Computational Creativity's Problems with Problem-Solving

Tarek R. Besold

Digital Media Lab, Center for Computing and Communication Technologies (TZI)
University of Bremen, Germany
Tarek.Besold@uni-bremen.de

Computational creativity (i.e., "the philosophy, science and engineering of computational systems which, by taking on particular responsibilities, exhibit behaviours that unbiased observers would deem to be creative", cf. [2]) is steadily gaining popularity and has become a recognised field of scientific activity. Still, while work in art-performing and artefact-producing computational creativity ("artistic creativity") has greatly advanced, research into the computerisation of other forms of creativity, such as general creative capabilities for computational cognitive agents or automated creative problem-solvers addressing real-world scenarios in a domain-independent way ("problem-solving creativity"), is lacking far behind and is receiving only limited attention from the community. In fact, over the last years neighbouring disciplines (such as cognitive science/cognitive modelling or machine learning) have been reporting developments which—when looked at from the perspective of computational creativity—could turn out to be crucial stepping stones for advancing towards closing this gap between artistic and problem-solving creativity, but which went mostly unnoticed by the majority of computational creativity researchers.

Coherent with the view expressed in [1], I suggest that this asymmetry in development and dedication between the two main strands of computational creativity research can be traced back to (at least) two self-created shortcomings within the field: the lack of a commonly acknowledged definition of the concept of creativity; and the tacit belief in (perceived or actual) differences in problem difficulty and complexity between artistic and problem-solving creativity. The absence of definition poses significant challenges from a theoretical/conceptual perspective—leaving the main objective of the research field undefined, and the discipline without decision criteria allowing to guide research, measure progress, or decide about success or failure of the entire program—which impact significantly stronger on work in problem-solving creativity than on efforts in artistic creativity, partially due to the commonly more output- than process-oriented nature of the latter.[1] This difference also is relevant for the second diagnosed ailment (which is more practical/empirical in nature), since a reasonably general

[1] Even when accepting that a single, all-encompassing definition of creativity might indeed be hard to achieve already for conceptual and historical reasons (cf., for instance, the arguments presented in [4]), the lack of even an approximation or an incomplete, but widely accepted working definition—together with the absence of serious efforts aiming to establish either—in my view constitutes a major methodological failure.

computational model of problem-solving creativity seems to require solutions to (often even in their disciplines of origin) still unsolved research questions relating to, for instance, the re-representation of concept or search spaces and the frame problem in classical AI, or processes and mechanisms driving human creativity in psychology and cognitive science.

These observations notwithstanding, I strongly suggest that problem-solving creativity has to form an equally important part of computational creativity research as does artistic creativity. Moreover, I argue that already now many open questions in problem-solving creativity could feasibly be addressed through active exchange and collaboration with other disciplines investigating creativity-relevant capacities and concepts (other subfields of AI, psychology, cognitive science, etc.). Additionally, I am confident that over the course of this process also the two described deficiencies within the field of computational creativity could and would be overcome, significantly advancing the field in its entirety. In order to provide additional support for these claims—and to specifically emphasise the one concerning joint efforts with related fields—two examples for research work outside of computational creativity shall be mentioned, each of which carries the promise to be highly relevant in systems aiming to solve (aspects of) problem-solving creativity: Meta-Interpretive Learning and predicate invention [3] in Inductive Logic Programming (ILP), and Bayesian theory learning applicable to the modeling of concept and theory formation in humans [5]. The former equips ILP systems with the ability to learn recursive logic programs, as well as to introduce new (i.e., previously unforeseen) predicates into learned logical theories, while the latter constitutes a (in certain variants of the overall approach strongly cognitively-inspired) modeling framework for processes of observation-driven theory formation.[2] And these are only two cases among what I am convinced is a surprisingly high number of recent revolutionary developments within the computation- and cognition-related fields of research. But while these breakthroughs—once identified and properly re-contextualised—also have the potential to turn out as revolutions from the computational creativity point of view, they hitherto simply have gone unnoticed by the big majority of researchers currently active in computational creativity.

Converting and re-applying approaches from related disciplines in problem-solving creative applications also offers another advantage. The transfer and adaptation of already tested techniques and implementations can be expected to accelerate development and reduce the risk of failure for operational proof of concept systems. While some might consider this scientifically less prestigious—an assessment I personally oppose—initial successes are more likely to happen in a timely manner, providing evidence that advances are also possible on the side of problem-solving creativity; which, in turn, I expect to play an important motivational role in putting work on the corresponding topics and questions back on the map.

[2] Cf. [1] for a short assessment of both lines of work from the perspective of computational creativity.

In closing, I want to emphasize that a focus on problem-solving creativity would also be desirable for computational creativity as scientific community. Creativity in its many forms unquestionably forms part of human cognition and of the way humans as cognitive agents interact with their environment. Still, this spectrum of interaction spans far beyond the realm of artistic activity and performance and basically covers all domains of our daily lifeworld. Work in artistic computational creativity is important and the achieved recognition is well-deserved, but this should not limit the focus of the community to certain forms of creativity whilst ignoring others. Looking at AI as a field, it has to be noted that—in parts simply due to some recent successes and the corresponding public attention, which also entices a still growing number of researchers to reinvest time and resources into the field—solutions to some longstanding problems have been presented, and promising approaches to others have been suggested. For example cognitive systems and cognitive robotics (but also other disciplines) are advancing at a pace unseen for the previous two decades, and the addressed questions while still quite basic are coming closer to also addressing (parts of) phenomena such as creativity. I am confident that these communities could greatly profit from interacting with the field of computational creativity and its practitioners, drawing on the latter's knowledge and experience when starting to address creativity-related questions in their respective research endeavours. And as argued before, the exchange would be reciprocal since also computational creativity—and especially the problem-solving creative line of investigation—certainly could greatly profit from developments in other disciplines of AI and related areas.

References

1. Besold, T. R. The Unnoticed Creativity Revolutions: Bringing Problem-Solving Back into Computational Creativity. In *Proceedings of the AISB 2016 Symposium on Computational Creativity.* (AISB 2016, April 4, University of Sheffield, UK), 2016.
2. Colton, S. and Wiggins, G. Computational Creativity: The Final Frontier? In *Proceedings of the 20th European Conference on Artificial Intelligence (ECAI)*, pp. 21–26. IOS Press, 2012.
3. Muggleton, S. H., Lin, D., Pahlavi, N., and Tamaddoni-Nezhad, A. Meta-interpretive learning: application to grammatical inference. Machine Learning, 94, pp. 25–49. 2014.
4. Still, A. and d'Inverno, M. A History of Creativity for Future AI Research. In Proceedings of the 7th International Conference on Computational Creativity. (ICCC 2016, June 27 - July 1, Paris, France), 2016.
5. Ullman, T. D., Goodman, N.D., and Tenenbaum, J. B. Theory Acquisition as Stochastic Search. Proceedings of the 32nd Annual Meeting of the Cognitive Science Society. (CogSci 2010, August 11 - 16, Portland, Oregon, USA), 2010.