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

Several discussions have emerged about the remarkable results that Learning Endowed Generative Systems are achieving in various well-accepted creative fields and their impact on jobs in those areas. But (how) can we study formally creative limitations? Computational Creativity has prolifically provided us with philosophical concepts 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. Attempting to create tools to formally and philosophically address questions involving both creativity and learning, we endeavoured to explore the parallels between these two currently disparate areas, Computational Creativity and Formal Learning Theory, through a transdisciplinary approach. This merged view is believed not only to spawn new studies in generative models, computability of learning, and computational creativity but mainly to bring new insights to some philosophical debates on the relationship between Artificial Intelligence, Machine Learning and Computational Creativity.

Motivation

For a long time, creativity was believed to be so hard to automatize that it was considered to be a human-exclusive capacity. This capacity in humans is constantly described as infinite, unlimited and unbounded, yet, in an opposite argument, human creations usually seem to result from an appropriation, transformation or combination of previously known things, i.e., that everything we create "is a remix" (Ferguson 2011). Assuming there are inherent limitations to what one can learn from the environment, it then seems reasonable to explore the boundaries to the infinity of both human and computer creativity. Given this goal, we ask the question: (How) can we systematically and formally study these philosophical limitations?

We believe that to answer this question, we need new formal tools that abstract and model characteristics of both creative and learning behaviours. Besides, we focus on formal tools, i.e., tools that can be used to draw fundamental conclusions about the limitations of creativity by making use of other previously developed formal paradigms. We start walking this path by comparing different frameworks and paradigms relating to creativity and learning, namely and Formal Learning Theory (FLT) à la Gold respectively, to answer the question: (how) can we formalize the relationship between Learning and Creativity?

Background

During the last decade, Machine Learning (ML) techniques allowed computers 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). Tools like DALL-E and ChatGPT came to plunge forward the state-of-the-art on image and text generation systems. Not only are these deep generative models such as GAN's, VAE's, Transformers and Diffusion Models (Goodfellow et al. 2014; Kingma and Welling 2013; 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 are also dominating the academic research and reaching the public. These new ML capabilities seem to challenge several traditional arguments against "creativity in machines" (Du Sautoy 2019).

The Computational Creativity (CC) area has explored creativity as a branch of AI since the late 1990s. The spectrum of CC approaches (y Pérez 2018; 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). One formal model deserves mentioning at this point: the Creative Systems Framework (CSF) (Wiggins 2006a;

*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

2006b). 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 formal 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 point of view of ML.

On learning, in the 1960s, Gold (1965) 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.

Methodology

1. Study FLT by exploring the book (Jain et al. 1999) and its results, while keeping notes might bring new insights to CC;
2. Re-read some proposals in CC that might be promising for expansion to include learning, namely (Wiggins 2006a) and (Ritchie 2006);
3. Follow a series of though experiments concerning LEGS and creative humans aiming at creating a set of formal definitions of Computational Creativity (CC) concepts and their counterparts in FLT;
4. Translate results from FLT to CC, or explore new results given this new set of definitions;
5. Illustrate the obtained definitions and results using some grounded examples in domains such as scientific discovery and music;
6. 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, for FLT à la Gold, motivated by the creativity context, that might lead to new results 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;
- 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.

Preliminary results

From our research, we developed a paper entitled "Towards a Formal Creativity Theory: Preliminary Results in Novelty and Transformativeness,"[1] submitted to ICCC'24 (the long paper was rejected). In this paper, we introduce the main formal concepts behind E. Mark Gold's (1965) learning paradigm. We suggest that one of the main relationships between learning and creativity is what Boden (1994) describes as transformational creativity, where learning happens during the exploration process. Based on these definitions, we are able to model the contextualized properties of artefacts such as novelty and transformativeness, i.e., the power of an artefact to change an agents mind after experiencing a certain sequence of other artefacts. We provide some simple preliminary results to illustrate what we can prove given this formalization: we prove that, for all possible behaviours we can have in an agent, novelty is not necessary nor sufficient for transformativeness. This suggests that if one aims at creating transformative artefacts then searching for novelty does not guarantee nice results, unless we constrain ourselves to a proper subclass of possible agents, that is actually strictly less powerful and for many cases can be an oversimplification of the human behaviour.

We are preparing a new batch of results in typicality and utility, in the context of CSF.

[1] Availalble in: https://arxiv.org/abs/2405.02148

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, and Blum. 1975. Toward a mathematical theory of inductive inference. *Infect. Control*.

Boden, M. A. 1994. Creativity: A framework for research. *Behav. Brain Sci.* 17(3):558–570.

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.

E. Mark Gold. 1965. Limiting recursion. *J. Symbolic Logic* 30(1):28–48.

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.".

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.

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

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

Ritchie, G. 2006. The transformational creativity hypothesis. *New Generation Computing* 24(3):241–266.

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.

y Pérez, R. P. 2018. The computational creativity continuum. In *ICCC*, 177–184.