

Creativity and Markov Decision Processes

Joonas Lahikainen¹, Nadia M. Ady^{1,2} and Christian Guckelsberger^{1,3}

¹Department of Computer Science, Aalto University, Espoo, Finland

²Helsinki Institute for Information Technology, Helsinki, Finland

³School of Electronic Engineering and Computer Science, Queen Mary University of London, London, UK
christian.guckelsberger@aalto.fi

Abstract

Creativity is already regularly attributed to AI systems outside specialised computational creativity (CC) communities. However, the evaluation of creativity in AI at large typically lacks grounding in creativity theory, which can promote inappropriate attributions and limit the analysis of creative behaviour. While CC researchers have translated psychological theory into formal models, the value of these models is limited by a gap to common AI frameworks. To mitigate this limitation, we identify formal mappings between Boden’s process theory of creativity and Markov Decision Processes (MDPs), using the Creative Systems Framework as a stepping stone. We study three out of eleven mappings in detail to understand which types of creative processes, opportunities for (aberrations), and threats to creativity (uninspiration) could be observed in an MDP. We conclude by discussing quality criteria for the selection of such mappings for future work and applications.

Introduction

Since the inception of Artificial Intelligence (AI), researchers have sought to model creativity in artificial systems (McCarthy et al., 2006; Boden, 2015). Until recently, most Computational Creativity (CC) research has been driven by relatively small subfields of the larger AI community (Cook and Colton, 2018). But now, AI at large has progressed greatly, and creativity is now attributed to systems developed outside CC subfields, e.g., *AlphaGo* (Bory, 2019; Natale and Henrickson, 2022), with increasing frequency.

Echoing, e.g., Besold (2016), we hold that many instances of AI already exhibit creativity to an extent. To what extent though is an open question, as attributing creativity to such systems is usually based on intuitive judgement and not theoretically grounded.¹ This may promote incorrect attributions of creativity, including failure to recognise creativity in a system entirely. Additionally, it limits our ability to distinguish distinct types of creativity exhibited by systems, which may inhibit system development for specific purposes.

¹Intuitive judgement remains a valuable type of creativity evaluation (Colton and Wiggins, 2012; Natale and Henrickson, 2022), but we contend that structured and theoretically grounded accounts are at least equally important for supporting scientific progress.

The required systematic reflection on the creativity of AI can be supported by insights into how a system’s components and dynamics, *captured formally*, relate to creativity theory. Translational research between Psychology and AI has recently gained more traction (e.g., van Rooij et al., 2023; Lintunen et al., 2024; Ady et al., 2022). It is notoriously hard though, in part due to the challenge of interpreting informal theory formally. Drawing on psychological theory, CC researchers have developed a few formal frameworks for evaluating systems’ (potential) creativity (e.g., FACE, Colton, Charnley, and Pease, 2011; DeVER, Aguilar and Pérez y Pérez, 2015; CSF, Wiggins, 2019). However, most AI systems are not originally formalised through creativity frameworks. Assessing a given system’s creativity through theory thus involves considerable interpretation, which requires effort and expertise and may introduce inconsistency. Specifically, AI systems that engage in sequential decision-making are most prevalently formalised with Markov Decision Processes (MDPs) (Sutton, 1997, p. 273). A formal mapping between a creativity framework and agents in interaction with MDPs would immediately allow for the standardised analysis of numerous AI systems in terms of how they might be considered creative. To fulfil this purpose, mapping would need to express how each mathematical object involved in one formal framework could be understood as one of the mathematical objects in the other.

Here, we propose formal mappings between AI systems modelled using MDPs to Margaret Boden’s (2004) process theory of creativity. We chose Boden’s theory because it allows to distinguish different types of, obstacles to, and opportunities for creativity, while not assuming specific cognitive faculties. Moreover, it has already been formalised—and, enabled by formal rigour, further differentiated—twice in the Creative Systems Framework (CSF; Wiggins, 2006a,b, 2019; Ritchie, 2012), allowing us to use the CSF as a stepping stone for mapping Boden’s theory to MDPs.

Colin et al. (2016) made a pioneering effort to map between the CSF and hierarchical MDPs (see Related Work). We here extend their work by proposing mappings to the more general framework of MDPs, and additionally illustrate ambiguity in our possible mappings. Our contributions are: (1) We argue for MDPs as a minimal framework from which to map Boden’s theory (2) We identify eleven potential mappings between MDPs and the CSF and evaluate three

in detail; (3) We discuss which types of creative processes, uninspiration, and aberration could be observed in an MDP under the three mappings detailed; (4) We propose quality criteria to surface mapping issues and to support selecting from mapping candidates for future analytical work and applications. We thus follow the call made by Ady and Rice (2023) for CC researchers to explain choices made in selecting and interpreting definitions.

This paper thus does not conclude, but opens up an interdisciplinary research agenda, establishing a new formal basis to strengthen the dialogue between CC and Psychology with AI research at large. As such, it benefits and is written for an audience of multiple stakeholders. *CC researchers* receive the means to better analyse AI systems beyond CC to further the field’s engineering and cognitive research continuum (Pérez y Pérez, 2018). *AI researchers* from other fields can better understand the (potential) creativity of their systems, and means to foster it. Boden’s original theory and the CSF have been widely used for CC evaluation and system design, and extended in several ways (e.g., Wiggins and Forth, 2018; Linkola and Kantosalo, 2019; Linkola, Guckelsberger, and Kantosalo, 2020; Kantosalo and Toivonen, 2016). Our mappings to MDPs promise to make this CC research heritage more relevant and accessible to AI at large. Finally, our mappings can allow *psychologists* to simulate process theories of creativity in more diverse systems towards alleviating the formalisation and replication crises in Psychology (Oberauer and Lewandowsky, 2019).

Background

Boden’s Process Theory of Creativity

Theories of creativity have been conceived from four distinct perspectives (Rhodes, 1961; Jordanous, 2016): the *person* or *producer* as originator of the work, the *process* as the steps the producer undertakes when being creative, the *product* as the outcome, and the *press* as the sociocultural environment which shapes our views on, and the assessment of, creativity.

We focus on process theories, which, by treating how creative products are made, have been deemed central to understanding and supporting the evaluation of creativity in natural (e.g., Walia, 2019) and computational domains (e.g., Colton, 2008), and have been promoted for including “many important ideas that can and should influence the design of a creative system” (Lamb, Brown, and Clarke, 2018).

Crucially though, from Wallas’ (1926) classic four-stage model to Green et al.’s (2023) recent definition, many theories of the creative process assume human cognitive features such as attention or unconscious reasoning, limiting their applicability to artificial systems (Lamb, Brown, and Clarke, 2018). A notable exception is the theory of Margaret Boden, who distinguishes three types of creative *processes*: combinatorial, exploratory, and transformational. These rest on what Boden denotes a *conceptual space*: a structured way of thinking which both constrains and makes possible a particular variety of thoughts (2004, p. 58). It can be conceived as a space of all complete and incomplete things (including both mental concepts and physical artefacts) that could be generated according to a set of (agreed) rules.

Combinatorial creativity refers to the creation of new concepts by combining features of existing ones. *Exploratory creativity* corresponds to exploring the conceptual space for new and valued concepts. Finally, Boden introduces *transformational creativity* as reaching new points only made accessible by altering the rules defining the space itself “so that thoughts are now possible which previously (...) were literally inconceivable” (Boden, 2004, p. 6). Many process theorists consider transformational creativity as the most profound (summarised by Lamb, Brown, and Clarke, 2018), including Boden herself, who notes that transformational creativity is “the most arresting of the three” and relatively rare (Boden, 2013, pp. 6–8). For this reason, an important focus in our paper will be on how systems, as viewed through our mappings, might exhibit transformational creativity.

The Creative Systems Framework

Abstaining from assumptions of human cognitive facilities, Boden’s theory has unsurprisingly become the most popular process theory of creativity in CC (Lamb, Brown, and Clarke, 2018). Wiggins (2006a,b, 2019) formalised the theory in set-theoretic terms, hereby providing a more concrete interpretation and adding to Boden’s account. Ritchie (2012) later re-formulated and extended Wiggins’ framework while retaining backward compatibility. We rely on Ritchie’s formulation because it simplifies the formalism, dropping Wiggins’s reliance on a universal language and interpreters. Moreover, it has been used by Colin et al. (2016) in the only instance of related work; using the same framework thus also eases comparison. The definitions that follow are taken from Ritchie (2012) and Wiggins (2019), modified only for brevity or clarity; most of our notation is that of Ritchie (2012, see p. 43 for relation to Wiggins).

Definition 1 (General notation).

1. For any sets A and B , B^A denotes the set of functions from A to B . In particular, $[0, 1]^A$ denotes the set of functions from A to real values between 0 and 1, inclusive.
2. For any set A , $\text{tuples}(A)$ denotes the set of all finite tuples of elements in A ; e.g., if $A = \{1, 2, 3\}$, then $\text{tuples}(A) = \{1, 2, 3, (1, 1), (1, 2), \dots\}$.
3. For any set of tuples X , we define the set of distinct elements in the tuples of X as $\text{elements}(X) := \{x \mid \exists (y_1, \dots, y_n) \in X \wedge \exists i \in \{1, \dots, n\} : x = y_i\}$. This flattens a set of tuples into a set of distinct elements: for set A and $X = \text{tuples}(A)$, we have $\text{elements}(X) = A$.
4. For any set A , function $f \in [0, 1]^A$, and threshold $\alpha \in [0, 1]$, we define the strong α -cut of A as $f^{>\alpha}(A) := \{a \in A \mid f(a) > \alpha\}$.

To formalise transformational creativity, Wiggins introduced a set representing all conceivable concepts.

Definition 2 (Universe). *The universe, \mathcal{U} , is a set (specifically a multidimensional space; Wiggins, 2019, p. 26) capable of representing anything, and all possible distinct concepts correspond with distinct points in \mathcal{U} .*

Axiom 1 (Universality). All possible concepts, including the empty concept, \top , are represented in \mathcal{U} .

Axiom 2 (Non-identity of concepts). All concepts represented in \mathcal{U} are non-identical, meaning $\forall c_1, c_2 \in \mathcal{U}, c_1 \neq c_2$.

Axioms 3 & 4 (Universal inclusion). All conceptual spaces (3) are strict subsets of \mathcal{U} and (4) include \top .

Ritchie (2012, p. 43) adapts Wiggins’ formulation to define an *exploratory creative system*² as a 4-tuple:

Definition 3 (Exploratory Creative System, ECS). We define an ECS \mathcal{E} as a 4-tuple $(\mathcal{P}, \mathcal{N}, \mathcal{V}, \mathcal{Q})$ consisting of:

1. $\mathcal{P} \subseteq \mathcal{U}$, sub-universe
2. $\mathcal{N} \in [0, 1]^{\mathcal{P}}$, acceptability function
3. $\mathcal{V} \in [0, 1]^{\mathcal{P}}$, evaluation function
4. $\mathcal{Q} : [0, 1]^{\mathcal{P}} \times [0, 1]^{\mathcal{P}} \rightarrow \text{tuples}(\mathcal{P})^{\text{tuples}(\mathcal{P})}$, traversal strategy

A *conceptual space*, defined below, is the set of acceptable concepts out of a subset, \mathcal{P} , of all concepts. Restriction to subset $\mathcal{P} \subseteq \mathcal{U}$ (a *sub-universe*) allows us to make distinctions between systems w.r.t. their access to the full universe (Ritchie, 2012, p. 42).

Definition 4 (Conceptual space). For acceptability function $\mathcal{N} \in [0, 1]^{\mathcal{P}}$, acceptability threshold $\alpha \in [0, 1]$, and sub-universe $\mathcal{P} \subseteq \mathcal{U}$, we define the conceptual space \mathcal{C} as

$$\mathcal{C} := \mathcal{N}^{>\alpha}(\mathcal{P}) = \{c \in \mathcal{P} \mid \mathcal{N}(c) > \alpha\} \quad (1)$$

While, according to Wiggins, evaluation, \mathcal{V} , does not play into the *definition* of a conceptual space (see Critical Assumptions), value is a core requirement of creativity (e.g., Runco and Jaeger, 2012). Thus, not only the acceptability function, \mathcal{N} , but also the evaluation function, \mathcal{V} , influence the traversal strategy, \mathcal{Q} . The ECS starts at an initial concept and “searches” through the conceptual space by means of its traversal strategy. In the CSF, *exploratory creativity* corresponds to exploring the conceptual space \mathcal{C} for new concepts valued in terms of \mathcal{V} .

Wiggins (2006a, p. 454) offered a mechanistic explanation of transformational creativity by formalising it as exploratory creativity, but on a *meta* level. A meta-level creative system, $(\mathcal{P}^{\text{meta}}, \mathcal{N}^{\text{meta}}, \mathcal{V}^{\text{meta}}, \mathcal{Q}^{\text{meta}})$, searches a conceptual space in which the concepts are triples $(\mathcal{N}, \mathcal{V}, \mathcal{Q})$, with the sub-universe, $\mathcal{P}^{\text{meta}}$, being the set of all such triples (Ritchie, 2012, p. 44). Recursively, a creative system could have many such levels. The lowest level is considered the *object-level* system.

Boden (2004, p. 58) originally explained transformation as modifying the existing rules of a conceptual space to make possible concepts that were not possible before. Wiggins extended the notion of transformation to changing the traversal strategy, noting that such a change might “make accessible concepts which were not previously available” to the agent (2019, p. 32). We therefore consider:

- *\mathcal{N} -transformation*: modifying the conceptual space via changes to the acceptability function.
- *\mathcal{Q} -transformation*: changing the traversal strategy.

²We adopt the name “exploratory creative system” by convention. We do not want to convey the impression that such a system exhibits creativity at any moment. Instead, we hold that it has the *potential* to exhibit creativity as defined by Boden (2004).

For the purpose of assessing transformational creativity, we follow Wiggins’ (2019, p. 33) suggestion that a transformation is valued if, given a fixed object-level evaluation function \mathcal{V} , the transformation admits new concepts valued under \mathcal{V} , either to the set of reachable concepts (defined below) or the conceptual space itself. Following Boden’s (2004, p. 10) requirements for creativity to involve both novelty and value, we assume a transformation can only be considered transformational creativity if it is valued.

Wiggins (2006a, p. 456) additionally introduced *aberrations*, characterising the traversal strategy \mathcal{Q} reaching a set of concepts \mathcal{B} which lie outside the conceptual space (cf. Ritchie, 2012, p. 45). Aberrations could be used to trigger and guide transformations as opportunities for creativity.

Definition 5 (Set of reachable concepts). Let ECS $\mathcal{E} = (\mathcal{P}, \mathcal{N}, \mathcal{V}, \mathcal{Q})$. Then, starting from an initial tuple of concepts B , we denote the set of reachable concepts within \mathcal{E} in m steps by

$$\mathcal{E}^m(B) := \text{elements} \left(\bigcup_{n=0}^m \mathcal{Q}(\mathcal{N}, \mathcal{V})^n(B) \right) \quad (2)$$

where exponent n denotes repeated applications of traversal strategy \mathcal{Q} . Notably, B can be given by \top , i.e. the empty concept, corresponding to a blank canvas.

An aberration occurs when $\mathcal{B} := \mathcal{E}^\infty(B) \setminus \mathcal{C}$ is non-empty. Aberrations fall into three categories depending on how this set \mathcal{B} is valued, given threshold $\beta \in [0, 1]$:³ *perfect* ($\mathcal{V}^{>\beta}(\mathcal{B}) = \mathcal{B}$), *productive* ($\mathcal{V}^{>\beta}(\mathcal{B}) \subset \mathcal{B}$) and *pointless* ($\mathcal{V}^{>\beta}(\mathcal{B}) = \emptyset$) aberrations.

Wiggins furthermore characterises three cases of *uninspiration*, causing the system to fail at being creative (notation adopted from Ritchie, 2012, pp. 43–46):

- *Generative uninspiration*. When the traversal strategy does not find any valued concepts, $\mathcal{V}^{>\beta}(\mathcal{E}^\infty(\top)) = \emptyset$.
- *Conceptual uninspiration*. When there are no valued concepts in the conceptual space, $\mathcal{V}^{>\beta}(\mathcal{C}) = \emptyset$.
- *Hopeless uninspiration*. When there are no valued concepts in the sub-universe, $\mathcal{V}^{>\beta}(\mathcal{P}) = \emptyset$.⁴

A threat to creativity, uninspiration complements aberration.

Markov Decision Processes

A *Markov Process (MP)* (Markov, 1906; Levin and Peres, 2017) models stochastic *transitions* between *states* of a system. As per the *Markov property*, the probability of moving to a specific next state depends only on the current state. The structure of MPs often remains useful even if the Markov

³Crisp sets ease discussion, but yield binary descriptions of conceptual spaces, *uninspiration*, and *aberration*, contradicting Boden’s philosophy (2004, p. 2): “Rather than asking ‘Is that idea creative, yes or no?’ we should ask ‘Just how creative is it, and in just which way(s)?’” However, modifying the CSF is out of scope.

⁴Ritchie’s (2012, pp. 43, 45) re-definition of hopeless uninspiration is contradictory, leaving unclear whether it refers to a lack of valued concepts in the *sub-universe*, \mathcal{P} , or *universe*, \mathcal{U} .

property does not exactly hold, and so MPs are used abundantly across many academic fields to model the dynamics of natural and artificial systems, and especially within AI.

A *Markov Decision Process (MDP)* extends MPs by introducing a notion of agency in the form of *actions* that influence the transition between states. Actions are chosen based on the current state via a *policy*. Moreover, MDPs define *rewards* as scalar feedback from state transitions to define the problem to be solved with a suitable policy.

Definition 6 (Markov Decision Process). *We define a discrete-time MDP as a 4-tuple (S, A, T, R) , consisting of state space S , action space A , stationary transition probabilities T , and stochastic reward function R .*

In this paper, we chose to focus on discrete-time MDPs for its fit with the use of discrete “steps” in the CSF. We denote the probability of transitioning from state s to the next state s' with action a as

$$T_a(s, s') := P(s' | s, a) \quad (3)$$

with the intuition that each action results in a unique transition matrix indexed by states. We assume both the reward function $R : S \times A \times S \rightarrow \mathbb{R}$ and policy $\pi : S \times A \rightarrow [0, 1]$ to be stochastic, without loss of generality. While the stochastic policy is characterised by a probability distribution, our stochastic reward function directly samples rewards from an underlying reward distribution. For generality, the output of $R(s, a, s')$ need only depend on a subset of these inputs. Excluding the policy from Definition 6 lets us discuss transitioning through the same MDP with different policies.

The goal for the agent is to maximize its *return*. Return, G_t , can be formalized in multiple ways, but is a function of the rewards observed after time t : $R(s_t, a_t, s_{t+1}), R(s_{t+1}, a_{t+1}, s_{t+2}), \dots$. Different formulations of return include discounted or undiscounted sums of rewards or as the sum of differences between received rewards and the average reward (Sutton and Barto, 2018, pp. 54–55, 249–250); we do not assume a particular formulation here. Under a given policy, π , the return allows us to define the *value* of each state (*ibid.*, p. 54):

$$v_\pi(s) := \mathbb{E}_\pi [G_t | s_t = s] \quad (4)$$

The definition of return changes the problem the agent is set to solve, but in each case, solving the problem means finding an optimal policy π^* such that $v_{\pi^*}(s) \geq v_\pi(s)$ for all states $s \in S$ and all policies π (*ibid.*, p. 62).

Such an optimal policy π^* as one solution to the sequential decision-making problem defined by the MDP can be found with a plethora of techniques, e.g., from reinforcement learning (Sutton and Barto, 2018). Some techniques maintain an estimate of some $v_\pi(s)$, which we denote by V .

MDPs are part of a family of extensions to MPs, including Partially Observable Markov Decision Processes (POMDPs), which extend MDPs with partial observations of the state, and Hierarchical MDPs, which extend MDPs with temporally extended actions and a hierarchical structure. We focus here on MDPs for their generality and widespread application (Sutton, 1997, p. 273), but other members of this larger family may be worth considering for other mappings (cf. Colin et al., 2016, Related Work).

Mappings

In this paper, we seek to construct mappings from agents in interaction with MDPs to the CSF. AI systems employ numerous different algorithms and architectures, but with MDPs, one way of abstracting at least part of their agency⁵ or individuality is via their policy. For this reason, we model an agent in interaction with an MDP as a pairing between an MDP and a (potentially non-stationary) policy. **Our goal, then, is to understand any given MDP-policy pair as an ECS.** We thus aim to construct mappings of the form:

$$\begin{aligned} \mathcal{M} : & \left\{ (S, A, T, R, \pi) \mid \begin{array}{l} (S, A, T, R) \text{ an MDP,} \\ \pi \text{ a policy} \end{array} \right\} & (5) \\ & \rightarrow \left\{ (\mathcal{U}, \mathcal{P}, \mathcal{N}, \mathcal{V}, \mathcal{Q}) \mid \begin{array}{l} \mathcal{U} \text{ a universe,} \\ (\mathcal{P} \subseteq \mathcal{U}, \mathcal{N}, \mathcal{V}, \mathcal{Q}) \text{ an ECS} \end{array} \right\} \end{aligned}$$

Ideally, each mapping should be a total function, capable of mapping any MDP-policy pair to an ECS. We can then use the CSF to analyse any system that is modelled as a policy over an MDP. By virtue of this choice, we can retain the expressivity of Boden’s framework. This in turn enables studying the extent to which different types of creative processes, opportunities for, and obstacles to creativity are expressed in specific MDP instances, or the MDP framework in general.

We map agents in interaction with MDPs to Boden’s theory, but usually only capture agents w.r.t. their policy; In some cases, though, we also used the agent’s estimated value function (extending the 5-tuple in the domain shown in Equation 5 to a 6-tuple to include V). We were initially concerned that this choice would be limiting, as many algorithms do not use estimated value functions (see, e.g., Sutton and Barto, 2018, p. 321). However, we only use V in defining the evaluation function \mathcal{V} . Since \mathcal{Q} only can, but does not have to be, dependent on \mathcal{V} (Wiggins, 2019, p. 29), with careful choice of \mathcal{Q} , V can be safely ignored if it is not implemented in a system of interest. This allows the mapping to retain the semantic similarity of the value function to *evaluation*. Below, we explore in one mapping the use of rewards (which “determine the immediate, intrinsic desirability of environmental states”) and in two others the use of values (which “indicate the long-term desirability of states”) for their semantic similarity to the *value* component of evaluation (Sutton and Barto, 2018, p. 6).

Under the paradigm of sequential decision-making and creative autonomy (Jennings, 2010; Saunders, 2012), we assume exploration, transformation, aberration, and uninspiration to be driven by the agent, rather than a separate cause (such as the system designer). Therefore, in our interpretations of these dynamics, we only consider elements that the agent can change, notably excluding S, A, T , and R .

MDPs as Mapping Domain

MDPs are very widely used (Sutton and Barto, 2018, Section 1.7), but are an extension of a simpler process model: MPs. If we mapped MPs, we could map any MDP by extension. However, we determined that mapping MPs is not

⁵Defining agency in AI systems is an ongoing effort, cf. e.g., Biehl and Virgo (2022); Kenton et al. (2023).

appropriate since the CSF requires a means to *evaluate* concepts; defining an ECS requires a choice of evaluation function \mathcal{V} . Some authors consider evaluation to be a minimal requirement for creative autonomy (Jennings, 2010) and transformational creativity (Wiggins, 2019). Neither states nor transition probabilities (nor combination of them) have similar meaning to the evaluation function in the CSF.

Intuitively, MDPs as extensions of MPs appear as the most natural next mapping domain in that the evaluation function in the CSF has clear analogues in MDP defined in terms of the reward function. We recognise that further generalisations of MDPs such as POMDPs and Hierarchical MDPs would offer interesting viewpoints for mappings; however, they are in less widespread use than MDPs and their use would further expand the space of potential mappings, resulting in more candidate mappings to select from. Further benefits over MDPs in terms of semantic or formal connections to Boden’s theory, as well as Wiggins’ and Ritchie’s CSF are not apparent.

Mapping Procedure

Within the limited scope of this paper, we only discuss a selection of mappings without being exhaustive. Mirroring two-phase models of creativity (e.g., Kleinmintz, Ivancovsky, and Shamay-Tsoory, 2019), our mapping procedure involved two phases: a generation phase followed by an evaluation phase. In the generation phase, we explored potential mapping candidates, beginning with the decision of what aspect of the MDP and agent might be mapped to concepts (elements of \mathcal{U}). Table 1 lists all candidates we considered. In this phase, we considered not only the components listed on the left-hand side of Equation 5, but also policy-learning algorithms and their hyperparameters.

Mapping	Concepts $c \in \mathcal{C}$
\mathcal{M}_l	Policy-learning algorithms $l \in L$
\mathcal{M}_λ	Hyperparameters $\lambda_l \in \mathbb{R}^n$
$\mathcal{M}_{l,\lambda}$	$(l, \lambda_l) \in L \times \mathbb{R}^n$
\mathcal{M}_π	Policies $\pi \in \Pi$
* \mathcal{M}_s	States $s \in S$
\mathcal{M}_a	Actions $a \in A$
\mathcal{M}_r	Rewards $r \in \mathbb{R}$
$\mathcal{M}_{s,a}$	Tuples $(s, a) \in S \times A$
$\mathcal{M}_{s,a,r}$	Tuples $(s, a, r) \in S \times A \times \mathbb{R}$
* $\mathcal{M}_{s,a,s'}$	Transitions $(s, a, s') \in S \times A \times S$
* \mathcal{M}_τ	Trajectories $(s, a, s', \dots, s_{last}) \in Tr$

Table 1: Possible conceptual spaces identified in the exploratory stage. Starred * options were examined further.

In the evaluation phase, we filtered our initial candidates, aiming to keep a small but diverse sample to discuss in detail. In this phase, we followed our intuition about which candidates best match Boden’s theory and its extension via the CSF. Further below, we provide a detailed reflection on our intuitive choices, informing potential quality criteria for assessing mappings. Here, we only briefly express why we excluded eight of our original candidates. We dropped

mapping \mathcal{M}_π , for which exploratory creativity would correspond to the discovery of new and valued policies. This is essentially what MDP solvers do, but it resembles more what Wiggins and Ritchie describe as meta-level exploration, leaving ambiguity at the object-level. The same applies to \mathcal{M}_l , \mathcal{M}_λ and $\mathcal{M}_{l,\lambda}$ which were consequently also excluded. We excluded \mathcal{M}_a since many applications of MDPs use discrete and small action spaces, which would be quickly exhausted. We moreover excluded \mathcal{M}_r and $\mathcal{M}_{s,a,r}$, as understanding a reward as representing a concept does not seem intuitive. This left us with groups of mappings from state-action tuples or trajectories to concepts, from which we picked one each for diversity. We chose $\mathcal{M}_{s,a,s'}$ over $\mathcal{M}_{s,a}$ because we saw very natural mappings for the acceptability and evaluation functions, as the transition-probability function and reward function both take (s, a, s') as input. This leaves us with the mappings \mathcal{M}_s , $\mathcal{M}_{s,a,s'}$ and \mathcal{M}_τ , which we detail in Table 2 and the following subsections, and which form the basis of our discussion.

Mapping \mathcal{M}_s

Universe & Sub-universe. We map both *universe* \mathcal{U} and *sub-universe* \mathcal{P} to the union of all conceivable state spaces $\bigcup_S S$. Consequentially, concepts are mapped to states.

Acceptability & Conceptual Space. We map *acceptability* \mathcal{N} to a membership function $\mu : \bigcup_S S \rightarrow \{0, 1\}$ which outputs one if the input state belongs to S and zero otherwise. That is, this function describes the logic by which states are included in state space S . As a result, *conceptual space* \mathcal{C} is mapped to state space S , as long as the acceptability threshold $\alpha < 1$; if $\alpha = 1$, then $\mathcal{C} = \emptyset$ (a consequence of Def. 4).

Evaluation. We map *evaluation* \mathcal{V} to a normalisation $\hat{V} : \bigcup_S S \rightarrow [0, 1]$ of the agent’s estimated value function V .

Traversal Strategy. *Traversal strategy* \mathcal{Q} maps to a one-step rollout of policy $\pi: s' \sim T_{a \sim \pi(s)}(s, \cdot)$ for every state s in its input. That is, \mathcal{Q} takes a tuple of current states as input, and, for each state s , samples action a from $\pi(s)$, and outputs the following state s' to compose an output tuple. \mathcal{Q} traversing from tuples of concepts to tuples of concepts is directly from the CSF (Wiggins, 2019, p. 35).⁶

Transformations. \mathcal{N} -*transformation* would require the conceptual space \mathcal{C} to change, but \mathcal{C} corresponds to the state space S , which we assume cannot be modified by the agent. \mathcal{Q} -*transformation* maps to a change in policy π .

Aberration & Uninspiration. *Aberrations* are left unmapped since finding states outside of state space S is theoretically impossible. *Generative uninspiration* here reflects policy π performing poorly with respect to normalised value function \hat{V} , as none of the reachable states exceed value threshold β when evaluated by \hat{V} . Generative uninspiration may reflect either a truly suboptimal policy or a poorly estimated value function. *Conceptual uninspiration* suggests an ill-formed value function \hat{V} or value threshold β where no

⁶Appealingly, potential to traverse multiple rollouts at once may be suited to more general agents including asynchronous (Mnih et al., 2016) and model-based agents (Sutton and Barto, 2018, Ch. 8).

CSF	\mathcal{M}_s	$\mathcal{M}_{s,a,s'}$	\mathcal{M}_τ
Concept c	State s	Transition $\delta = (s, a, s')$	Trajectory $\tau = (s, a, s', \dots, s_{\text{last}})$
First-order definitions			
Universe \mathcal{U}	$\bigcup_S S$	$\bigcup_{S,A} S \times A \times S$	$\bigcup_{S,A} \text{tuples}(S \times A)$
Sub-universe $\mathcal{P} \subseteq \mathcal{U}$	$\bigcup_S S$	$S \times A \times S$	$\text{tuples}(S \times A)$
Acceptability \mathcal{N}	Membership function μ of S	$p_\delta(\delta) := T_a(s, s')\pi(a s)$	$p_\tau(\tau) := P(\tau s, \pi, T)$
Evaluation \mathcal{V}	$\hat{V} := \text{normalised } V$	$\hat{R} := \text{normalised } R$	$\bar{V} := \text{normalised } V^{(s_{\text{last}})}$
Traversal strategy $\mathcal{Q}(\mathcal{N}, \mathcal{V})$	$f : \text{tuples}(S) \rightarrow \text{tuples}(S)$	$f : \text{tuples}(\mathcal{P}) \rightarrow \text{tuples}(\mathcal{P})$	$f : \text{tuples}(\mathcal{P}) \rightarrow \text{tuples}(\mathcal{P})$
Higher-order definitions			
Conceptual space \mathcal{C}	$\mu^{>\alpha}(\mathcal{P})$	$p_\delta^{>\alpha}(\mathcal{P})$	$p_\tau^{>\alpha}(\mathcal{P})$
\mathcal{N} -transformation	—	$\pi \rightarrow \pi', \text{ s.t. } \mathcal{C} \neq \mathcal{C}'$	$\pi \rightarrow \pi', \text{ s.t. } \mathcal{C} \neq \mathcal{C}'$
\mathcal{Q} -transformation	$\pi \rightarrow \pi'$	$\pi \rightarrow \pi'$	$\pi \rightarrow \pi'$
Aberration	—	Reaching δ s s.t. $p_\delta(\delta) \leq \alpha$ and	Reaching τ s s.t. $p_\tau(\tau) \leq \alpha$ and
Perfect	—	all δ s are rewarding.	all τ s are valued.
Productive	—	some δ s are rewarding.	some τ s are valued.
Pointless	—	no δ s are rewarding.	no τ s are valued.
Uninspiration	No $s \in \hat{V}^{>\beta}(\mathcal{P})$	No $\delta \in \hat{R}^{>\beta}(\mathcal{P})$	No $\tau \in \bar{V}^{>\beta}(\mathcal{P})$
Generative	can be found with π .	can be found with π .	can be found with π .
Conceptual	exists in S .	has $p_\delta(\delta) > \alpha$.	has $p_\tau(\tau) > \alpha$.
Hopeless	exists in $\bigcup_S S$	exists in \mathcal{P} .	exists in \mathcal{P} .

Table 2: Mappings $\mathcal{M}_s, \mathcal{M}_{s,a,s'}$ and \mathcal{M}_τ . Notation: apostrophe $'$ indicates succession (state s' follows state s), “—” denotes no mapping, the union of all conceivable sets (excluding the union itself) is denoted as $\bigcup_X X$, $\mu : \mathcal{P} \rightarrow \{0, 1\}$ is the membership function for state space S , f is defined by outputting the result of a one-step rollout of policy π on each concept in its input. Higher-order definitions are derived from first-order definitions.

states in S are sufficiently valued. *Hopeless uninspiration* means no state in any MDP would be sufficiently valued.

Mapping $\mathcal{M}_{s,a,s'}$

Universe & Sub-universe. *Universe* \mathcal{U} comprises all conceivable state transitions in all potential MDPs, $\delta \in \bigcup_{S,A} S \times A \times S$. *Sub-universe* \mathcal{P} then maps to a narrowed down version $S \times A \times S$ defined by a particular S and A in the mapped MDP.

Acceptability & Conceptual Space. *Acceptability* \mathcal{N} maps to probability p_δ , where $p_\delta(\delta) := T_a(s, s')\pi(a|s)$. *Conceptual space* \mathcal{C} then contains transitions with $p_\delta(\delta) > \alpha$.

Evaluation. We map *evaluation* \mathcal{V} to the normalised reward function $\hat{R} : S \times A \times S \rightarrow [0, 1]$.

Traversal Strategy. *Traversal strategy* \mathcal{Q} maps transitions to the next transition triple, given a one-step rollout of policy π : that is, transition (s, a, s') maps to (s', a', s'') where $s'' \sim T_{a' \sim \pi(s')}(s', \cdot)$. As in \mathcal{M}_s , inputs and outputs can be tuples.

Transformations. Both \mathcal{N} -transformation and \mathcal{Q} -transformation map to changes in policy π . However, while any change to π corresponds to a \mathcal{Q} -transformation, \mathcal{N} -transformation additionally requires that the conceptual space changes. That is, the probability of some transition previously $\leq \alpha$ must now exceed it, or one previously $> \alpha$ must drop below or meet it.

Aberration & Uninspiration *Aberration* here refers to experiencing transitions which have at most probability α of occurring. These aberrations are then further categorised into *perfect*, *productive*, and *pointless*, depending on

whether all, only some, or none of these aberrant transitions produce rewards exceeding value threshold β .

Uninspiration here refers to a complete lack of transitions with rewards exceeding value threshold β , either due to a flawed policy π (*generative* and *conceptual*), or ill-defined state and action spaces S, A (*hopeless*).

Mapping \mathcal{M}_τ

Universe & Sub-universe. The *universe* \mathcal{U} includes finite trajectories $\tau = (s, a, s', \dots, s_{\text{last}})$ across all conceivable state and action spaces. As in $\mathcal{M}_{s,a,s'}$, the *sub-universe* \mathcal{P} comprises only trajectories from the mapped MDP.

Acceptability & Conceptual Space. We map *acceptability* \mathcal{N} to a conditional probability distribution p_τ over variable-length trajectories: $p_\tau(\tau) := P(\tau|s, \pi, T) = \prod_{s,a,s' \in \tau} \pi(a)T_a(s, s')$. Our *conceptual space* \mathcal{C} comprises trajectories with this probability larger than α .

Evaluation. We map *evaluation* \mathcal{V} to some normalisation \bar{V} of the agent’s estimated value of the final state in the trajectory, $V^{(s_{\text{last}})}$ such that $\bar{V} : \text{tuples}(S \times A) \rightarrow [0, 1]$.

Traversal Strategy. *Traversal strategy* \mathcal{Q} takes tuples of trajectories as input and outputs the same trajectories appended with one new action and state resulting from a one-step rollout of policy π . During traversal, the observation of each new state rules out some potential trajectories and changes the probabilities of others. Consequently, conceptual space \mathcal{C} changes during traversal in accordance with α .

Transformations. \mathcal{N} -transformation refers to a change to policy π such that conceptual space \mathcal{C} changes. When we change our policy, the set of possible trajectories can also change. \mathcal{Q} -transformation maps to any change to policy π .

Aberration & Uninspiration *Aberration* maps to experiencing unlikely trajectories given the trajectory so far, policy π , and transition probabilities T . A trajectory can have a low likelihood for two reasons: our policy can assign low probabilities to certain actions, or certain transition probabilities might be low according to T . Different types of aberration are categorised similarly as in mapping $\mathcal{M}_{s,a,s'}$: *perfect*, *productive*, and *pointless* aberration correspond to situations where all, only some, or none of the aberrant trajectories are valued by the normalised value function \bar{V} past the value threshold β . *Uninspiration* is mapped similarly: *generative*, *conceptual*, and *hopeless* uninspiration match with no valued trajectory reachable by policy π , with probability larger than α , or expressible with given state and action spaces S, A . Additionally, the value function \bar{V} may be flawed.

Quality of Mappings

Our mappings resulted from an iterative and intuitive process, gradually incorporating insights about how our mapping choices interplay and whether they retain compatibility with Boden’s original theory, the CSF, and MDPs. Based on this process, we formulated quality criteria to guide researchers’ choice of mappings. Our formulation was retrospective, non-exhaustive, and independent of additional theory or related work. We demonstrate below that each criterion can surface issues of the quality of a given mapping via application to \mathcal{M}_s , $\mathcal{M}_{s,a,s'}$, and \mathcal{M}_τ .

Semantic Similarity to Boden’s Original Theory

A mapping should preserve as much of the meaning established in Boden’s theory, once cast to the elements and dynamics of MDPs. Due to space limitations, we reflect on the central notion of *transformations* only. Boden (2004, p. 6) describes (\mathcal{N} -)transformations as the “deepest cases of creativity.” However, in \mathcal{M}_s , mapping the state space S to the conceptual space \mathcal{C} renders \mathcal{N} -transformations impossible, as the agent cannot modify S , contradicting Boden’s theory. To support generality in the design of our mappings, we imposed only minimal requirements (a policy) on our agent, and in this case found no natural way to allow changes to \mathcal{C} . Relaxing this design choice, a small change to \mathcal{M}_s could overcome the issue, mapping between the conceptual space and an agent’s current *model* of the state space (e.g., Lolos et al., 2017), rather than the MDP’s “objective” state space. In $\mathcal{M}_{s,a,s'}$ and \mathcal{M}_τ , \mathcal{N} -transformation refers to a change in the policy that changes the conceptual space. Thus, a conceptual space is partly determined by the policy. This association is semantically similar to Boden’s description of conceptual spaces as a *style* of thinking: “actions might be totally mental” (Sutton, 1997, p. 275), and a policy reacts to the situation and past experiences to decide what to do next.

Functional Similarity to the CSF

In the CSF, acceptability \mathcal{N} and evaluation \mathcal{V} are parameters to traversal strategy \mathcal{Q} . This gives \mathcal{Q} “awareness” of conceptual space \mathcal{C} and enables decision-making based on evaluation \mathcal{V} . In \mathcal{M}_s , acceptability \mathcal{N} is formalised as a membership function μ to state space S . \mathcal{Q} , then, does not make use of \mathcal{N} , as all possible outputs will themselves be states, and therefore elements of the conceptual space. In contrast,

evaluation \mathcal{V} , mapped to value function \hat{V} , quite reasonably informs \mathcal{Q} , as policies are commonly inferred from value functions. A problem regarding \mathcal{N} as a parameter to \mathcal{Q} also exists in $\mathcal{M}_{s,a,s'}$ and \mathcal{M}_τ : acceptability \mathcal{N} is rarely available (Sutton, 1997, p. 275) to the agent, which often only implicitly “experiences” the transition probabilities T while rolling out a policy. Another concern with $\mathcal{M}_{s,a,s'}$ and \mathcal{M}_τ is the dependency between \mathcal{N} - and \mathcal{Q} -transformation: the former forces the latter, excluding the possibility of changing our acceptability \mathcal{N} without change to traversal strategy \mathcal{Q} . In the CSF, these two are independent components. Even though the behaviour of \mathcal{Q} will be impacted by a change in \mathcal{N} or \mathcal{V} , the traversal strategy itself will not change.

Compatibility Between Frameworks

Rewards and values are arguably the most semantically aligned choice with the CSF evaluation function. However, in contrast to the latter, they are not meant to be squashed into the unit interval; their potential to extend the real line allows them to flexibly represent both real-world quantities and challenges. We leave open *how* to normalise these quantities, introducing undesirable ambiguity. As another compatibility issue, in \mathcal{M}_s , *aberrations* are left unmapped due to policies operating strictly inside state spaces.

Applicability Across Diverse AI Systems

In \mathcal{M}_s , acceptability \mathcal{N} is defined as the membership function μ , outputting binary values. For some systems, a continuous μ may be more appropriate; e.g., fuzzy definitions of state space S have been used (Buşoniu et al., 2010). Alternatively, acceptability \mathcal{N} could correspond to a probability of a state belonging to state space S , perhaps combined with a notion of confidence (Grace and Maher, 2015).

Identifying Types of Creativity

We aim to identify *exploratory* and *transformational creativity* in our mappings. For brevity, “sufficiently valued” concepts here mean concepts that pass value threshold β .

Exploratory Creativity (finding novel valued concepts)

- \mathcal{M}_s : finding novel states with sufficient estimated value.
- $\mathcal{M}_{s,a,s'}$: finding novel transitions with sufficient reward.
- \mathcal{M}_τ : finding novel trajectories with sufficient estimated value for the final state.

Transformational Creativity

(finding \mathcal{N} and \mathcal{Q} that admit novel valued concepts)

- \mathcal{M}_s : \mathcal{N} -transformational creativity (\mathcal{N} -TC) requires redefining the membership function μ such that the state space S admits new states with sufficient estimated value.
- \mathcal{Q} -transformational creativity (\mathcal{Q} -TC) requires updating policy π such that new valued states are found.
- $\mathcal{M}_{s,a,s'}$: \mathcal{N} -TC here is updating policy π such that new sufficiently rewarded transitions are found, but with the additional requirement that the set of transitions that pass probability α , i.e. the conceptual space \mathcal{C} , changes. \mathcal{Q} -TC works similarly but without said requirement.
- \mathcal{M}_τ : \mathcal{N} -TC and \mathcal{Q} -TC are the same as in mapping $\mathcal{M}_{s,a,s'}$ but instead of transitions we have entire trajectories.

Critical Assumptions

This research required a very close reading and comparison of Boden’s original account and Wiggins’/Ritchie’s formalisations, highlighting several inconsistencies.

Boden (2004, p. 4) defines a conceptual space as “any disciplined way of thinking that is familiar to (and valued by) a certain social group”, suggesting that all concepts in the space are considered by default valuable. Wiggins and Ritchie in contrast assume that membership in the conceptual space is only conditional on the *typicality* of the concept/artefact, and that value is evaluated separately through \mathcal{V} (e.g., Wiggins, 2019, p. 28). We adopt this view for its useful distinction between typicality and value.

We also consider Wiggins (2019, p. 28) and Ritchie (2012, p. 44) describing \mathcal{V} as defining “value” problematic; it suggests the evaluation is *only* of value, leaving open how the other core components of creativity, most importantly *novelty* but also *surprise* (Boden, 2004, p. 1), are assessed. Here, we assume that \mathcal{V} may or may not evaluate these core components. In neither mapping do we specify which components the reward or value function mapped to evaluation capture; consequently, we leave it open whether traversal under this mapping actively promotes exploratory creativity.

Related Work

We relate our contributions to Colin et al.’s (2016) work as the only close predecessor in terms of goals and methods. They mapped the CSF as formalised by Ritchie (2012) to hierarchical MDPs as instantiated in the options framework proposed by Sutton, Precup, and Singh (1999). Their object-level conceptual space comprises policies, evaluation is mapped to a discounted return function and traversal is mapped to a policy update function. Via online learning (e.g. standard temporal difference learning), an agent traverses the space of possible policies. On the meta-level, it traverses pairs of discounted return and policy update functions, with meta-level evaluation assessing the value of a given pair for solving the problem at hand. The policy update function is constrained to evolving policies within a single option only, i.e. policies that start from a specific set of initiation states, and end on a specific termination condition. Traversal over these pairs changes which option and underlying space of potential policies is used to tackle the problem at hand.

Colin et al.’s object-level thus matches our proposed mapping \mathcal{M}_π , which we did not pursue further as we considered it a meta-level interpretation. Crucially though, the CSF allows for many meta-levels and Colin et al. thus complement the present work by contributing a candidate mechanism to facilitate transformational creativity on the space of policies by extending MDPs to include behavioural hierarchies. Crucially though, the focus on hierarchical MDPs and the options framework as only one (albeit popular) formalisation thereof limits the wider applicability of their findings. We, in contrast, map MDPs as a more general but less expressive framework to increase the applicability of our findings while understanding MDPs’ limitations w.r.t. modelling creativity.

As a second distinction, Colin et al. propose only a single mapping, not motivated beyond its relevance to robot con-

trol. In contrast, we discuss three of eleven potential mappings in depth, complemented with domain-agnostic quality criteria for probing and comparing different mappings.

As a third and final distinction, Colin et al. focus on computational models of insight, corresponding to transformational creativity, but do not integrate uninspiration as threats to, and aberration as opportunities for creativity. Related, they do not consider \mathcal{N} part of the explored meta-level triplets, hence providing a mapping for \mathcal{Q} -transformation and revisions to \mathcal{V} only. Despite being interested in insight, they thus do not explicitly address \mathcal{N} -transformation as the most widely known and accepted type of transformational creativity, and only instance presented by Boden (2004) originally. We interpret aberration, uninspiration and transformational creativity under each mapping, thus embracing Boden’s theory and the CSF more comprehensively.

Conclusion & Future Work

We have set out to provide theoretically grounded tools for the assessment of creativity in AI systems at large. To this end, we put forward eleven, and discuss in detail three, formal mappings between agents in interaction with Markov Decision Processes (MDPs) and Margaret Boden’s process theory of creativity, using the Creative Systems Framework (CSF) as a stepping stone. We leveraged these three mappings to reflect on the types of, opportunities for, and threats to creativity conceivable in a system formalised on an MDP.

Our findings motivate exploration of further mappings as the most imminent future work, which will be supported by our critical reflection on quality criteria, and our discussion on the constraints on potential mappings imposed by formal features of MDPs and the CSF. The influence of hyperparameters α and β must also be investigated. This object-level effort should be complemented with formalising the corresponding meta-levels to highlight not only which forms of creative transformations are possible, but also how they can be brought about. This can for instance shed more light on the mechanisms behind policy changes and, in consequence, different forms of transformational creativity. Following this exploratory research, the best mapping candidates should be further tested through application to existing systems that have been attributed creativity, or that realised major AI milestones for which creativity is commonly considered necessary. Our present focus on MDPs allowed us to define potential mappings by trading off simplicity and widespread applicability; as another avenue for future work and prerequisite for evaluating the creativity of more constrained but realistic systems, e.g., with partial observability and a need to learn models of the world, we recommend leveraging this foundation to integrate creativity theory and more general—but also more complex—sequential decision-making frameworks such as POMDPs.

The synthesis of established creativity theory and AI frameworks can significantly enhance our understanding of, and consequently the potential for, creativity in AI. We invite researchers from Psychology and AI more generally and CC specifically to join this interdisciplinary effort.

References

- Ady, N. M., and Rice, F. 2023. Interdisciplinary methods in computational creativity: How human variables shape human-inspired AI research. In *Proceedings of the International Conference on Computational Creativity*.
- Ady, N. M.; Shariff, R.; Günther, J.; and Pilarski, P. M. 2022. Five properties of specific curiosity you didn't know curious machines should have. *arXiv preprint arXiv:2212.00187*.
- Aguilar, W., and Pérez y Pérez, R. 2015. Dev er: A computational model of early cognitive development as a creative process. *Cognitive Systems Research* 33:17–41.
- Besold, T. R. 2016. The unnoticed creativity revolutions: Bringing problem-solving back into computational creativity. In *Proceedings of the AISB International Symposium on Computational Creativity*, 1–8. CRC Press, Taylor & Francis Group.
- Biehl, M., and Virgo, N. 2022. Interpreting systems as solving POMDPs: a step towards a formal understanding of agency. In *International Workshop on Active Inference*, 16–31. Springer.
- Boden, M. A. 2004. *The Creative Mind: Myths and Mechanisms*. London, UK: Routledge. First published in 1990.
- Boden, M. A. 2013. Creativity as a neuroscientific mystery. In Vartanian, O.; Bristol, A. S.; and Kaufman, J. C., eds., *Neuroscience of Creativity*. MIT Press. 3–18.
- Boden, M. A. 2015. How computational creativity began. In Besold, T. R.; Schorlemmer, M.; and Smaill, A., eds., *Computational Creativity Research: Towards Creative Machines*. Atlantis Press. 5–13.
- Bory, P. 2019. Deep new: The shifting narratives of artificial intelligence from Deep Blue to AlphaGo. *Convergence* 25(4):627–642.
- Buşoniu, L.; Ernst, D.; De Schutter, B.; and Babuška, R. 2010. Approximate dynamic programming with a fuzzy parameterization. *Automatica* 46(5):804–814.
- Colin, T. R.; Belpaeme, T.; Cangelosi, A.; and Hemion, N. 2016. Hierarchical reinforcement learning as creative problem solving. *Robotics and Autonomous Systems* 86:196–206.
- Colton, S., and Wiggins, G. 2012. Computational creativity: The final frontier? In *Proc. ECAI*, volume 12, 21–26.
- Colton, S.; Charnley, J. W.; and Pease, A. 2011. Computational creativity theory: The face and idea descriptive models. In *Proceedings of the International Conference on Computational Creativity*, 90–95. Mexico City.
- Colton, S. 2008. Creativity versus the perception of creativity in computational systems. In *AAAI Spring Symposium: Creative Intelligent Systems*, volume 8, 7. Palo Alto, CA.
- Cook, M., and Colton, S. 2018. Neighbouring Communities: Interaction, Lessons and Opportunities. In *Proceedings of the International Conference on Computational Creativity*, 256–263. Association for Computational Creativity.
- Grace, K., and Maher, M. L. 2015. Specific curiosity as a cause and consequence of transformational creativity. In *Proceedings of the International Conference on Computational Creativity*, 260–267.
- Green, A. E.; Beaty, R. E.; Kenett, Y. N.; and Kaufman, J. C. 2023. The process definition of creativity. *Creativity Research Journal* 1–29.
- Jennings, K. E. 2010. Developing creativity: Artificial barriers in artificial intelligence. *Minds and Machines* 20:489–501.
- Jordanous, A. 2016. Four PPPerspectives on computational creativity in theory and in practice. *Connection Science* 28(2):194–216.
- Kantosalo, A., and Toivonen, H. 2016. Modes for creative human-computer collaboration: Alternating and task-divided co-creativity. In *Proceedings of the International Conference on Computational Creativity*, 77–84.
- Kenton, Z.; Kumar, R.; Farquhar, S.; Richens, J.; MacDermott, M.; and Everitt, T. 2023. Discovering agents. *Artificial Intelligence* 322:103963.
- Kleinmintz, O. M.; Ivancovsky, T.; and Shamay-Tsoory, S. G. 2019. The two-fold model of creativity: the neural underpinnings of the generation and evaluation of creative ideas. *Current Opinion in Behavioral Sciences* 27:131–138. Creativity.
- Lamb, C.; Brown, D. G.; and Clarke, C. L. 2018. Evaluating computational creativity: An interdisciplinary tutorial. *ACM Computing Surveys (CSUR)* 51(2):1–34.
- Levin, D. A., and Peres, Y. 2017. *Markov Chains and Mixing Times: Second Edition*. American Mathematical Society.
- Linkola, S., and Kantosalo, A. 2019. Extending the creative systems framework for the analysis of creative agent societies. In *Proceedings of the International Conference on Computational Creativity*. Association for Computational Creativity.
- Linkola, S.; Guckelsberger, C.; and Kantosalo, A. 2020. Action selection in the creative systems framework. In *Proceedings of the International Conference on Computational Creativity*.
- Lintunen, E. M.; Ady, N. M.; Guckelsberger, C.; and Deterding, S. 2024. Advancing self-determination theory with computational intrinsic motivation: The case of competence. osf.io/preprints/psyarxiv/4cbuv.
- Lolos, K.; Konstantinou, I.; Kantere, V.; and Koziris, N. 2017. Adaptive state space partitioning of Markov decision processes for elastic resource management. In *Proceedings of the IEEE International Conference on Data Engineering*, 191–194.
- Markov, A. A. 1906. Extension of the law of large numbers to dependent events. *Izvestiia Fiz.-Matem. Obsch. Kazan Univ.* 2:135–156. (In Russian).
- McCarthy, J.; Minsky, M. L.; Rochester, N.; and Shannon, C. E. 2006. A proposal for the Dartmouth summer

- research project on artificial intelligence. *AI Magazine* 27(4):12–12. First published in 1955.
- Mnih, V.; Badia, A. P.; Mirza, M.; Graves, A.; Lillicrap, T.; Harley, T.; Silver, D.; and Kavukcuoglu, K. 2016. Asynchronous methods for deep reinforcement learning. In Balcan, M. F., and Weinberger, K. Q., eds., *Proceedings of the International Conference on Machine Learning*, volume 48, 1928–1937.
- Natale, S., and Henrickson, L. 2022. The Lovelace effect: Perceptions of creativity in machines. *New Media & Society* 14614448221077278.
- Oberauer, K., and Lewandowsky, S. 2019. Addressing the theory crisis in psychology. *Psychonomic Bulletin & Review* 26:1596–1618.
- Pérez y Pérez, R. 2018. The computational creativity continuum. In *Proceedings of the International Conference on Computational Creativity*.
- Rhodes, M. 1961. An Analysis of Creativity. *The Phi Delta Kappan* 42(7):305–310.
- Ritchie, G. D. 2012. A closer look at creativity as search. In *Proceedings of the International Conference on Computational Creativity*, 41–48. Open University Press.
- Runco, M. A., and Jaeger, G. J. 2012. The standard definition of creativity. *Creativity Research Journal* 24(1):92–96.
- Saunders, R. 2012. Towards autonomous creative systems: A computational approach. *Cognitive Computation* 4:216–225.
- Sutton, R. S., and Barto, A. G. 2018. *Reinforcement Learning: An Introduction*. Massachusetts, USA: MIT Press.
- Sutton, R. S.; Precup, D.; and Singh, S. 1999. Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial intelligence* 112(1-2):181–211.
- Sutton, R. S. 1997. On the significance of Markov decision processes. In Gerstner, W.; Germond, A.; Hasler, M.; and Nicoud, J.-D., eds., *Artificial Neural Networks — ICANN’97*, 273–282. Berlin, Heidelberg: Springer Berlin Heidelberg.
- van Rooij, I.; Guest, O.; Adolphi, F. G.; de Haan, R.; Kolokolova, A.; and Rich, P. 2023. Reclaiming AI as a theoretical tool for cognitive science. osf.io/preprints/psyarxiv/4cbuv.
- Walia, C. 2019. A dynamic definition of creativity. *Creativity Research Journal* 31(3):237–247.
- Wallas, G. 1926. *The Art of Thought*, volume 10. Harcourt, Brace.
- Wiggins, G. A., and Forth, J. 2018. Computational creativity and live algorithms. In *The Oxford Handbook of Algorithmic Music*. Oxford University Press.
- 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 Generation Computing* 24:209–222.
- Wiggins, G. A. 2019. A framework for description, analysis and comparison of creative systems. In *Computational Synthesis and Creative Systems*. Springer International Publishing. 21–47.