From Computational Creativity to Creative Problem Solving Agents

Evana Gizzi[*,1], **Lakshmi Nair**[*,2], **Jivko Sinapov**[1], **Sonia Chernova**[2]
[1] Department of Computer Science, Tufts University, MA, USA
[2] College of Computing, Georgia Institute of Technology, GA, USA
Evana.Gizzi@tufts.edu, lnair3@gatech.edu, Jivko.Sinapov@tufts.edu, chernova@cc.gatech.edu

Abstract

Creative problem solving (CPS) is a skill that can greatly improve resourcefulness and adaptability of existing artificial intelligence (AI) systems. In this paper, we discuss how CPS leverages theoretical aspects from Computational Creativity (CC) and planning in AI. We present a definition of CPS, and discuss how CPS is achieved using aspects of CC and problem solving in AI.

Introduction

The Apollo 13 incident of 1970 is an example of how human ingenuity and creativity saved the lives of the three astronauts on-board. In order to combat the increasing carbon dioxide levels in the spacecraft, the astronauts crafted a carbon dioxide filter using available objects (Cass 2005). Similar capabilities are currently beyond the scope of existing artificial agents. In this paper, we focus on Creative Problem Solving (CPS) - a skill that can greatly improve resourcefulness of existing artificial intelligence (AI) systems. We discuss how CPS adapts theoretical aspects from Computational Creativity (CC) and general problem solving in AI, thus combining the two. We focus specifically on agents that plan and learn over states and actions, in the context of creative problem solving. We adapt terminologies frequently used in the planning and learning literature in AI, and link them to theoretical aspects in CC. While there have been existing efforts at formalizing CPS, the formalizations have focused specifically on *either* AI (such as classical planning (Sarathy 2018; Erdogan and Stilman 2013)), or perspectives in CC (such as concept re-representation (Olteţeanu 2015; 2014)). In contrast, we present a formalization of CPS by adapting aspects of problem solving from *both* AI and CC. This interdisciplinary approach allows us to take a holistic perspective on CPS, to stimulate further research in the area.

Definition of Creative Problem Solving

We begin by defining the components of a problem to be solved by an agent acting in its environment through planning or learning. The planning or learning problem specification in AI typically consists of a task goal G to be accomplished, given a set of environment states S and agent actions A. The

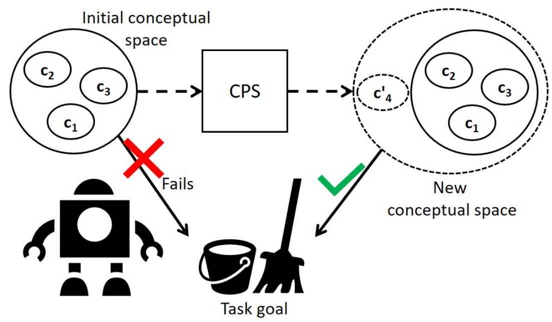

Figure 1: Creative Problem Solving (CPS) occurs when the initial conceptual space of the agent is insufficient to complete the task, and the agent needs to expand its conceptual space to achieve the task goal. Traditional planning or learning in AI would return a failure in such scenarios.

agent produces a task plan or policy Π over the states and actions, in order to accomplish a goal. Thus, $\Pi : S \rightarrow A$, represents a mapping from the set of environment states, to the set of actions.

In creative problem solving, we broadly define the notion of a *concept*, as a state, action, or a task plan or policy. More specifically, depending on the problem formulation, concepts could refer to the actions that an agent can perform (e.g., “open”, or “close”), the task plan or policy generated via planning or learning approaches, or the state space including states of objects in the agent’s environment (e.g., “clean”, “dirty”), and the state of the agent (e.g., configuration of the agent’s joints). While existing problems in AI focus on one or more of these aspects (actions, states, task plan or policy), grouping them under the term “concepts” allows us to unify the range of problem formulations within a single definition.

More generally, we denote concepts as c^X, and *conceptual space* C^X as the set of all concepts c^X, where X is one of $\{\Pi, S, A\}$. More specifically, C^Π denotes the set of all task plans or policies c^π, C^S denotes the set of all states c^S, and C^A denotes the set of actions c^A. Furthermore, let C^{X^*} denote the universal set of concepts c^X, such that C^{X^*} contains every possible concept that the agent could know, e.g., the set of all possible actions. In this work, we assume that the initial conceptual space $C^X \subset C^{X^*}$, i.e., the agent’s

[*]Both authors contributed equally

Proceedings of the 11th International Conference on Computational Creativity (ICCC'20)
ISBN: 978-989-54160-2-8

initial knowledge is limited. Note that $C^X = C^{X^*}$ is often not a practical assumption for real-world agents, since it implies that the agent knows every concept possible.

A crucial aspect of CPS that differentiates it from general problem solving is that the initial conceptual space C^X known to the agent is *insufficient* to accomplish the task goal. Traditional planning or learning approaches in AI often yield a failure in these circumstances. For example, the agent might need to learn about a completely new state or action to be able to accomplish the task. Thus, CPS is characterized by its *flexibility or adaptability* to handle novel problems (Guilford 1967). We formally define CPS as (See figure 1):

Definition 1 *Creative problem solving is defined as the process by which the agent discovers new concepts that were not in the initial conceptual space of the agent, allowing it to accomplish a previously impossible goal. Formally, CPS refers to the process by which the agent discovers new concepts $c'^X \in C'^X$, where $C'^X \nsubseteq C^X$, such that C'^X enables the agent to solve a task that was previously unsolvable when only C^X was known.*

In other words, the space of concepts that is explicitly represented by the agent defines the boundaries of what the agent can plan to accomplish. Creativity arises when the agent uses what it already knows, to discover something new. In the context of problem solving, the newly discovered knowledge is applied to solve a previously impossible task. In the following sections, we expand upon our definition, highlighting the theoretical aspects in CC that apply to CPS.

Adaptations of Aspects in Computational Creativity to Creative Problem Solving

Although there is no widely accepted singular definition for computational creativity (Jordanous 2012), there do exist generally accepted aspects. In this section, we review four major aspects of CC, and their inheritance and adaptations to creative problem solving. The aspects listed below are grouped into two categories; *output-based aspects* and *process-based aspects*. These categorizations are not meant to divide types of systems, but rather, to group key aspects.

Output-based Aspects

In output-based aspects, the focus is on evaluating the creativity of a system by determining whether the output produced in a task is considered to be creative. These systematic outputs, referred to as *artefacts* in the CC community, may take physical and/or non-physical form (e.g. paintings, songs, recipes). The first aspect (*Novelty* and *Value*) describes two key characteristics of a creative output, whereas the second aspect (*Evaluative Methods*) describes methods for making the evaluations.

Novelty and Value: This first aspect stems from the work of Margaret Boden, who proposed that creativity necessitates both novelty and value (Boden 1998). Novelty guarantees that the generated outputs of a creative process are in some way original, whereas the value criteria ensures that the generated outputs are not just random generations, but in some way geared towards a creative goal. Both novelty and value have contextual considerations. For example, an agent may produce a novel painting, but in the context of a scenario which calls for a creative recipe, the novel painting would not be considered valuable (Sosa and Gero 2016; Varshney, Wang, and Varshney 2016).

Evaluative Methods: In evaluative methods, the creativity of a system is evaluated by judging the output of its processes, deeming them to be creative or not creative. Similar in nature to the Turing test, these methods focus on using the judgement of an observer on the product of a creative process. The evaluation can either happen computationally (Colin et al. 2016; Colton, Wiggins, and others 2012; Varshney, Wang, and Varshney 2016), from a human evaluator (Bishop and Boden 2010; Guckelsberger, Salge, and Colton 2017), or from a social group (Varshney, Wang, and Varshney 2016). The nature of these evaluations vary, where in some cases the output is compared to a human's creative output, and in others, the output is judged in a social context. Additionally, the evaluation of a creative output can be determined by the agent's ability to explain its own intentions and motivations to a human (Cook et al. 2019).

CPS Adaptation of Output-based Aspects: In the context of problem solving, the novelty criteria is fully inherited. Creative solutions may not be completely original themselves, but rather in their *application* to the problem. For example, using Tupperware as a container may not be original in itself, but using Tupperware as a replacement for a soap dish may be considered a creative solution to a problem. Formally, a concept c'_X is said to be novel when it is not contained within the initial conceptual space C^X of the agent for that problem, i.e., $c'_X \notin C_X$. The second criteria in CC is that creativity necessitates value. In the context of CPS, this criteria is inherited as *usefulness* or *utility*. That is, does the solution actually solve the problem? A conceptual space C'_X is said to be useful, when the goal state G can be accomplished via concept(s) $c'_X \in C'_X$.

CPS does not directly inherit evaluative methods. This is because the output of a CPS process is simply evaluated as either successful or not successful, based on its ability to solve the problem. As such, a successful solution to a problem which necessitates CPS is inherently creative. Thus, evaluation in CPS involves evaluating whether the new conceptual space C'_X is sufficient to accomplish the current goal.

Process-based Aspects

Process-based aspects focus on the method by which creative output is generated. Contrary to output-based, process-based aspects are concerned with the question of *how* outputs are produced as opposed to the evaluation of *what* is produced. The first aspect (*Procedural Methods*) reviews existing methods for synthesizing the creative process, whereas the second aspect (*Boden's Types of Creativity*) reviews three ways of implementing procedural methods.

Procedural Methods: In procedural methods, the focus lies on evaluating creativity based on the method by which it systematically generates its creative output. These methods have been explored across many processes, ranging from machine learning approaches (Toivonen and Gross 2015), to associative algorithms (Varshney, Wang, and Varshney 2016), and autonomous evaluation algorithms (Jennings 2010). A popular approach in CC is a two part method, consisting of an *expansion* phase where the agent synthesizes a large set of possible outputs for a creative process, and a *contraction* phase where the agent processes the candidate outputs in order to select valuable output. Analogous conceptualizations of the expansion phase include *divergent thinking*, *generative thinking*, and *defocused attention*. Analogous conceptualizations of the contraction phase include *convergent thinking*, *evaluative thinking*, and *focused attention* (Guilford 1967; Pereira and Cardoso 2002; Zhang, Sjoerds, and Hommel 2020; Sarathy 2018).

Boden's Types of Creativity: Boden proposed three categories that describe the process of generating creative outputs, namely, *combinational creativity*, *transformational creativity*, and *exploratory creativity* (Boden 1998). Combinational creativity involves taking known or familiar information, and combining it in a way that generates a novel output (Pereira and Cardoso 2002; Lieto et al. 2019). Transformational creativity involves transforming one or more dimensions of the solution/output space to provide the means for new structures to emerge in the transformed space. Lastly, exploratory creativity involves an exhaustive search of a solution/output space to find a novel solution.

CPS Adaptation of Process-based Aspects: Creative problem solving directly utilizes process-based approaches. CPS is typically triggered by an *impasse* moment, where the agent detects that nominal problem solving techniques are insufficient for accomplishing the goal (Knoblich et al. 1999). Impasse is followed by a period of *incubation*, where the agent generates the solution space, synthesizing possible ways of solving the problem using a relaxed representation of the problem and domain. Once a viable solution is found in this space, the agent is said to reach its *insight* or "Aha!" moment (Colin and Belpaeme 2019), wherein the agent proceeds to use the solution to solve the problem. We call this process the *impasse-incubation-insight* process. While there exist other general formalizations of the creative process (Mumford et al. 1991; 1997), we use the *impasse-incubation-insight* paradigm to facilitate our adaptations. The impasse-incubation-insight process can be implemented using the two part method of expansion and contraction in the following manner – the impasse moment triggers incubation, where the agent enters the expansion phase and generates a new conceptual space C'_X from C_X. This is followed by the contraction phase, wherein the agent applies the newly discovered concepts $c'_X \in C'_X$ to generate a plan for accomplishing the goal (insight moment).

Boden's types of creativity are inherited into creative problem solving by providing three ways to expand an agent's initial conceptual space. Thus, Boden's types of creativity can be applied to manipulate the initial conceptual space C^X of the agent, in order to come up with a new conceptual space $C'^X \nsubseteq C^X$ for solving the problem.

Combinational methods involve combining existing information in an agent's conceptual space to generate a novel conceptual space for solving the problem. The agent creates new concepts $c'^X \in C'^X$ by combining existing concepts in C^X. We define a function f that combines concepts in C^X to create the new conceptual space, such that $f(c_i^X) = c_i^X$, when more than one concept is not combined:

$$f : C^X \to C'^X \mid c'^X = f(c_1^X, ...c_k^X);$$
$$c_1^X, ...c_k^X \in C^X, C'^X \nsubseteq C^X$$

If $c_1^X, ...c_k^X \notin C^X$, we can redefine a new conceptual space of $\bigcup_{i=1}^k c_i^X$, where combinational creativity applies. In CPS, combinational creativity can be observed when learning new behaviors or skills as a composition of previously known behaviors (Hangl et al. 2020), or constructing new tools by combining objects (Nair, Balloch, and Chernova 2019).

Transformational methods involve transforming the problem representation in some way to generate a novel and previously unknown representation of the same problem, i.e., the agent transforms the initial conceptual space C^X into a new conceptual space C'^X. The set of concepts $c'^X \in C'^X$ can be represented as follows:

$$f : C^X \to C'^X \mid c'^X = f(c^X)\, \forall\, c^X \in C^X, C'^X \nsubseteq C^X$$

Thus, f denotes a surjective function that maps every concept in $c^X \in C^X$ to a new concept $c'^X \in C'^X$. Transformational creativity involves a mapping from the initial conceptual space to a new conceptual space, via an appropriate transform, e.g., rotations or translations (Fitzgerald, Goel, and Thomaz 2017), and segmentations (Gizzi, Castro, and Sinapov 2019).

Exploratory methods involve searching the universal conceptual space C^{X^*}, to discover a novel solution. The agent may discover a new conceptual space $C'^X \subset C^{X^*}$ either via random exploration of its environment (i.e., babbling), or guided exploration using heuristics or loss or reward functions. If the agent uses a loss function, the concepts $c'^X \in C'^X$ can be represented as follows:

$$\{c'^X = argmin_{c^{X^*}} \mathcal{L}(c^{X^*})\ s.t.\ c^{X^*} \in C^{X^*}\},$$

where $\mathcal{L}$ denotes an appropriate loss function, and C'^X contains novel concepts from the universal conceptual space such that $C'^X \nsubseteq C^X$. In general, approaches that explore the state space to derive a solution, e.g., reinforcement learning (with reward functions), search through planning spaces (Erdogan and Stilman 2013), and motor babbling (Sinapov and Stoytchev 2007), fall into this type. In large conceptual spaces, this form of creativity can be prohibitive.

Conclusion

In this paper, we presented a formal definition of creative problem solving as the intersection of computational creativity and problem solving in AI. We specified key aspects of CC systems, and formalized their adapted inheritance into CPS. Research in creative problem solving has taken place mostly within the confines of the artificial intelligence community, and we believe that highlighting this problem in the CC community will enable necessary, and more aggressive advancements in developing computational methods for CPS.

Acknowledgements

This work is supported in part by NSF IIS 1564080, ONR N000141612835 and DARPA award W911NF-19-2-0006. The authors would also like to thank Dr. Daniel C. Dennett and Dr. Christopher G. Atkeson for their valuable feedback.

References

Bishop, M., and Boden, M. A. 2010. The Turing test and Artistic Creativity. *Kybernetes* 3(39):409—-413.

Boden, M. A. 1998. Creativity and Artificial Intelligence. *Artificial Intelligence* 1-2:347–356.

Cass, S. 2005. Apollo 13, We Have a Solution. *IEEE Spectrum On-line, 04* 1.

Colin, T. R., and Belpaeme, T. 2019. Reinforcement Learning and Insight in the Artificial Pigeon. In *CogSci*, 1533–1539.

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.; Wiggins, G. A.; et al. 2012. Computational Creativity: The Final Frontier? In *Ecai*, volume 12, 21–26. Montpelier.

Cook, M.; Colton, S.; Pease, A.; and Llano, M. T. 2019. Framing in Computational Creativity – A Survey and Taxonomy. In *The 10th International Conference on Computational Creativity*, 156–163.

Erdogan, C., and Stilman, M. 2013. Planning in Constraint Space: Automated Design of Functional Structures. In *2013 IEEE International Conference on Robotics and Automation*, 1807–1812. IEEE.

Fitzgerald, T.; Goel, A.; and Thomaz, A. 2017. Human-Robot Co-Creativity: Task Transfer on a Spectrum of Similarity. In *2017 International Conference on Computational Creativity*, 104–111.

Gizzi, E.; Castro, M. G.; and Sinapov, J. 2019. Creative Problem Solving by Robots Using Action Primitive Discovery. In *2019 Joint IEEE 9th International Conference on Development and Learning and Epigenetic Robotics*, 228–233. IEEE.

Guckelsberger, C.; Salge, C.; and Colton, S. 2017. Addressing the "Why?" in Computational Creativity: A Non-Anthropocentric, Minimal Model of Intentional Creative Agency. In *The 8th International Conference on Computational Creativity*.

Guilford, J. P. 1967. Creativity: Yesterday, Today and Tomorrow. *The Journal of Creative Behavior* 1(1):3–14.

Hangl, S.; Dunjko, V.; Briegel, H. J.; and Piater, J. 2020. Skill Learning by Autonomous Robotic Playing using Active Learning and Exploratory Behaviour Composition. *Frontiers in Robotics and AI*.

Jennings, K. E. 2010. Developing Creativity: Artificial Barriers in Artificial Intelligence. *Minds and Machines* 20(4):489–501.

Jordanous, A. 2012. A Standardised Procedure for Evaluating Creative Systems: Computational Creativity Evaluation Based on What it is to be Creative. *Cognitive Computation* 4(3):246–279.

Knoblich, G.; Ohlsson, S.; Haider, H.; and Rhenius, D. 1999. Constraint Relaxation and Chunk Decomposition in Insight Problem Solving. *Journal of Experimental Psychology: Learning, memory, and cognition* 25(6):1534–1555.

Lieto, A.; Perrone, F.; Pozzato, G. L.; and Chiodino, E. 2019. Beyond Subgoaling: A Dynamic Knowledge Generation Framework for Creative Problem Solving in Cognitive Architectures. *Cognitive Systems Research* 58:305–316.

Mumford, M. D.; Mobley, M. I.; Reiter-Palmon, R.; Uhlman, C. E.; and Doares, L. M. 1991. Process analytic models of creative capacities. *Creativity Research Journal* 4(2):91–122.

Mumford, M. D.; Supinski, E. P.; Baughman, W. A.; Costanza, D. P.; and Threlfall, K. V. 1997. Process-based measures of creative problem-solving skills: V. overall prediction. *Creativity Research Journal* 10(1):73–85.

Nair, L.; Balloch, J.; and Chernova, S. 2019. Tool macgyvering: Tool construction Using Geometric Reasoning. In *2019 International Conference on Robotics and Automation (ICRA)*, 5837–5843. IEEE.

Olteţeanu, A.-M. 2014. Two General Classes in Creative Problem-Solving? An Account Based on the Cognitive Processes Involved in the Problem Structure-Representation Structure Relationship. In *The International Conference on Computational Creativity*.

Olteţeanu, A.-M. 2015. "Seeing as" and Re-Representation: Their Relation to Insight, Creative Problem-Solving and Types of Creativity. In *The Workshop on Computational Creativity, Concept Invention, and General Intelligence*, 105.

Pereira, F. C., and Cardoso, A. 2002. Conceptual Blending and the Quest for the Holy Creative Process. In *The 2nd Workshop on Creative Systems, Approaches to Creativity in Artificial Intelligence and Cognitive Science, European Conference on Artificial Intelligence*.

Sarathy, V. 2018. Real world problem-solving. *Frontiers in human neuroscience* 12:261.

Sinapov, J., and Stoytchev, A. 2007. Learning and Generalization of Behavior-Grounded Tool Affordances. In *2007 IEEE 6th International Conference on Development and Learning*, 19–24. IEEE.

Sosa, R., and Gero, J. S. 2016. Multi-dimensional Creativity: A Computational Perspective. *The International Journal of Design Creativity and Innovation* 4(1):26–50.

Toivonen, H., and Gross, O. 2015. Data Mining and Machine Learning in Computational Creativity. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 5(6):265–275.

Varshney, L. R.; Wang, J.; and Varshney, K. R. 2016. Associative Algorithms for Computational Creativity. *The Journal of Creative Behavior* 50(3):211–223.

Zhang, W.; Sjoerds, Z.; and Hommel, B. 2020. Metacontrol of Human Creativity: The Neurocognitive Mechanisms of Convergent and Divergent Thinking. *NeuroImage* 210:116572.