



DRADEK

Studies in Philosophy of Literature, Aesthetics
and New Media Theories

Vol. III Num. 1-2 2017

ISSN 2465-1060
[online]

Remediation:
art, technology and humanity

Edited by Ester Fuoco

powered by

ZETESIS
RESEARCH GROUP

<http://zetesisproject.com/>

Scientific Board:

Prof. Leonardo Amoroso (Università di Pisa), Prof. Christian Benne (University of Copenhagen), Prof. Andrew Benjamin (Monash University, Melbourne), Prof. Fabio Camilletti (Warwick University), Prof. Luca Crescenzi (Università di Trento), Prof. Paul Crowther (NUI Galway), Prof. William Marx (Université Paris Ouest Nanterre), Prof. Alexander Nehamas (Princeton University), Prof. Antonio Prete (Università di Siena), Prof. David Roochnik (Boston University), Prof. Antonietta Sanna (Università di Pisa), Prof. Claus Zittel (Stuttgart Universität).

Executive Board:

Matteo Bensi, Danilo Manca (coordinator), Lorenzo Serini, Valentina Serio, Marta Vero

Review Board:

Alessandra Aloisi, Pia Campeggiani, Ester Fuoco, Annamaria Lossi, Cathrin Nielsen, Francesco Rossi

ODRADEK. Studies in Philosophy of Literature, Aesthetics and New Media Theories.
ISSN 2465-1060 [online]

Edited by Associazione “Zetesis-Progetto di studi e Dialoghi Filosofici”,
via Paoli, 15 - 56126 Pisa. Registered by Agenzia delle Entrate di Pisa, n. 3705, serie III,
23.10.2014




License Creative Commons

Odradek. Studies in Philosophy of Literature, Aesthetics and New Media Theories
di Zetesis is licensed under a Creative Commons attribution, non-commercial 4.0
International.

Further authorization out of this license terms may be available at <http://zetesisproject.com> or writing to: zetesis@unipi.it.

Layout editor: Stella Ammaturo

Volume Editor: Ester Fuoco



DRADEK

Studies in Philosophy of Literature, Aesthetics
and New Media Theories

Vol. III Num. 1-2 2017

ISSN 2465-1060
[online]

Remediation:
art, technology and humanity

Edited by Ester Fuoco

powered by

ZETESIS
RESEARCH GROUP

<http://zetesisproject.com/>

Creativity and the Machine

How Technology Reshapes Language

Fabio Fossa

Abstract

Recent advancements in AI and robotics have provided the means to design, program and assemble machines which appear to be able to produce artworks as different as portraits, songs, poems and movie scripts. Such technologies are known as creative machines. In scientific communications and journal articles the functioning of creative machines is frequently described with words which usually apply to human aesthetic experiences. This may result in blurring the line that separates human and machine creativity. Indeed, it sometimes leads to wonder if human artists are bound to be replaced by creative machines.

Such scenario, however, appears to be an illusion generated by the way we adapt ordinary language to speak effectively of creative machines. By referring to the philosophical debate on machine aesthetics I try and develop a critical standpoint which may help clarify what happens, when words pertaining to the domain of artistic creativity are used to describe the functioning of machines. In sum, I argue that interpreting creative machines as mediums of human artistic experience may prove useful to see through the ordinary use of language and avoid confusion, misconceptions and irrational expectations.

1. *Technology, language and artistic creativity*

The impact of robotics and AI on society can hardly be overestimated and it is bound to become even more significant in the next future. Many sides of human life—for instance, the organisation of industrial work—have already been entirely re-shaped by AI and robotic technologies. What intelligent machines will bring and take in the years to come, nobody can say for sure. What is instead safe to say is that robots are going to enter the social stage and share spaces and activities with us. We are currently standing on the edge of a robotic revolution which is about to massively affect both our everyday practices and worldviews.

The revolutionary potential of robotics and AI has been promptly recognized and explored by many forms of artistic expression, most notably in literature and movies (Kakoudaki 2014). Intelligent machines have actually proven to be a very rich source of artistic inspiration and poetic reflection. Besides, artificial persons have been an object of mythical and poetic imagination long before they became an object of technological science (Wiener 2013, 39-40; Henry 2013). Even the very word ‘robot’ was firstly introduced in a drama—*R.U.R. Rossum’s Universal Robots* by Karel Čapek (2004). The robotic world was firstly and mostly discovered by means of poetic expression¹.

1 As a consequence, sometimes robots are invested with expectations and fears based more on science fiction than hard science. Needless to say, hard science has its own enthusiasts too (who reject any limitations to what technology

Recent technological advancements, however, have now opened a new scenario in the relation between art and intelligent machines. In fact, some interesting results have been achieved in the mesmerizing field of machine *artistic* creativity²: technologies have been designed which paint, compose music, write movie scripts and even poems. Robots and AI have already been objects of poetic expression; now, it seems that they are about to become *subjects* of poetic expression as well. Artistic creativity, which has long been considered a distinctive human trait (Jonas 1962), seems to have been finally welcomed to the lab.

Such astonishing experiments are frequently presented to the public by using the same words we normally use to speak of human artistic creativity. As I will show in more detail (§2), machines are said to create artworks, and these artworks are called paintings, poems, songs. As a result, words which used to indicate forms of human experiences and their objects are now applied to the functioning of machines and its objects.

This process of semantic extension, so to speak, is not only intuitive, but also rhetorically very effecti-

can achieve, thus contributing to blurring the line between imagination and reality) just as some science-fictional works exhibit more scientific awareness than others (thus contributing to putting the debate on the right track). That being said, it is in my opinion safe to claim that, as long as robots and AI are concerned, the influence poetic imagination has been having on the public is dominant compared to that of accurate scientific popularization. It may be the case that many of the fears and worries concerning robots and AI must be traced back more to artistic representations than to scientific data.

- 2 Creativity is a notion that extends beyond the boundaries of artistic expression and, as such, it has been a goal of AI since the very beginning (and still is). On creativity as a traditional goal of computer science, see for example Bringsjord et. al. (2001), Lipson (2007), Northon et al. (2013), Colton & Wiggins (2012). See also Coeckelbergh (2017), p. 289.

ve. Since artistic creativity is a form of experience which belongs utterly to human beings, it appears natural (if not inevitable) to describe machines which imitate artistic creativity through human-related words. To this extent, the word ‘creativity’ has already been reshaped to fit both human beings and machines³—or, more precisely, *creative machines*⁴.

The case of machine creativity is not an isolated one. In fact, a similar linguistic attitude can be observed in many fields of robotic research⁵. The use of biologically-inspired words to describe the functioning of machines is a ubiquitous rhetorical habit. The secret of its success, after all, is not too hard to guess. In a sense, most of the time roboticists and AI scientists reproduce organic behaviours and phenomena by way of inventive imitation⁶. Even if design

3 From now on, unless specified differently, the term ‘creativity’ stands for ‘artistic creativity’.

4 I think that the label ‘creative machines’, which in this case stands for ‘artistic creative machines’, fits the purposes of the present research better than other options such as, for example, ‘artificial artistic agents’. The label *creative machine* immediately puts forth the main issue at stake (that of creativity) without suggesting any definite stance on its features. On the contrary, the expression *artificial artistic agents* seems too strong to me, since it appears to imply that machines display artistic agency just like human artistic agents do. For ‘machine’, as many do, I mean basically a technological product which executes functions without requiring continuous human supervision (Wallach & Allen 2009, Steiner 2017). More on this in Section 3.

5 Another interesting example, I think, is that of *artificial phronesis*. Once it is realized that we need some technological imitation of phronetic judgement in order to build machines which behave in an ethically satisfying way, it is natural to frame the problem in terms of the need for artificial phronesis. However, since the concept of phronesis has traditionally been related to non-formalizable human experiences, it is unclear whether the expressions *artificial phronesis* and *natural phronesis* stand for identical, analogous or entirely different phenomena. On this see, for example, Gerdes (2016) and Ess (2016).

6 By saying this I do not mean to claim that robotics and AI slavishly follow Mother Nature’s instructions, for they evidently do not. Robotics and AI are forms of inventive imitation. As Norbert Wiener (1964, 30-32) brilliantly

is not strictly bio-inspired, the functions carried out by machines have often an organic counterpart or archetype—be it real or idealised⁷. Machines substantially carry out functions, and since organisms are the only entities existing which carry out functions, we tend to apply to machines the vocabulary of living things. Accordingly, we build machines that fly, walk, run, crawl, see, hear, think, learn, do specific jobs and so on.

Robotics and AI, thus, put a huge pressure on ordinary language, which is expected to offer appropriate and comprehensible words to talk about artificially intelligent machines. In this situation, the convenience of the process of semantic extension is almost impossible to resist. Besides, any attempt to linguistically signal the difference between copies and archetypes would be, from a rhetorical perspective, less effective, less intuitive, and ultimately inconsistent with the ordinary use of language⁸. As a result, machines are commonly spoken of through

explains, machines are not designed through a pictorial imitation of a particular entity, but through an operative imitation of a particular behaviour which may be naturally embodied by one or more entities. “Besides pictorial images”, Wiener (1964, 31) writes, “we may have operative images. These operative images, which perform the functions of their original, may or may not bear a pictorial likeness to it. Whether they do or not, they may replace the original in its action, and this is a much deeper similarity”. Once a behaviour is described in a satisfying fashion, technological creativity is free to devise its embodiment as it appears best.

7 For example, military drones look nothing like birds, and yet for convenience’s sake we say they fly and give them names such as Global Hawk, Grey Eagle, Snowgoose.

8 Sometimes in scientific papers authors signal this linguistic need by adding prefixes to words—as the letter ‘R.’ in Veruggio & Abney (2012), 356–357—or by writing words in block letters—as TRUST in Grodzinsky, Miller, & Wolf (2011). Although this may be a viable solution for academic writings, it is not equally functional in the case of ordinary language.

the very same words which describe their archetypes (Coeckelbergh 2017, 286). ‘Intelligence’, ‘agency’, ‘learning,’ and ‘morality’ as machine attributes are all suitable example of how technology is reshaping language.

This does not mean, however, that the process of semantic extension is to be regarded as justified by its own momentum⁹. On the contrary, it may sometimes contribute to blurring the line between the organic and the technological. Using the same words to describe both the archetype and the copy of technological imitation may in fact implicitly suggest that no significant difference separates the two objects. Yet, the evidence that some machines successfully reproduce specific behaviours does not necessarily imply that they do exactly what their archetypes do. Therefore, when machines are described in biological terms, the meaning of words cannot but be stretched.

This may result in two undesired situations. First, the process of semantic extension may engender irrational or misguided expectations towards the technology itself, which would be based more on the ‘organic-like’ vocabulary used than on the actual features of the technology. At the same time, since new phenomena are equated to old ones, we might risk losing the possibility to develop a specific understand-

9 On this point Coeckelbergh (2017, 296) seems to disagree: “if more people were to speak about what machines do in terms of ‘artistic creations’ and ‘works of art’, than would we really have an objective basis for saying that they are wrong? Even if today we might be opposed to the very idea of machine art, in the course of time, our language might change and let the machines in through the backdoor”.

ding of the technology and the particular way it is connected to our experience. Secondly, the revised, techno-related meaning of the words used may feedback on their original meaning, thus reducing their richness and adequateness. This may induce to frame biological and human activities by reference to technological criteria, which are usually easier to measure and control, thus forcing phenomena into conceptual schemes to which they do not belong.

This is to say that the process of semantic extension may implicitly lead to treat machines as organisms and organisms as machines. For the same reasons, when machines are supposed to reproduce typically human experiences, such as artistic creativity, this linguistic process may implicitly support the humanization of machines and the mechanization of human beings.

Developing critical awareness of how words change in light of technological advancement may help see through the opaqueness of ordinary language and exploit the effectiveness of its flexibility without falling prey of its illusions. In order to do so, it is necessary to find a way to measure how words evolve alongside technology and how qualitative differences dilute in the process.

The case of machine creativity is in my opinion a good example of the issue I just sketched. In the next pages I focus on how creative machines are ordinarily spoken of in newspaper articles or in scientific communications addressed to the general public. Particular attention will be paid to how the process of semantic extension blurs the line between human

and machine creativity. I then turn to the philosophical debate on machine aesthetics in search of some guidance which may help see through the linguistic illusions of semantic extension and understand machine creativity by its own specific terms and logic.

2. *Creative machines*

In order to get a general idea of the linguistic process by which creativity has become a machine attribute, let's first consider how creative technologies are presented to the public. What words are called upon, when such machines make their entrance into the realm of ordinary language?

To answer this question, let's see how *The Next Rembrandt* (an artificial painter—an artificial Rembrandt, actually), Shimon (an artificial musician), Benjamin (an artificial movie script writer), and Google's AI artificial poet have been presented to the public both in newspapers and on dedicated websites.

2.1. *The Next Rembrandt*

The Next Rembrandt is a software system sponsored by ING and developed among others¹⁰ by TU Delft and Microsoft. The aim of the project was to

¹⁰ With contributions by Mauritshuis and Museum HET Rembrandthuis. The project was appointed for development by ING to the advertising agency J Walther Thompson, based in Amsterdam.

bring the great Dutch painter back to life to create one more portrait, as it reads on the homepage of the highly informative dedicated website¹¹. Two things already cannot pass unnoticed. First, the technology is rhetorically indicated as a substitute for a human artist perfectly capable of taking his place—which encourages to talk about it using human-related words. Secondly, it is written that the software has the purpose to *create* a new painting; thus, we are nudged to assume that we are reading about a machine which is at least as creative as its archetype was.

A lot can be learned on the components of The Next Rembrandt and the way it works by exploring the many pages of the site. The language used is as accessible as it can be, which requires a massive use of rhetorical and metaphoric expressions. Thanks to that, the reader is successfully talked through the different stages of the complicated design process: gathering the data, determining the subject, generating the features, bringing it to life¹². In few words—which do not do justice to the complexity of the undertaking but suit the informative purpose they pursue—the team explains how a database of Rembrandt’s works was created and specified through the selection of some parameters which led to the identification of a possible new portrait subject and its realization. Previously existent data concerning

11 <https://www.nextrembrandt.com>.

12 Note the choice of words for the last step of the process: “bringing it to life”. The connection between organic life, robotics and AI I referred to in note 6 comes now to surface in a particularly evident fashion. Both the points of the technological imitation—substitution and purposiveness as life-likeness—strongly influence how TheNextRembrandt is presented to the public.

Rembrandt's artistic traits were fed to a software which combined them in unprecedented ways, ultimately generating a new painting (first in 2D, then in 3D) similar to all the portraits considered, whilst being the same as no one. Moreover, Rembrandt's painting style was carefully analysed and replicated, so that the final result would look exactly as a work of the great master. Just as happened in Rembrandt's own experience, one may think, skills and inventiveness intertwine to beget a new masterpiece. The result, needless to say, is breath-taking.

On the official website, claims on what has actually been achieved by The Next Rembrandt and what differentiates the new painting from those of the flesh-and-blood Rembrandt are carefully avoided. However, the video presentation posted both on the website and on YouTube¹³ ends with the ambiguous sentence: "the Next Rembrandt makes you think about where innovation can take us. What's next?". Moreover, the verb 'create' is used in reference both to Rembrandt's paintings and the Next Rembrandt's portrait¹⁴.

Identification between the human artist and his technological copy seems to be rhetorically encouraged. In what sense, however, is the machine *creative*? What is the difference between the old paintings and the new one? Is there a difference at all, if nobody

13 <https://www.youtube.com/watch?v=IuygOYZ1Ngo>

14 One example. In section 3 of the website, *Determining the subject*, the team writes that the purpose of the software systems is "to create new artwork using data from Rembrandt's paintings". Below this, referring to Rembrandt, the authors write: "Then we found the period in which the majority of these paintings were (*sic*) created".

would notice it just by looking at the new artwork? Why should not the machine be as creative as the painter? Although these questions are critical, the process of semantic extension seems to bypass them.

A post on the news section of Microsoft.com faces these doubts in a more direct way¹⁵. The heading is already revealing: “The Next Rembrandt. Blurring the lines between art, technology and emotions”. The significance of the software system is now also unveiled:

“Blurring the boundaries between art and technology, The Next Rembrandt project is intended to fuel the conversation about the relationship between art and algorithms, between data and human design and between technology and emotion”.

Moreover, with a very interesting choice of words, the painting is defined as “a visualization of data in a beautifully creative form”. Acknowledging that thanks to The Next Rembrandt “an unprecedented opportunity to do new and great things creatively has been revealed, challenging the way we think about art, creativity and emotion”, not only the issue is stated in its exact terms, but its main questions are also spelled out: “Is this a piece of art? Can a programmer be considered an artist? Can this approach be applied to other artists, to music?”.

Given the way by which The Next Rembrandt has been presented, the temptation to frame the issue it poses as a matter of competition, substitution, and

¹⁵ <https://news.microsoft.com/europe/features/next-rembrandt/>

eventually human obsolescence can hardly be resisted. As Tim Nudd writes on Adweek¹⁶, “we have a harder time, especially those of us in the creative industries, entertaining the question of whether machines could ever be as creative as humans. Creativity is supposed to be our exclusive province, the spark that makes us special, the thing computers could never dream of mastering”.

As reported on the Windows post, art historian Gary Schwartz is convinced that nobody would ever think that Rembrandt “can be reduced to an algorithm”. In an article on the Guardian, Mark Brown reports that Bas Korsten¹⁷, who proposed the project in the first place, believes that “only Rembrandt could create a Rembrandt” and that “we are creating something new from his works”¹⁸. Tim Nudd notes that “The Next Rembrandt is a long way from a computer exhibiting true creativity”¹⁹. Still, the reasons supporting such beliefs are left to the reader to come up with. Even if the words are the same, in what sense The Next Rembrandt is creative and why it is not as creative as Rembrandt himself are questions yet to be answered.

16 Nudd (2016).

17 Bas Korsten is advertising executive of the agency J Walther Thompson, which developed the project for ING.

18 Brown (2016).

19 Nudd (2016).

2.2. *Shimon*

Let's talk music. Shimon is a four-armed, one-rhythmically-shaking-headed robotic marimba player which went on tour with its backup human band from May 2015 to July 2017. According to the dedicated website, Shimon “can listen to, understand, collaborate with, and surprise his (*sic*) human counterparts”. The robot was developed by the Georgia Tech’s Robotic Musicianship group at the Georgia Tech Center for Music Technology, and it implements “artificial intelligence and creativity algorithms”²⁰. On stage, Shimon plays the marimba in real time with its friend musicians and programmers, performing mixed human-robot jam sessions which are fascinating both to hear and to see²¹.

The activity of Shimon the Robot has been commonly framed in terms of creativity. In the website video section, a TEDEd Lesson written by Gil Weinberg²², Director of the Georgia Tech Center, illustrates in plain language and cool animations how music, artificial intelligence, genetic algorithms, and creativity overlap. The title of the lesson is “Can robots be creative?”. No answer is offered at the end of the video, but a bunch of bewildering questions are raised on what creativity is per se, what human creativity is, and what happens, when robots appear to act creatively. Eventually, it is also asked why it

²⁰ <https://www.shimonrobot.com>.

²¹ If you cannot attend to Shimon’s band live performances, videos are available online at <https://www.shimonrobot.com/video>.

²² <https://ed.ted.com/lessons/can-robots-be-creative-gil-weinberg>.

matters who or what creates a piece of art if it successfully triggers aesthetic experiences. Even though the questions are still open, the word “creativity” is used to describe the technologies involved without much hesitation.

Gil Weinberg is not only a programmer, but also a drummer. He has no intention to put his drumsticks away and leave the stage to robotic music performers. The whole point of the project is to enhance human creativity by endowing it with robotic partners which may stimulate the human experience of playing music in unforeseeable ways. However, the fact that Shimon can play along and improvise music cannot but make people wonder about what differentiates its musical creativity from that of human performers.

Furthermore, as Matt Burgess in an article for *Wired*²³ and Evan Ackerman in a post on *IEEE Spectrum*²⁴ explain, thanks to deep learning neural networks last summer (Deep) Shimon was able to compose two 30-second-long new songs²⁵ “all by itself”, blending classical music and jazz tunes in a “delicately soothing” vibe. The verb both Burgess and Ackerman use, in order to describe what Shimon does, is ‘to create’²⁶.

²³ Burgess (2017a).

²⁴ Ackerman (2017).

²⁵ You can listen to Deep Shimon’s songs on Youtube: <https://www.youtube.com/watch?v=j82nYLOnKtM>, <https://www.youtube.com/watch?v=j82nYLOnKtM>.

²⁶ Burgess (2017a): “it was able to use deep learning techniques to create two 30 second pieces of original music.” Ackerman (2017): “Now, Shimon has leveraged deep learning to create structured and coherent and totally unique compositions of its very own.”

The stages Shimon went through in order to compose original music are analogous—by and large at least—to those The Next Rembrandt went through: data collection, identification of patterns, insertion of parameters, composition and, finally, performance. Shimon does not just play along anymore: it makes its own music.

The conceptual similarity between the two creative machines is striking, and they put a strikingly similar pressure on language as well. By the same token, the scope and meaning of machine creativity remains blurred as compared to the single form of creativity normally available, that of human beings.

2.3. Benjamin

Benjamin, an AI developed by Ross Goodwin (AI researcher at NYU) and Oscar Sharp (director), writes extremely odd and funny movie scripts which defy even the wildest imagination. To be true, its original name was Jetson, until it suddenly introduced itself as Benjamin. Some short, completely absurd scripts Benjamin conceived are available on its Tumblr²⁷ page, which unfortunately offers no information about the science involved or the project in general.

Despite its eccentric style, Benjamin has managed to have two of its screenplays actually shot, both directed by Oscar Sharp. The first movie based on

²⁷ <http://benjamin-ai.tumblr.com>.

Benjamin's work is *Sunspring*²⁸, a sci-fi short film with Thomas Middleditch as main actor which was released during summer 2016. The second short film, *It's no game*²⁹, was presented in spring 2017 and, more or less, depicts the breakdown of its main character, masterfully played by David Hasselhoof. I better not dare to give a resume of what happens in these films: apparently, AI scriptwriting does not easily lend itself to synopsis.

In *It's no game* the part actually written by Benjamin is forerun by a satirical introduction³⁰. Two representatives of the Hollywood writer's league, which is about to call a general strike, are faced with Benjamin's work and the supposed threat it poses to their jobs. Of course, they complain about the fact that nothing makes sense in Benjamin's writing. As a result, however, they are momentarily transformed in actors reciting nonsensical lines from Benjamin through nanotech added in secret to their mugs of tea. After they regain consciousness and flee, the intro ends in a dystopian and, indeed, funny film industry delirium, followed by a quite disturbing dance sequence the moves and figures of which were selected and joined together by another algorithm. Last, but not least at all, Hasselhoff impersonates a man—or a cyborg?—on the edge of a breakdown and recites a heart-breaking dialogue entirely written by Benjamin.

Although Benjamin's writing is not good enough to make any sense, the question is put forth, even if

28 <https://www.youtube.com/watch?v=LY7x2lhqjmc>

29 https://www.youtube.com/watch?v=5qPgG98_CQ8.

30 On this, see Newitz (2017).

ironically: can a machine be as creative as a human scriptwriter? For the time being, the answer is clear: absolutely not. Benjamin's creators are not ashamed to compare their pupil's work to "monkey writing". But the technology is evolving, and that 'no' may soon become a 'not yet'.

Theoretically, Benjamin is not very different from *The Next Rembrandt* and *Shimon*. As explained in few screenshots before *Sunspring* begins, in order to produce the script a long short-term memory recurrent neural network was fed with science fiction screenplays from the 80s and the 90s. Then, a set of sci-fi prompts³¹ was also fed to the software: a title, a line of dialogue, some scenic hints and a general theme. The AI produced the screenplay moving from these starting points. Benjamin was also fed with thousands of pop songs, so that it would be able to output a song lyric added as soundtrack to the short film.

Annalee Newitz of arstechnica.com (which hosted both Benjamin's online debut and its second oeuvre) wrote two accurate articles on Benjamin, presenting it to the general public and discussing its main features and possible impacts with Sharp, Goodwin, and Hasselhoof. In their opinion, Benjamin's most interesting feature is the possibility it offers to address the study of screenplay by means of a new tool which is able to detect recurrent patterns, elements of style, and even the viewers' tritest expectations,

³¹ As explained by Newitz (2016), the prompts came from Sci-fi London, a movie contest that includes a challenge by which participants are given a set of prompts and 48 hours to shot a short film in which such prompts have to appear.

thus enhancing human scriptwriters' creativity and knowledge.

This, however, does not mean that the status of Benjamin as a creative machine did not have a confusing effect on the people who worked with it. On the one hand, as Newitz notes, it is hard not to anthropomorphize Benjamin, since it is impossible to describe what it does without using human-related words; and from language to reality it is but a short step. Oscar Sharp, Newitz reports, felt responsible for Benjamin's work, tried not to betray the script and grew to see the AI more as a co-author, whose choices he felt an obligation to respect, than just a machine void of intentions and meaning. On the other hand, he would not address Benjamin as a true colleague since, in his opinion, it lacks authenticity³². As Goodwin points out, even if Benjamin is a tool, the word 'tool' is not entirely fit for representing it: "we need a new word for it". Newitz ends: "Benjamin exists somewhere in between author and tool, writer and regurgitator". Up to a certain degree, Benjamin is an author, and there is no authorship without creativity. But what does "creativity" means here?

32 This is a pivotal point. In Newitz's (2016) words: "The answer is complicated, because the filmmakers felt as if Benjamin was a co-author, but also not really an author at the same time. Partly this boiled down to a question of authenticity. An author, they reasoned, has to be able to create something that's some kind of original contribution, in their own voice, even if it might be cliché. But Benjamin only creates screenplays based on what other people have written, so by definition it's not really authentic to his voice—it's just a pure reflection of what other people have said".

2.4. *Google's AI Poetry*

Last, but not least, the artificial poet. As sometimes happens with people as well, nobody suspected that the extension of the standard recurrent neural network language model (RNNLM), at which a team of scientists from Google, Stanford, and UMass³³ were working at, cherished poetic ambitions. In fact, it does not, but its outputs may very well be (mis) taken for poems—good or bad, judge for yourselves. The scientists' intention was to develop a neural network able to connect two previously fed sentences through a list of plausible phrases. In order for the machine to carry out the task, it had to have some grasp of both the general topic underlying the two initial sentences and plausible chances of connection between them. To achieve this aim, the neural network was trained on romance novels³⁴. Here is one of the results (the first and last sentences were fed to the machine, the sentences in between were automatically generated):

he was silent for a long moment.
 he was silent for a moment.
 it was quiet for a moment.
 it was dark and cold.
 there was a pause.
 it was my turn.³⁵

33 Samuel R. Bowman (Stanford), Luke Vilnis (UMass), Oriol Vinyals, Andrew M. Dai, Rafal Jozefowicz & Sami Bengio (Google). See Bowman et al. (2016).

34 Bowman et al. (2016), p. 11. The Google's team which authored the paper is applying the same strategy to multiple Google's services. On this, see Kantrowitz (2016).

35 Bowman et al. (2016), p. 8.

The scientists presented their results in a paper, in which they steer carefully clear of any reference to poetry or creativity. The poetic resemblance of such lines was later noticed by Samuel Gibbs, a journalist at *The Guardian*, who missed no opportunity to highlight the AI's lack of poetic talent³⁶. However, when he had to choose a word to describe what the machine did, the verb 'create' was used—as the subheading reads³⁷. The poems, as dull as they are, are framed as machine creations nonetheless. Matt Burgess of *Wired* did not share Gibb's point of view and found the AI's poems "amazingly mournful"³⁸. Despite the difference in matter of taste, there is agreement on the choice of words to describe the functioning of the neural network: "It was then up to the AI to create its own poetry based on the available information"³⁹.

Due perhaps to the inner strangeness of poetic expression, it is not always easy to say whether a poem was written by a human being or a machine. On Benjamin Laird and Oscar Schwartz's website *Bot or Not*⁴⁰ everyone can test her ability in distinguishing human-written poems from machine-generated sentences by taking the Turing Test for poetry. Laird and Schwartz define computer-generated poetry as "text, that is generated through an algorithm, which is executed by a digital, electronic computer,

36 Gibbs (2016).

37 See Gibbs (2016): "Inspired by thousands of dramatic novels, technique creates verse that rivals that of Douglas Adams' Vogons".

38 Burgess (2017b).

39 Burgess (2017b).

40 <http://botpoet.com/about/>

which is intended, by whoever it may be, to be read as poetry”⁴¹; their research is aimed at clarifying “what constitutes a human poem, and what constitutes a computer poem”⁴². After a couple of mistakes in ascribing poems to computers or human beings, the question as to what differentiates human artists from creative machines cannot but arise. As Laird and Schwartz’s definition of computer-generated poetry suggests, the point of view of the observer must not be lost. That being said, it seems that some form of creativity must still be acknowledged to the machine.

2.5. Machine creativity and ordinary language

As clearly emerges from the last sections, human-related words pertaining to the domain of creativity are commonly used to describe how creative machines work. The rhetorical potential of this semantic extension is great, since it allows describing what such machines do in a non-technical language which many can understand. Nonetheless, in a sense, the process runs too fast and notions struggle to keep up. A gap opens between the inclusive evolution of language and its semantic power. As a consequence, it becomes hard to grasp what the same word means when applied to entities as different as human beings and machines. In our specific case, it becomes

41 <http://botpoet.com/what-is-computer-poetry/>

42 <http://botpoet.com/about/>

hard to realize what ‘creativity’ means as applied to machines.

It seems, then, that the flexibility of language generates confusion if not kept under control. By the same token it may be said that the flexibility of language is exactly what fuels the evolution of our notions and concepts. Either way, the process of semantic extension must be carefully criticized; otherwise, words will have the effect of diluting differences, thus losing expressive power and ultimately leading to deception.

The analysis of ordinary language sufficiently indicates the issue at stake. The semantic extension of the word ‘creativity’ to technological products seems to imply that human creativity and machine creativity are one and the same thing or, at least, two points along a continuum. Language is the mirror looking into which human beings and machines reflect one another. However, many intuitively claim that the mirror is enchanted and the reflection is illusory: The Next Rembrandt is not Rembrandt himself, Benjamin is not entirely an author, Google AI’s poet is not actually a poet, and so on. To back up such claims, either intuitions are introduced or concepts such as ‘authenticity’ or ‘true creativity’ are evoked.

The differences covered by the process of semantic extension cannot but come to surface in the form of the already mentioned questions. All these doubts point at the very gap ordinary language jump across but cannot bridge. Modern machines blur the line between selves and tools. The linguistic attribution of creativity to AI and robotic technologies intro-

duce them to a domain of meaning which used to be exclusively human, thus generating the impression that machines should be interpreted more as selves than tools. In this sense, ordinary language must be counted among the many factors which lead people to anthropomorphize machines. Nevertheless, the intuition remains strong that creative machines, being technological products, are to be conceived as tools and not as agents themselves. However, how can we make sense of this form of tool-like creativity?

Tools execute functions the ends of which are already set: this is the main principle of instrumental behaviour (Johnson 2011). ‘True creativity’, on the other hand, seems to require ‘authenticity’, spontaneity, something beyond a mere ‘visualisation of data’. This is the content of the main intuition against the identification of creative machines and human artists. And yet, such visualisation is described as creative and such machines are described as creative as well. This is the dead-end in which ordinary language seems to lead: the mutual relations between data elaboration and creativity, tools and selfhood, instrumental and spontaneous behaviour appear to be no longer intelligible.

3. Machine aesthetics

As the previous section shows, we already talk and write about some technological products in terms of creative machines. What still remains unclear is

whether the word ‘creativity’ applies indiscriminately to human beings and machines or not. At first sight, it appears not: many maintain that machines are not and cannot be as creative as human artists. Machine creativity seems to be of a different kind. However, no detailed answer is usually offered to the question as to what differentiates the two forms of creativity. As a result, the meaning of the word becomes fuzzy, may lead to anthropomorphism and ultimately induce irrational expectations towards both machines and human artists. Hence, it is necessary to clarify what ‘creativity’ means when it is attributed to technological products. In this section I try to shed light on what is distinctive of machine creativity by reviewing some of the studies which deals with the many issues creative machines raise from the point of view of philosophical aesthetics (Gunzel 2017; Coeckelbergh 2017, p. 296).

In order to analyse machine creativity, a good place to start consists in focusing on the entities which are actually under scrutiny. First and foremost, machines are technical artefacts, i.e., “purposefully created or modified physical structures that serve a practical function”, where “a function is the *intended* effect of a technical artifact” (Steiner 2016, §3). Technological products are essentially tools human beings conceive and use in order to attain ends more easily, more efficiently, less dangerously and so on (Bryson & Kime 2011). However, modern machines are not simple tools. Whilst traditional tools continuously require operators’ intervention, machines carry out functions without requiring constant hu-

man supervision. Machines autonomously mediate the achievement of given ends.

Although machines execute functions autonomously, the purpose of their functions is external to them and set by human beings. No machine can be built without an end in mind. This is why, from a conceptual perspective, machines cannot be adequately understood apart from the purpose they are built to serve and the function they are intended to carry out. Since machines autonomously achieve given ends, they can be properly understood only if inscribed in the human intentional context to which they belong (Johnson 2011). So, creative machines can be understood only by reference to the intentional context in which they are conceived and deployed (Steiner 2016, §5).

At first sight, this seems to exclude machines from the realm of creativity. Creativity, in fact, is often interpreted as the autonomous expression of inner states, emotions and beliefs artists infuse in their works for the audience to experience (Coeckelbergh 2017, 289-290). The same spontaneity is a key element to the development of a distinctive style (Crowther 2015, pp. 18-24). Creativity, thus, appears to be inseparable from the unique expression of ideas, feelings, emotions, meaning, and inner states which pertains exclusively to the artist's sensitivity and selfhood (Crowther 2015, pp. 9-12).

According to this common viewpoint, creativity is substantially a form of self-expression. Since machines have neither inner states to express nor the possibility to grasp what is fundamental in the expe-

rience of life, very little room is left for attributing creativity directly to them. Only selves can be creative, tools cannot.

The claim according to which machines, being “purpose-built artifacts” (Bryson & Kime 2011), cannot be as creative as human artists are, appears to be solid. Still, this is not enough to settle the question on machine creativity. In fact, the claim not only shows the negative side of the issue, but also points at the positive conditions by which machine creativity can be explored. Machine creativity cannot be exclusively approached by way of an abstract comparison with human creativity. On the contrary, its peculiar meaning must be searched starting from its intentional context. Machine creativity must *also* be understood by its own terms and logic (Coeckelbergh 2017, p. 297).

The very purpose by which creative machines are designed and built is exactly the reproduction of artistic creative experience by means of technology (Steiner 2016)⁴³. In other words, creative machines are always embedded in artistic contexts in which ends and aesthetic values are implicitly or explicitly set. Now, what characterises creative machines is that they carry out their artistic function auto-

43 This is the main claim in Steiner (2016). Moving from this thesis, the author contends that the products of creative machines are artworks, since they are intended to be so. Since creative machines are designed and built to output artworks, if they work correctly they must be acknowledged to do so—even if, of course, this says nothing on the quality of the artworks produced. Although his viewpoint is very interesting and, in a sense, compelling, it does not say much about what machine creativity is and how it differs from human creativity. We need to ask further: since creative machines produce what are intended to be works of art, which is their function, are they creative as we are? And if they are not, in what sense they are, and why we call them ‘creative’?

mously, i.e., without requiring human intervention. This peculiar process of mediation distinguishes creative machines from all other artistic tools. So, even if “the non-human creator is created by human creators, the work created by the non-human agent is not directly created by the humans”; in a sense, then, “the algorithm, not the human, is the ‘artistic’ agent”⁴⁴ (Coeckelbergh 2017, p. 286).

The functional autonomy by which machines mediate creative experiences is an unprecedented phenomenon which adds a new and confusing degree of mediacy between human creators and artworks. The capacity displayed by creative machines of carrying out artistic functions autonomously is exactly what blurs the line between artists and tools, thus indicating the need of new concepts and words⁴⁵. Functional autonomy, then, is key to address machine creativity: it allows establishing a new domain of creativity, proper to machines, which flows directly from what machines actually do without losing track of the intentional context to which such machines belong.

According to this approach, then, machines cannot be taken as mere substitutes for human artists, since the way by which they function is both essentially different and partially dependent from human artistic experience. As functional artefacts, machines

44 In my opinion, the attribution of agency to machines is too strong a stance—and a not necessary one, after all. In this I agree with Steiner (2016). *See infra*.

45 For example, functional autonomy is the ground on which several scholars brought forth the issue concerning intellectual property in the case of machine-generated content. On this see for example Davis (2011), Bridy (2012), Guadamuz (2017).

are *mediums*: they autonomously mediate the expression of given artistic ends and values. So, machine creativity must be understood by reference to a relational framework which concentrates on how machines and human beings interact in the experience of creating art. Machines and human beings must be framed as partners in the co-creation of artworks (Coeckelbergh 2017, p. 297) or partakers in the creative “human-machine system” (Sandry 2016).

By assuming a cooperative stance, it is possible to develop an analysis of the hybrid experience of creating art with machines which acknowledges the role that the “incommensurable differences between humans and machines” plays “in the creative collaboration” (Sandry 2016, §1). It must be noticed, however, that such differences—as incommensurable as they might be—are not at variance with the notion that creative machines mimic or imitate human creative experiences. However deep the ‘otherness’ of machines may be, they can only be understood against a human background. If the difference between human and machine creativity were absolute, we would not be talking about machine creativity at all; we would have named it differently. In reference to human beings, machines exhibit difference in similarity. In order to clarify in what sense creativity applies to machines, this difference in similarity needs to be addressed.

Some hints in this direction can be found in Crowther (2015). The author assumes the perspective I have presented in the last paragraphs and sums it up in what he calls the Cohen’s Principle:

Unless computers acquire selfhood then they will always follow a different order of creativity from that of humans, but the difference at issue here can actually be used to extend the scope of human creativity itself⁴⁶.

In order to clarify what machine creativity is, then, it is better to start from what differentiates it from human creativity. This difference provides the key to understanding why machines are said to be creative:

Human selfhood and creativity is clearly of an entirely different qualitative order from that of machines—however sophisticated they might be. Can't we just leave it at that? The answer is no. For if we fail to explain what these differences centre on, then we fail to do justice to the other half of the principle—namely that it is the very difference between computer and human creativity that allows the former to assist in developing the latter⁴⁷.

So, in order to properly understand machine creativity, it is necessary to frame creative machines neither as agents, artists, selves nor as simple tools, but as mediums: entities which autonomously execute artistic functions and, thus, are always embedded in (human) intentional contexts. Creative machines mediate human creativity through partially independent creative behaviour. As such, creative machines offer new and unprecedented means of artistic expression. This is the ground on and the extent to

46 Crowther (2015), p. 9.

47 Crowther (2015), p. 12.

which it makes sense to talk about creative machines.

It seems, then, that the proper domain of machine creativity is the autonomous execution of artistic functions. Hence, the specific form of creativity machines display may be named *functional creativity*. Functional creativity appears to be the form of creativity proper to artistic machines as mediums.

The 'autonomous' way by which creative machines execute their functions adds a unique element to the picture, which partly lies beyond human control. This is why, although the general context is known, machines can still surprise and yield 'unexpected' results—which is perhaps the most effective psychological evidence of machine creativity. Sure enough, artistic surprise does not pertain to the context in which the machine is embedded. If the output fails to be in line with the expectations set by context, then the machine is not working appropriately and simply need to be fixed. A machine, which does not execute the task it is programmed to carry out, is not creating but malfunctioning. When it comes to the elements of context, machines cannot artistically surprise us, but only betray our technological expectations. Artistic surprise, on the other hand, originates from the ability the machine displays to elaborate data in unpredictable and strange ways which, therefore, yield artistically unexpected outcomes.

Much more, however, remains to be said on what precisely constitutes machine creativity and whether expressions as 'creative elaboration of data' or 'creative execution of functions' are not oxymorons to begin with. In what sense the process creative ma-

chines carry out is creative? Does it add something new to the data they have collected or been fed with? Or perhaps is there a meaning of creativity which does not stress the addition of something new, but is related to the original disposition and combination of contents? Is the autonomous execution of artistic functions enough to justify the recourse to the word 'creativity'? How are the relations between programmers, artists, creative machines and their outputs to be intended? A full-blown theory of machine creativity must address these and other questions. The notion of creative machines as mediums, perhaps, may prove to be a good starting point in this direction.

4. *Conclusion*

The concept of functional creativity offers the means to a critical discussion of the process of semantic extension active in ordinary language. Not only this framework shows how irrational it is to fear that machines will take the place of artists; it also offers a general standpoint from which to determine the meaning of the words we use to talk about machine creativity. In fact, the notion of creative machines as technological mediums of human artistic experience allows thinking both the connections and the differences between human and machine creativity. The analogy between the two forms of creativity, which supports the process of semantic extension, is not denied; however, it is contextualized and de-

terminated. Offering an analytic method to untangle the multiple meanings of terms which are used in metaphorical or analogical sense, this conceptual approach provides a way to adequately think and speak of the grey area between authors and tools where creative machines find their place.

There is no point in demanding ordinary language to stick rigidly to the result of conceptual analysis. It would be a mistake not to acknowledge the great communicative potential the process of semantic extension has in bridging technological research and shared social knowledge. However, it is also important to develop the means to clarify the meaning of the words we use, when needed. If this task is neglected, the specific qualities which distinguish one form of creativity from the other get lost, thus generating confusion, impoverishing the expressive power of our words and ultimately leading to misconceptions and irrational expectations. This is why, in my opinion, the re-shaping of language in light of technology must not only be acknowledged, but also critically assessed.