory.
- Schmidhuber, J. 2019. "Unsupervised Minimax: Adversarial Curiosity, Generative Adversarial Networks, and Predictability Minimization." arXiv preprint arXiv:1906.04493.
- Shannon, C. E. 1948. "A mathematical theory of communication." Bell system technical journal, 27(3), 379-423. <https://doi.org/10.1002/j.1538-7305.1948.tb01338.x>.
- Smith, G. W. "An Interview with Frieder Nake." Arts, 8(2), 69. <https://doi.org/10.3390/arts8020069>.
- Wiener, N. 2019. Cybernetics or Control and Communication in the Animal and the Machine. MIT press. <https://doi.org/10.7551/mitpress/11810.001.0001>.

CV



Bruno Caldas Vianna

University of the Arts, Helsinki

bruno.caldas@uniarts.fi

Bruno Caldas Vianna lives in Barcelona and is pursuing a PhD from Uniarts in Helsinki in Visual Arts and Machine Learning. He studied Computer Engineering but graduated in Film Studies. He has a master's from NYU's Interactive Telecommunications Program. He creates visual narratives using classical and innovative supports, having directed short and feature films, as well as working in live cinema, augmented reality, mobile applications and installations. From 2011 until 2016 he ran Nuvem, a rural art laboratory and residency space, located between Rio and São Paulo, and he worked as a teacher at Oi Kabum! art and technology school in Rio until 2018.