Theoretical Learning Creators and Creative Scientists

Luís Espírito Santo[*]
CISUC, DEI
University of Coimbra, Portugal
lesanto@dei.uc.pt

Amílcar Cardoso
CISUC, DEI
University of Coimbra, Portugal
amilcar@dei.uc.pt

Geraint Wiggins
AILab, DINF
Vrije Universiteit Brussel, Belgium
geraint@ai.vub.ac.be

Abstract

Can creative jobs be performed by machines? Such questions are in debate, as Learning Endowed Generative Systems threat to invade creative areas by recently achieving great results in several widely-accepted creative tasks. Computational Creativity has prolifically provided us with formal tools to address such argument, systematically leaving "learning" out of the equation. Before that, Formal Learning Theory, also informally known as "learning in the limit", allowed to study some of the limits of learning, yet mainly pinning these results to the language acquisition and scientific discovery problems, with no known example of generalized analogies to other more widely accepted creative domains. We will endeavour to explore the parallels between these two currently disparate areas, Computational Creativity and Formal Learning Theory, by identifying points of contact and clear differences and expanding both in a convergent joint transdisciplinary direction. This merged view is believed not only to spawn new studies in generative models, computability of learning, and computational creativity but also bring new insights to some philosophical debates on the relationship between Artificial Intelligence and Computational Creativity and the nature of human creativity.

Motivation

Throughout my life, as a researcher and an amateur artist, there were several times I ambitiously endeavoured to create something truly new, something that everyone could agree that was completely revolutionary in an helpful way, or in other words, that was deeply creative. This never-ending quest to find the next big thing that will wonder crowds and help us moving forward as a society is far from being exclusively mine since it has been discussed for generations ever since humanity started studying its own evolution by looking into the past as a way to understand the future.

Nowadays, revolutionary advances in both arts and science happen at a fast pace when compared with decades ago and it seems to keep on accelerating (Oliveira 2017). This hastening allowed everyone in a period of a lifespan to witness several advances pushing for a better understanding of the nature of intelligence and creativity while also deeply impacting society by forcing a continuous anxiety for creating things that results in a self-fed spiral of innovation acceleration. Technology has been playing a great role on this whole feedback loop and, more recently, Artificial Intelligence (AI) has kept on threatening to revolutionize our daily lives while sub-fields such as Machine Learning (ML) and sub-sub-fields such as Deep Learning (DL) have been providing great solutions for tasks that were exclusively performed by humans.

Creative tasks are some of the tasks that many believed to be so hard to automatize that were considered to need human exclusive capacities. Human creativity has been several times said to be infinite, unlimited and unbounded, yet human creations seem always to be an appropriation, transformation or combination of previously known things. Knowing that there are certain limits to what we can learn from our environment, and if we believe that everything we create "is a remix" (Ferguson 2011), there might be some limits on the infinity of human and computer creativity which seem to be directly linked to what one can learn.

Background

During the last decade, we witnessed the rising of machine learning techniques that allowed them to autonomously create new realistic looking things, those being faces (Karras et al. 2020), illustrations (Chen et al. 2020) or music (Dong et al. 2018; Espírito Santo 2019). Even this month, DALL-E2 (Ramesh et al. 2022) came to plunge forward the state of the art on image generation systems, by creating astonishing contextualized images when provided with a single text prompt. This revolution was pioneered in 2014 by the presentation of two very powerful generative deep models: GAN's (Goodfellow et al. 2014) and VAE's (Kingma and Welling 2013). Not only are these and other deep generative models such as Transformers and Diffusion Models (Vaswani et al. 2017; Dhariwal and Nichol 2021; Foster 2019; Goodfellow, Bengio, and Courville 2016) becoming the backbone of very powerful Learning Endowed Generative Systems (LEGS) but also are populating the academic research while also reaching the general pub-

[*]This work is funded by national funds through the FCT - Foundation for Science and Technology, I.P., within the scope of the project CISUC - UID/CEC/00326/2020 and by European Social Fund, through the Regional Operational Program Centro 2020. The first author is funded by Foundation for Science and Technology (FCT), Portugal, under the grant 2021.05529.BD

lic. These new machine learning capabilities seem to challenge several traditional arguments against "creativity in machines" (Du Sautoy 2019).

On creativity, the Computational Creativity (CC) area is a branch of AI which began the late 1990s. The CC-continuum (Ackerman et al. 2017) encompasses opposing approaches: at one extreme, those that simulate creative behavior with machines, at the other, theoretical approaches that focus on modeling general creativity. Concerning modelling creativity, although authors have agreed that creativity involves novelty and usefulness (Mumford 2003), there are still several kinds of creativity to take into account (Boden 2009; Kaufman and Beghetto 2009) and several different complex components (Rhodes 1961):

- the Product, that might be a theory or a song, for example;
- the Person, the author;
- the Press, that refers to the cultural and social environment;
- and the Process, the method by which the Person achieved the Product while imbued in the Press.

From the several interpretations on the creative process, two models proposed deserve mentioning: the Creative Systems Framework (CSF) (Wiggins 2006b; 2006a) which four main components are a universe, a conceptual space, an evaluation function, and a strategy to explore that universe; and Information Dynamics of Thinking (IDyOT), contextualized in (Wiggins 2020). There are many other references of neurological, psychological, educational, philosophical, and cognitive debates around creativity and its domains (Koestler 2014; Pigrum 2009; Wallas 1926; Hill and Monroy-Hernández 2013). However, and even though adaptability has been constantly considered a way to implement creativity (Ackerman et al. 2017), there is still no known text explicitly and methodically discussing the role of "learning" in creativity and most of the works that try to shed light on how creativity allows learning lack the very insightful view provided by the most recent developments in AI and ML.

On learning, in the 1960s, Gold (1967) formalizes a model for inductive learning, motivated by language acquisition in infants, and proves that not all classes of languages can be identified by every kind of "learner". On top of this framework, Blum and Blum (1975), Case (1983; 2012) and several other authors contributed to create a theory with definitions and criteria for "scientists" identifying both recursive functions (scientist for functions) and recursively enumerated sets (scientists for sets) on a provided text, i.e., a sequence of positive information such as any enumeration of a function's points or any enumeration of sentences in a language. We refer to this theory as the Formal Learning Theory (FLT), which is addressed and compiled in (Kelly 1996; Martin and Osherson 1998; Jain et al. 1999). FLT has been used to prove some results that confront some widely spread ideas about learning and knowledge while also, according to Costa (2013; 2017; 2019), providing a better understanding on the "large scale limitations of scientific discovery". We believe that these insights can and should be applied to other creative domains such as music, a unique domain according to Wiggins (2020), in order to also grasp the limitations of creativity. Moreover, we believe the formal foundations of FLT deserve to be known in a broader context for its generality, elegance and interesting yet controversial results and their implications in other fields of research.

Research question

How can we formalize the relationship between Learning and Creativity?

Methodology

1. Study in depth the area of FLT and its results, while keeping notes on how those results might bring new insights to CC;
2. Collect a set of formal models of CC that might be promising for expansion to include learning;
3. Follow a series of though experiments aiming at creating a mapping between the several CC models and the formalized components of FLT and translate the results from one to the other area, while keeping in mind the current state of art of LEGS;
4. Detail how the new formalization take form on different domains such as scientific discovery and music;
5. Document and share our results at different conferences, workshops and communities to obtain feedback and different opinions as well as disseminate the new insights the spawn from this investigation.

Expected results

- Proposals of new learning paradigms, criteria, or strategies, motivated by the creativity context, that might lead to new results on FLT and might be leveraged to better understand the limitations of learning;
- Expansions for some CC models in order to better formalize and include Learning as part of the process while explicitly identifying where this component has already been considered;
- Parallels and relationships between FLT and CC models, considering the new definitions and results achieved on the previous point, possibly resulting in a new formal model for general creative concepts, and producing new viewpoints on the relationship between learning, creativity, and intelligence;
- Some views on how these new formal models take form on several learning and creativity related domains such as scientific discovery (usually more associated to learning) and music (undoubtedly related to creativity), possibly studying the theoretical limitations and highlighting the clear interdisciplinarity of both FLT and CC;
- Potentially, a taxonomy for LEGS based on the way they can be formalized using the new formal tools developed;
- New attention towards the new insights brought by the two rather disjoint and underrated areas of FLT and CC.

Significance

The expected results might not only shed some light on the increasingly blurred line between machine and human creation, but also would help to bring both ML and CC closer together. We believe that joining of forces is one of the paths to achieve greater and more general AI in a way to keep both public and private investors interested in the area in an joint effort to avoid another speculatively prophesied AI winter.

References

Ackerman, M.; Goel, A.; Johnson, C.; Jordanous, A.; León, C.; Perez y Perez, R.; Toivonen, H.; and Ventura, D. 2017. Teaching computational creativity. In Sony CSL Paris, F., ed., *Proceedings of the Eigth International Conference on Computational Creativity (ICCC 2017)*, 9–16.

Blum, L., and Blum, M. 1975. Toward a mathematical theory of inductive inference. *Information and control*.

Boden, M. A. 2009. Computer models of creativity. *AIMag* 30(3):23–23.

Case, J., and Smith, C. 1983. Comparison of identification criteria for machine inductive inference. *Theor. Comput. Sci.* 25(2):193–220.

Case, J. 2012. Algorithmic scientific inference. *International Journal of Unconventional Computing* 8(3).

Chen, M.; Radford, A.; Child, R.; Wu, J.; Jun, H.; Luan, D.; and Sutskever, I. 2020. Generative pretraining from pixels. In Iii, H. D., and Singh, A., eds., *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, 1691–1703. PMLR.

Costa, J. F. 2013. Incomputability at the foundations of physics (a study in the philosophy of science). *J. Logic Comput.* 23(6):1225–1248.

Costa, J. F. 2017. The unity of science as seen through the universal computer. *International Journal of Unconventional Computing* 13:59–81.

Costa, J. F. 2019. On discovering scientific laws. *International Journal of Unconventional Computing* 14:285–318.

Dhariwal, P., and Nichol, A. 2021. Diffusion models beat gans on image synthesis. *Adv. Neural Inf. Process. Syst.* 34.

Dong, H.-W.; Hsiao, W.-Y.; Yang, L.-C.; and Yang, Y.-H. 2018. MuseGAN: Multi-track sequential generative adversarial networks for symbolic music generation and accompaniment. *AAAI* 32(1).

Du Sautoy, M. 2019. *The Creativity Code: Art and Innovation in the Age of AI*. Harvard University Press.

Espírito Santo, L. 2019. Automatically generating novel and epic music tracks. Master's thesis, Instituto Superior Técnico, University of Lisbon.

Ferguson, K. 2011. Everything is a remix.

Foster, D. 2019. *Generative Deep Learning: Teaching Machines to Paint, Write, Compose, and Play*. "O'Reilly Media, Inc.".

Gold, E. M. 1967. Language identification in the limit. *Information and Control* 10(5):447–474.

Goodfellow, I.; Bengio, Y.; and Courville, A. 2016. *Deep Learning*. MIT Press.

Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial networks.

Hill, B. M., and Monroy-Hernández, A. 2013. The remixing dilemma: The Trade-Off between generativity and originality. *Am. Behav. Sci.* 57(5):643–663.

Jain, S.; Osherson, D.; Royer, J. S.; and Sharma, A. 1999. *Systems that Learn: An Introduction to Learning Theory*. MIT Press.

Karras, T.; Laine, S.; Aittala, M.; Hellsten, J.; Lehtinen, J.; and Aila, T. 2020. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8110–8119.

Kaufman, J. C., and Beghetto, R. A. 2009. Beyond big and little: The four C model of creativity. *Rev. Gen. Psychol.* 13(1):1–12.

Kelly, K. T. 1996. *The Logic of Reliable Inquiry*. Oxford University Press, USA.

Kingma, D. P., and Welling, M. 2013. Auto-Encoding variational bayes.

Koestler, A. 2014. *The act of creation*. One 70 Press.

Martin, E., and Osherson, D. N. 1998. *Elements of Scientific Inquiry*. MIT Press.

Mumford, M. D. 2003. Where have we been, where are we going? taking stock in creativity research. *Creat. Res. J.* 15(2-3):107–120.

Oliveira, A. 2017. *The Digital Mind: How Science Is Redefining Humanity*. MIT Press.

Pigrum, D. 2009. *Teaching Creativity: Multi-mode Transitional Practices*. A&C Black.

Ramesh, A.; Dhariwal, P.; Nichol, A.; Chu, C.; and Chen, M. 2022. Hierarchical Text-Conditional image generation with CLIP latents.

Rhodes, M. 1961. An analysis of creativity. *Phi Delta Kappan* 42(7):305–310.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Adv. Neural Inf. Process. Syst.* 30.

Wallas, G. 1926. *The art of thought*, volume 10. Harcourt, Brace.

Wiggins, G. A. 2006a. A preliminary framework for description, analysis and comparison of creative systems. *Knowledge-Based Systems* 19(7):449–458.

Wiggins, G. A. 2006b. Searching for computational creativity. *New Gener. Comput.* 24(3):209–222.

Wiggins, G. A. 2020. Creativity, information, and consciousness: The information dynamics of thinking. *Phys. Life Rev.* 34-35:1–39.