Transformation Is All You Need

Sean Luke
Department of Computer Science
George Mason University
Washington, DC, USA
sean@cs.gmu.edu

Matthew Cronin
Costello College of Business
George Mason University
Washington, DC, USA
mcronin@gmu.edu

Abstract

We argue that the computational creativity field has had a minor impact on AI, in large part because of its emphasis on exploratory creativity, which maps primarily to already well understood algorithms in search and optimization. We argue that although it is relatively less studied, transformational creativity may yet instead provide a richer vein for inspiration of AI algorithms due to its meta-nature and the opportunities it affords in combination with multiagent systems. We offer a new framework for multiagent transformational creativity, and discuss possible application to AI fields.

Introduction

A major motivation for the field of computational creativity is its potential to give rise to useful AI algorithms that are in some sense artificially creative. But "creative" is a vague and slippery concept with many weak or ill-considered definitions over the years. When it is applied to AI, computational creativity lacks a concrete or rigorously quantifiable task or goal. Measurements of computationally creative processes or products are fraught with bias, context sensitivity, and disagreement as to what constitutes a creative product (Colton and Wiggins 2012). Compare this to machine learning, which has a specific task: given examples drawn from the world, produce a generalizable model which can perform discriminative or generative operations; and this generalization is straightforwardly measurable with validation examples.

Lacking a specific task for computational creativity, we can still attempt to understand the *process* of human creativity to see if this understanding can *inspire* us to develop useful artificially intelligent algorithms for whatever purpose. This is similar to historical efforts in evolutionary computation, neural networks, computer vision, soft computing, and swarm robotics, which all borrowed liberally from biology or cognitive science to inspire competitive AI algorithms. However in our opinion the output from the computational creativity community has not contributed very much new to AI, resulting in a multi-decade research effort by this community with little to show for it. Algorithms deemed "creative" after the fact have come from fields such as deep learning: but not often from computational creativity.

We think that this is because the study of the creative process, at least in the computational creativity literature, has been dominated by the wrong thing: exploring of the space of novel and valuable concepts, that is, *exploratory creativity*. Models of exploratory creativity largely boil down to straightforward search, optimization, and iterative revision, and these topics have been well studied in engineering and AI for over a century. We think that not only has exploratory creativity not offered much new to AI, it is not really a useful notion of *creativity in humans* either: it's just exploration.

But there is a contending notion of creativity whose process is well justified. Humans have limited working memory, and so we get stuck in local optima in the value and novelty functions of exploratory creativity's "concept space". This is manifested in writer's block, for example, or getting stuck on a problem. We also may be blinded by our biases as to the existence of certain fruitful avenues to explore. We can can overcome these problems by looking for a new angle on the task, or a new way to organize the space, to bring formerly distant concepts into view. That is, we engage in what Boden (1992) calls *transformational creativity*.

In his 2000 AAAI Presidential Address, Bruce Buchanan (2001) argued for the importance of transformational creativity to AI, claiming "that the key to creativity is at the metalevel." Transformational creativity may be modeled as a meta-exploration through the space of possible topologies which organize the concept space: it changes how likely we are to reach concepts, and thus jumbles up the local optima, revealing new pathways to make progress. This is often mistakenly thought of as a phenomenon of high intelligence and with global impact, but in fact it happens all the time in everyday life. Meta-level search, optimization, and representation have been studied in AI only to a limited degree, and so this opens avenues of study. And transformational creativity is potent in combination with *multiagent* scenarios, where multiple creators spread to one another not only news of their creations, but *new ways to think about the problem itself.*

For these reasons, we think that transformational creativity deserves to be the primary research thrust of computational creativity. And yet it has seemed to take a back seat in the literature. Out of the 755 papers from the 2010–2025 ICCC conferences, the word *transformational* appeared in 79 papers, and transformational creativity was discussed in a meaningful (non-passing) manner in only 37. The proportion of the discussion in recent ICCC conferences has dropped by half relative to the earlier ones.

Previous Work

Creativity has been studied since Greek antiquity, and informal models of creativity have been proposed at least since von Helmholtz (1896) and Wallas (1926), but the seminal study of *computational* creativity may have been by Newell, Shaw, and Simon (1958). They noted that according to the psychological literature of the time, "creative thinking" was characterized by at least one of four features, including "the product of the thinking has novelty and value", and "the thinking is unconventional, in the sense that it requires modification or rejection of previously-accepted ideas." Novelty, value, and (to a lesser extent) unconventionality have since come to dominate later definitions of creativity and computational creativity (Sternberg and Lubart 1999).

Margaret Boden's theory of the process of creativity continued along these lines (Boden 1992). Boden proposed three major kinds of creativity: *exploratory, combinatorial,* and *transformational*. In exploratory creativity, the creator is exploring the space of *concepts* (potential, perhaps incomplete, creative artifacts), in search of ones which are *valuable* and *novel*. A *valuable* concept is one that someone cares about in some way or that satisfies some aesthetic or performance criterion, and a *novel* concept is one that has not been discovered before. Combinatorial creativity expands on exploratory creativity by defining creative acts as the combination of two concepts from different and often unrelated domains, perhaps drawn from different concept spaces.

Transformational creativity is touted as an "advanced" form of creativity, as the person's creative act *changes the topology, boundary, or nature of the concept space* and thus enables her to rapidly explore new areas in the space. Boden's oft quoted canonical example is Friedrich August Kekulé's discovery of the cyclic nature of benzene: before it, his (and others') concept space for molecules in chemistry was restricted to trees or to crystalline meshes, but afterwards, organic chemistry was forever changed to consider cycles.

Creativity is also studied in a multiagent context. Early work in this area was based on Csikszentmihalyi's Domain-Individual-Field Interaction (DIFI) model (Feldman, Csikszentmihalyi, and Gardner 1994). Here a creative *individual*, inspired by the existing historical body of work (the *domain*), produces new creative works which are then assessed for novelty and value by *the field*: critics, audiences, etc. If his works are deemed sufficiently creative, they are added to the domain and the process continues.

Semiformal Models of Computational Creativity Many computational models have stemmed from Boden (Wiggins 2006b; Ventura 2011; Ritchie 2012). These have largely built off of Wiggins's seminal work extending Boden's concept space into a kind of state-space to search (Wiggins 2006a). Collectively this body of work is known as *Creativity as Search*. Wiggins describes an agent engaged in exploratory creativity as traversing a concept space. He further defines a set of rules which constrains which concepts in this space are of interest to the agent (or *valid* (Wiggins 2006b)); a second set which describes how the space may be traversed (perhaps via heuristic search); and a third set which assesses the value of concepts in the space. Wiggins extends this model to include transformational creativity by defining two *meta-rulesets*: one which defines the space of validity rulesets, and one which defines the space of traversal rulesets. By selecting a new ruleset from the validity space, the agent would thus modify the boundaries defining the valid concepts to search. By selecting a new traversal ruleset, the agent would change the transition topology over the concept space itself.

Ritchie also developed a framework for transformational computational creativity (Ritchie 2005; 2006), which has since been extended to consider biases in search (Demke and Ventura 2024).

Other work has tried to combine Boden with Csikszentmihalyi or other multiagent models. Sauders (2019; 2001) describes an agent-based model where agents assess and provide feedback for other agents' creative works. Linkola (2019) redefine Wiggins's rulesets as ones concerning themselves with the personal concept space of the agent, then create global validity, traversal, and value rulesets corresponding to the personal ones, but which are shared collectively among agents at a society-wide level.

Why Transformational Creativity Is Important

As pointed out by Buchanan (2001), there is a critical difference between search guided by the assumptions embedded in a problem representation, and search for the right representation (the meta-level). This was originally called the distinction between understanding and solving a problem (Simon and Hayes 1976), but has also been described as problem finding versus problem solving (Mumford and Gustafson 1988), as exploration versus exploitation (in the psychological sense, not the optimization sense), or most recently as the distinction between perspective use and perspective change (Cronin and Loewenstein 2018). In all cases, one kind of search applies known search operators (exploratory creativity, combinatorial creativity), and the other kind of search seeks to change the very operators themselves.

Exploratory and combinatorial creativity discover new things that may expand current capabilities, but they do not fundamentally change the way one thinks about a topic, only the current actions one might take to solve it. Transformational creativity changes how one thinks about a topic, and as such opens up more possibilities for the topic itself. Its potential to change things in a significant and global way is why transformational creativity is often thought as "Big-C" (Kaufman and Beghetto 2009) creative or "disruptive" (Schumpeter 1942). But transformational creativity happens all the time in real life. We might transform the way students think about topics (e.g., "think of negotiation as problem solving") or the way trainees think about tactics (e.g., "hold a golf club like you are holding a bird"). These, too, alter the way people approach even rudimentary topics and thus open up the potential for advanced understanding and action.

We also argue that, from the perspective of an AI researcher, exploratory creativity maps largely to search or optimization with a diversity criterion, and this has already been long studied. Thus we do not think that the development of exploratory creativity-inspired algorithms will benefit AI much more than optimization and search already have. But transformational creativity maps to *meta-level* optimization,

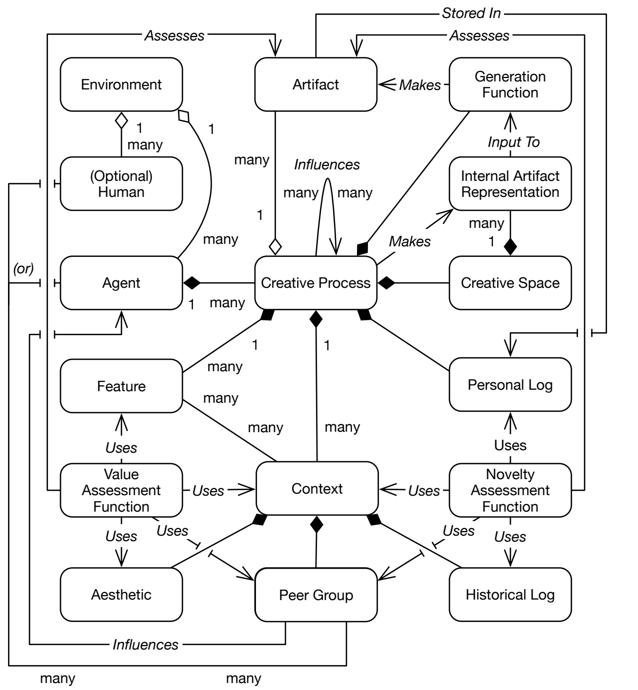

Figure 1: UML diagram of model classes and relationships, from Luke (2023).

and while there are a few examples of meta-level optimization in AI, particularly in stochastic optimization, it is still a poorly studied area rich with research opportunities where computational creativity could contribute.

Proposed Model

In prior work on exploratory creativity we criticized models which focused on search rather than on multiobjective optimization, and for not situating creativity in a dynamic and multiagent context (Luke 2023). We developed a model of exploratory creativity which connected these elements. Our model was similar to Wiggins in that it described a space of concepts, rules for value and novelty, and the notion of traversal: but we rejected the use of validity rules to constrain the space. The goal was to create a framework in order to define and compare computational creativity algorithms. We now extend this model to consider transformational creativity.

The original model defined an "exploratory" optimization space of artifacts called a *creative space*, the rough analog of Boden's concept space. We assume that this space has a *topology*, or neighborhood function, approximately defining the probability that some new prospective artifact r' would be discovered given past discovered artifacts $\langle r_1, ...r_n\rangle$. Note that this is different from the distance function used to assess novelty in Luke (2023). There are many things that could affect this probability, such as similarity, but also whether r' is just too exotic. Thus we can do away with a Wiggins-style constraint or validity function: for example, benzene was not *invalid* to Kekulé, but just too unusual a solution for him to likely see. Also note that while an agent could "transform" a space by changing the exploratory value function rather than

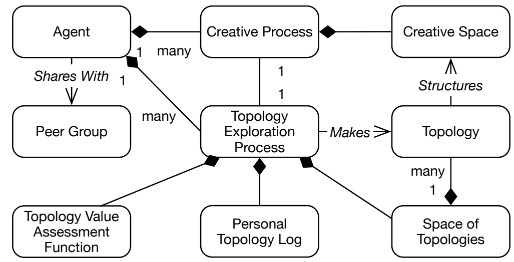

Figure 2: UML diagram of model extension to accommodate multiagent transformational creativity.

the topology, this would be more akin to formerly *rejecting* ideas rather than simply not *realizing* them due to biases.

We augment this model by giving each agent a meta-level "transformational" space: the space of topologies. The agent can, when appropriate, search this space to attempt to discover a new, high-quality topology, that is, one which encourages the exploratory optimizer to produce valuable and novel artifacts. As with Luke (2023), we define this space to be an optimization space, with its own neighborhood function, rather than as a state-space to search as in Demke and Ventura (2024). Wiggins (2006a) calls this just "exploratory creativity at the meta-level," but we note that this meta-level optimization process is likely much slower than exploratory-level optimization; and also the objective function is likely not the same. Perhaps a meta-level objective function would assess the quality of a topology as a maximum or weighted mean value over artifacts created so far in the exploratory space applying this topology. Or would it consider novelty?

Multiagent Transformational Creativity In our original model, an agent could receive discovered artifacts from a *peer group* of other agents or humans, which informed his sense of novelty and also potentially seeded the agent's exploratory optimization function. Conversely discoveries by an agent could spread to his peers, and their peers, inspiring them with new possible concepts or artifacts to consider.

But transformational creativity offers a very different additional seeding opportunity: agents may communicate *topologies* (in the form of problem representations, ways of viewing the problem, changes in bias) to one another. As an agent discovers a new and fruitful way of conceptualizing the problem, he can spread this to others. We model this as agents communicating topologies at the meta-level, thus seeding agents' meta-level topological optimization processes.

Model Details

First we introduce certain necessary concepts from our original model in (2023), but omit many of its other details. We use the nomenclature from the original model. The model describes an *environment* E consisting of one or more *agents*. The state of the various agents at time t may be described as $\mathbf{A}^t = \{A_1^t, ..., A_a^t, ...\}$. At time t each agent A_a^t employs a set of one or more *creative processes* $\mathbf{P}^{a,t} = \{P_1^{a,t}, ..., P_p^{a,t}, ...\} \subseteq \mathbb{P}$, which are optimization algorithms attempting to produce high-quality, novel *artifacts*.

The space of artifacts over which some optimization process $P_p^{a,t}$ operates is called a *creative space* $S^{a,p,t} \in \mathbb{S}$, and effectively serves as Boden's Concept Space for exploratory creativity. The creative process $P_p^{a,t}$ maintains $L^{a,p,t}$, a *personal log* of artifacts it has produced or is aware of up until time t. An agent shares concepts with other agents $\mathbf{A}^t$, and possibly with other humans $\mathbf{M}^t = \{M_1^t, ..., M_m^t, ...\} \subseteq \mathbb{M}$. The group with which it shares concepts is called its *peer group* $G^{a,p,c,t} \subseteq (\mathbf{A}^t \vee \mathbf{M}^t)$. This peer group is part of the *context* $\mathbf{C}^{a,p,t} = \{C_1^{a,p,t}, ..., C_c^{a,p,t}, ...\} \subseteq \mathbb{C}$ in which artifacts are assessed. Another part of the context is an *aesthetic* $Z^{a,p,c,t}$ which encapsulates all information used to assess the value of an artifact. Contexts and their aesthetics may change over time as an agent develops new opinions about artifacts.

Creative Space Topologies We augment the model by defining the notion of a space of *neighborhood topologies* $\{T_1, ..., T_i, ...\} \subseteq \mathbb{T}$. At time t a creative space $S^{a,p,t}$ is currently applying some topology T_i. The topology defines the neighborhood distribution of any given artifact in the concept space. A topology T_i might be modeled as a distribution $P(r|L^{a,p,t}), r \in S^{a,p,t}$ indicating the probability that artifact r would be discovered given the concepts in $L^{a,p,t}$.

The function *Quality*$(T_i, S^{a,p,t}, L^{a,p,t}) \rightarrow \mathbb{R}$ provides a quality assessment of the topology T_i as it has performed so far, based on the value and novelty of concepts in $L^{a,p,t}$ discovered using that topology. If need be, *Quality* could instead be multiobjective.

Topology Exploration For each creative process $P_p^{a,t}$ the agent maintains a parallel meta-level *topology exploration process* $P_p'^{a,t}$. This process explores through T, starting at T_i, until it finds some new topology T_j. It does this by considering T_i and its current *Quality*, plus all the topologies stored in the process's *log*, $L'^{a,p,t}$, of previously used topologies, along with their *Quality* assessments.

An agent is normally engaged in its primary exploration process: executing $P_p^{a,t}$ to find high-quality and novel artifacts. When the agent feels it is stuck, or has exhausted the space, or that it cannot make further useful progress, it may engage the meta-exploration process $P_p'^{a,t}$ to produce the new topology T_j. It then replaces T_i with T_j as the topology for $S^{a,p,t}$, adds T_j to $L'^{a,p,t}$, and continues executing $P_p^{a,t}$.

Topology and Aesthetic Sharing In Luke (2023) agents may *influence*, or be influenced by, other agents in their peer group by sharing artifacts with them. These new artifacts, and their value and novelty, can bias an agent's creative process.

In this revised model, agents may additionally influence other agents by sharing new *topologies*. If a topology T_i has proven fruitful to explore, an agent may share it with other members of its peer group $G^{a,p,c,t}$. Any other agent $A_{a'}^t \in G^{a,p,c,t}$ may then choose to apply T_i to its own creative space $S^{a',p',t}$, assuming that p' is a creative process similar to p.

Finally, we amend the revised model to allow an agent to share an aesthetic $Z^{a,p,c,t}$ with its peer group $G^{a,p,c,t}$, and to revise it based on ones shared from members of the group, thus influencing how the agent assesses the value of artifacts.

Discussion

Our purpose in developing this framework is to serve as a first step in developing useful algorithms inspired by transformational creativity, which we hope may ultimately prove more fruitful than exploratory creativity has. Classically many areas of artificial intelligence have drawn from processes in the sciences for inspiration in the development of useful algorithms and models. But what about computational creativity? Beyond basic models of creativity, there exist some actual algorithms which borrow from exploratory creativity principles to engage in iterative refinement (for example, Venkatesh et al (2025)). But we are struck with the paucity of such algorithms as opposed to, for example, the fields of evolutionary computation (EC) or neural networks.

Before we examine where transformational creativity might make an impact in AI, we first must acknowledge the possibility that computational creativity principles could in fact largely *be useless* in the development of AI algorithms. If creativity is primarily a tool to get around the restricted ability of humans to explore concept spaces due to our biases and limited working memory, and if computers *do not have these limitations nearly to the degree that humans do*, then the creative process may not be very useful to AI algorithms.

Even so we might ask if there are places where approaches inspired by transformational creativity may prove useful: for example, in overcoming limitations caused by restricted distributed communication in limited memory, in multiagent systems, in dynamically changing value functions, in environments with many local optima, or in meta-level optimization. Buchanan (2001) identified multiple examples of meta-level state-space search in the literature as possible targets for transformational creativity research. We think another obvious target would be stochastic optimization fields such as EC, which are rife with meta-optimization and multiagent optimization in a variety of guises. These include meta-evolutionary algorithms which optimize the parameters of other EC algorithms; iterated local search, which seeks optimize in the space of local optima; island models, in which multiple optimization processes seed one other over distributed networks; and grammatical evolution, in which solutions are discovered by interpreting them through a grammar (which can itself can be optimized and modified).

Finally, we might also examine how transformational multiagent creativity can be used in a *co-creative context*. For example, could a computer identify when a human is stuck, or ignoring obvious solution avenues due to bias, and suggest ways to reframe the problem to help him out? Can a human do the same for the computer, acting as a guide to nudge the computer as it performs an exploratory search?

Conclusion

Transformational computational creativity has been neglected, but it should be at the forefront of the computational creativity research area. Transformational creativity's meta-level optimization it is well positioned to have an impact on AI applications, and the notion of spread of topologies poses valuable insights in a multiagent context. We hope to see transformational creativity applied to real AI applications and to have a greater effect on the AI research field.

References

Boden, M. 1992. *The Creative Mind*. Abacus.

Buchanan, B. 2001. Creativity at the metalevel: AAAI-2000 presidential address. *AI Magazine* 22(3):13–28.

Colton, S., and Wiggins, G. 2012. Computational creativity: the final frontier. In *European Conference on Artificial Intelligence (ECAI)*.

Cronin, M. A., and Loewenstein, J. 2018. *The Craft of Creativity*. Stanford Business Books.

Demke, J., and Ventura, D. 2024. Overcoming algorithmic bias as a measure of computational creativity. In *International Conference on Computational Creativity (ICCC)*, 245–253.

Feldman, D. H.; Csikszentmihalyi, M.; and Gardner, H. 1994. *Changing the World: A Framework for the Study of Creativity*. Praeger.

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

Linkola, S., and Kantosalo, A. 2019. Extending the creative systems framework for the analysis of creative agent societies. In *International Conference on Computational Creativity (ICCC)*, 204–2111.

Luke, S. 2023. Computational creativity as dynamic, multiobjective, multiagent optimization. In *International Conference on Computational Creativity (ICCC)*.

Mumford, M., and Gustafson, S. B. 1988. Creativity syndrome: Integration, application, and innovation. *Psychological Bulletin* 103(1):27.

Newell, A.; Shaw, J. C.; and Simon, H. A. 1958. The processes of creative thinking. Technical Report P-1320, The RAND Corporation.

Ritchie, G. 2005. On transformational creativity. In *International Joint Conference on Artificial Intelligence (IJCAI)*.

Ritchie, G. 2006. The transformational creativity hypothesis. *New Generation Computing* 24:246–266.

Ritchie, G. 2012. A closer look at creativity as search. In *International Conference on Computational Creativity (ICCC)*.

Saunders, R., and Gero, J. 2001. The digital clockwork muse: A computational model of aesthetic evolution. In *Society for the Study of Artificial Intelligence and Simulation of Behavior (AISB)*.

Saunders, R. 2019. Multi-agent-based models of social creativity. In Veale, T., and Cardoso, F. A., eds., *Computational Creativity: The Philosophy and Engineering of Autonomously Creative Systems*. Springer. 305–326.

Schumpeter, J. A. 1942. *Capitalism, Socialism and Democracy*. New York: Harper and Brothers.

Simon, H. A., and Hayes, J. R. 1976. The understanding process: Problem isomorphs. *Cognitive Psychology* 8(2):165–190.

Sternberg, R. J., and Lubart, T. I. 1999. *The Concept of Creativity: Prospects and Paradigms*. Cambridge University Press. chapter 1.

Venkatesh, K.; Dunlop, C.; and Yanardag, P. 2025. CREA: A collaborative multi-agent framework for creative image editing and generation. In *NeurIPS*.

Ventura, D. 2011. No free lunch in the search for creativity. In *International Conference on Computational Creativity (ICCC)*.

von Helmholtz, H. 1896. *Vorträge und Reden*. Friedrich Vieweg.

Wallas, G. 1926. *The Art of Thought*. Harcourt, Brace and Company.

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

Wiggins, G. A. 2006b. Searching for computational creativity. *New Generation Computing* 24:209–222.