Studi di estetica, anno XLVI, IV serie, 3/2018
ISSN 0585-4733, ISSN digitale 1825-8646, DOI 10.7413/18258646060

Mario Verdicchio[1]

The digital in digital art

Abstract

Theorising over the relation between art and digital technology is challenging, because a new layer of analysis on the artistic use of computers was added in the 1960s over a debate on art that was far from over. Some technological and sociotechnical aspects of computers must be taken into account to form a more complete picture on what is going on in digital art. Technological characteristics of computers depend on the physical properties of their components, while their sociotechnical aspects derive from the fact that these artefacts are conceived, designed, built, and deployed in society. Dealing with electronics on the one hand and with companies on the other may not look relevant for a discourse in aesthetics, but computers are fundamentally dependent on these aspects of reality and, thus, an enquiry on what digital art is must take off from this standpoint.

Keywords
Art, Computers, Digital technology

1. *Introduction*

Both art and digital technology are very vast areas of human culture, intersecting in so many different ways, all of which deserving of philosophical analysis. Theorizing over the relation between art and digital technology is not at all easy, because not only the millennia-old debate on art is far from over, but also a new layer of analysis on the artistic use of computers was added in the 1960s when the works of the first algorists (experts in algorithms who were also artists) were shown to the public. While still debating on the question “what is art?” we started asking, “what is digital art?”.

[1] mario.verdicchio@unibg.it.

In this work, I will tackle the question by shedding light on some issues that I identified as critical on the path towards an answer, because they point to the characteristics of digital technology that may determine or hinder its role within the context of art. These characteristics have already been discussed in other philosophical works on digital art, but some technological and sociotechnical aspects have been neglected that are necessary to form a more complete picture on what is going on in this field. By "technological" I mean characteristics of computers that depend on the physical properties of their components, which can be analysed in isolation. Instead, by "sociotechnical" I refer to those aspects of digital technology that derive by the fact that these artefacts are conceived, designed, built, and deployed in society, and thus their working is based on agreements, conventions, contracts and, more in general, all sorts of relations holding within the human society.

I will try to show that computers are fundamentally dependent on these aspects of reality and, thus, any discourse on digital art that aspires to answer the question must move along these paths.

This work is organised as follows. In paragraph 2, I will argue against the traditional chasm between the "analog" and the "digital" by showing that not only the two technologies are not in opposition, but they coexist, and that digital technology is predominant because it allows for easier storage and transmission of data, which lead to two key aspects: memory and connectivity. In paragraphs 3 and 4, I will elaborate on these aspects: memory allows for automated and iterative processes, the key ingredients of generative art; connectivity enables geographically-distributed computers to work in synchronization, making net art possible. In paragraph 5, I will focus on speed, which is not exclusive to digital technology but is needed for real-time responses from computers, which are fundamental for interactive art. In paragraph 6, I will analyse the latest trend in digital art, which encompasses all the previously analysed aspects of digital technology, namely, machine learning applied to art. In paragraph 7, I will propose a general way to analyse digital art and build a case against a particular way of using computers for creative endeavours.

2. *Analog and digital*

Philosophical analyses on digital art typically begin with juxtaposing digital technology with analog technology, presenting the former as a newer, more powerful evolution of the latter, with characteristics that are indicated as creating a chasm between the digital-based endeavours and their analog counterparts. Thomson-Jones's account on digital art indeed begins with an analysis of the distinctions that philosophers have drawn between these two different realms (Thomson-Jones 2015). Such distinction, an exercise in ontology, is not straightforward, and the involved philosophers are not in agreement (Lewis 1971, Goodman 1976). One take in particular is dismissive of such contraposition: Haugeland states that the digital is "a mundane engineering notion [...]. It only makes sense as a practical means to cope with the vagaries and vicissitudes, the noise and the drift, of earthy existence" (Haugeland 1981).

Indeed, if we observe digital technology up-close, we notice that analog and digital instruments have more shared features than differences. In electronics, although analog signals are treated as a continuous and dense series of voltage values and digital signals as a discrete sequence of voltage pulses, from a physical perspective they are the same phenomenon, that is, they are all electromagnetic waves. They do have a difference in shape: if drawn in a diagram of voltage over time, digital signals look like staircases because they only assume discrete values, whereas analog signals present a more sinuous shape because they have no restrictions on their values. Apart from this distinction, the two types of signals behave in the same way and, in particular, they are both subject to physical phenomena like noise and drift.

This is where, as Haugeland rightfully indicated, engineering plays a role: by having a limited range of possible voltage values (i.e. only discrete ones), digital signals are much easier to store and transmit over long distances than analog signals. Let us consider a digital binary system that works with a high voltage of 5 Volts and a low voltage of 0 Volts. Imagine some disturbance that affects a signal in a point of high voltage, by bringing the voltage down to 4.9 Volts. Since there are only two possible voltage values, the risk of error is minimal: the system elaborating the signal can incorporate detection and compensation mechanisms to bring the signal back to its original value. On the contrary, in an analog system where all voltage values are possible (at least within the range of operation), the relevant system would not able to establish whether a 4.9 Volt value is the result of a disturbance or the correct value of the original signal.

This is the key difference that determined the success of digital technology over the analog. It is a matter of practicality rather than an actual ontological distinction: most systems now rely on digital signals because they are less affected by disturbances and this makes them easier to store and to transmit.

This does not mean that analog systems are out of the picture at all. For instance, even if we call it “digital photography”, the very first point of contact between a camera and the external world is based on analog technology: the light sensor in CCD (charged-coupled device) cameras registers the light waves from the outside in the form of continuous electrical charges. After an electrical phenomenon has been created that is analogous to the light phenomenon captured by the camera, digital technology intervenes and those recorded electrical charges can be turned into discrete values to be stored, elaborated, transmitted, etc. Digital photography is just one of many examples of technological endeavours where not only the analog and the digital are not exclusive alternatives, but indeed coexist and cooperate within one system.

Digital imaging is a context that enables me to illustrate one of the caveats of digital systems: using only discrete values to describe a signal instead of a whole continuous range implies a loss of information. To show an image on a digital LCD monitor, electric pulses are sent to command the emission of coloured light by the electronic components that form the pixels on the screen. These pulses represent numerical values that are put in correspondence with different colours. A correspondence between colours and the numerical values of the frequencies of the relevant light waves was first introduced by the International Commission on Illumination, or CIE (Commission Internationale de l'Eclairage) in 1931 (CIE 1932), and later such convention was adapted to digital technology by IBM in 1981 for their first Personal Computer (IBM 1981). Nowadays most digital monitors use the RGB (red, green, and blue) scheme for displaying colour at each pixel: a colour is specified in terms of its red, green, and blue components, each of which is expressed with an integer value that can vary between a minimum of 0 and a maximum of 255. For instance, the triple (0, 0, 0) corresponds to black (no colour at all), (255, 0, 0) to pure red, (255, 0, 255) to magenta (mix of red and blue), (255, 255,

255) to white (all colours present), and so on. Given that there are 256 × 256 × 256 = 16,777,216 different possible triples, this means that each pixel of a digital monitor can produce as many different colours.

This is a significant technological achievement; however, it comes with a limit. A monitor can produce the colour sage green (136, 179, 120) and also the unnamed shade of green corresponding to the RGB code (136, 180, 120), which differs from the previous because of a +1 difference in the green component. In the digital colour space, these two shades of green are adjacent, but in the physical world it is possible to have a light wave with a frequency that is an intermediate value between sage green and the unnamed green shade. This means that the digital colour palette is a subset of what is physically possible.

After all, we might really need to establish an ontological distinction between analog and digital systems, not based on their intrinsic physical characteristics but on their scope. However, there is no such distinction when it comes to colour phenomenology: most, if not all, human beings are not able to distinguish the abovementioned two shades of green. This is why we can afford to have digital technology for our images: its intrinsic limitations lie below the threshold of our perception. Now let us move on to the strengths of digital technology.

3. *Memory*

Before dealing with the memory and connectivity devices that digital technology allows for, I need to make a clarification. The fact that current digital systems use 256 different values to indicate the level of a colour component does not mean that these values are stored and transmitted in and between electronic circuits as instances of 256 different voltage levels. It is much more practical to exploit the versatility of binary numerical systems and have only two levels with which to implement binary sequences that correspond to the 256 different values. We call the components of these binary sequences “bits” and we usually indicate them with “0” and “1”, referring to the low voltage level and to the high voltage level in the circuit, respectively. This is the concept underlying current digital systems, which are often described as storing, elaborating, and transmitting bits. This makes digital systems more complex and expensive than analog systems, because instead of having one analog device (e.g. a capacitor) to store one physical, non-integer voltage value, we must have an array of digital devices, each storing one bit that is part of the binary digital approximation of that value. Digital memories have nevertheless been extremely successful because, as explained before, their contents are simpler to maintain.

Another success factor for digital memories derives from the versatility of the binary code, which enables computer designers to easily create encodings, that is, mathematical correspondences between finite sequences of 0s and 1s and entities in the physical world. We have already seen the RGB scheme, where the encoded entities are colour descriptions, but the system can be straightforwardly extended to encode not only all sorts of quantifiable data, but also commands that a computer is to execute. This was the great intuition that brought digital memories to the centre stage of computer science mid-20th century: the possibility to store not only the data to elaborate, but also the instructions by which such data were to be elaborated. The adoption of the so-called "stored program" concept is still considered the most significant single step in the development of modern computing (Haigh 2013). This invention has two fathers, the greatest computer scientists of all time, Alan Turing and John von Neumann, and the debate on the one who conceived it first is not over yet (Pelaez 1999).

This may sound like an exquisitely technical detail, but the stored program concept is what made digital computers stand out not only among electronic devices, but also among all inventions of humanity, because it expanded the possibilities of automation to an unprecedented extent. Automation on a large scale wasn't born with computers, but with the industrial revolution, and it was then based on mechanical machines. The physical design of the machine would determine the one operation that could be automated by means of that machine (e.g. a wheel would rotate, a piston would move up and down). However, with the advent of digital memories, there was the new possibility of writing different commands inside the electronic machine itself, in the form of bits in the memory. Thus, digital memories with the stored program concept allowed, for the first time in the history of technology, for the storage of data and the operations to perform on those data. This was the birth of automated iteration, that is, the possibility to program a machine to perform complex sequences of different operations, which paved the way for the first examples of digital art.

The beginning of digital art goes back to the 1960s, when three computer scientists independently experimented with their computers to create geometrical designs: George Nees at Siemens in Erlangen in Germany, Michael Noll at the Bell Labs in New Jersey, and Frieder Nake at the University of Stuttgart, Germany. Some works of Nake and Nees were shown in the gallery Wendelin Niedlich in Stuttgart in November 1965, in what is universally considered the first digital art exhibition (Bense 1965). The works that these algorists proposed are all extraordinarily similar, to the point that it is hard to believe that they were developed independently. They all consist of graphical compositions with a number of randomly oriented lines that form closed or open polygons. Nake provides a possible explanation, quoting Nietzsche, who wrote in 1882 to his secretary Köselitz "our writing instrument attends to our thought", about a typewriter with only upper case letters. This is to state that even a free mind like an artist's ends up following the affordances provided by the instrument chosen for their creative activity. In the case of Nees, Noll, and Nake such instrument was a digital computer of the 1960s with very limited graphical operations (including at least a function to trace segments between two points) but with the abovementioned possibility to automate the iteration of such operations. According to Nake, anybody with some artistic ambitions and such a machine at their disposal would have obtained results like his (Nake 2012).

I need to remind the reader that, although they appear to be randomly oriented, no lines traced by means of a digital computer can be considered really random, since everything that happens inside these machines is deterministically programmed. We should rather refer to pseudorandomness, that is, the illusion of randomness provided by a sequence of outputs of complex mathematical functions stored in the computer's memory. These functions are designed in a way to provide results that cannot be easily recognized by a human observer as belonging to an arithmetic scheme but, as every value inside a digital computer, they actually do. A digital artist who exploits pseudorandom numbers to create an artwork has a rather precise idea of the drawing they are going to make, but cannot exactly foresee the positions at which the geometrical elements will be placed because they are not able nor willing to do the math the computer is going to do to establish such positions. Thus, once the work is completed, the artist is looking at a partially unexpected result. Since this kind of creative process has a significant generative character, the field of digital art that relies on iteration and pseudorandom numbers is called "generative art".

4. *Connectivity*

While the first digital artists were showing their works to the public, in a very different setting, other computer scientists were looking for ways to exploit digital technology in the field of telecommunications. In 1969, the first version of the Internet was born: ARPANET (Advanced Research Projects Agency Network) connected three universities and one research centre in the USA (University of California Los Angeles, University of California Santa Barbara, University of Utah, and Stanford Research Institute) to enable the sharing of the computational power of the machines on all these premises (Roberts 1978). The idea was rather simple but extremely clever: since digitized data are a sequence of electric impulses, instead of sending them from origin to destination preserving the sequence, it is possible to spread them over different channels, and then rebuild the sequence once all impulses reach the destination. This wouldn't be possible with analog signals, because one continuous wave cannot be broken down in parts, and the advantage is that the transmission is more robust, since if one route is not working, the data packets can be sent over other alternative routes. Moreover, since digital data are easy to regenerate in case of decay, noise, and drift, the signal at destination is identical to the signal initially sent.

After 1969, more universities, research centres, and military bases joined the initiative and got connected, but this new level of telecommunications did not have an impact on creative practices around the world until three decades later. If digital computers had almost an immediate impact on art in the form of generative art, why didn't the same happen with digital telecommunications? Firstly, the Internet requires a great amount of cables, and those cables require time to be placed all around the planet, including under the sea. Moreover, since we are dealing with digital telecommunications, we need digital computers at both ends of these cables for the system to work, that is, people needed to have a computer to get access to the Internet. Given that computers were, at least in the early years of computer science, only devoted to scientific computation and extremely expensive, the diffusion of this digital technology needed further time. However, this didn't stop the slow but steady expansion of the Internet, which reached tens of millions of people in the 1990s, when the online population formed a critical mass that allowed for the first Internet-based artistic experiments, also known as "net art".

Net art began when the Internet and, in particular, its hyperlinked page-based content (i.e. the Web) became a widespread commodity in the Western world, and the works in this context present a very tight relation with the evolution of the Internet and the way people are enabled to use it. First and foremost, this kind of endeavour exploits the connectivity of the Internet: net art works are created by artists who set up a computer-based content that viewers can access through the Internet, most frequently by means of their Web browser. Some readers might recall what websites were like in the early days of the Internet and, especially, how limited the role of the user was. The communication paradigm was one-way: net artists created Web-based content, and users at home viewed it on their browsers. All of the most significant net art works (Greene 2004), differences in content aside, can be traced back to this paradigm: for instance, Yael Kanarek's *World of awe* presents a fictional online diary comprised of love letters to cyborgs and 3D models and landscapes (Kanarek 1995), whereas Olia Lialina's *My boyfriend came back from the war* uses black and white images and short bits of text to tell the story of an unnamed boyfriend coming back from an unspecified war (Lialina 1996).

These works can be viewed as a new way to perform story telling. Although the users at home are relegated to a very simple role of viewers, net art provides something different from the more traditional TV or movie viewing experience: by means of the hyperlinks in these webpages, the users can determine the order with which they view the different contents of the artwork and, given that the computers providing such contents are constantly on and connected (unless some technical issue rises), users can experience the artists' narratives at any time, and anywhere there is a computer online at their disposal. For this to be possible, the above-mentioned digital memory plays a fundamental role in enabling the retention of the net art content in the form of bits on the artists' computers, to be sent around the Internet at the users' requests.

The early net art works, however, still shared the characteristics of the one-to-many broadcast model with traditional TV: there was one content producer, the artist, and many content consumers, the users at home on their computer. This changed in the second half of the 2000s, not because of a radical technological innovation (although an evolution of the programming languages for the Web helped), but rather for a change in the role of computer users at home. Facebook was born in 2004, and along with it a number of other social network websites, which provided standardized, ready-to-use platforms for users to publish their own multimedia content on the Web, for others to see, listen to, and comment on. The enormous success of these endeavours did not make all Web users digital artists, obviously, but opened the gates to a flood of digitally shared personal contents, which inspired net art works with an unprecedented level of personalisation, only this time with the Web user, and not the artist, at centre stage.

The most notable work of personalised net art came from a collaboration between director Chris Milk and a team of programmers at Google, headed by Aaron Koblin (Koblin 2010). *The wilderness downtown* is a browser-based experience (specifically, Google's browser Chrome) that exploits the data of Street View, the company's image database of the streets around the world. At the beginning, the user is asked to insert the home address from their childhood, and then multiple browser windows open, to a song by Canadian rock band Arcade Fire. The windows show previously recorded videos of a teenager running on a street, alternated with computer-generated flocks of birds in the sky, moving aerial photographs and street views of the area of the address inserted by the user. Towards the end, when the song reaches its climax, computer-generated trees start quickly rising up from the ground, covering the streets of the user's childhood, as if nature is reclaiming the whole area. Once the song is over, a new browser window opens in the centre of the screen, showing a blank canvas, and inviting the user to write a message to themselves as a child, by typing on the keyboard and also drawing with the mouse. The drawn lines graphically develop into branches and, after the message is complete, a number of birds fly in and rest on them.

The technical advancements between the 1990s and the 2010s in terms of the contents that a browser can show are obvious: in little more than a decade we go from text and digital photographs to full-fledged videos, superpositions of computer-generated graphics and photos, computer-generated graphics interacting with user-generated drawings on the fly, and so on. These enhancements, which are theoretically made possible by the digitisation of the contents, are made practically feasible by the technological evolution of digital devices, comprised of circuits that are every year more miniaturised and denser with transistors, which increases the number of operations that a computer is able to perform per unit of time. I will get back to operational speed in the next section. Here I want to focus on the role of the connectivity provided by digital technology: not only it eliminates the distances between a net artist and their audience at home, but it also enriches the contents included in their work to a point where they can be based on data selected by the users and retrieved from a remote repository. The experience of *The wilderness downtown* is built upon code written by a net artist, Google’s Street View data, and a selection by the user. The locations of these three sources do not matter: digital connectivity provides a virtual space where they can coalesce and create a unique, personalised aesthetic experience.

5. *Speed*

Speed is often associated with computers and digital technology and it is necessary to make many current works in digital art possible. However, it is not with computers that speed entered the realm of human culture. Speed was there already decades before the first digital computers. For instance, it played a significant role in the manifesto of Futurism, where a new kind of beauty was sung, the beauty of speed, embodied in the manifesto by a racing car (Marinetti 1909). Almost one century before that, the first practical photographic process was perfected by Luois Daguerre, who proposed to cover copper plates with light-sensitive silver iodide, to enable an instantaneous impression of the light waves of a scenery onto a small, rectangular surface. I did not choose these pre-digital examples of fast technology randomly: they are two manifestations that present a significant difference that can be found in digital technology and digital art still today. I am referring to the chasm between the microscopic and the macroscopic, a contraposition that is much more obvious and important than the analog/digital dichotomy.

A car is fast: faster than a running person, faster than a horse. Even faster than a car is a jet plane, which can take us across the Atlantic in a few hours, compared to the days needed on a ship at the times of the manifesto of Futurism. However, the speed of these mechanic artefacts cannot compare to the speed at which even the most rudimentary photographic apparatus operates in capturing the light waves on film. Photography literally works at the speed of light, because its fundamental process intrinsically depends on the movements of photons, the wave-particles responsible for light. No matter the evolution of mechanical engineering, it is a given of the laws of physics that moving macroscopic objects with a mass requires a significant amount of energy that grows with the mass that needs to be moved and the speed at which it is supposed to be moving. Instead, endeavours like photography, telecommunications or digital computing, whose operations rely on microscopic particles with little to no mass, can function at a much higher speed, following the same laws of physics, but on a very different scale. The speed of microscopic phenomena is exploited nowadays by digital technology, but it has long been a fundamental part of analog technology, not only in the above-mentioned photography, but also in telecommunications: the first transatlantic phone call was placed between New York and London in 1923 (Young 1991).

Memory devices in computers are a good example of how the micro/macro divide is more determinant a factor for speed than the analog/digital. Think of mechanical hard disk drives (HDD): piles of magnetic platters whose surfaces are used to store bits in the form of oriented magnetic domains, that is, microscopic natural compasses that are interpreted as “0” if oriented in one direction, as “1” if oriented in the other direction. Although the single data components may be stored in microscopic form, this technology relies on macroscopic mechanics: a motor that rotates the discs, and another that moves an arm with a magnetic point that reads and writes on the platters.

HDDs are digital technology with a mechanical component, whereas the latest storage technology, called Solid State Drive (SSD), is digital and completely electronic, without moving parts. In SSDs, the selected data need not be retrieved by a mechanical arm: their digital values are moved along the circuits in the system in the form of electric pulses. Currently, SSDs are able to read and write data 5 times faster than the best HDDs.

Speed is not a definitory characteristic of digital technology, but it is there whenever a system operates with elementary particles like electrons and photons. A system that operates at light speed can respond to users' requests very quickly, and this is why computers were so successful in the first place. However, if a system can also receive such requests quickly, then a new paradigm of interaction between users and machines can come into play. The evolution of photographic technology in the form of light sensors, combined with the operating speed of computers, allowed for another branch of digital art, called "interactive art".

American artist Scott Snibbe is one of the first artists to work with interactivity by means of computer-controlled sensors and projectors. One of his most successful efforts in this context is *Boundary functions*, shown at the Ars Electronica festival in Linz, Austria, for the first time in 1998 and, after that, at several other venues around the world until 2008 (Snibbe 1998). The work is an installation including a flat, square platform on the floor, on which people are to walk around. The presence of people on the platform triggers a projection from above, so that straight geometric lines appear on the floor, tracing boundaries that separate each person on the platform from the others. The projected image is dynamic, and the lines move in accordance with the positions of the participants: as people move, so does the image to always keep a line between any pair of persons on the platform. The artist's intention is to provoke a discussion on "personal space", which he wants to show as the result of a negotiation between individuals rather than an area developed independently.

One of the stand-out characteristic of this work is interactivity: the system is comprised of devices that enable participants to move around freely and receive a response in accordance with the artist's vision in real time. The speed of operations is fundamental for the artistic statement of *Boundary functions*, because if the system lags and the lines are not updated according to the participants' movements, then the spell would be broken and the aesthetic experience would be affected negatively. The artwork relies on the sensors in the camera on the ceiling to create, in a matter of microseconds, a snapshot of the platform from above. Such image is then processed by a computer to determine, by means of an algorithm that analyzes the colours of the image's pixels, the positions of the participants. Mathematical rules are then applied to compute the coordinates of the lines delimiting the areas surrounding the participants. Such partitioning rules were first discovered by Russian mathematician Georgy Feodosevich Voronoy between the end of the 19th and the beginning of the 20th century. Finally, once the so-called Voronoy diagram of the plane divided into different areas (each surrounding exactly one participant) is ready, it is projected down onto the platform by means of a projector. All this process must be completed before the participants make the next step, because once one of them moves, their new positions must be captured, and a new diagram must be calculated and projected.

It takes a lot of imagination to think of realising a similar kind of interactive experience with any other technology. We could build a mechanically and electrically enhanced platform, comprised of a matrix of small scales that on the basis of the participants' weight distribution can activate tiny lamps to create the lines of the Voronoy diagram. Another solution could be a highly skilled group of people with aptly masked flashlights with which they could create the diagram on the spot, according to the participants' movements. Naturally, not only both these solutions are extremely hard to implement, but also their results could not compare to the speed and precision of the experience that Snibbe's work offers.

Indeed, the need for computers in the endeavour to create interactive artworks like *Boundary functions* is so stark that there are philosophers of art, like Dominic McIver Lopes, who consider interaction as a necessary property: a work of computer art[2] is such that it is interactive, and it is interactive by means of a computer (McIver Lopes 2010). This is a very restrictive definition that put the status of more "traditional" works in doubt: is the print of an abstract graphic work of generative art still computer art in this perspective? The centrality of interaction and, more specifically, interaction made possible thanks to a computer seems to require the necessity of digital technology for the realization of a work of computer art. If there is a feasible alternative method, such as manual drawing with the help of a ruler and a compass, is it still computer art? This is a discourse that involves the concepts of craftsmanship and authorship to which I will return later.

[2] McIver Lopes speaks of "computer art" and, in a world where analog computers are not produced anymore, this concept can be considered synonymous with "digital art".

Of course, even if we followed McIver Lopes and considered interactivity as a necessary condition for digital art, it is not sufficient: there are aesthetic experiences that are interactive but not digital, like interactive theatre, where actors break the fourth wall and involve members of the audience in their performances. If interactions based on mathematical concepts like *Boundary functions* can exist only thanks to computers, others may rely on this technology for more pragmatic reasons of feasibility and practicality. *Carne y arena* is a mixed-media installation heavily relying on Virtual Reality (VR) technology by Mexican director Alejandro González Iñárritu, who designed it with the aim to give spectators a first-person perspective on the hardships that an immigrant has to face while trying to cross the Mexico-US border. It was the first ever VR-based project presented at the Cannes Film Festival in 2017, and since then it has been shown around the world in a tour that included Italy, Mexico and, currently, the USA (Iñárritu 2017). Once the spectator reaches the VR-based stage of this installation, they are helped by two assistants, who make them wear a head-mounted device (HMD) that provides an immersive experience set in the Mexican desert. The 3D graphics is extremely detailed, and the fact that the spectator is barefoot in a room full of sand increases the realism even more by means of tactile sensations. At the beginning, their role is that of an observer, accompanying the Mexican immigrants on their journey through the desert. Even when the immigrants are caught by border patrol with a helicopter, Hummers and dogs, the spectator is still free to wander around and observe the scenes, the people and the objects in it from any perspective they want. In the final part of the immersive experience, however, there is a sudden change: border patrol agents detect the presence of the spectator, and start treating them as one of the immigrants. I will stop here not to spoil *Carne y arena* for those who will have the chance to experience it. I was lucky enough to try it and it was an extremely intense emotional journey. Indeed, the producers of the installation warn that it is not recommended for individuals with heart conditions.

Carne y arena might be implemented in the form of interactive theatre, with several actors, a sound stage, dogs and vehicles, but it is clear that having an immersive 3D model of the scene that users can navigate through with an HMD allows for multiple repetitions of the experience for a great number of spectators without any fatigue of the actors and the animals involved. Moreover, graphically simulating a hovering helicopter by means of digital technology is much cheaper and practical than having a real one in a sound stage or, to avoid the complexities of governing an aircraft indoors, setting up the installation in an actual desert area.

VR artworks like *Carne y arena* are the latest stage of the evolution of interactive art that started 20 years ago with Snibbe's pioneering work. Now like then, speed is the key characteristics of technology that makes interaction possible. Let me remind you once again that such speed does not rise from the digital, but from the microscopic: whether with projectors or with HMDs, we are dealing with photons that move at the speed of light, whose trajectories are computed by circuits in which electrons move as quickly. Attach these computers to a macroscopic peripheral, and the seamless, real-time interaction is lost, like when Roman Verostko attached a paintbrush to a plotter for some of his paintings (Verostko 2004), or when Benjamin Dillenburger and Michael Hansmeyer sand-printed their algorithmically generated geometries (Dillenburger 2014). Processes based on mechanical machines may be quicker than handmade ones, but they still take a significant amount of time, even if they are controlled by computers. In some cases, speed makes art quicker to create. In interactive art, speed simply makes it possible.

6. *Machine learning*

We are currently living another wave of Artificial Intelligence (AI). This subfield of computer science, born in 1956 at a summer research retreat (Moor 2006), aims at modelling, reproducing, and improving via computers what is traditionally attributed exclusively to humans, like thinking, reasoning, being creative. During the last 60 years the discipline had its ups, in terms of promising new computational techniques and technologies many got overexcited about, and downs, the so-called "AI winters", when the disappointment of unfulfilled promises redirected interest and funds towards other research endeavours, relegating AI to small academic circles.

Now AI is in popular demand again, in a great number of fields, ranging from finance to the automotive industry, mainly thanks to technological advancements that made theories from the 20th century feasible to implement on actual computer systems. The basic idea dates back to the 1940s, when Warren McCulloch, a neuroscientist, and Walter Pitts, a logician, proposed a simplified computational model of a neuron, in their quest to understand how the connections of many basic cells in the human brain could produce complex patterns (McCulloch and Pitts 1943). A fundamental stepping-stone was given in the 1980s by collaboration between cognitive scientists from the University of California San Diego and computer scientists at the Carnegie-Mellon University: a new procedure, called back-propagation, that repeatedly adjusts the numerical numbers associated with the connections in a network of computational neurons (Rumelhart *et al.* 1986). Very roughly inspired by the human brain, these neural networks are mathematical models of interconnected units that, like neurons, fire a signal out if they receive stimulation in input above a certain threshold. Every connection between two neurons is characterized by a multiplying factor, called "weight", that indicates the strength of that connection. This groundbreaking work proposed a technique to modify those weights in accordance with how far the network was from a desired behaviour (e.g. sending a "yes" in output when fed the digital data of a photo of a cat in input, in order to create an automatic image classifier) so as to minimize that distance. The authors showed that, as a result of these weight modifications, units inside the network come to represent fundamental features of the task domain in numerical terms (in the example, the common traits that characterise images of cats), and the regularities in the task are captured by the interactions of these units (i.e. the network "learns" what cats look like and stores that knowledge in the weights of its connections). In the 1980s computers were not fast enough to compute networks of significant complexity, so the tasks they could deal with were extremely limited.

In the 2010s, advancements in digital technology enabled Google to design and implement AlphaGo, a system based on neural networks that not only learned how to play the game of Go, but also became the strongest player in the history of the game (Chen 2016). The strength of AlphaGo was also based on a vast database of past games compiled by master players hired by Google for the project, but the most striking aspect of the initiative was indeed the use of neural networks to make a computer learn how to play well. This is a kind of AI that is profoundly different from the discipline that was born in the 1950s: back then logic was the foundation of a research effort aimed at specifying explicit formal rules of the task the computer scientists were trying to automate, whereas now the knowledge is formed within the weights of a neural network, in a form that cannot be understood or analysed by human beings, if not through the results that the network is producing, in terms of closeness to the desired results, whether they are about recognising images of cats or playing games of Go. This is the subfield of AI called "machine learning".

Machine learning is the newest effort in computer science to be applied to art. AI and art already intersected in the past, prominently thanks to Harold Cohen, a London-born artist who spent decades, starting from the 1970s, writing code for his AARON project, a computer program comprised of formal rules, just as envisioned by the AI pioneers, to create visual art with a particular style (Cohen 1995). The current endeavours in machine learning and creativity are based on a different methodology, as follows: since neural networks learn the features and the regularities of the data they are trained with, what if we train them with artistic images? Neural networks were able to find and learn regularities in the works of different artists. A group of computer scientists based in Tübingen, Germany, shed some light on how neural networks store information on images in their units. In particular, they discovered that the layers of the network that are closer to the input layer (the array of neurons receiving the image) store more information on the details of the pixels and how they form visual textures or "style", whereas the layers closer to the output (where the network gives the results of its computation) are dedicated more to abstract information about objects and colour blobs in the image, their contours and position, that is, what they call the "content" of the image (Gatys *et al.* 2015). The researchers used neural networks to isolate the "artistic style" of a painting like Van Gogh's *The starry night*, separate the content of a photograph, in terms of object contours, and mesh the two together, obtaining a photograph that presents a visual appearance that strongly resembles that of Van Gogh's masterpiece, although with a different subject.

Google's role with neural networks has been very strong not only with games, but also in the visual arts. In trying to understand better how the different features of an image were learned by the layers of neural networks, a team led by engineer Alexander Mordvintsev came up with a technique to feed random images to neural networks already trained to detect a specific kind of pattern (Mordvintsev 2015). After a number of iterations, the images manipulated by the neural network showed a very particular psychedelic aesthetics, in which the original image was deformed to show many instances of the pattern the network was trained to look for. The most famous example of this technique, which Google named Deep Dream, is a music video in which a trip along the aisles of a supermarket becomes a hallucination with dogs all around (Pouff 2015).

Another innovation with neural networks proposed by Google is to pit two networks against each other in a technique called Generative Adversarial Networks (GANs). A "generative" network trains over a set of data to learn their features and generate new data that share those features but are not from the training set (i.e. they are original instances of a given kind); a "discriminative" network has the task to recognize which outputs from the generative network are original and which are from the training set. The two networks work against each other in the sense that the generative network aims at increasing the error rate of the discriminative one, while the latter tries to increase the accuracy of its detection process (Goodfellow *et al.* 2014). German artist Mario Klingemann is one of the most prominent contemporary artists who make use of GANs. His latest efforts are devoted to create portraits of faces and bodies of entities which look like real people, but they are actually the result of a generative neural networks trained over hundreds of thousands of photographs retrieved on the Web. For his 2017 series titled *Pose-to-picture*, Klingemann made use of pornographic pictures from the blogging platform Tumblr, because they were a good source of full-figure imagery. The dataset was fed into the GAN to produce images exploiting the speed of the computer running the system. The generative process can obviously go on until the artist finds an output that he considers aesthetically interesting. In case all outputs show the same kind of figurative defects, then it means that the parameters of the neural network need some tweaking. Although the results are still far from realistic portraits, even the initial attempts consist of very interesting quasi-figurative works, with uncanny deformations of the human figure that remind us of Francis Bacon's works (Elliott 2017).

Machine learning applied to art synthesizes all the most important characteristics of digital art that we have seen so far: it relies on the speed of connected computers, whose memories store mathematical models of images in the form of digital data and the algorithmic operations to further elaborate them. Is this art? Are the people who make use of such sophisticated technology artists? The scrutiny that started with Nake's pioneering works is still there and, in the light of the latest computing devices that automatize most part of the creative process, that scrutiny is closer than ever.

7. *Digital and human*

The very first criticism early digital artists had to face in the 1960s went along these lines: people who are able to use computers are mathematicians and physicists; mathematicians and physicists are not artists; hence, whatever is conceived, implemented and produced by them is not art. Reflecting on this critique almost 50 years later, digital art pioneer Frieder Nake fully agrees with the first point: back then, only a very specific kind of scholars had access to computers and were able to program them to create visual works. He also admits that the aesthetics of the early digital artworks may have been rather simple due to his and his colleagues' professional background: if people with a more solid artistic background had taken on programming to create their works, digital art could have been born with much more engaging visuals (Nake 2012). However, Nake is adamant about the artistic value of digital art's pioneers: if the value of a work is established also by the audience's appreciation, then whether the first digital artworks were made by mathematicians or more traditional artists, the spectators surely appreciated the undeniable revolutionary aspect of these geometric lines plotted on paper, for the first time in the history of human culture created by means of a programmable computing machine. As a playful reminder of this novelty, Nake used to sign his early works not only with his name, but also with the model numbers of the computer and the plotter that he had used.

In pointing at the audience as a source of artistic value, Nake openly embraced avant-garde artist Marcel Duchamp's beliefs on the role of the audience in art: "the creative act is not performed by the artist alone; the spectator brings the work in contact with the external world by deciphering and interpreting its inner qualifications and thus adds his contribution to the creative act" (Duchamp 1957). Nake made Duchamp his champion because a "this is not art" criticism was also moved against the French-American artist when he introduced the "readymades" in his art in the first decade of the 20th century. With the concept of the "readymade", an object that becomes an artwork by an act of selection of the artist, Duchamp shed light on choice as a determining factor in art. I consider this aspect of his artistic theory more important for the enquiry on digital art than the role of the audience.

Nake's argument on the groundbreaking adoption of computers as artistic tools is historically circumscribed: it is meant for the 1960s, and it wouldn't work for the contemporary world, where computers are part of our everyday life in so many aspects. There is nothing revolutionary in trying to employ computation to deal with a task, nowadays. We need to look elsewhere to find a significant characteristic of digital artists, and choice is one. "When an artist paints using a palette he is choosing the colours" said Duchamp (Powell 1966), but think of the choices that digital artists have at their disposal when they are using a computer. The choices that digital technology provides are not a superset of those connected with a canvas and colours: we are talking about two different domains, although both digital art and traditional painting can produce visual artworks. The materiality and physicality of traditional painting, in terms of gestures by the painter, is something that cannot be found in digital art: a digital artist can design and 3D print an artefact that looks like a painted canvas, but the process to get to that result is completely different from what a painter does. However, digital technology offers something that traditional art does not have: a programmable machine.

Simply put, a paintbrush can be used in many different ways, but all the acts that a person can perform using a paintbrush are of a specific kind. The same goes for a chisel, or a violin. In a very general sense, computers are no different: we can only perform computation with them, but thanks to all the technology that converts electric impulses representing numbers into physical phenomena and vice versa, digital artists can use that computational power to draw, sculpt, play music, and so on. Given that computers are connected through the Internet, those operations can be distributed over geographical distances; since computers operate at extremely high speed, interaction with users in real time is possible. Moreover, computers' internal representation of physical phenomena in the form of numbers, and this is where the digital really makes a difference in aesthetic terms, digital artists are given the possibility to manipulate those numbers with arithmetic operations, and then convert them back to new instances of phenomena that cannot be easily created in the physical world, like synthesized timbres in music, or shape morphing in animation.

This very significant potential for creativity is what made British philosopher Paul Crowther argue that digital art extends the structural scope of visual representation, involving "semantic, syntactic and broader aesthetic features that build on – and then exceed – what traditional representation and abstraction can offer" (Crowther 2008). He is also very careful to point out that creativity in digital art goes beyond issues of mathematical or technological competence: the artist must adjust the output of computers to the structure and conditions of the artwork they are creating, keeping in mind the traditional questions on spatial composition and narrative construction. We are back to the issue of choice again: we can rely on computers for the generation of a variety of outputs, but it is still up to the human artist to organize them into an aesthetically meaningful ensemble.

Herein lies one of the most critical issues in digital art: there is a trade-off between the possibilities for aesthetic choices that an artist gains in adopting digital technology, in terms of unprecedented visuals and interactivity, and the freedom that they give up in having to adapt to a complex set of devices that work only thanks to an extremely long series of decisions made by other people. Every aspect of digital technology, from the most basic and physical (e.g. the electric pulses inside a circuit) to the most elaborated and abstract (e.g. the visual interface of a Web browser), are based on agreements between organizations. Makers of computer circuitry and storage devices need to agree on the voltage of the electrical pulses that travel between these devices to ensure interoperability; software houses that create operating systems need to know from computer chip makers what kind of commands, in terms of arrays of 0s and 1s, control those chips; all monitor makers need to adapt to the convention established between numbers and colours, so that when their screens receive orders from a computer they show the expected colours; telecommunication companies need to be authorized by governments to install cables for the Internet and use specific wavelengths and frequencies for their cellular networks. The list could go on and on.

Computers are different from paintbrushes, chisels and violins not only because they are programmable machines rather than tools, but also because those programs are based on a series of conventions that attribute a specific meaning to each physical phenomenon happening inside of them. Those conventions ensure that the different components of a computer work together, and also that different computers can connect and cooperate. The more technological layers are added to computers, in terms of new software, online services, peripherals, the more conventions a digital artist is bringing into their creative process.

All this raises an issue on authorship in digital art that is an evolution of the early criticism. Now computers are accessible to and used by many more people than in the 1960s, and we are no longer talking only about mathematicians and physicists. The users have an enormous computational power at their disposal that they can exploit at their will. Given the very particular nature of the devices they are using, are their choices significant enough to make them full-fledged creators of their artworks, or are they navigating within a restricted operational space determined by all the people and organizations that conceive, design and produce digital technology?

I do not have a definitive answer to this question, but I propose a way to conceptualize digital art that follows the guidelines traced by Duchamp about choice, but in a different way than what Nake did in the early days of computer-based creativity. Back then using computers was a choice radical enough to give the pioneers' act an artistic meaning. Today, artists need new space for their choices, because computers have become mainstream, and so it makes sense to ask why an artist should use digital technology for their endeavour. This question underlies all philosophical enquiries on digital art, and many scholars try to focus on what essentially defines it. I suspect that the extraordinary development that has characterized physics, electronics, and computer science in the last decades has led many to focus too much of their attention on the relevant technological accomplishments. The issue that I have with this drift, especially when it comes to aesthetics and philosophy of art, is that digital technology is very versatile and is employed in an ever-growing number of facets of human society, including stock markets, healthcare, mechanical industry, warfare, which means that employing a computer to make art not only does not seem to give art a special status, but it might actually dilute it into the myriads of human activities that are nowadays computer-based. For instance, while there are only so many uses for material pigments, the colour scheme regulating the monitor a digital artist uses for their visual artworks is the same employed by air traffic control personnel to track planes.

With Duchamp himself and his peers, who disrupted the art world with what we call today "conceptual art" in a broad sense, content, in the form of concepts and ideas, took centre stage at the expenses of medium-based craftsmanship and technique. In Sol LeWitt's words: "the idea or concept is the most important aspect of the work" (LeWitt 1967). I argue that a philosophical discourse on digital art should borrow from this fundamental definition of conceptual art. This may appear to be a contradiction, since digital art is essentially dependent on the microscopic physicality of electronic circuits on the one side, and on the geographically widespread connections of Internet cables on the other, so the creative method, rather than the idea, seems to define digital art. However, given that these material apparatuses and the relevant protocols and software are used mostly for non-artistic purposes, we need to abstract away from them to look for a criterion to isolate digital art or, since we are excluding media from this discourse, simply art. Moving from technique and technology to pure content, we go back to the traditional "what is art?" question.

Following this path, when we are in front of, or immersed in, a computer-based work, I would like to propose to analyze the "is this digital art?" question into two: "is this based on an artistic idea?" and "does this need digital technology to exist?" I do not claim that this separation will be crystal clear in all possible cases: often, if not always, we will have to take into account historical and cultural contexts surrounding the analyzed works. For instance, in Nake's early works the artistic idea and digital technology are tightly related and cannot be considered in isolation. However, the task may be easier for those works that came after that singularity at the beginning of digital art history. In particular, I would like to isolate the technological aspects of machine learning applied to art. The problematic aspect of this kind of endeavour, which is nothing short of extraordinary from a technological point of view, is that the neural networks that churn out image after image and adjust their connections to reach the goals set by their programmers perform their training over enormous databases of existing images taken in some cases from the Web, in other cases from the artworks of one particular artist or style from the past.

Does this need digital technology to exist? Certainly. However, what is the idea behind such an effort? What is the artist trying to express? What is the viewer to think? So far, machine learning has only produced visual artworks that are instances of already known aesthetics: Klingemann has modified technology implemented by Google researchers to create works that happen to be in the style of Francis Bacon; the Amsterdam branch of the advertising company J. Walter Thompson in collaboration with Microsoft trained a neural-network-based system with all the portraits by Rembrandt to generate a new painting in the style of the Dutch master (Nudd 2016). What is the message that we are supposed to take home from these works? Is there one? Perhaps we are to admire the technology that made such results possible. Indeed, the people involved in these projects are always very keen to discuss the techniques they conceived and implemented.

At the risk of sounding like the detractors in the 1960s, I consider this kind of endeavour detrimental for digital art, because it appears to be a self-congratulatory technological exploit that distracts from the real, significant support that digital technology can give contemporary artists.

If art is about ideas, and we are looking for new ideas about this ever-changing world, we should let digital technology be an incredibly sophisticated backdrop to the artists who are able to use it to express themselves more effectively, and not the other way around.

Bibliography

Bense, M., *Computergrafik bei Wendelin Niedlich. Speech at Stuttgart, Germany, November 5, 1965*, available at http://dada.compart-bremen.de/-item/publication/481.

Chen, J.X., *The evolution of computing: AlphaGo*, "Computing in Science and Engineering", n. 18/4 (2016), pp. 4-7.

CIE. *Commission Internationale de l'Éclairage Proceedings, 1931*, Cambridge, Cambridge University Press, 1932.

Cohen, H., *The further exploits of AARON, painter*, "Stanford Humanities Review", n. 4/2 (1995), available at http://www.cs.uml.edu/~fredm/courses/91.548-spr04/-papers/furtherexploits.pdf.

Crowther, P., *Ontology and aesthetics of digital art*, "The Journal of Aesthetics and Art Criticism", n. 66/2 (2008), pp. 161-70.

Dillenburger, B., *Arabesque wall (with Michael Hansmeyer)*, 2014, available at http://benjamin-dillenburger.com/arabesque-wall/.

Duchamp, M., *The creative act. Talk at the convention of the American Federation of Arts*, Houston, Texas, 3-6 April 1957, "Art News", n. 56/4 (1957).

Elliott, L., *AI through the technologist's eye*, "Flash Art", n. 316 (2017), available at https://www.flashartonline.com/article/mario-klingemann/.

Gatys, L.A., Ecker A.S., Bethge M., *A neural algorithm of artistic style* (2015), available at: https://arxiv.org/pdf/1508.06576.pdf.

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y., *Generative adversarial nets*, in *Proceeding of advances in neural information processing systems 27 (NIPS 2014)*, available at https://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf.

Goodman, N., *Languages of art. An approach to a theory of symbols*, Indianapolis, Hackett Publishing, 1976.

Greene, R., *Internet art*, London, Thames and Hudson, 2004.

Haigh, T., *Stored program concept considered harmful: history and historiography*, in P. Bonizzoni, V. Brattka, B. Löwe (eds.), *The nature of computation. Logic, algorithms, applications. CiE 2013*, "Lecture Notes in Computer Science", n. 7921 (2013), pp. 241-51.

Haugeland, J., *Analog and analog*, "Philosophical Topics", n. 12 (1981), pp. 213-26.

IBM. Personal computer announced by IBM (1981), available at https://www-03.ibm.com/ibm/history/exhibits/pc25/pc25_press.html.

Iñárritu, A.G., *Carne Y arena* (2017), available at https://carneyarenadc.com.

Kanarek, Y., *World of awe* (1995), available at http://www.worldofawe.net-/thejournal/landing/.

Koblin, A., *Street view and the wilderness downtown* (2010), available at https://maps.googleblog.com/2010/08/street-view-and-wildernessdowntown.html.

Lewis, D., *Analog and digital*, "Noûs", n. 5/3 (1971), pp. 321-27.

LeWitt, S., *Paragraphs on conceptual art*, "Artforum", n. 5/10 (1967), pp. 79-83.

Lialina, O., *My boyfriend came back from the war* (1996), available at http://www.teleportacia.org/war/.

Marinetti, F.T., *Manifeste du Futurisme*, "Le Figaro", February 20, 1909.

McCulloch, W.S., Pitts, W.H., *A logical calculus of the ideas immanent in nervous activity*, "Bulletin of Mathematical Biophysics", 5 (1943), pp. 115-33.

McIver Lopes, D., *A philosophy of computer art*, London, Routledge, 2010.

Moor, J., *The Dartmouth College artificial intelligence conference. The next fifty years*, "AI Magazine", n. 27/4 (2006), pp. 87-91.

Mordvintsev, A., Olah, C., Tyka, M., *Inceptionism: going deeper into neural networks* (2015), available at https://ai.googleblog.com/2015/06/inceptionism-going-deeper-into-neural.html.

Nake, F., *Construction and intuition: creativity in early computer art*, in McCormack, J., d'Inverno, M. (eds.), *Computers and creativity*, Berlin, Springer, 2012, pp. 61-94.

Nudd, T., *Inside "The next Rembrandt": how JWT got a computer to paint like the old master*, "AdWeek" (2016), available at https://www.adweek.com/-brand-marketing/inside-next-rembrandt-how-jwt-got-computer-paint-old-master-172257/.

Pelaez, E., *The stored-program computer: two conceptions*, in "Social Studies of Science", n. 29/3 (1999), pp. 359-89.

Pouff, *Grocery trip* (2015), available at https://youtu.be/DgPaCWJL7XI.

Powell, T., *Rebel ready-made Marcel Duchamp* (1966), available at https://youtu.be/ZGkFa8lf3A0.

Roberts, L.G., *The evolution of packet switching*, "Proceedings of the IEEE", n. 66/11 (1978), pp. 1307-13.

Rumelhart, D.E., Hinton, G.E., Williams, R.J., *Learning representations by back-propagating errors*, "Nature", n. 323 (1986), pp. 533-6.

Snibbe, S., *Boundary functions*, (1998), available at www.snibbe.com/-projects/interactive/boundaryfunctions/.

Thomson-Jones, K., *The philosophy of digital art*, in Zalta, E.N. (ed.) *The Stanford Encyclopedia of philosophy*, available at https://plato.stanford.edu/-archives/spr2015/entries/digital-art/.

Verostko, R., *Imaging the unseen. A statement on my pursuit as an artist* (2004), available at http://www.verostko.com/archive/statements/statement04.html.

Young, P., *Person to person. The international impact of the telephone*, London, Granta, 1991.