Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence (IJCAI-23)
Journal Track

Creative Problem Solving in Artificially Intelligent Agents: A Survey and Framework (Extended Abstract)*

Evana Gizzi[1] , **Lakshmi Nair**[2] , **Sonia Chernova**[2] and **Jivko Sinapov**[1]
[1]Department of Computer Science, Tufts University
[2]College of Computing, Georgia Tech
evana.gizzi@tufts.edu, {lnair3, chernova}@gatech.edu, jivko.sinapov@tufts.edu

Abstract

Creative Problem Solving (CPS) is a sub-area within artificial intelligence that focuses on methods for solving off-nominal, or anomalous problems in autonomous systems. Despite many advancements in planning and learning in AI, resolving novel problems or adapting existing knowledge to a new context, especially in cases where the environment may change in unpredictable ways, remains a challenge. To stimulate further research in CPS, we contribute a definition and a framework of CPS, which we use to categorize existing AI methods in this field. We conclude our survey with open research questions, and suggested future directions.

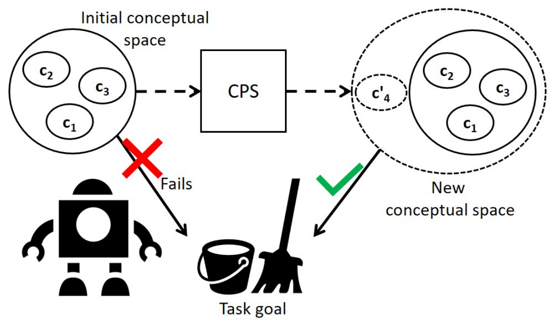

Figure 1: Creative problem solving occurs when the initial conceptual space of the agent is insufficient to complete the task, and the agent needs to expand its conceptual space to achieve the task goal. Traditional planning/learning approaches in AI would often return a failure in such scenarios.

1 Introduction

Creative problem solving (CPS) focuses on using creative processes in the context of problem solving. While numerous real-world scenarios have highlighted the importance of CPS skills in humans [Cass, 2005; Turner *et al.*, 2020], such skills are currently beyond the scope of AI. CPS capabilities can greatly improve the resourcefulness of existing AI systems, but it currently remains an under-explored research area. Our survey article [Gizzi *et al.*, 2022b] seeks to address this gap in AI and to do so, we combine theoretical aspects from planning and learning in AI with theoretical aspects from Computational Creativity (CC) – an active area of research that seeks to develop computational methods that are capable of generating a creative output, reminiscent of the creative processes in humans. To the best of our knowledge, this is the first survey that is specifically focused on creative problem solving in AI, leveraging the literature from *both* CC and AI. We believe that a comprehensive discussion of CPS, combining CC and AI principles, is vital for encouraging future work in the area.

The goals of our article are primarily to: a) define creative problem solving, in order to create a common understanding of what constitutes creative processes in AI systems; b) establish a framework for designing CPS systems in AI by leveraging theoretical aspects from both Computational Creativity and planning/learning in AI; c) organize and categorize existing CPS work within our framework; and d) lay the ground for future work by identifying open research questions and the gaps to be addressed in the future based on our survey.

2 Defining Creative Problem Solving (CPS)

Towards the first goal of our survey, we present a novel definition of creative problem solving (CPS) and describe the relevant formalisms. To begin, we broadly define a *concept* as a state (of the environment and/or agent) or action; and *concept space* as the set of all concepts known to the agent. The *universal conceptual space* represents a theoretical conceptual space (not realizable in practice), containing every possible concept that the agent could possibly know about.

In practice, the agent often encounters problems that it is unable to solve given the initial information (i.e., concepts) available to it, thus requiring it to adapt, e.g., robots operating in unstructured environments often have to improvise to effectively solve the task [Atkeson *et al.*, 2018]. Hence, *a crucial aspect of CPS that differentiates it from general planning or learning problems in AI is that the initial conceptual space known to the agent is insufficient to accomplish the task goal*. We refer to such task goals as "*un-achievable goals*". Thus, CPS is characterized by its *flexibility or adaptability* to handle novel problems where traditional AI approaches often

*This paper is an extended abstract of an article in the Journal of Artificial Intelligence Research [Gizzi *et al.,* 2022b]

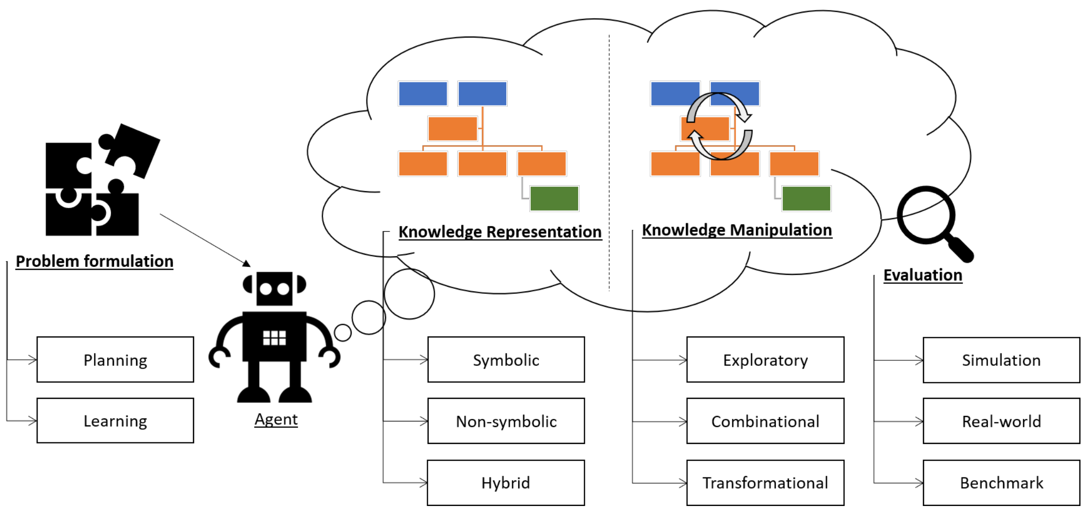

Figure 2: Creative problem solving framework, beginning with the problem formulation, followed by representation of the initial conceptual space (knowledge representation). The agent then operates on the initial conceptual space to derive a new conceptual space for solving the task (knowledge manipulation), and evaluates the solutions generated from the new conceptual space for their success (evaluation).

fail. On this basis, we define CPS as follows (See Figure 1):

Definition 1. *Given an un-achievable goal due to an insufficient conceptual space, creative problem solving is defined as the process by which the agent manipulates its currently known conceptual space in order to discover new concepts that are not in its current conceptual space, thus allowing the agent to accomplish the previously un-achievable goal.*

In other words, the space of concepts (i.e., states and actions) that is explicitly represented by the agent defines the boundaries of what the agent can accomplish. Creativity arises when the agent uses what it already knows to discover something new. In CPS, the newly discovered knowledge is applied to solve a previously impossible task. Our full article further expands on this definition and provides running examples to clarify the different formalisms.

3 Establishing a Framework for Designing Creative Problem Solving Systems in AI

Towards the second goal of our survey, we leverage theoretical aspects from both Computational Creativity and planning/learning in AI to develop a novel framework that captures the process of CPS. The full article details the process of *impasse-incubation-insight* that forms the basis for the CPS framework, that we briefly preview here (See Figure 2).

Given a task that is currently unsolvable (i.e., impasse), the first step in our framework involves formulating the problem, either as a *planning* or *learning* problem, with a few exceptions. Once the problem is formulated, the next step involves appropriately representing the relevant information (i.e., the concepts) to form the initial conceptual space that defines the problem. Third, since the initial conceptual space is insufficient for completing the task, the agent must expand it in order to "derive" a new conceptual space for accomplishing the goal (i.e., incubation). Lastly, the framework involves evaluating the new conceptual space for its effectiveness in solving the problem, by generating a solution from the new conceptual space (i.e., insight). In summary, we organize existing work in CPS through the following questions:

- *How is the problem formulated?* - We discuss the two primary problem formulations found within the CPS literature, namely, a) *planning* and, b) *learning*, specifically reinforcement learning, with a few exceptions that follow alternate paradigms;
- *How are the concepts represented?* - We discuss the different modes of representing information within the conceptual space of the agent, namely *symbolic*, *non-symbolic* and *hybrid* representations;
- *How is the new conceptual space derived?* - We discuss three ways for the agent to "discover" or "derive" a new conceptual space, based on Boden's three levels of creativity [Boden, 1998]. Specifically, the agent can discover new concepts for solving the task by: a) exploring the universal conceptual space beyond the constraints of the agent's initial conceptual space (*Exploratory*); b) combining existing concepts within the initial conceptual space (*Combinational*), or c) transforming the initial conceptual space to a new one where the solutions to the task become apparent (*Transformational*);
- *How is the new conceptual space evaluated?* - We discuss the different modes of evaluation adopted by existing CPS approaches, to distinguish theoretical models from models that have been tested in real-world settings. This ranges from evaluating the CPS approaches in simulation to evaluating them on physical robotic platforms.

In the full article, we discuss the above steps in-depth, and also present algorithmic formulations for implementing the three ways of deriving the new conceptual space within an AI system. We further expand on the different types of conceptual spaces in CPS depending on whether they capture actions or states (states of the agent, states of the terrain or states of the objects in the environment).

4 Organizing Existing Work in CPS

To capture the current status of work in creative problem solving, the third goal of our survey seeks to organize existing literature guided by our CPS framework. We review 51 papers that perform CPS in AI and describe how they fit within the steps of our framework. Specifically, we organize them based on how they formulate the problem, represent the conceptual space, derive the new conceptual space, and evaluate the result. Additionally, we categorize them based on the types of conceptual spaces that they operate on (states vs. actions). Lastly, we present four existing CPS architectures, namely, *CreaCogs* [Olteţeanu and Falomir, 2016], *Robogyver* [Nair *et al.*, 2020; Nair, 2020], *ICARUS* [Choi *et al.*, 2018], and *DIARC* [Muhammad *et al.*, 2021] that cover the different aspects of our framework. We hope that this taxonomy will provide readers with a full review and characterization of the existing CPS literature in AI.

5 Laying the Ground for Future Work

Towards the final goal of our survey, we review our major findings and present open research questions to be addressed in the future. Our article also provides a statistical data summary of the papers reviewed in our survey to offer insights regarding highly explored vs. under-explored sub-areas in CPS based on our taxonomy. Here, we briefly present some of the research questions discussed in our article.

5.1 Hybrid Representations of Concepts

Hybrid representations that combine symbolic and non-symbolic modes of information could be beneficial for the development of improved CPS systems. Future work should consider what information should be represented as symbolic vs. non-symbolic, and how we can design hybrid systems that effectively leverage the relative strengths of each representation. Most planning-based CPS methods use purely symbolic representations, whereas learning-based CPS methods tend to utilize non-symbolic representations. Current methods that fall under the "hybrid" category include hybrid planning and reinforcement learning [Strens and Windelinckx, 2004], and neurosymbolic AI [d'Avila Garcez and Lamb, 2020]. Both these methods have already been shown to greatly enhance the reasoning capabilities of AI systems and could be significantly helpful in CPS. Only a few CPS systems covered in our review have looked at such hybrid representations.

5.2 Generalizability of CPS Methods

A key consideration in CPS research is the flexibility to handle inherently novel tasks. While existing research has looked at generalizability within the individual problem domains (e.g., generalizing from one tool to another within the domain of cooking), there is yet to be an extensive research and testing of cross-domain generalizability of CPS systems (from one task domain to another, e.g., cooking to woodworking). Despite the domain-agnostic theoretical formulations of CPS methods which exist in current research, testing the specific implementations for cross-domain generalizability remains a largely unexplored challenge.

5.3 Lifelong CPS

Most of the CPS methods reviewed in our survey focuses on creating flexible CPS methods for solving the "problem at hand". However, there has not been extensive research into methods for using past CPS encounters to improve future CPS performance, thus employing "lifelong" learning for CPS. For example, many CPS methods surveyed involve environment exploration, used for optimizing a specific task solution. An open question here is, *how can the agent learn from, and adapt its prior experiences to effectively solve CPS tasks in the future?* The agent's prior interactions could be used to support lifelong CPS, by minimizing the amount of environment interaction and effort required by the agent in a different, future CPS task. While there does exist preliminary work in the area of lifelong CPS [Gizzi *et al.*, 2022a], future research should continue consideration of how agents can improve their own CPS abilities as they are continually put into scenarios where they need to discover new information.

5.4 Computational Creativity in CPS

In our previous work [Gizzi *et al.*, 2020] as well as our survey article [Gizzi *et al.*, 2022b], we emphasized the value of utilizing methods from Computational Creativity in CPS, which has been largely unexplored. Some of the works explored in this survey consider key features of CC, such as "novelty" and "value" when evaluating creative solutions, and methods that enable agents to find solutions to CPS problems. As a future direction, the CC method of "conceptual blending" [Falomir and Plaza, 2020; Schorlemmer and Plaza, 2021] could be an effective combinational method of deriving a new conceptual space for CPS. Similarly, the CC method of "framing" [Colton *et al.*, 2011; Guckelsberger *et al.*, 2017; Cook *et al.*, 2019] where an intelligent agent attempts to "explain" their own creativity can potentially be used as a mode of transformational method for deriving the new conceptual space. However, algorithmic implementations of these methods in AI systems for CPS remains an open question.

5.5 Universal Metric for CPS Solutions

The lack of a concise and universally accepted definition of CPS has made it challenging to develop a universal measure of the "quality" of a CPS solution. Developing such a metric for CPS would be highly beneficial, as it would help streamline and benchmark existing and future CPS methods. For instance, existing metrics for assessing whether a solution is "creative" involves subjective novelty-based measures. In contrast, an objective metric for evaluating approaches would be helpful when characterizing the outputs of CPS methods, such as the outputs of reinforcement learning, which often depend on objective measures, like loss or reward.

We hope that this survey will encourage research into the relatively unexplored field of CPS, bridging the gap between CC and AI to build more resourceful and adaptable AI systems.

Acknowledgments

We would like to note that the two first authors, Evana Gizzi and Lakshmi Nair, both contributed equally to all parts of the survey.

References

[Atkeson *et al.*, 2018] Christopher G Atkeson, PW Babu Benzun, Nandan Banerjee, Dmitry Berenson, Christoper P Bove, Xiongyi Cui, Mathew DeDonato, Ruixiang Du, Siyuan Feng, Perry Franklin, et al. What happened at the darpa robotics challenge finals. In *The DARPA Robotics Challenge Finals: Humanoid Robots to the Rescue*, pages 667–684. Springer, 2018.

[Boden, 1998] Margaret A. Boden. Creativity and artificial intelligence. *Artificial Intelligence*, 1-2:347–356, 1998.

[Cass, 2005] Stephen Cass. Apollo 13, we have a solution. *IEEE Spectrum On-line, 04*, 1, 2005.

[Choi *et al.*, 2018] Dongkyu Choi, Pat Langley, and Son Thanh To. Creating and using tools in a hybrid cognitive architecture. In *2018 AAAI Spring Symposium Series*, 2018.

[Colton *et al.*, 2011] Simon Colton, John William Charnley, and Alison Pease. Computational creativity theory: The face and idea descriptive models. In *The International Conference on Computational Creativity (ICCC)*, pages 90–95, 2011.

[Cook *et al.*, 2019] Michael Cook, Simon Colton, Alison Pease, and Maria Theresa Llano. Framing in computational creativity–a survey and taxonomy. In *The 10th International Conference on Computational Creativity (ICCC)*, pages 156–163, 2019.

[d'Avila Garcez and Lamb, 2020] Artur d'Avila Garcez and Luis C Lamb. Neurosymbolic ai: The 3rd wave. *Arxiv preprint*, arXiv:2012.05876, 2020.

[Falomir and Plaza, 2020] Zoe Falomir and Enric Plaza. Towards a model of creative understanding: deconstructing and recreating conceptual blends using image schemas and qualitative spatial descriptors. *Annals of Mathematics and Artificial Intelligence*, 88(5):457–477, 2020.

[Gizzi *et al.*, 2020] Evana Gizzi, Lakshmi Nair, Jivko Sinapov, and Sonia Chernova. From computational creativity to creative problem solving agents. In *International Conference on Computational Creativity (ICCC)*, pages 370–373, 2020.

[Gizzi *et al.*, 2022a] Evana Gizzi, Wo Wei Lin, Mateo Guaman Castro, Ethan Harvey, and Jivko Sinapov. Toward life-long creative problem solving: Using world models for increased performance in novelty resolution. In *13th International Conference on Computational Creativity*, 2022.

[Gizzi *et al.*, 2022b] Evana Gizzi, Lakshmi Nair, Sonia Chernova, and Jivko Sinapov. Creative problem solving in artificially intelligent agents: A survey and framework. *Journal of Artificial Intelligence Research*, 75:857–911, 2022.

[Guckelsberger *et al.*, 2017] Christian Guckelsberger, Christoph Salge, and Simon Colton. Addressing the" why?" in computational creativity: A non-anthropocentric, minimal model of intentional creative agency. In *The 8th International Conference on Computational Creativity (ICCC)*, 2017.

[Muhammad *et al.*, 2021] Faizan Muhammad, Vasanth Sarathy, Gyan Tatiya, Shivam Goel, Saurav Gyawali, Mateo Guaman, Jivko Sinapov, and Matthias Scheutz. A novelty-centric agent architecture for changing worlds. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*, pages 925–933, 2021.

[Nair *et al.*, 2020] Lakshmi Nair, Nithin Shrivatsav, and Sonia Chernova. Tool macgyvering: A novel framework for combining tool substitution and construction. *Arxiv preprint*, arXiv:2008.10638, 2020.

[Nair, 2020] Lakshmi Velayudhan Nair. *Robogyver: Autonomous Tool Macgyvering for Inventive Problem Solving*. PhD thesis, Georgia Institute of Technology, 2020.

[Olteţeanu and Falomir, 2016] Ana-Maria Olteţeanu and Zoe Falomir. Object replacement and object composition in a creative cognitive system. towards a computational solver of the alternative uses test. *Cognitive Systems Research*, 39:15–32, 2016.

[Schorlemmer and Plaza, 2021] Marco Schorlemmer and Enric Plaza. A uniform model of computational conceptual blending. *Cognitive Systems Research*, 65:118–137, 2021.

[Strens and Windelinckx, 2004] Malcolm Strens and Neil Windelinckx. Combining planning with reinforcement learning for multi-robot task allocation. In *Adaptive Agents and Multi-Agent Systems II*, pages 260–274. Springer, 2004.

[Turner *et al.*, 2020] MC Turner, LV Duggan, BA Glezerson, and SD Marshall. Thinking outside the (acrylic) box: a framework for the local use of custom-made medical devices. *Anaesthesia*, -, 2020.