RIFL (2020) Vol. 14, n. 2: 35-46
DOI: 10.4396/AISB201904

Learning through creativity: how creativity can help machine learning achieving deeper understanding

Caterina Moruzzi
Universität Konstanz
caterina.moruzzi@uni-konstanz.de

Abstract In this paper, I address the difficult task of analysing the nature of creativity by suggesting a more objective way of defining it. In particular, I propose a minimal account of creativity as autonomous problem-solving process. This definition is aimed at providing a baseline that researchers working in different fields can agree on and that can then be refined on a case by case basis. Developing our insight on the nature of creativity is increasingly necessary in the light of recent developments in the field of Artificial Intelligence. In the second part of the paper, I discuss how an investigation on the main features of human creativity can support the advancement of machine learning models in their current areas of weakness, such as intuition, originality, innovation, and flexibility. I suggest how methods such as modelling the human brain or simulation can be useful to extract the main mechanisms underlying creative processes and to translate them to machine learning applications. This can eventually aid both the development of machine learning systems that achieve a deeper and more intuitive understanding and our exploration of human creativity.

Keywords: creativity, artificial intelligence, autonomy, problem-solving, machine learning

Received 20/01/2020; accepted 10/03/2020.

0. Introduction

The notion of creativity has been a subject of discussion in many research fields for centuries. Indeed, the diversity of the debates concerning creativity is one of the reasons why they usually end leaving participants even more confused on what the nature of creativity is. This is partly due to the fact that creativity is defined differently by people working in different fields. Artists normally deem creativity as a property of *genius* that only some people seem to possess, at least at its more distinguished level (Albert, Runco 1999). Computer scientists, for their part, consider creativity as something that can be recreated through an algorithm and they usually place more relevance on the final product rather than on the process that leads to it (Bridy 2012: 1-28; Dartnall 1994). Cognitive scientists and neuroscientists look for neural correlates of creativity in the brain and try to understand the mechanisms underlying it (Goldman 1993; Thagard 2006). Lastly, philosophers conduct a meta-study of what researchers in the other fields say about creativity. As a consequence, they describe creativity in a variety of ways: as a property of a process, a property of a product, an emergent entity, etc. (Paul, Kaufman 2014).

One of the main aims of this paper is to contribute to this debate by proposing a definition of creativity that people working in different fields of research can agree on. In section 1 I am, thus, going to discuss some of the major difficulties that scholars encounter when facing the necessity of finding an objective definition of creativity. As a solution to this conundrum, I will suggest a minimal account of creativity which defines it as an autonomous problem-solving process. In the second part of the paper, I discuss the relevance of investigating the nature of creativity in the light of recent technological developments in the field of Artificial Intelligence (AI). In particular, I will reflect on how an enhanced understanding of human creativity can aid the development of more efficient and flexible machine learning systems. In section 2, I am going to describe one of the main struggles that current machine learning systems experience: the difficulty in performing generalisations from unknown distributions. In section 3, I will argue that the analysis of the concept of creativity can be crucial in aiding this technological development and I will suggest some of the methods that can be used for achieving this. I conclude the paper by pointing at some of the difficulties that we may encounter when embarking in this quest and by presenting the main benefits that emerge from this research.

1. What is creativity?

Numerous definitions of creativity have been proposed in the literature on the topic. The definition that has set the baseline for many of the ones that would have followed is the one proposed by Margaret Boden: «Creativity is the ability to come up with ideas or artefacts that are new, surprising and valuable» (Boden 1990: 1). For Boden, a creative activity produces something which has never been done before. Also, the novelty which is brought about by a creative action should necessarily be accompanied by a feeling of surprise. Finally, this novel and surprising idea or artefact needs to be *valuable*. It should, that is, have a recognised quality which makes us value it for some reason. The attribute of *valuable* is a very important addition to the other two attributes of *novel* and *surprising*. Indeed, in this paper I might write a few sentences which have never been written before and that, for some people, result as surprising. Yet, for other people they may not count as valuable and, thus, neither as creative. There are other available definitions of creativity but they generally agree on the features that Boden identifies (Colton 2008; Gaut 2010: 1034-1046; Wiggins 2005: 209-222).

Boden's analysis on the theme of creativity is widely valued for having brought the topic to the attention of philosophical analysis. However, the definition she gives is not without flaws. A controversial aspect of Boden's definition is the fact that it is too strong. Indeed, Boden's inclusion of the attribute *valuable* among the requirements for creative acts has the unwelcome result that, for example, a song which we do not like nor find valuable is not a product of creativity[1].

Nevertheless, Boden's description of an act of creativity as producing something *valuable* highlights a characteristic which our common sense essentially attributes to the notion of creativity: its being *subject-dependent*. It is this subjective nature of creativity that leads to the difficulty in finding an objective definition of it. Indeed, creativity is not a univocal and determined concept but can instead be identified with the property of a

[1] See Wiggins 2006: 449-458 for a formalisation of Boden's definition of creativity that aims at solving some of its inconsistencies.

RIFL (2020) Vol. 14, n. 2: 35-46
DOI: 10.4396/AISB201904

product or of a process (Wiggins 2005), as a single or communal activity (Still, D'Inverno 2016) or, ultimately, as something that cannot be defined and that, probably, does not even exist (Minsky 1982).

From this quick overview, it is clear that there is no consensus on what creativity is. The first step we need to take in the investigation on the nature of creativity is, thus, trying to shed light on the issue. In this paper I do not have the presumption of providing a definite answer on what creativity is, nor to indicate its neural correlates or the way in which to replicate it on an artificial substratum. Rather, my aim is to pave the way for one interpretation of creativity which may hopefully lead to interesting results if pursued further and, moreover, which can be an objective baseline that the majority of people researching creativity can agree on.

We, arguably, share the intuition that creative processes present something special in respect to other, more mundane, activities. This special feature of creativity has been interpreted in different ways, as value, originality, intentionality, etc. (Ariza 2009). I am not in the position of stipulating what these special features are. Instead, I aim to provide a minimal account of creative processes which is intended as a starting point, a way to determine the primary nature of creativity in order to proceed in the investigation and refine our picture of creativity as we gain more insights on it.

I suggest that creativity can be minimally defined as an autonomous process of problem-solving. The characterisation of creativity as a problem-solving process emphasises how creativity is essential to many areas of application, not only the arts (Colin *et al.* 2016; Kralik *et al.* 2016; Simon 1985). Problem-solving abilities subsume a certain degree of creativity in the exploration and analysis of the different routes that can be taken in finding a solution to the problem we are confronted with. Processes of scientific and technological discovery certainly require this aspect as much as the creation of an artistic artefact does.

The choice of focusing on the creative process instead of the product is motivated by the need to try and achieve a view on creativity which is as objective as possible. Addressing our attention to the product that originates from a creative process, in fact, does not provide us with any insight regarding how creative thought originates or what it entails (Van Der Schyff *et al.* 2018: 1-18). The consideration of a product as creative is not objectively measurable but, rather, it depends on how we perceive it and on a series of contextual factors[2].

For example, we generally consider the notorious *Fountain* by Marcel Duchamp creative because of the context in which it is displayed and of the process of creation that lies behind it. We would not consider it as creative if, instead of an art gallery, it was installed in a public toilet. It is, thus, Duchamp's thinking process that is responsible for our evaluation of *Fountain* as creative, not the intrinsic properties of the product itself. The consideration of *Fountain* as a product is dependent on that and it is relative to the context in which it is perceived.

I believe that the exploration of the mechanisms underlying a creative process may instead be more liable to an objective investigation. If we individuate the features that characterise the creative process underpinning the production of an outcome, in fact, we may be able to justify the essential properties of it and of generalising them in order to categorise processes as creative or non-creative.

I defined the process of creativity as *autonomous*. Autonomy is widely recognised as an essential property of creativity (Briot, Hadjeres, Pachet 2012; Luck, D'Inverno 2012:

[2] See Dewey 1934; Elton 1995: 207–213. The subject-dependence of the notion of creativity as applied to a product is also the reason why the Turing Test is deemed by many a non-reliable measure of creativity. See Ariza 2009: 48-70; Briot, Hadjeres, Pachet 2017; Wiggins, Pearce, Müllensiefen 2011.

RIFL (2020) Vol. 14, n. 2: 35-46
DOI: 10.4396/AISB201904

332-346). However, autonomy can be variously defined as the generation of inner goals (Luck, D'Inverno 2012), the ability to respond to known and unknown inputs (Bown 2011: 73-85), or the freedom to generate ideas (Dong *et al.* 2017). A possible concern is that autonomy is too strong a requirement for a process to be creative. Also humans, in fact, are not completely autonomous but instead constrained by their body, social context, the tools they use, etc. In this paper, I intend autonomy not as complete freedom but as not being completely reliant on a given set of data (D'Inverno, McCormack 2012). Once the learning stage is complete, the process of creation needs to lead to directions which depart from the training set. This is the kind of autonomy I refer to when analysing the possibility for natural and artificial agents to be creative.

I already mentioned how, with the rapid development of AI systems and the diffusion of their fields of application, there is the urgency of addressing questions regarding how this new kind of intelligence differs from human intelligence, including questions regarding its creative potential. In the following section, I will present one of the principal areas where machine learning systems need to improve: achieving flexible and intuitive understanding. In this respect, I will argue that the investigation of the mechanisms underlying creativity can help us develop machine learning systems that overcome this obstacle[3].

2. Machine learning and understanding

Machine learning is an application of AI that allows artificial systems to learn how to perform some tasks without being explicitly programmed (Alpaydin 2009). Machine learning is increasingly used in many application areas: image recognition and classification, natural language processing, spam and malware filtering, and even for performing creative tasks such as writing a poem or composing a melody (Moruzzi 2018).

The structure of many machine learning models takes inspiration from the processes going on in the human mind. Particularly relevant for applications in the scientific field is the process of extracting knowledge by generalisation from a set of data, i.e. the process of induction. However, even if machine learning models have learnt to replicate some processes of the human mind quite successfully (the process of induction is one of them) they struggle to reproduce other processes, such as common sense, plasticity, and tacit knowledge, which help humans adapting to different environments and solving different kinds of problems (Dartnall 1994; Goldman 1993; Jiang, Thagard 2014; Thagard, Stewart 2011).

In particular, machine learning systems currently exhibit a poor generalisation performance across unknown distributions of data. A task that might seem trivial to humans, for example learning to distinguish a Segway from cars, bicycles, motorcycles, scooters, and other vehicles just by seeing it once (what is known as one-shot learning), seems instead insurmountable to machines (Lake *et al.* 2011). To put in other, and more philosophical terms, what machines are struggling with is finding the *type* of different *tokens*.

Systems of classification have fixed categories and they are not able to manage concepts with low granularity. Many systems of image recognition are trained on ImageNet which has a large number of different categories[4]. Among these categories there are a lot of different breeds of dogs: Labrador, poodle, retriever, etc. What these systems lack, though, is the general category, or type, *dog*, that generalises from all the individual

[3] For more recent research on this topic, see my papers (Moruzzi 2020 and Moruzzi 2020a).
[4] See http://www.image-net.org.

RIFL (2020) Vol. 14, n. 2: 35-46
DOI: 10.4396/AISB201904

breeds, as the net is not capable to elaborate a univocal concept of dog with a coarse granularity. Thus, the two tasks that machine learning systems currently need to improve on are the extrapolation of coarse concepts from fine-grained ones and the capacity to distinguish new fine-grained concepts from similar ones.

An intelligent system should be able to learn to adapt to different environments and to solve the problems that arise in each situation. This is the challenge that machine learning faces in many application areas, such as speech and robotics. In order to work towards finding a solution to this limitation, I propose that it would be beneficial to investigate creativity as a fundamental process of the human brain. Understanding how the brain carries out creative thought is necessary in order to have a full understanding of human thinking and to replicate this process in machine learning (Dartnall 1994; Thagard 2010). I believe that the exploration of creativity can help us understand how the different models of the brain are applied to different and complex tasks and it can contribute to the development of better and more adaptable machine learning models. In particular, the study of abduction, an essential element of the creative process, can help developing algorithms that go beyond mere generalisation and achieve a deeper understanding (Halford, Levinson 1994; Harney 1994; Thagard 1990).

Creativity is essential to intelligence, understood as a problem-solving system, and it involves the discovery of methods in the arts, sciences, and technology (Jiang, Thagard 2014; Thagard, Stewart 2011; Thagard 2012). Understanding creativity is, thus, essential to understand the human capacity of insight and intuition and, as a consequence, it is essential also to develop machine learning models which perform tasks more efficiently.

3. Learning through creativity: methodology

A possible question that may be raised is why we should consider creativity in order to improve machine learning systems and not, instead, another concept. Two are the main reasons why I deem the investigation of creativity beneficial in this respect: (1) the involvement of both the conscious and the unconscious side of the mind during creative processes, and the fact that (2) creativity subsumes most of the aspects which machine learning systems currently lack.

As I mentioned, machine learning models struggle in reproducing processes pertaining to the more intuitive side of the human brain. Investigating a process of the human brain which stems from the combination of both a conscious (explicit) and a subconscious (intuitive) approach such as creativity can, thus, be beneficial (Wallas 1926; Chella, Manzotti 2012).

In addition, creativity presents a series of features that are essential for the development of areas in which machine learning is currently less strong: intuition, flexibility, originality, innovation, etc. Creativity and intuition are traditionally described as defining features of natural intelligence[5]. This is one of the reasons why AlphaGo and its 'move 37' have raised so much enthusiasm and bewilderment. Unlike chess, in fact, the game of Go requires intuition, and the (apparent) intuition behind move 37 of AlphaGo surprised everyone, Lee Sedol included (Silver *et al.* 2017).

In order to improve machine learning systems through the study of creativity, though, we first need to understand better the processes of the human brain and how these processes are applied in different situations. Once we have a better insight into how the

[5] Either human or animal. See «What Computers Can't Create», MIT tech review, accessed at https://www.technologyreview.com/s/612913/a-philosopher-argues-that-an-ai-can-never-be-an-artist/, 21/02/2019.

human brain works, we can implement these findings for the development of more adaptable and intuitive machine learning models.
In what follows, I suggest two possible ways that might be followed in order to achieve the final goal of improving the capacities of machine learning systems: (1) modelling the human brain and (2) performing a simulation of the mechanisms underlying creative processes in order to translate these findings to machine learning applications.

3.1. Modelling the human brain

Researchers in machine learning look at the human brain to be inspired to build new models and validate algorithms. In the case of creativity, we could use techniques such as fMRI, EEG, neuroimaging, and others to individuate the underlying mechanisms of creativity in the human brain. The main characteristics of these processes can be extracted and can prove themselves to be useful for creating more intuitive machine learning systems. For example, by looking at the patterns of activation in the human brain we can detect whether creativity is a hard-wired, an emergent, or a reducible property of the brain.
It should be said here that neuroscientific studies have not managed to pin down exactly where does creativity lie in the brain and which mechanisms are concerned, given the plasticity and complexity of the system and of the connections involved. Notwithstanding this difficulty, there are research groups which try to reproduce the structures and functions of the human brain. For example, it is the case of spiking neural networks: artificial neural network models that try to closely mimic organic neural networks by incorporating the concept of time in their model (Tavanaei *et al.* 2018). It is the case also of various other projects such as the Blue Brain project[6] or the Human Connectome Project[7]. The Blue Brain project, for example, aims at recreating the human brain and its connections. By using these techniques it may be possible to recreate the process of creativity in an artificial substratum.
However, trying to model the human brain to extract the main features of creativity may be not the ideal way of obtaining the necessary information needed to develop better machine learning systems. It is, in fact, possible that the brain gives us hints only at the high level about how creativity and other processes work. For instance, as we did not create planes by copying the exact mechanism of birds' wings but instead by understanding their overarching system, we could be able to recreate artificial creativity not by mapping natural creativity but by understanding how it works at a higher level. For this reason, we may decide to take a different route than modelling the human brain and use simulation instead as a method.

3.2. Creativity and Simulation

Simulation can act as a connection between thought experiments and actual measurements, providing empirical grounding to theoretical questions. In addition, simulation is a useful tool to test the consequences of changes in the hypotheses of the model without the need to build the hardware of the model itself. In the discussion at issue, simulations can be used as a substitute of an actual recreation of all the connections existing in the organic brain. For example, it could be possible to simulate models of creative processes and test how different components of creativity may affect the overall performance of the system. Competing models can then be tested and their

[6] See https://www.epfl.ch/research/domains/bluebrain/.
[7] See http://www.humanconnectomeproject.org/.

performance analysed, thus raising our understanding of the substratum of creativity and of its essential features that can then be implemented in machine learning systems.
As I mentioned, creativity is a feature currently investigated in many fields (robotics, machine learning, neuroscience, healthcare, animal studies, etc.) and simulation is used as a method by researchers working in these areas. For example, simulations are used in machine learning in self-supervised learning and multi-environment agent simulation (Le 2019). Here the concept of creativity plays a relevant role to create models which are more autonomous through techniques like interaction with the environment and evolutionary algorithms[8].
Simulations can also help to explore different ways in which machine learning systems can learn, hopefully leading to systems with a more intuitive and deeper understanding. The kinds of learning which machine learning has mostly focused on so far are: rote learning, learning from instructions, learning by analogy, learning from examples, and learning from observation (Thagard 1990). A kind of learning which is essential in humans and which, I believe, needs to be addressed more in machine learning research is learning from experience.
The latter requires an embodiment of the agent in order to interact with the world. In this respect, it has been argued that an essential part of the creative process is its perceptual, emotional, and social dimension (Stanciu 2015). The ability to perceive and interact with the environment may be beneficial also for the implementation of more usable and transparent human-machine interfaces[9]. As a consequence, the investigation of the role of the body in the process of creative discovery and problem-solving may help us understand whether an artificial embodied system would be more effective in the collaboration with humans (Fabisch *et al.* 2019; Lesort *et al.* 2019).
Lastly, I conclude by mentioning the benefits that can emerge from the integration of machine learning and simulation systems. Machine learning and simulation are both aimed at understanding and predicting the behavior of complex systems. It has been shown that, when these two approaches are working together, they are exponentially more efficient (Abdurahiman, Paul 1994; Parker 2008). As mentioned, one of the weaknesses of machine learning is that it struggles to make predictions from unknown distributions. Simulation, instead, can allow us to predict things that have never happened before and, thus, it can be used to train machine learning systems in order to improve their predictive power. Moreover, simulations can allow machine learning to be more adaptable in response to the changing behavior of other systems. At the same time, machine learning can be useful to perfectionate the mechanism of simulations.
Understanding creativity through simulation is essential to appreciate the process of problem-solving and discovery and, as a consequence, it is essential also to develop more adaptable and transparent machine learning models. Achieving a better transparency of the process that the model undertakes is indeed necessary to the efficiency of the interface between humans and the algorithm that collaborates with them in generating hypotheses or pruducts. The ideal outcome of this research would then be the implementation of machine learning models that can be more usable, adaptable, and transparent in helping humans pursue scientific experiments and discoveries of various types (Castellano 2018).

[8] See Fabisch *et al.* 2019, Chella *et al.* 2016 and the research done by the Berkeley Artificial Intelligence Lab at https://bair.berkeley.edu/blog/?refresh=1.
[9] For 'transparency' I intend the possibility for a model to be interpretable, «given the input data and the model parameters, it should be possible to step through the calculations that lead to a prediction.» (Silva *et al.* 2019: 3). A transparent model is a model whose process through which an output is produced can be understood by its users.

4. Difficulties and benefits in the application of creativity

In the first sections of the paper, I already called attention to the subject-dependent nature of the notion of creativity. This feature, which is one of the reasons for the difficulty in achieving an objective definition of creativity, can also cause some complications when trying to apply the principal features of creativity to the development of understanding in machine learning systems.

Indeed, creativity is not a definite and binary concept but it is instead graded. Thus, we can have different levels of creativity: from the minimal definition of autonomous problem-solving that I proposed in this paper, to extreme unpredictability. When talking about AlphaZero, Demis Hassabis - co-founder of DeepMind – explains the graded nature of creativity by using the notions of interpolation, extrapolation, and out of the box[10]. When doing interpolation we try to understand the shape of the given data that we have. When doing extrapolation we understand the general pattern and predict the shape of the distribution outside of the given data. The out-of-the-box step is to find out that there is another additional dimension that cannot be predicted through the data in our possession. The latter identifies the extreme kind of creativity: coming up with things that are not predictable from data. Despite the fact that this kind of creativity may be the more interesting and fruitful one, it however may lead to some difficulties in handling it.

Creativity in fact is a double-edge sword. It is normally interpreted as novelty and unpredictability, just as a black-box where, given some input, we cannot be completely sure about the output that will come out of it. However, too much unpredictability may lead to problems of interpretability. Interpretability can be understood in two ways: (1) as knowing exactly how the system internally works, or (2) as knowing which input variables have an impact on the final output. While the first case is the kind of interpretability that we can find in systems of shallow learning - more interpretable than deep learning systems (Wu *et al.* 2019) - the interpretability that researchers mostly focus on is the second type.

Research in the field of machine learning is concentrating on interpretability as understanding how a difference in the initial variables can affect the final decision of the system. This is done through models that try to peer inside the 'minds' of deep learning systems, in order to prevent adversarial attacks and in general to tackle the issue of fairness and transparency (Carter *et al.* 2019). Indeed, the issue of interpretability in machine learning leads to the consideration of another challenge of machine learning systems that I do not have the space here to address in full details: the problem of transparency. Solving the problem of transparency is extremely relevant for tackling issues of accountability, fairness, and biases. Achieving a better transparency of the process that the model undertakes is indeed necessary to ensure an ethical use of AI tools and for the efficiency of the interface between humans and AI.

However, I believe that reaching an absolute optimum of transparency is improbable. And maybe achieving complete interpretability is not even needed. We just need to learn to interpret the behaviour of machines and how it changes given different variables. Indeed, it works the same for humans: if I develop an argument and present it to you, for example in a research paper, you validate it or not independently from the knowledge of the inner functions of my brain.

What we need instead is to find a balance. Human behaviour is not completely predictable nor completely random, but somewhere in between. Many human properties are graded and the properties of flexibility and vagueness is something that machine learning systems need to improve on if they want to reach intuitive understanding.

[10] See https://cbmm.mit.edu/video/power-self-learning-systems.

Finding the balance between accuracy and interpretability and between interpretability and intuition may lead to benefits in the development of machine learning systems which achieve a deeper understanding. Once more, the investigation of a graded concept such as creativity may help to improve this balance and to convey to machine learning systems the same flexibility exhibited by humans.
In conclusion, I wish to point out some of the benefits and positive outcomes that may come out from this approach. Now that the error rates in programming machine learning systems are lower, we are assisting at a shift in the focus of research on machine learning. From the hard sciences, it is opening up to the humanities, and this is because we need humanities and social sciences to inform the development of research in this field, addressing its societal and ethical consequences, but also supporting practitioners through the analytical and problem-solving skills that the field of humanities provide researchers with.
I believe that this research can help us understand the relevance of creativity as a feature of human intelligence and, I believe, this will constitute a step forward in our capacity of predicting the development of future artificial general intelligence. Creativity, in fact, seems to sum up the features that artificial intelligence currently lacks and that it needs in order to achieve human-like levels of understanding. The exploration of creativity can help us interpret the processes of the human brain that machine learning models should strive to replicate. Achieving a better insight on the process of creative discovery will foster the implementation of systems which could help humans in generating new hypotheses, methods, and products (Edmonds 1994; Thagard 2012).
The research I presented in this paper and the further examinations that can be conducted will hopefully contribute to answering questions regarding creativity in AI systems. The relevance of addressing the questions which this paper focuses on does not only stand on the implications that they bear in relation to finding an objective definition of creativity and individuating how its analysis can foster machine learning applications. Considering how AI can fare in the creative sector in respect to humans is beneficial also to address the concern shared by many that AI may in the future outperform and replace humans. It is easier, and less alarming, to conceptualise and program a machine able to work in a warehouse or to build cars and other commodities. It is instead more challenging and disturbing to think of a machine that could equate or even surpass humans in one of the abilities which distinguish us from other beings: creativity.

References

Abdurahiman, Paul (1994), «Machine Learning and Simulation Model Specification» in *Simulation Practice and Theory*, vol. 2, pp. 1-15.

Albert, Robert, Runco, Mark (1999), *The History of Creativity Research*, in Sternberg, Robert (ed.) (1999), *Handbook of Human Creativity*, Cambridge University Press, New York.

Alpaydin, Ethem (2009), *Introduction to Machine Learning*, MIT Press, Cambridge (MA).

Ariza, Christopher (2009), «The Interrogator as Critic: The Turing Test and the Evaluation of Generative Music», in *Computer Music Journal*, vol. 33, pp. 48-70.

Boden, Margaret (1990), *The Creative Mind: Myths and Mechanisms*, George Weidenfeld and Nicolson, London.

Bown, Oliver (2011), «Experiments in Modular Design for the Creative Composition of Live Algorithms», in *Computer Music Journal*, vol. 35, pp. 73-85.

Bridy, Annemarie (2012), «Coding Creativity: Copyright and the Artificially Intelligent Author», in *Stanford Technology Law Review*, vol. 5, pp. 1-28.

Briot, Jean Paul, Hadjeres, Gaetan, Pachet, Francois (2017), «Deep Learning Techniques for Music Generation - A Survey», arXiv:1709.01620 [cs.SD].

Carter, Shan *et al.* (2019), «Activation Atlas», from https://distill.pub/2019/activation-atlas/.

Castellano, Federico (2018), «What can Philosophy Teach Machine Learning?», in *Medium*, 7/12/2018, from https://towardsdatascience.com/what-can-philosophy-teach-machine-learning-4ff091d43de6.

Chella, Antonio, Griffiths, Sascha, Wiggins, Geraint (eds.) (2016), «Robotics and Creativity», special issue in *Robotics and Autonomous Systems*, vol. 86.

Chella, Antonio, Manzotti, Riccardo (2012), «Jazz and Machine Consciousness: Towards a New Turing Test», proceedings of AUSB/IACAP 2012 Symposium.

Colin, Thomas *et al.* (2016), «Hierarchical Reinforcement Learning as Creative Problem Solving», in *Robotics and Autonomous Systems*, vol. 86, pp. 196-206.

Colton, Simon (2008), «Creativity Versus the Perception of Creativity in Computational Systems», in *AAAI Spring Symposium: Creative Intelligent Systems*.

Dartnall, Terry (ed.) (1994), «Artificial Intelligence and Creativity», in *Studies in Cognitive Systems*, vol. 17.

Dewey, John (1934), *Art as Experience*, Minton, Balch & Co, New York.

D'Inverno, Mark, McCormack, Jon (eds.) (2012), *Computers and Creativity*, Springer, London.

Dong, Hao-Wen *et al.* (2017), «MuseGAN: Multi-track Sequential Generative Adversarial Networks for Symbolic Music Generation and Accompaniment», arXiv:1709.06298.

Edmonds, Ernest (1994), «Introduction: Computer-based Systems that Support Creativity» in Dartnall (1994), pp. 327-334.

Elton, Matthew (1995), «Artificial Creativity: Enculturing Computers», in *Leonardo*, vol. 28, pp. 207–213.

Fabisch, Alexander *et al.* (2016), «A Survey of Behavior Learning Applications in Robotics: State of the Art and Perspectives», in arXiv:1906.01868.

Gaut, Berys (2010), «The Philosophy of Creativity», in *Philosophy Compass*, vol. 5, n. 12, pp. 1034-1046.

Goldman, Alvin I. (ed.) (1993), *Readings in Philosophy and Cognitive Science*, MIT Press, Cambridge (MA).

Halford, Graeme S., Levinson, Robert (1994), «Introduction: Creativity and Cognition» in Dartnall (1994), pp. 139-144.

Harney, Maurita (1994), «Clues to Creativity» in Dartnall (1994), pp. 195-208.

Jiang, Mingming, Thagard, Paul (2014), «Creative Cognition in Social Innovation», in *Creativity Research Journal*, vol. 26, n. 4, pp. 375-388.

Kralik, Jerald *et al.* (2016), «Modeling Incubation and Restructuring for Creative Problem Solving in Robots», in *Robotics and Autonomous Systems*, vol. 86, pp. 162-173.

Lake, Brenden *et al.* (2011), «One Shot Learning of Simple Visual Concepts», from https://cims.nyu.edu/~brenden/LakeEtAl2011CogSci.pdf.

Le, Nam (2019), «Evolving Self-taught Neural Networks: The Baldwin Effect and the Emergence of Intelligence», arXiv:1906.08854v1.

Lesort, Timothée *et al.* (2019), «Continual Learning for Robots», arXiv:1907.00182v1.

Levinson, Robert (1994), «Experience-based Creativity», in Dartnall (1994), pp. 161-179.

Luck, Michael, D'Inverno, Mark (2012), «Creativity Through Autonomy and Interaction» in *Cognitive Computation*, vol. 4, pp. 332-346.

Minsky, Marvin (1982), «Why People Think Computers Can't», in *AI Magazine*, vol. 3, n. 4, DOI: https://doi.org/10.1609/aimag.v3i4.376 .

Moruzzi, Caterina (2018), «Creative AI: Music Composition Programs as an Extension of the Composer's Mind», in Müller, Vincent C. (ed.), *Philosophy and Theory of Artificial Intelligence 2017*, Springer, Berlin, DOI: 10.1007/978-3-319-96448-5_8.

Parker, Wendy S. (2008), «Computer Simulation», in Curd, Martin, Psillos, Stathis (eds.) *The Routledge Companion to Philosophy of Science*, Routledge, New York, pp. 135-145.

Paul, Elliot, Kaufman, Scott (eds.) (2014), *The Philosophy of Creativity*, Oxford University Press, Oxford.

Silva, Vivian *et al.* (2019), «On the Semantic Interpretability of Artificial Intelligence Models», https://arxiv.org/abs/1907.04105.

Silver, David *et al.* (2017), «Mastering the Game of Go Without Human Knowledge», in *Nature*, vol. 550, n. 7676, p. 354.

Simon, Herbert (1985), «What We Know About the Creative Process», in Kuhn, Robert (ed.), *Frontiers in Creative and Innovative Management*, Ballinger, Cambridge, pp. 3–22.

Stanciu, Marius (2015), «Embodied Creativity: A Critical Analysis of an Underdeveloped Subject», in *Procedia - Social and Behavioral Sciences*, vol. 187, pp. 312-317.

Still, Arthur, D'Inverno, Mark (2016), «A History Of Creativity For Future AI Research», in Pachet, F. *et al.* (eds.), *Proceedings of the 7th International Conference on Computational Creativity*, (ICCC), pp. 147-154.
Tavanei, Amirhossein *et al.* (2018), «Deep Learning in Spiking Neural Networks», arxiv:1804.08150.

Thagard, Paul (1990), «Philosophy and Machine Learning», in *Canadian Journal of Philosophy*, vol. 20, pp. 261-76.

Thagard, Paul (2006), *Philosophy of Psychology and Cognitive Science*, Elsevier, Amsterdam.

Thagard, Paul (2010), *The Brain and the Meaning of Life,* Princeton University Press, Princeton.

Thagard, Paul (2012), *The Cognitive Science of Science: Explanation, Discovery, and Conceptual Change*, MIT Press, Cambridge (MA).

Thagard, Paul, Stewart, Terrence (2011), «The AHA! Experience: Creativity Through Emergent Binding in Neural Networks», in *Cognitive Science*, vol. 35, pp. 1-33.

Van Der Schyff, Dylan *et al.* (2018), «Musical Creativity and The Embodied Mind: Exploring The Possibilities of 4E Cognition And Dynamical Systems Theory», in *Music and Science*, vol. 1, pp. 1-18.

Wallas, Graham (1926), *Art of Thought*, Harcourt-Brace, New York.

Wiggins, Geraint (2005), «Searching for Computational Creativity», in *New Generation Computing*, vol. 24, pp. 209-222.

Wiggins, Geraint (2006), «A Preliminary Framework for Description, Analysis and Comparison of Creative Systems», in *Knowledge-Based Systems*, vol. 19, pp. 449-458.

Wiggins, Geraint, Pearce, Marcus, Müllensiefen, Daniel (2011), «Computational Modelling of Music Cognition and Musical Creativity», in *The Oxford Handbook of Computer Music*, Oxford University Press, Oxford.

Wu, Felix *et al.* (2019), «Simplifying Graph Convolutional Networks», arXiv:1902.07153.