

Creative Problem Solving in Artificially Intelligent Agents: A Survey and Framework

Evana Gizzi

*Department of Computer Science, Tufts University
Medford, MA 02155*

EVANA.GIZZI@TUFTS.EDU

Lakshmi Nair

*College of Computing, Georgia Tech
Atlanta, GA 30332*

LNAIR3@GATECH.EDU

Sonia Chernova

CHERNOVA@GATECH.EDU

Jivko Sinapov

*Department of Computer Science, Tufts University
Medford, MA 02155*

JIVKO.SINAPOV@TUFTS.EDU

Abstract

Creative Problem Solving (CPS) is a sub-area within Artificial Intelligence (AI) that focuses on methods for solving off-nominal, or anomalous problems in autonomous systems. Despite many advancements in planning and learning, resolving novel problems or adapting existing knowledge to a new context, especially in cases where the environment may change in unpredictable ways post deployment, remains a limiting factor in the safe and useful integration of intelligent systems. The emergence of increasingly autonomous systems dictates the necessity for AI agents to deal with environmental uncertainty through creativity. To stimulate further research in CPS, we present a definition and a framework of CPS, which we adopt to categorize existing AI methods in this field. Our framework consists of four main components of a CPS problem, namely, 1) problem formulation, 2) knowledge representation, 3) method of knowledge manipulation, and 4) method of evaluation. We conclude our survey with open research questions, and suggested directions for the future.

1. Introduction

Creativity is often described as a hallmark of sophisticated intelligence. The Oxford English Dictionary defines “creativity” as “*Inventive, imaginative; of, relating to, displaying, using, or involving imagination or original ideas as well as routine skill or intellect, esp. in literature or art*” (Dictionary, 1989). Despite our familiarity with the notion of creativity, understanding and implementing creativity in artificially intelligent systems continues to be a challenge. Computational Creativity (CC) is 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. The CC research community includes a diverse body of researchers, spanning the fields of psychology, neuroscience, philosophy, and computer science. The goal of CC research, as described by the Association for Computational Creativity¹, is “[to gain] the ability to model, simulate or replicate creativity using a computer, to achieve one of several ends, including the construction of a program or computer capable of human-level

1. The Association of Computational Creativity is a nonprofit organization dedicated to the advancement of CC, and organizing body of the International Conference on Computational Creativity (ICCC)




Figure 1: Examples of CPS in human and non-human species: The jury-rigged filter constructed by the astronauts on Apollo 13 (left, image credit: NASA). Rook extracting a bucket by bending a piece of wire to make a hook (right) (Bird & Emery, 2009).

creativity, to better understand human creativity and to formulate an algorithmic perspective on creative behavior in humans, and to design programs that can enhance human creativity without necessarily being creative themselves” (ICCC, 2022). While there has been extensive work in the area of Computational Creativity in Artificial Intelligence (AI), these works are primarily focused on the generation of creative artifacts, e.g., paintings, poems etc. In contrast, there is very limited focus on creativity that is specifically task-oriented, i.e., creativity in problem solving.

Creative problem solving (CPS) focuses on using creative processes in the context of problem solving. Both human and non-human species have been shown to creatively solve problems (Boesch & Boesch, 1990; Baker, Kanitscheider, Markov, Wu, Powell, McGrew, & Mordatch, 2019) (Figure 1), e.g., crows have been shown to spontaneously modify tools by shaping hooks out of wire, and using the modified tools in the correct sequence of actions required to retrieve food (Bird & Emery, 2009). Various models of CPS in humans have also been proposed by researchers over the years (Isaksen & Treffinger, 2004), beginning as early as 1952 with the work of Alex Osborn, who presented a comprehensive description of a seven-stage CPS process (Osborn, 1952). Prior work by Mumford et al. (1991, 1997) has also investigated CPS processes in humans by conducting human participant studies and evaluations. They define creative problems as problems that arise in ill-defined situations, thus eliciting creativity in humans. Further, their work highlights core cognitive processes involved in CPS, including problem construction, information encoding, idea evaluation, implementation, and monitoring. While these works primarily explore CPS in humans and other animals, there is very limited work focusing on CPS in artificial agents.

What makes CPS important for AI? Numerous real-world examples demonstrate the practical importance of CPS, particularly when dealing with crises or time-constrained scenarios. For instance, in the Apollo 13 incident of 1970, astronauts on board the spacecraft creatively constructed a makeshift CO₂ filter using unconventional materials, enabling them to safely return home (Cass, 2005) (See left, Figure 1). More recently, makeshift ventilators built using low-cost 3D printed parts and off-the-shelf items such as manual resuscitator PVC bags and motors, have been used to combat the widespread equipment shortages during COVID-19 (Kitchen, 2020; Turner, Duggan, Glezerson, & Marshall, 2020). However, similar skills are currently beyond the scope of AI. Developing artificial agents with similar capabilities, can greatly improve the resourcefulness and adaptability of existing AI systems.

These capabilities will be especially useful for robots that explore, as well as work in space, underwater, remote locations on land, and disaster sites, where the robots are highly likely to face unprecedented circumstances, requiring them to adapt (Atkeson, Benezun, Banerjee, Berenson, Bove, Cui, DeDonato, Du, Feng, Franklin, et al., 2018).

In this survey, we describe how CPS can leverage concepts from CC and planning/learning in AI to improve the adaptability of existing AI systems to novel scenarios. Similar surveys have previously focused specifically on CC by reviewing interdisciplinary work in CC along with evaluation techniques for CC systems (Jordanous, 2013; Lamb, Brown, & Clarke, 2018). Rowe et al. (1993) surveyed CC works explicitly in an AI context, in which they suggest five key aspects of CC systems. These include a) flexibility of knowledge representations, b) tolerance to ambiguity in the knowledge representations, c) avoiding functional fixity, d) assessing the usefulness of the creative output, and e) the capacity to elaborate on the creative output to find out their consequences. The special issue journal on “*Problem-solving, Creativity and Spatial Reasoning*” by Falomir et al. (2019) compiled selected works in existing CPS research for a multi-disciplinary perspective on problem solving in CC, focused specifically on highlighting the synergies between the traditionally separate research areas. In contrast to prior surveys, 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. Our goal in this survey is to contribute a taxonomy of research in CPS, and to provide organization and clarity into what has already been achieved as well as open research questions. We believe that a comprehensive discussion of CPS, combining CC and AI principles, is vital for encouraging future work in this area.

Our paper is organized as follows. We begin in Section 2, by discussing existing definitions of CPS in AI along with presenting our definition, and contrasting it to the existing work. We then describe relevant aspects in CC in Section 3, and describe how they can be adapted to CPS. In Sections 4 and 5, we present a framework that highlights four key processes involved in CPS. We further classify existing research in the context of the presented framework. In Section 6, we discuss examples of CPS architectures in the current literature. Finally, we conclude our survey in Section 7 with open research questions.

2. Defining Creative Problem Solving in AI

We begin by presenting existing definitions of creative problem solving. Derived from the formulation of planning and learning problems in artificial intelligence (which we explain in this section), we introduce our definition for CPS, and highlight its contribution and differences from existing definitions.

2.1 Existing Definitions of Creative Problem Solving

There have been efforts within the research community to formalize a general definition and process for CPS. Closely related to our work, prior research has formulated CPS based on classical planning and concept re-representation.

Prior work by Oltețeanu (2014) defines CPS from an object affordance perspective. Affordances broadly refer to action possibilities for objects, e.g., cups are pour-able and doors are open-able (Gibson, 1977). Formally, affordances are defined as relationships between objects, actions, and the effects of applying the actions on the objects (Şahin, Çakmak,

Doğar, Uğur, & Üçoluk, 2007). Olteteanu (2014) distinguishes creative problems from normal problem solving tasks in that CPS problems have poor representational structure, in particular the representation of objects (initial states, relations between objects etc.). Collectively, these entities are referred to as *concepts*, and their representations consist of affordances, visuospatial features, or semantic tags. From this, Olteteanu defines the process of CPS as “*the ability of a cognitive, natural, or artificial system to use new objects to solve a problem, other than the ones that have been stored in its memory as tools for that specific purpose (if any), or to create those objects by putting together objects or parts of objects the system has access to. Depending on the problem, objects can be either physical or abstract/informational (concepts, problem templates, heuristics or other forms of representations)*” (Olteteanu, 2015). However, their framework specifically refer to objects, and do not cover other concepts relevant to problem solving in AI, e.g., actions and environment states.

In a contrasting definition, Sarathy and Scheutz (2018) define the notion of “MacGyver-esque” creativity as embodied agents that can “*generate, execute, and learn strategies for identifying and solving seemingly unsolvable real-world problems*”. They approach CPS as a planning problem, introducing the notion of a *MacGyver Problem* (MGP) as a *planning problem in the agent’s world that has a goal state that is currently unreachable by the agent* (Sarathy & Scheutz, 2018). They formalize MGP with respect to an agent t , as a planning problem in the agent’s world \mathbb{W}_t , that has a goal state g currently unreachable by the agent. In order to solve an MGP, the agent modifies its domain knowledge (through domain expansion or contraction) by sensing and perceiving its environment and its own self, enabling the agent to discover previously unknown information needed for accomplishing the goal. However, the proposed definition of MGPs do not cover learning approaches in AI for CPS and does not describe potential methods for modifying domain knowledge of the agent.

2.2 Proposed Definition of Creative Problem Solving

In this section, we present a novel formalization of CPS. We begin by first defining the components of a traditional planning or learning problem to be solved by an agent acting in its environment.

2.2.1 COMPONENTS OF TRADITIONAL PLANNING AND LEARNING IN AI

The planning or learning problem specification in AI typically consists of a task goal G to be accomplished, given a set of environment states \mathcal{S} and agent actions \mathcal{A} . The agent then produces a solution Π for accomplishing the task goal. Depending on whether the initial problem is formulated as planning or learning, the generated solution corresponds to either a task plan or policy respectively, over the states and actions. Thus, $\Pi : \mathcal{S} \rightarrow \mathcal{A}$, represents a mapping from the set of environment states, to the set of actions. In the case of planning, a full task plan is typically computed prior to the agent acting in the environment. Hence, CPS may be observed before any action executes physically. In the case of learning, the policy is typically learned on-the-fly as actions are executed in the environment and CPS may be observed concurrently.

To highlight the definitions presented in this paper, we introduce the following running example in a planning domain. Consider a scenario with a robot in a 2-dimensional 3x3 ‘grid-world’ room with 9 locations denoted by $t_{(1,1)}, t_{(1,2)}, t_{(1,3)}, t_{(2,1)}, t_{(2,2)}, \dots, t_{(3,3)}$. The room

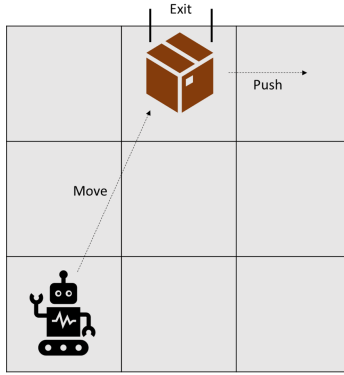


Figure 2: Nominal case where creative problem solving is not required, given the knowledge in Table 1. The robot can push the box onto a nearby empty space and proceed to the exit.

C_S	$c_s^1 \dots c_s^9 =$ robot at one of $t_{(1,1)}, t_{(1,2)}, t_{(1,3)}, t_{(2,1)}, t_{(2,2)}, t_{(2,3)}, t_{(3,1)}, t_{(3,2)}, t_{(3,3)}$ $c_s^{10} \dots c_s^{18} =$ box ₁ at one of $t_{(1,1)}, t_{(1,2)}, t_{(1,3)}, t_{(2,1)}, t_{(2,2)}, t_{(2,3)}, t_{(3,1)}, t_{(3,2)}, t_{(3,3)}$
C_A	$c_a^1 =$ move to a designated location $c_a^2 =$ push box to a location that is <i>empty</i>

Table 1: Initial concept space of the grid-world agent, shown partially.

consists of one box in front of an exit door. The robot may have the goal of exiting the room (e.g., $G =$ robot not in room), with two types of available actions, which allow the robot to move around the room, and to move boxes around the room (e.g., $A = \{a_1, a_2\}$, where $a_1 =$ move to a location, and $a_2 =$ push box to an *empty* location). This agent could then generate a solution to get from its current state in the room to the goal state, by using the actions that are known to it (e.g., $\Pi =$ move to the exit, push the box covering the exit door to the empty location on the side, move through exit door). This is shown in Figure 2.

2.2.2 COMPONENTS OF CREATIVE PROBLEM SOLVING IN AI

Given the formalisms from planning and learning problems in AI, we now describe the components of CPS. In creative problem solving, we broadly define the notion of a *concept*, as a **state** (of the environment and/or agent) or **action**. More specifically, depending on the problem formulation, concepts could refer to the actions that an agent can perform (e.g., “move”), or the state space including states of objects in the agent’s environment (e.g., “box at location $t_{(1,2)}$ ”), and the state of the agent (e.g., “agent at location $t_{(2,3)}$ ”). Grouping states and actions under the single term “concepts” allows us to unify the broad range of problem formulations within a single definition. More generally, we denote concepts as c_x , and a *conceptual space* C_X as the set of all concepts c_x . We use c_x only as notational convenience since c_x can refer to either a state s or an action a . This allows us to explain our CPS framework across different concepts. Here, X relates to either the set of environment states \mathcal{S} or actions \mathcal{A} . Hence, C_S denotes the set of all states where each state is denoted by c_s , and C_A denotes the set of actions, where each action is denoted by c_a . For our grid-world example, the concept space involved is partially highlighted in Table 1.

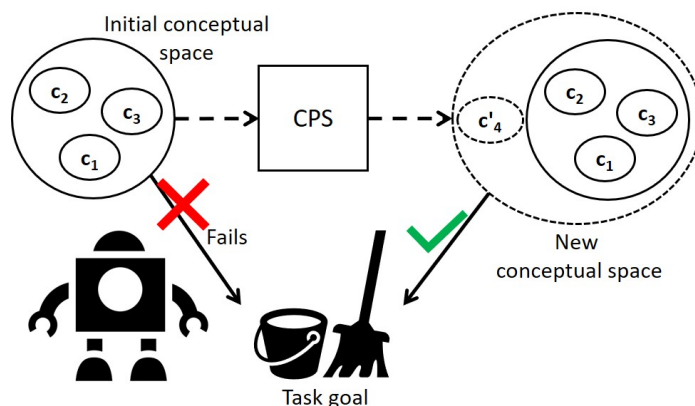


Figure 3: Creative problem solving (CPS) occurs when the initial conceptual space of the agent is insufficient to complete the task, and the agent needs to expand its conceptual space to achieve the task goal. Traditional planning/learning approaches in AI would return a failure in such scenarios.

Furthermore, let \check{C}_X denote the universal set of the concepts c_x , such that \check{C}_X represents a theoretical conceptual space containing every possible concept that the agent could potentially know about, e.g., \check{C}_A as the set of all possible actions that the agent can perform. In this work, we assume that the initial conceptual space $C_X \subset \check{C}_X$, i.e., the agent’s initial knowledge is limited. Note that $C_X = \check{C}_X$ is not a practical assumption for real-world agents, since it implies that the agent knows every concept possible. In practice, the agent often encounters problems that it is unable to solve given the initial information available to it. Similar scenarios are commonly observed in robots operating in unstructured environments where they often have to improvise to effectively solve the task.

2.2.3 KEY DEFINITION

A crucial aspect of CPS that differentiates it from general planning or learning problems in AI is that the initial conceptual space C_X known to the agent is *insufficient to accomplish the task goal*. We refer to such task goals as “*un-achievable goals*”. Consider the original grid-world example, with an added switch that needs to be *kept pressed* in order to turn off the lights in the room. The new goal for the robot in this case is, $G = \text{robot not in room AND lights off}$. The initial concept space of the robot does not include any action allowing it to accomplish this goal (i.e., Table 1 does not contain *switch press* actions, or *switch is_pressed* states). Traditional planning approaches in AI often yield a failure in these circumstances, since the initial set of actions available to the agent is insufficient for completing the task. In traditional learning approaches, the initial action space that the agent is allowed to explore is limited and often not extensible to include new actions. Thus, CPS is characterized by its *flexibility or adaptability* to handle novel problems (Guilford, 1967a). In particular, CPS seeks to enable the agent to discover new concepts for accomplishing the task, by modifying the agent’s initial conceptual space. Here, given an initial conceptual space the agent must generate a creative solution Π for accomplishing the task. We now present our definition of CPS as follows (also shown in Figure 3):




Figure 4: Example of CPS in the grid-world planning domain. In the left image, the robot has the goal of exiting the room with the lights turned off. Note that the blue switch needs to be kept pressed in order to keep the lights off. A creative solution in this case involves pushing the brown box over the light switch in order to accomplish the task goal, as shown to the right. The agent here has to somehow discover that the box can be pushed onto the switch to keep it pressed.

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. Formally, CPS refers to the process by which the agent discovers a new conceptual space $C'_X \not\subseteq C_X$, such that $C'_X = f(C_X)$ is the result of applying some function f on the current conceptual space, enabling the agent to solve the previously unsolvable task by using C'_X .*

In other words, the space of concepts 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. In our grid-world example with the light switch, the initial conceptual space of the agent (shown in Table 1) is insufficient for accomplishing the goal of $\langle G = \text{robot not in room AND lights off} \rangle$. However, if the agent is able to somehow discover (via modification of its initial conceptual space) that the box can be pushed onto the light switch to keep the switch pressed, rather than only being able to move it to empty locations, then the agent has exhibited CPS since it can now exit the room while the lights are off (See Figure 4). More specifically, the agent would have to discover a new action $c_a^3 = \text{push box onto switch}$, which has the effect of $is_pressed(\text{switch})$. In a learning context, exploration is one possible approach for the agent to discover this new concept, i.e., the agent may push the box around until it happens to be on the switch. This can be computationally prohibitive in larger state spaces. In this paper, we will present three classes of approaches (including *informed* exploration) by which the agent can manipulate the initial conceptual space more efficiently to discover new concepts for accomplishing the task (Section 4).

Our definition of CPS differs from previous definitions (Section 2.1) in the following ways. Firstly, in contrast to the prior definition by Olteteanu where the concepts focused primarily on object affordances, we describe the notion of concepts in terms of all the core entities involved in a planning or learning problem, i.e., the actions and states (including, but not limited to the states of objects). This allows us to capture CPS approaches that manipulate

non-object related concepts (e.g., the agent itself) as well. Secondly, in contrast to the prior definition by Sarathy and Scheutz that formulated macgyvering problems specifically as a planning problem, we present CPS as a planning *or* learning problem, and further connect problem solving in AI to relevant aspects of CC. In particular, we describe i) how CPS can be performed by efficiently manipulating the initial conceptual space via the function $f(\cdot)$ that correlates to existing methods developed in CC, and ii) how the output of CPS can be evaluated by leveraging existing notions of output evaluation in CC. In the following sections, we expand upon our definition, highlighting the theoretical aspects in CC that apply to CPS.

3. Aspects of Computational Creativity for Creative Problem Solving

In this section, we review four major aspects of CC, and their inheritance and adaptations to creative problem solving. The four aspects include: novelty and value, evaluative methods, procedural methods, and Boden’s types of creativity. These aspects 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. We leverage these aspects when presenting the components of CPS in Section 4.

3.1 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 creative. These systematic outputs, referred to as *artefacts*, may take physical and/or non-physical form (e.g. paintings, songs). The first aspect (*Novelty* and *Value*) describes two key characteristics of a creative output, whereas the second aspect (*Evaluative Methods*) describes methods for evaluating the output.

NOVELTY AND VALUE: Prior work by Boden (1998) proposed that creativity necessitates both novelty and value. Novelty guarantees that the generated outputs of a creative process are original, whereas the value criteria ensures that the generated outputs are not random, but targeted to accomplishing a task goal. Both novelty and value have contextual considerations. 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 & Gero, 2016; Varshney, Wang, & Varshney, 2016).

EVALUATIVE METHODS: In evaluative methods, the creativity of a system is evaluated by judging the output of its processes, if it is creative or not. Similar in nature to the Turing test, these methods focus on using the judgement of an external observer. The evaluation can either happen computationally (Colton, Wiggins, et al., 2012; Varshney et al., 2016), from a human evaluator (Bishop & Boden, 2010; Guckelsberger, Salge, & Colton, 2017), or from a social group (Varshney et al., 2016). The nature of these evaluations vary, the output may be compared to a human’s creative output, judged in a social context, or evaluated based on the agent’s ability to explain its own intentions to a human (Cook, Colton, Pease, & Llano, 2019).

ADAPTATION OF OUTPUT-BASED ASPECTS IN CREATIVE PROBLEM SOLVING: In creative problem solving, novelty is important with contextual considerations. 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. 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?

CPS does not directly inherit evaluative methods, because the output of a CPS process is simply evaluated by the agent 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, because problem solving in this case requires the discovery of concepts which are novel relative to the agent. Thus, evaluation in CPS involves evaluating whether the new conceptual space is sufficient to accomplish the current goal.

3.2 Process-based Aspects

Process-based aspects are concerned with the question of *how* creative outputs are 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: Procedural methods of generating creative outputs consist of two phases: 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/contraction phases include *divergent thinking/convergent thinking* (Guilford, 1967a; Zhang, Sjoerds, & Hommel, 2020), *generative thinking/evaluative thinking* (Ellamil, Dobson, Beeman, & Christoff, 2012), and *defocused attention/focused attention* (Sarathy, 2018).

BODEN'S TYPES OF CREATIVITY: Boden proposed three ways 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. 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.

ADAPTATION OF PROCESS-BASED ASPECTS IN CREATIVE PROBLEM SOLVING: CPS directly utilizes process-based approaches. CPS is triggered by an *impasse* moment, where the agent detects that nominal problem solving techniques are insufficient for accomplishing the goal (Knoblich, Ohlsson, Haider, & Rhenius, 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 & 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 for our CPS framework.

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. Upon

generating new concepts, the agent enters the contraction phase, wherein the agent applies the newly discovered concepts to generate a plan for accomplishing the goal (insight moment). Boden’s types of creativity provides three ways to generate the new conceptual space during the expansion phase (referred to as “Knowledge Manipulation” in the following sections). We formalize each of these methods in detail in Section 4, describing how each method operates on the initial conceptual space.

4. Components of a Creative Problem Solving Framework

In this section, we introduce a novel computational framework for creative problem solving (also shown in Figure 5), leveraging the aspects discussed in Section 3. Our CPS framework is motivated by two key ideas: First, we develop our CPS framework based on the *impasse-incubation-insight* approach in computational creativity, described in Section 3. Specifically, we contextualize and formulate the three steps based on how planning and learning problems are formulated in the AI literature. We believe this is a key feature of our framework that bridges the gap between *impasse-incubation-insight* in computational creativity and planning and learning methods in AI. Secondly, we recognize the need to organize the vast breadth of existing CPS literature through common formalisms. This would benefit future endeavors in the area, and help identify existing gaps in the research. We begin by describing our framework in detail, and in the following sections we demonstrate how existing research fits within our framework.

Given a task that is currently unsolvable (i.e., impasse), the first step within our framework involves appropriately formulating the problem. In this case, the problem may be formulated as a *planning* and/or *learning* problem, with a few exceptions. Once the problem is formulated, the agent must appropriately represent the relevant information (i.e., the concepts) in order to form the initial conceptual space. In CPS, the initial conceptual space is insufficient for accomplishing the task and as a result, the agent must expand its conceptual space (i.e., incubation) to discover a new conceptual space for accomplishing the goal. The final component of 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: a) **How is the problem formulated?** (Section 4.1: Problem Formulation); b) **How are the concepts represented?** (Section 4.2: Knowledge Representation); c) **How is the new conceptual space derived?** (Section 4.3: Knowledge Manipulation); and d) **How is the new conceptual space evaluated?** (Section 4.4: Evaluation).

4.1 Problem Formulation

How is the problem formulated? There are primarily two problem formulations within the CPS literature, namely, a) planning problem, and b) learning problem. In particular, learning refers to reinforcement learning. We categorize and discuss each paper in terms of their predominant methodology. A small subset of the papers in our review do not fall clearly within either problem formulation, which we discuss in detail.

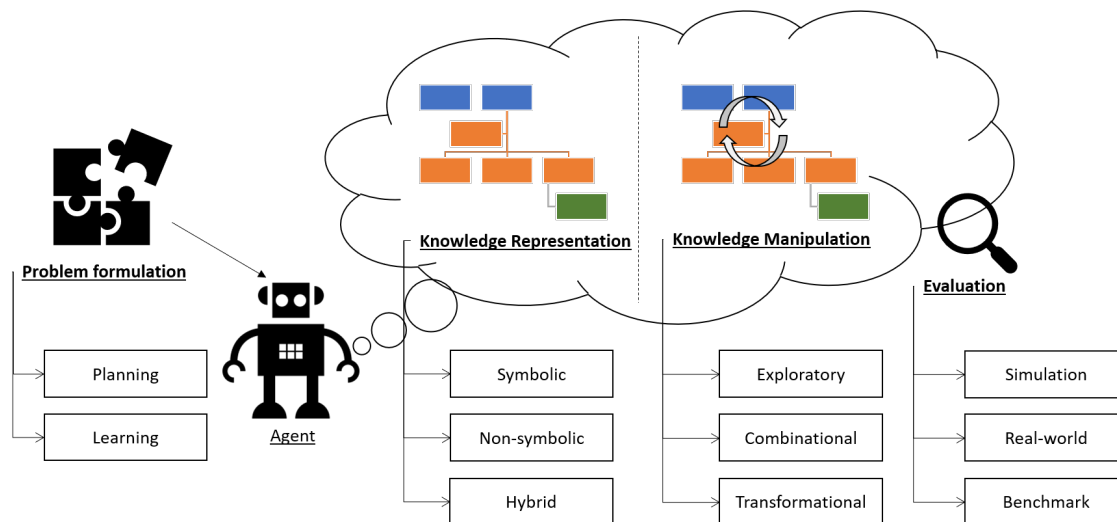


Figure 5: Creative problem solving framework, beginning with the problem formulation, followed by representation of the initial conceptual space (knowledge representation). The agent then operates on the initial conceptual space to derive a new conceptual space for solving the task (knowledge manipulation), and evaluates the solutions generated from the new conceptual space for their success (evaluation).




Figure 6: Examples of planning in CPS: Task planning within ICARUS (Choi et al., 2018) (left) where the specifics of the motion trajectory is not considered; Task and Motion Planning (TAMP) (right) for sequential manipulation that considers both the task and motion trajectories (Toussaint et al., 2018).

4.1.1 PLANNING

We begin by defining a planning problem in AI. A planning problem consists of a set of states S , a set of actions A and state transitions γ . Further, the formulation consists of an initial state denoted as s_i , and a goal state denoted as s_g . Most commonly in classical planning, a problem specification consists of a domain definition $\mathcal{P}_D = (S, A, \gamma)$, and a problem definition $\mathcal{P}_T = (\mathcal{P}_D, s_i, s_g)$. Given the domain and problem definitions, planning involves identifying a sequence of actions that can get the agent from the initial state to the goal state (i.e., a task plan). Within CPS, planning formulations may be integrated into a larger architecture or framework, or alternatively presented as standalone approaches.

ARCHITECTURAL INTEGRATION

As an example of architectural integration of task planning, Choi et al. (2018) introduce planning in the context of a cognitive architecture called *ICARUS*, for the creative construction of navigational structures, e.g., ramps and bridges. In similar work, Freedman et al. (2020) introduce analogical reasoning capabilities into a classical planning architecture called *CIRCA*, to enable a robotic agent to reason about analogies when identifying substitute objects for building navigational structures. In contrast to the construction of navigational structures, prior work by Nair et al. (2020, 2020) introduce the *Robogyver* architecture that extends classical planning through supervised learning, to enable a robot to create or “macgyver” novel tools from available objects. In similar work, Wicaksono and Sammut (2017, 2020) integrate planning within the *CREATIVE* architecture, to enable a robot to craft novel tools through 3D printing, as opposed to using available objects. Lastly, in the cognitive architecture of *SOAR*, Lieto et al. (2019) show that concept representation in a knowledge base can be used as a means for “subgoal” resolution (or plan repair) within planning.

STANDALONE TASK PLANNING

In addition to planning in the context of cognitive architectures, several standalone planning approaches have also been proposed for CPS. The standalone approaches often focus on either task planning, or motion planning. Task planning involves generating high-level action sequences for accomplishing a task, whereas motion planning focuses on generating sequences of valid joint configurations (of the robot) for performing different actions. In these cases, creativity arises either at the level of the task plan, or at the level of the motion trajectory.

Within task planning for construction of navigational structures, prior work by Erdogan and Stilman (2013) incorporate constraint optimization as a means of evaluating candidate states during planning. More specifically, they search for specific object configurations (when combining objects to construct structures such as bridges), in a convex continuous domain. For each abstract action, they partition the convex space, and evaluate whether a feasible solution exists in the partitioned spaces. Levihn and Christensen (2015) extend constraint relaxed planning (using A^*) with inverse affordances to enable a robot to navigate its environment in novel ways, e.g., cross gaps by making bridges. Inverse affordances are a mapping from a failed action to object properties that are required to make the action feasible. The agent then locates objects that satisfy the desired properties, and incorporates them into the planner. Prior work by Saboia et al. (2019) leverage mathematical descriptions of elevation in the terrain to extend planning, and incorporate the construction of ramps to navigate uneven terrains. In closely related work, Tosun et al. (2018) use a specialized planner to enable robots to create ramps for navigational tasks. The planner itself serves two functions: 1) it synthesizes a robot controller for achieving the task, and 2) it executes the controller. In contrast to navigational tasks, Boteanu et al. (2015), focus on using Hierarchical Task Networks (HTNs) for CPS by incorporating novel uses of objects within the planner, e.g., using a bowl instead of a basket. Prior work has also focused on discovery of new actions in the context of planning. Suarez-Hernandez et al. (2020) introduce a novel approach for the unsupervised synthesis of new action primitives. Gizzi et al. (2019, 2021) discover novel actions through action segmentation and behavior babbling, respectively, which they use as a method for knowledge expansion. In this way, the agent can then re-plan toward a goal in a novel scenario with the new knowledge. In similar work, Sarathy et al.




Figure 7: Examples of learning in CPS: Agent policy optimization for solving tool-based puzzles (Allen et al., 2019) (left) where the agent design is not considered, and joint optimization of agent policy and design, e.g., morphology of the legs of the agent (Ha, 2019) (right).

(2020), discover action operators through RL as a way to expand the knowledge base for CPS through re-planning.

MOTION PLANNING

While the approaches described above focus on task planning, several CPS approaches have also been introduced for creative motion planning. Fitzgerald et al. (2017, 2014) focus on the transfer of skills (i.e., motion trajectories) from a source object to a target object. For example, their work introduces approaches for adapting motion trajectories for novel uses, such as adapting trajectory for a ladle with a short handle to one that has a much longer handle. In similar work, Gajewski et al. (2019) focus on adapting trajectories from a source object to a target object by reasoning about the geometric similarities between the source and target objects. Their work differs from that of Fitzgerald et al. in terms of the underlying representations used. In contrast to trajectory adaptation, Qin et al. (2020) introduce an approach for enabling robots to manipulate novel objects as tools for different manipulation tasks. Here, the robot reasons about tool use based on the underlying object representations, as opposed to adapting a known trajectory from a source to a target object. In contrast to tool manipulation, Murooka et al. (2019) focus on incorporating the dynamics of screws and screwdrivers into motion planning, in order to enable a robot to self-tighten loose screws on its body to augment its physical capabilities.

Some CPS approaches leverage high-level task constraints for low-level motion planning. For example, Toussaint et al. (2018) introduce a novel approach for “Task and Motion Planning” (TAMP) to address the problem of sequential manipulation for tool use. They impose explicit task constraints on actions (in a continuous space) regarding the physical dynamics of objects, and optimize over the constraints using motion planning. Similarly, Silver et al. (2021) show how bottom up relational learning can support learning new probabilistic operators in a TAMP paradigm. Low-level transitions are converted into high level state representations of lifted effects (predicates with argument placeholders), which are then processed via greedy/best-first search to discover new preconditions. Hence, all of the approaches described in this section enable robots to adapt to novel and unforeseen task environments through planning.

4.1.2 LEARNING

Within the CPS literature, learning is used in two ways: a) to learn a solution or policy for accomplishing a task goal (i.e., through reinforcement learning), and b) to learn representations that can then be combined with planning or reinforcement learning (RL) techniques. In this section, we focus on the former since it relates specifically to the problem formulation. In section 4.2, we focus on the latter, and discuss how learning is applied to learn representations for the conceptual space.

We begin by presenting the components of an RL problem. An RL problem is typically formalized as a Markov Decision Process (MDP) consisting of the tuple (S, A, P, γ, R) that represents a set of states S , a set of actions A , and transition probabilities $P(s'|s, a) = Pr(s_{t+1} = s' | s_t = s, a_t = a)$, i.e., the probability of transitioning to a state s' , given an initial state s and action a . The MDP further specifies a discount factor γ , and a *reward function* $R(s, s', a)$ that indicates a reward for transitioning from s to s' via action a . The goal of the agent is to generate a policy (mapping from states to actions) that maximizes the expected reward. In contrast to planning, the agent explores its environment to gather information (via the reward function) to eventually settle on an appropriate policy or solution. In CPS, existing RL approaches have been used to generate creative solutions to tasks. While most of the approaches focus on optimizing the agent’s policy alone, a subset of the approaches focus on jointly optimizing for the agent policy *and* the agent’s physical design.

AGENT POLICY OPTIMIZATION

When optimizing for the agent’s policy, Xie et al. (2019) use model-based RL to enable a robot to improvise using tools, e.g., using unconventional tools to grab out-of-reach objects. Similar work by Allen et al. (2019) introduce the Sample-Simulate-Update-Predict (SSUP) model for solving tool-based puzzles that improve upon more traditional RL techniques. Their work samples actions that operate close to objects (in Cartesian space), and simulates the potential outcomes of the actions using a physics simulator. Based on the simulations and the actual outcomes of executing the actions, the agent eventually generates novel policies for completing tool-based puzzles. Prior work by Baker et al. (2019) demonstrated emergent creative tool use by agents playing a game of hide-and-seek. In their work, the agent policies were generated by combining two separate networks: 1) a policy network that learns an action distribution, and 2) a critic network that predicts the discounted feature returns. The policies were optimized by using Proximal Policy Optimization (PPO), and Generalized Advantage Estimation (GAE). In contrast to creative tool use, Bapst et al. (2019) generate policies for structure construction via two approaches: 1) Multi-layer Perceptron (MLP) Policy that uses a multi-layer perceptron-based algorithm to output actions or Q-values, and 2) Graph Network (GN) Policy that uses a stack of three Graph Networks to output a policy.

At the intersection of CPS and psychology, cognitively inspired RL approaches for CPS have also been proposed. Chitnis et al. (2021) use model-based RL to learn a goal/action relational model, where exploration goals are selected based on a novelty measure to optimize the exploration of un-visited state spaces. Here, there is no extrinsic reward function, but rather an intrinsic “novelty” motivation. They call this approach a “goal-literal” babbling (GLIB) approach. In similar work, Oddi et al. (2020) use the notion of intrinsic motivation for skill learning based on self-generated goals, where the measures driving goal generation are based on competence (composed of novelty, curiosity, exploration, and surprise). Goals

are selected based on those which have the highest competence improvement rate. Kralik et al. (2016) apply Q-Tree learning, a variant of Q-learning wherein a hierarchical state representation is learned simultaneously with the action policy. They demonstrate that the Q-Tree model simulates CPS observed in monkeys, thus hypothesizing potential underlying cognitive mechanisms. Similarly, Colin et al. (2019) develop RL-based algorithms for modeling insight in pigeons, motivated by a real life CPS experiment. They use a basic MDP, along with a CNN to support an actor-critic model for agent behavior, suggesting basic RL as a viable possible mechanism for cognitive CPS. In the context of discovering new actions for accomplishing a task, Kroemer et al. (2014) use RL-based approaches to learn low-level movements corresponding to primitives. In similar work, Jain et al. (2020) learn novel actions through unsupervised methods and apply PPO for learning task policies.

JOINT AGENT POLICY AND DESIGN OPTIMIZATION

In contrast to optimizing for agent policies alone, existing work in CPS has also focused on jointly optimizing for agent policy and design. These approaches demonstrate the emergence of novel morphologies that are well-suited to specific environments. Prior work by Pathak et al. (2019) focus on self-assembling morphologies wherein individual agents (“limbs”) can combine to form new morphologies. They adopt an RL framework wherein the policy parameters are optimized to jointly maximize the reward for each limb. Prior work by Ha (2019) uses the REINFORCE algorithm to enable an agent to jointly optimize for its policy and physical design. In closely related work, Schaff et al. (2019) use Proximal Policy Optimization (PPO) as opposed to REINFORCE. In these CPS approaches, the agent learns novel designs that enable effective locomotion in its environment.

Note that most RL approaches represent information in a strictly non-symbolic manner whereas classical planning often represents information in a symbolic manner. In some cases, the representations used in RL are themselves learned using supervised and semi-supervised learning techniques (Xie et al., 2019). We will describe these in more detail in Section 4.2.

4.1.3 OTHER PARADIGMS

The papers described in this section conform to our definition of CPS, but do not fall strictly under planning or learning categories, since they do not use the newly discovered concepts to accomplish some task goal. A notable observation is that the majority of the papers in this section pertain to tool-use. They involve the selection of tools that *can be* improvised for a particular function but is not demonstrated in the context of accomplishing some task goal. Hence, we include them in a separate section.

Within tool use, prior work has focused on identifying creative alternate uses for objects, through various representations such as graphs (Zhu, Zhao, & Chun Zhu, 2015; Yang, Lan, Zhang, & Zheng, 2020), semantic networks (Oltețeanu & Falomir, 2016), geometric representations (Schoeler & Wörgötter, 2015; Abelha, Guerin, & Schoeler, 2016; Shrivatsav, Nair, & Chernova, 2019; Nair, Srikanth, Erickson, & Chernova, 2019; Nair, Balloch, & Chernova, 2019), and perceptual functions that capture the effects of interacting with objects using tools (Sinapov & Stoytchev, 2007, 2008). In contrast to accomplishing a task goal, the approaches discussed here are evaluated by comparing to ground truth data. For instance, the approach may identify a rock as a good substitute for a hammer, but it is not used to actually perform a hammering task. Rather, it is evaluated against a ground truth label

indicating how well rocks can be used for hammering. In this sense, these approaches do not strictly fit a planning or RL paradigm.

Apart from tool use, the other papers in this section focus on discovering new skills (Hangl, Dunjko, Briegel, & Piater, 2017; Xu, Nair, Zhu, Gao, Garg, Fei-Fei, & Savarese, 2018) and discovering new agent designs (Zhao, Xu, Konaković-Luković, Hughes, Spielberg, Rus, & Matusik, 2020). Hangl et al. (2017) combine already known skills to discover new behaviors which are evaluated for their success at small sub-tasks such as grasping and placement of objects, without a planning or RL formulation. Xu et al. (2018) leverage Neural Task Programming (NTP) to learn skills by decomposing demonstrated tasks into generalizable substructures. While they demonstrate learning of new skills, they do not conform to the typical planning/RL formulation. In the context of learning new agent designs, Zhao et al. (2020) use Model Predictive Control (MPC) to evaluate and optimize for agent designs for traversing different kinds of environments. As a hybrid approach, Sarathy et al. (2020) combine planning and RL to discover and execute new actions/policies to acquire an environmental states in which the previously unachievable goal can be attained.

4.2 Knowledge Representation

How is the conceptual space represented? In the previous section, we described the components of a problem formulation (for learning or planning) including states and actions. These specifications form a part of the conceptual space of the agent. In this section, we discuss how the conceptual space (i.e., information regarding actions and states) is represented. In existing CPS works, there are two broad classes of representations that are most commonly used, namely symbolic and non-symbolic. More recent approaches have also sought to combine the two, to develop hybrid approaches that can leverage the relative strengths of both representations. Note that, as described in Section 4.1, symbolic representations are typically associated with planning problems, whereas non-symbolic representations are typically associated with learning. In the case of hybrid representations, symbolic and non-symbolic formulations are combined and used for planning and/or learning.

4.2.1 SYMBOLIC REPRESENTATIONS

Symbolic representations involve explicitly modeling all known and newly discovered concepts in a declarative form, via facts and rules. These representations encode high-level information as predicates and/or fluents, and the agent often learns relationships between concepts as rules or facts. In this section, we review approaches that use symbolic representations for CPS, including automated planners and large-scale semantic networks.

PLANNING LANGUAGES

Automated planning methods in CPS depend on symbolic representations of states, actions and transitions in the form of planning languages. Two commonly used planning languages include the Stanford Research Institute Problem Solver (STRIPS) language (Fikes & Nilsson, 1971), and the Planning Domain Description Language (PDDL) (McDermott, 2000; McDermott, Ghallab, Howe, Knoblock, Ram, Veloso, Weld, & Wilkins, 1998). Planning languages encode information using logical predicates describing states and actions. States are represented as a list of logical predicates that hold true when the agent is in that particular state. Actions are encoded with *preconditions*, i.e., a list of logical predicates that must hold

true before action execution, and *postconditions*, i.e., logical predicates that are expected to hold true after execution. Given an initial state and goal state, planners then use forward or backward chaining of predicates and postconditions to yield a solution. Apart from PDDL and STRIPS, examples of planning languages include PDDL 2.1, which handles temporal planning domains (Fox & Long, 2003), PDDL 1.0, which is able to capture domains with probabilistic effects (Younes & Littman, 2004), and PDDL 3, a constraints based PDDL planner (Gerevini, Haslum, Long, Saetti, & Dimopoulos, 2009). Other examples include the Action Description Language (ADL) (Pednault, 1987) and Hierarchical Task Networks (HTNs) (Erol, Hendler, & Nau, 1994).

Erdogan et al. (2013) break down problems in high dimensional continuous spaces into discrete, symbolic PDDL actions used in a constraint-based planner. Choi et al. (2018) use a STRIPS-inspired planning language extended with numeric representations in order to capture quantitative properties of objects such as their dimensions and weight. Planning and execution is then performed through a modified planner in the ICARUS architecture, which take these quantitative attributes into consideration. Prior work by Wickasono and Sammut (2017, 2020) capture properties of tools within a hierarchical planning language that is then used to create novel tools. Their hierarchy captures aspects such as length, width, and shape of the tools. They capture the state representation at two levels: primitive and abstract. Primitive states contain quantitative values, such as pose of objects, whereas abstract captures qualitative relationships between objects. Suarez-Hernandes (2020) introduce a novel algorithm for the synthesis of STRIPS actions from execution traces (sequences of actions), within a planning framework. In Sarathy et al. (2020), the agent uses RL in a fully symbolic 2D 'grid-world' domain to resolve failures in PDDL plans. Chitnis et al. (2021) utilize a fully symbolic problem representation, with grounded and un-grounded literal representations of states, actions, and goal states (which they refer to as "goal-literals"). Using this information, agents "plan to learn", where goal-literals drive exploration of novel goal/action combinations to learn a symbolic transition model.

SEMANTIC NETWORKS

In contrast to planning languages, semantic networks are symbolic encodings of information with their inter-connections (e.g. semantic similarity), usually taking the form of a graph structure (also referred to as "knowledge graphs" (Hélie & Sun, 2010)). The connections and their corresponding strengths can be hand-coded, or learned over time. In Lieto et al. (2019), CPS is performed by exploring conceptual blending (Fauconnier & Turner, 2008) in a semantic network of concepts, represented using a non-monotonic logic system which allows for concept composition. This representation accounts a measure for "common-sense applicability" on its semantic concepts. In the work of Oltețeanu (2014, 2016), semantic information (referred to as "concepts") is captured based on affordances, visual features, and explicit semantic tags. This collective knowledge base is called a "semantic map," and can assume many different structures. The semantic map is used for object replacement and object composition (OROC) for creating new objects. Boteanu et al. (2015) use large-scale semantic networks, particularly ConceptNet, that capture object affordances in the context of Hierarchical Task Networks (HTNs), to enable a robot to reason about object replacement. In similar work, Freedman et al. (2020) encapsulate knowledge about various object features (such as geometry, length, width, rigidity etc.) within a graphical semantic network to enable a robotic agent to perform analogical reasoning. Prior work by Zhao et al. (2020)

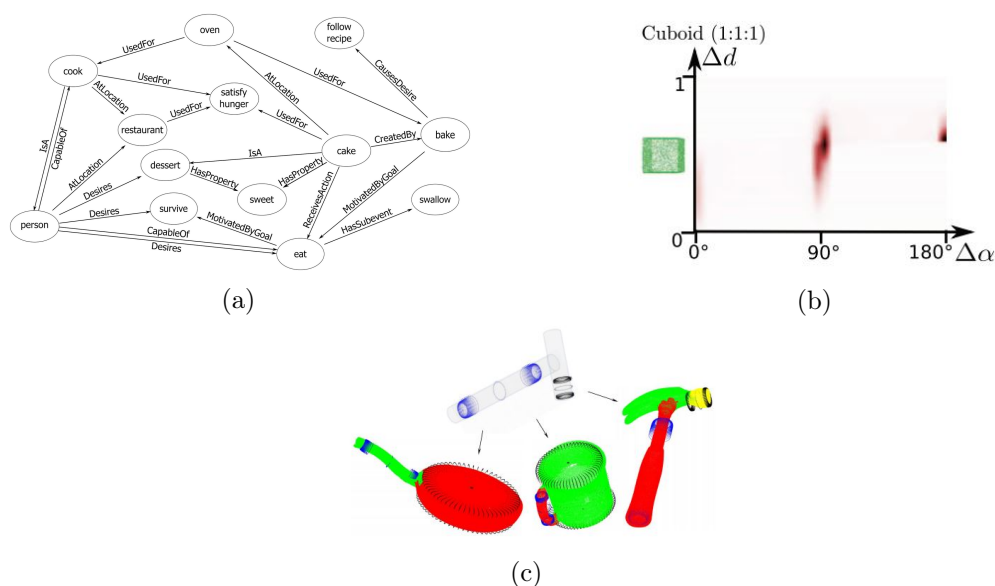


Figure 8: Examples of representations in CPS: a) ConceptNet 5 (Speer & Havasi, 2013): Symbolically represented semantic networks expressing object-action relationships; b) Non-symbolic non-parametric representation of a cuboid as a histogram (Schoeler & Wörgötter, 2015); and c) Non-symbolic parametric representations of objects using superquadrics (e.g., cylinders representing the handles of containers) (Abelha et al., 2016).

introduced the notion of a robot “grammar” (called “RoboGrammar”) for generating novel robot designs for environment traversal. They represent the robot’s morphology using a graph that symbolically describes the different parts of the robot, such as *head*, *body*, *joints*, *connectors*, and *limbs*. Each node within the graph denotes a particular part, and the edges denote the relationship between the parts.

A key advantage of symbolic approaches (planning languages or semantic networks) is that they provide a vehicle for explanatory reasoning that is human readable and allows the reasoning process to be easily understood and interpreted. However, a major disadvantage of symbolic approaches is that in many cases, knowledge must be encoded a-priori, requiring expertise and domain knowledge which can be difficult for non-expert users. Additionally, although knowledge acquisition in symbolic systems is possible (as in the case of the CPS), it can be challenging to acquire complete and useful information in symbolic systems.

4.2.2 NON-SYMBOLIC REPRESENTATIONS

In contrast to symbolic representations, non-symbolic representations do not explicitly model “rules” or “facts”. Instead, non-symbolic representations use *parametric* or *non-parametric* means of representing information. Parametric representations are characterized by the use of numerical parameters that encode some physical meaning. For instance, spheres are characterized by the value of the radius, and the height of a terrain is characterized by the numerical value of its elevation. In CPS, parametric representations are often mathematical models that are pre-specified by the user/developer. In contrast, non-parametric

representations cannot be characterized by physically meaningful parameters and are not pre-specified, but rather learned from observed data (called *representation learning*). While pre-defined mathematical models such as models of physics, are used to parametrically represent concepts such as the behavior of objects, non-parametric representations are learned from data via representation learning. In the following paragraphs, we describe these two types of non-symbolic representations.

MATHEMATICAL MODELS (PARAMETRIC)

The mathematical models used in CPS literature are used to generate outputs, and can be thought of as a function f applied to an input x . A key distinction between mathematical models and representation learning is that the function $f(x)$ is pre-specified by the user rather than learned from observed data. These mathematical models often capture the underlying dynamics of a system, such as the physics of the environment in order to predict how objects will behave in a given scenario. Within CPS, these models include 3D-transformations, models of the 3D scene or environment, geometric models used to represent objects, mathematical models of environment physics, and inverse kinematic models.

Prior work by Fitzgerald et al. (2019) use 3D transformations (rotations and translations) to represent the relationships between tool use trajectories in a source and target environment in order to enable a robot to adapt tool use trajectories learned from human demonstrations in a source environment to a new target environment. In similar work, Abelha et al. (2016) and Gajewski et al. (2019) model tools using Superquadrics (SQs) to identify creative substitutes for missing tools. SQs are geometric shapes similar to other quadrics, but raised to arbitrary powers as opposed to the power of two. SQ representations use 13 parameters that characterize various geometric properties of the tools they represent, such as length, width etc. Prior work by Allen et al. (2019) demonstrate CPS using a Sample-Simulate-Update-Predict (SSUP) model. They model environment physics in order to simulate object behavior in a given environment. The simulations are then used to update a policy for predicting actions that can enable the agent to solve novel puzzles.

Within navigation, prior work has focused on deriving parametric representations of the environment by modeling various attributes such as elevation and gaps. Saboia et al. (2019) model the navigability of an environment by describing mathematical models that capture the height difference between points on the terrain, the pose stability of the robot and, reachability of a particular point on the terrain. These representations are used to enable the robot to construct ramps using available objects in order to traverse the terrain effectively. Closely related work by Tosun et al. (2018) introduce a probabilistic template-based terrain characterization algorithm that uses feature templates to identify regions of elevation, given a 2.5D elevation map of the robot's environment. The robot then uses objects that are available to it in order to construct ramps to scale the elevation. Levihn and Stilman (2015) focus on using objects as simple machines, for enabling a robot to traverse its environment in new ways, such as using a bar of wood to prop open doors. Their work uses an A^* planner defined over a set of constraints in a discretized configuration space of the robot. The constraints model environment physics, such as kinetic energy, momentum, and mass.

More recent work by Murooka et al. (2019) use parametric representations modeled through inverse kinematics to enable a robot to perform augmentation and self-repair by tightening screws on its body. They incorporate the screwdriver and screw dynamics into the inverse kinematic models, including physical concepts such as force, moment, and coefficients

of static friction. In some creative results, the robot attaches hooks to its body to enable it to carry grocery bags. Prior work by Hangl et al. (2017) parameterize the notions of “curiosity” and “boredom” using Shannon entropy, in order to guide the learning of new robot behaviors.

LEARNED REPRESENTATIONS (NON-PARAMETRIC)

Representation learning refers to a set of techniques for learning representations or “features” that are useful for encoding a given set of inputs (Bengio, Courville, & Vincent, 2013). Here, the term “useful” is used to indicate features that are capable of differentiating between classes of inputs. When applied to CPS, non-parametric representations of the initial conceptual space of the agent is learned through observed data. Within the new representation, new concepts are discovered for solving tasks. The vast majority of models used for learning non-parametric representations in CPS include different forms of neural networks, such as Long Short Term Memory networks (LSTMs), Conditional Variational Auto-Encoders (CVAEs), and Feedforward Neural Networks (FNN). In addition to neural networks, few approaches in CPS have also used Support Vector Machines (SVM), k-Nearest Neighbors (kNN), and Decision Trees (DTs). While we do not delve into the details of each machine learning algorithm, we briefly describe the ML approaches used in each paper.

Supervised Learning: Among existing CPS approaches that learn non-parametric representations using neural networks, Xie et al. (2019) use LSTMs to learn skill representations for improvisational tool use. In this case, the robot uses data collected from human demonstrations in order to train the LSTMs to represent tool use trajectories for any given tool, often adapting the trajectory to demonstrate improvisational skills. Similar work by Qin et al. (2020) use the PointNet architecture (2017) for learning keypoint (i.e., points of interest) representations of tools. The network predicts novel keypoints, given an input tool, e.g., Grasping locations for unconventional uses of the tool. Baker et al. (2019) use LSTMs within their policy network architecture to enable a set of agents to learn creative policies that involve using objects to succeed in a game of hide-and-seek. Within the space of tool construction, prior work by Nair et al. (2019, 2019, 2019) have used FNNs and DNNs to perform tool substitution and tool construction. These networks are used to represent objects in terms of their shape and material, in order to identify appropriate objects for constructing or substituting a missing tool. In closely related work, Yang et al. (2020) use Gated Graph Neural Networks (GGNNs) to represent object shapes and reason about pairs of objects for tool construction. Bapst et al. (2019) combine two types of internal representations, namely, Recurrent Neural Networks (RNNs) and Convolutional Neural Networks (CNNs) for the construction of physical structures in simulation. Prior work in tool substitution by Schoeler et al. (2015) uses SVMs to predict alternate uses for given objects. They formulate their work as a part-based affordance learning problem, wherein the models learn shape representations of the tools as histograms (called part signatures) through supervised learning. These representations are then used for identifying tool substitutes.

Apart from representing tools and structures, neural networks have also been successfully applied to represent the agent. In these cases, the non-parametric representations model the agent’s physical design or morphology. Prior work by Pathak et al. (2019) have used Dynamic Graph Networks (DGN) in order to enable an agent to reason about its morphology. Here, the limbs and joints of the agent are represented in the form of a graph learned using DGN,

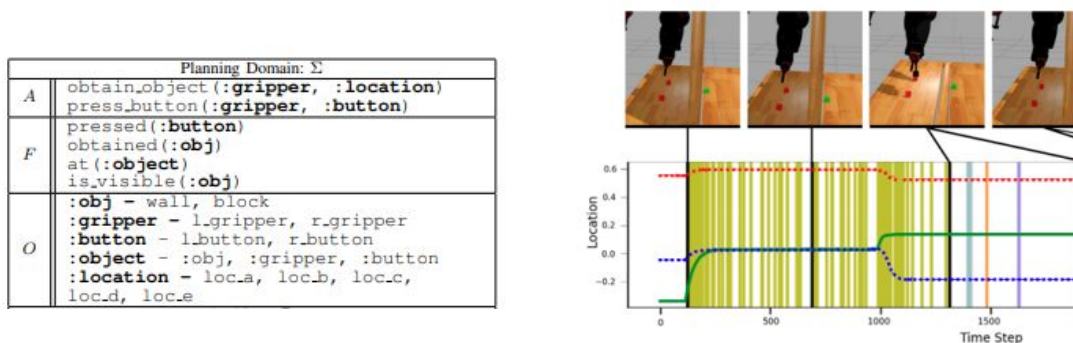


Figure 9: Example of hybrid symbolic and non-symbolic representation in CPS: The symbolic PDDL representation of actions (left) is combined with their trajectory representations (action executors) highlighting change points in the trajectory (right) (Gizzi et al., 2019).

and combined with RL to learn motion policies consistent with the graph structure. Similar work, has also used FNNs to represent agent designs, in conjunction with RL approaches such as REINFORCE (Ha, 2019) or Proximal Policy Optimization (PPO) (Schaff et al., 2019).

Unsupervised/Semi-Supervised Learning: Existing approaches in CPS have also used unsupervised/semi-supervised learning methods. Prior work by Sinapov et al. (2007, 2008) represent tools using k-nearest neighbors (kNN) and Decision Trees (DTs). In their work, the robot gathers observations through repeated interactions with various tools and the collected observations are used to cluster and learn affordance representations for new tools. In Kralik et al. (2016), external state representations are clustered into hierarchical state representations, thus discretizing a continuous non-symbolic state space for RL. Work in generating reusable dynamic movement primitives (DMPs) commonly employ unsupervised learning techniques to extract motion structures for accomplishing tasks. Kroemer et al. (2015) sample non-symbolic state spaces comprised of robot end effector data (end effector torque/position, joint angles, and contact features) and environment data (position, orientation, and relational distances of objects) to learn new representations.

A key advantage of non-parametric representations is that they do not require significant engineering on the part of the user, as they can be learned from data. However, they often lose any meaningful interpretation. In contrast, parametric representations carry some physical interpretation (such as elevation or shape), but require some engineering. In order to leverage the advantages of interpretability and to alleviate hand-coding, several CPS approaches introduce hybrid symbolic and non-symbolic representations.

4.2.3 HYBRID SYMBOLIC/NON-SYMBOLIC REPRESENTATIONS

Hybrid approaches seek to leverage the strengths of both symbolic and non-symbolic approaches. In most cases, the symbolic representations are used to capture domain knowledge, that is then used to improve learning via non-symbolic approaches. Additionally, the framework may switch between symbolic and non-symbolic representations depending on the environment and the nature of the problem. In Zhu et al. (2015), videos of tool-use

demonstrations are decomposed into non-symbolic representations of objects and trajectories and composed to form a symbolic concept graph, containing spatial, causal, and temporal information of objects. Next, SVM clustering and weight-based ranking algorithms are used to capture the “essence” of the task, thereby providing a generative basis for problem representation. As a result, when faced with a new task, the agent is able to use the generative representations to “imagine” the use of candidate tools, eventually choosing an appropriate object. The task transfer work presented by Fitzgerald et al. (2017) uses a previously developed Tiered Task Abstraction (TTA) (Fitzgerald, Goel, & Thomaz, 2015) framework for abstracting actions using low-level object and motion information, and re-grounding them in novel CPS tasks using both trajectory information and high level object and task descriptors. In Xu et al. (2018), actions are called at a symbolic level, where underlying representations exist as low-level object position trajectories relative to the robots gripper frame. In Toussaint et al. (2018), authors propose “Task and Motion Planning” (TAMP) that combines symbolic and non-symbolic representations in classical search problems with optimization as a way to accomplish a task. Specifically, they ground a set of symbolic predicates describing dynamic and kinematic contact as physics-based constraint rules, using non-symbolic processing methods. They define a set of modes in terms of these constraints, and perform optimization using a decision tree. Lastly, these constraints are mapped to the action level using a pre-defined knowledge based on symbolic actions, which have constraints encoded in their end effects. In similar work, Silver et al. (2021) start with using low-level state representations, which are converted into symbolic transitions used for operator learning. This conversion happens through a deterministic parse function, which outputs lossy grounded predicate values (predicates with set arguments). In Gizzi et al. (2019), non-symbolic methods are used when symbolic planning fails to execute or plan for a creative task, respectively. The CPS agent, upon incurring a plan execution failure, attempts to discover new actions (represented both symbolically and non-symbolically, Figure 9) through segmenting formerly known actions. In similar follow-up work, Gizzi et al. (2021) utilize a framework for action discovery that applies low-level parameter variations to discover new actions. Low-level parameter variations change symbolic level predicates by either, a) generating a novel effect, or b) generating a set of effects equivalent to the original action, which are then added back into the knowledge base. The representations used here exist in both high-level symbolic form, and a low-level action controller form. Similarly, in Oddi et al. (2020), a set of 10 continuous low-level state representations are abstracted into predicate form descriptors, to populate the preconditions and effects of newly discovered PDDL operators. This form of skill abstraction is shown to enable problem solving in novel scenarios. Prior work by Nair et al. (2020) combine symbolic planning with non-symbolic representations learned through supervised learning techniques, to enable a robot to construct novel tools. They use feed-forward neural networks to predict “scores” for object-related symbols in the planner, that indicate visual and material fitness of objects for constructing tools. Here, the tool and object shapes are non-symbolically represented, whereas the actions and states are symbolically represented in the planner through PDDL.

The benefit of implementing hybrid approaches is that they limit the amount of expert knowledge that needs to be provided by the user, while also minimizing the data required to learn non-symbolic representations.




Figure 10: Three methods of knowledge manipulation based on Boden’s levels (Boden, 1998): a) Exploratory involving exploration in the universal conceptual space \check{C}_X ; b) Combinational where initial concepts are combined to discover new concepts; and c) Transformational where initial concepts are transformed into new concepts.

4.3 Knowledge Manipulation

How is the initial conceptual space manipulated to create a new conceptual space that enables creative problem solving? In this section, we elaborate on Boden’s three levels of creativity (described in Section 3). Specifically, the agent can discover new concepts by: a) exploring the universal conceptual space (Exploratory), b) combining concepts within the initial conceptual space (Combinational), or c) transforming the initial conceptual space (Transformational) (Boden, 1998). We formalize these approaches (also visualized in Figure 10) in terms of the initial conceptual space C_X , and the new conceptual space C'_X , and also categorize existing CPS literature accordingly.

4.3.1 EXPLORATORY METHODS

In exploratory methods of knowledge manipulation (Figure 10a), the agent searches the universal conceptual space \check{C}_X , in order to discover a new conceptual space $C'_X \subset \check{C}_X$ and $C'_X \not\subset C_X$, that enables the agent to solve the task. Since unguided exploration can be prohibitive, the search is informed by loss/reward functions or heuristics that “guide” the agent to explore specific regions of the universal conceptual space. However, this raises an important question, “*Is all of search or reinforcement learning an example of creative problem solving?*”. In Section 3.1, we noted the criteria of novelty applied to evaluate the generated output. Re-iterating our “Tupperware” example; if the search or exploration yields the solution: “Use Tupperware as food container”, it is not considered novel, and hence not an example of CPS. Whereas, if it yields the solution: “Use Tupperware as a soap-dish”, it is considered as an example of CPS. In this survey, we apply this criteria subjectively to identify the subset of the search-RL literature that qualifies as examples of CPS, and we focus our discussion on these papers, highlighting examples of how they satisfy the criteria.

We begin by formalizing exploratory CPS. In the formalization below, the agent uses a loss function, although it can be extended to heuristics and reward functions as well. A newly discovered concept $c'_x \in C'_X$ can be represented as follows:

$$\{c'_x = \operatorname{argmin}_{\check{c}_x} \mathcal{L}(\check{c}_x) \text{ s.t. } \check{c}_x \in \check{C}_X\},$$

Here, \mathcal{L} denotes an appropriate loss function, and C'_X contains novel concepts from the universal conceptual space such that $C'_X \not\subseteq C_X$ (Figure 10a). From our grid-world example, the agent may discover how to turn off the lights by pushing boxes around the room guided by an appropriate reward function, in order to eventually discover that it can accomplish the goal by pushing the box on top of the off button. Here, the agent discovers that in addition to moving objects to empty locations, they can be moved on top of some other objects.

LOSS/REWARD FUNCTIONS

Within existing CPS literature, prior work has focused on using cost and reward functions that guide the agent towards creative behavior. Xie et al. (2019) use a cost function that is the expected Euclidean distance between current and goal positions of a target object, in order to guide a robot to use tools in an improvised manner to move the target object from the initial to the goal position. The per-time step costs are summed together and used to select an appropriate tool-use trajectory with the lowest cost. Their work yields novel, i.e., non-prototypical uses of objects previously unknown to the robot, such as using a knife to pull objects closer. Within multi-agent systems, Baker et al. (2019) use reward functions for policy learning, wherein the reward functions seek to maximize the total expected discounted return for each agent within the multi-agent population. Their work demonstrates emergent creative tool use, such as using boxes to build a fort to win at hide-and-seek. Prior work by Bapst et al. (2019) use reward functions that enable the a set of agents to construct novel physical structures that achieve different tasks such as connecting multiple structures or covering existing structures.

In the context of discovering creative agent designs, Pathak et al. (2019) use reward functions that guide multi-agent systems to assemble into a single agent, leading to the discovery novel agent morphologies. Each self-assembling agent can be considered as a primitive limb with linking actions to join other limbs, and the reward function seeks to maximize the reward for each limb within the joint morphology. The reward itself measures the capability of the joint morphology to perform locomotion tasks. Similar prior work use reward functions within the REINFORCE algorithm (Ha, 2019) or Proximal Policy Optimization (PPO) (Schaff et al., 2019) to enable an agent to optimize for novel agent designs to traverse its environment in new ways.

HEURISTIC SEARCH

Apart from using reward and loss functions, the agent’s search may also be guided by heuristics. Heuristic search is commonly observed in the planning literature, wherein the agent uses informed search over a planning space, guided by heuristics, to output a task plan. Prior work by Erdogan and Stilman (2013), use a pre-defined action ordering (an uninformed heuristic) to decide which child nodes to inspect first during the search. Their work enables a robotic agent to construct novel structures for navigation by effectively searching through the configuration space of the objects. In similar work, Choi et al. (2018) enable robots to construct makeshift structures (e.g., ramps and bridges) using a unified forward-chaining planner that applies a numeric heuristic to guide a best-first search process. Their heuristic favors states that satisfy more elements in the goal formulae. Levihn and Christensen (2014), introduce an approach that samples a set of contact points between objects, and sorts the contact points using a custom scoring function (heuristic). Their work enables robots to use unconventional objects such as planks of wood, or loaded carts to open jammed doors.

In Chitnis et al. (2021), the agent builds a world transition model through state-based goal exploration, where selected goals are optimized on a novelty measure. In the space of tool creation, prior work by Wickasono and Sammut (2017, 2020) use heuristics to guide autonomous tool creation by a robot, where the heuristic prioritizes new tools that are similar to an existing one, where the similarity between tool representations is computed as the number of edit operations needed to transform one representation into the other. Nair et al. (2020) use supervised learning techniques to compute object fitness scores that are then incorporated into existing planning heuristics within A^* planning to enable a robot to perform tool construction.

In the context of other high-level tasks, prior work that uses heuristics have focused on high-level task planning (Boteanu et al., 2015), sequential manipulation (Toussaint et al., 2018; Allen et al., 2019), action discovery (Suárez-Hernández et al., 2020), and agent design optimization (Zhao et al., 2020). Prior work by Boteanu et al. (2015) use fitness functions to evaluate candidate objects that can serve as substitutes for a missing object within a hierarchical planning framework (HTN). Their approach to identifying substitute objects proceeds in three steps: generating candidates for the target, extracting contextual information from the HTN, and evaluating the fitness of each candidate within the context. Prior work by Toussaint et al. (2018) use the Multi-Bound Tree Search (MBTS) approach to enable a robot to perform creative sequential manipulation tasks using tools. MBTS acts as a best-first search approach, and is applied to a hybrid symbolic and non-symbolic Logic Geometric Program (LGP) to solve for the tool manipulation trajectory. Allen et al. (2019) use heuristics derived from physics simulators to evaluate agent actions based on their simulated outcomes. The heuristics guide the agent to discover creative policies for solving tool-based puzzles. Suarez-Hernandes et al. (2020) introduce cost functions to guide the search for the discovery of new PDDL actions. Their cost function encodes the number of editions required for the new action, such as the number of additions and deletions in effects, as well as changes in the number of preconditions for the actions. Similarly, Gizzi et al. (2021) apply behavior babbling to discover new actions, by varying action controller parameters by partitioning their values evenly along a range of permissible values. For example, given a min and max value of a parameters, of 0 and 100 respectively, a user selected partition value of 5 would cause babbling experimentation for 1, 25, 50, 75 and 100. Oddi et al. (2020) leverage a “competence”-based intrinsic heuristic for skill learning, where goals are generated to facilitate exploration. Prior work by Zhao et al. (2020) use Graph Heuristic Search over a conceptual space of graphs representing robot configurations. The heuristic function is learned as the search progresses, using ground-truth data from MPC-based (Model Predictive Control) evaluations of the robot’s performance.

4.3.2 COMBINATIONAL METHODS

In combinational methods of knowledge manipulation (Figure 10b), the agent discovers a new concept $c'_x \in C'_X$ by *combining* existing *distinct* concepts in C_X . Hence, the newly discovered concept c'_x can be thought of as a function of distinct concepts c^i_x in the initial conceptual space C_X . A key point to note is that, once the new concept c'_x is discovered through combinational manipulation, search or exploration may be used to evaluate or search through the newly discovered conceptual space in order to identify an appropriate solution

for completing the task (analogous to the “Aha” or insight moment described in Section 3.2). However, the discovery of new concepts itself does not happen through exploration, as is the case with exploratory methods of Section 4.3.1.

Combinational methods operate as a function over a set of concepts $c_x \in C_X$. Formally, we define a function f that combines k distinct concepts in C_X to discover new concepts², where a single concept $c'_x \in C'_X$ is represented as follows:

$$f : C_X \rightarrow C'_X \mid c'_x = f(c_x^1, \dots, c_x^k); c_x^1, \dots, c_x^k \in C_X, C'_X \not\subseteq C_X$$

Within our grid-world example, combinational methods involve the agent reasoning about combinations of objects, e.g., combination of the box and switch in order to discover that the box can keep the switch pressed. Here, the box-switch combination is a “composite” object that is newly discovered, enabling the agent to solve the task. The states of the box-switch composite can be thought of as a combination of the initial concepts regarding the states of the switch and states of the box, wherein the switch “*is_pressed*”, if the box is “*on(switch)*”.

PAIR-WISE CONCEPT COMBINATION

Within CPS, some existing works have focused on the use of pair-wise combinations, wherein *two* concepts are combined to discover new concepts. Nair et al. (2019, 2019) introduce the “Robogyver” framework for the construction of tools by combining pairs of available objects. They reason about visual properties of objects to output novel tool constructions that combine the properties of the individual objects, e.g., coin + pliers = screwdriver. Here, the visual “flatness” of the coin serves as the head of the screwdriver, while the “handle” of the pliers serves as its handle. The novel objects (concepts) are a combination of the visual properties of objects in the initial conceptual space. In closely related work, Yang et al. (2020) reason about object shapes using Gated Graph Neural Networks, that model relationships between pairs of objects for tool construction. The network takes depth maps of the available objects, and a reference tool that needs to be constructed. It outputs the pair of objects that can be combined to best match the provided reference tool. The resultant objects are a combination of the depth maps of the objects in the initial conceptual space. In contrast to reasoning about objects, prior work by Colin et al. (2019), show how RL can be used to combine and generalize behavior in agents. In their simulated experiment, they show that after a period of *shaping* (wherein the agent is exposed to a-priori reinforcement-based training), the agent is able to combine two actions ($a_1 =$ jumping on a box to peck a banana for a food reward, $a_2 =$ pushing a box to a green dot) and generalize their combination to a CPS task ($a' =$ pushing box to location under the banana in order to reach it and peck it for a reward).

MULTI-CONCEPT COMBINATION

In contrast to combining two concepts, existing works have also looked at combining more than two concepts. Prior work by Olteteanu and Falomir (2016) focus on composing available objects to create new objects for accomplishing a task (object composition), not limited to pair-wise combinations alone. They introduce the Object Replacement and Object Composition (OROC) framework that combines semantic tags or features associated with objects. Thus the novel concepts are a combination of the semantic features in the initial

2. In cases where more than one distinct concept is not combined, i.e., only a single input is provided to f , we define the function to be an identity function, $f(c_x^i) = c_x^i$

conceptual space. Closely related is the work of Lieto et al. (2019), where object substitution is found through combining existing object “concepts”, through conceptual blending methods (Fauconnier & Turner, 2008). Solution objects are selected from a set of candidate objects which are evaluated based on their rank of object similarity to the object being substituted. Prior work by Hangl et al. (2017) focus on learning new skills or behaviors as a composition of previously known behaviors. They introduce an approach that enables a robot to discover new behaviors through behavior composition. Given a set of behaviors B , new behaviors can be defined as a composition of behaviors $b_i \in B$, as $b_l \circ \dots \circ b_2 \circ b_1 \circ b^\sigma$. The goal of their approach is to extend the domain, where the domain is the set of states in which a skill can be applied successfully. The composite behaviors extend domain applicability if they can be applied successfully, i.e., $success(b_l \circ \dots \circ b_2 \circ b_1 \circ b^\sigma(e)) = true$.

4.3.3 TRANSFORMATIONAL METHODS

In transformational methods, the agent transforms the initial conceptual space C_X into a new conceptual space $C'_X \not\subseteq C_X$ via some function or “transform”. The agent then derives the creative solutions from the newly transformed conceptual space C'_X . Here, exploratory methods may be used to evaluate or search through the newly discovered conceptual space $c'_x \in C'_X$. However, the discovery of new concepts itself does not occur through exploration.

We formalize transformational manipulation using a transformation function or transform f . A concept $c'_x \in C'_X$ can then be represented as follows:

$$f : C_X \rightarrow C'_X \mid c'_x = f(c_x) \forall c_x \in C_X, C'_X \not\subseteq C_X$$

Thus, f denotes a surjective function that maps every concept $c_x \in C_X$ to a new concept $c'_x \in C'_X$. Transformational creativity thus involves a mapping from the initial conceptual space to a new conceptual space, via an appropriate transform. In our grid-world example, the agent may transform the initial conceptual space of actions into a new conceptual space that captures the forces applied by the actions. Within this new conceptual space, the agent may discover that pushing a box onto the switch applies a downward force, thus keeping the switch pressed. However, it is a challenging problem to identify the appropriate function for transforming the initial conceptual space into one where the solution for the task becomes evident. This closely relates to prior work by Oltețeanu et al. (2015) (presented in Section 2), discussing re-representation within CPS, i.e., thinking about the different representations of concepts wherein particular representations yield the solution to the task. A key point to note that there is a strong correlation between transformational methods and parametric mathematical representations as discussed in Section 4.2.2. Examples of transformations in CPS mostly include geometric and graphical transformations, with some other paradigms.

GEOMETRIC TRANSFORMATIONS

Geometric transformations manipulate geometric aspects of the concepts in the initial conceptual space, e.g., applying rotations and translations. Thus, the concepts are re-represented in terms of their geometrically transformed versions. Prior work by Fitzgerald et al. (2017, 2019) use 3D rotations and translations in order to re-represent skill trajectories in a new conceptual space where the skills can be re-applied to novel tools. Within the new conceptual space, the robot would be able to adapt previously known skills to new tools that it has not seen before. In similar work, Abelha et al. (2016), and Gajewski et al.

(2019) introduce techniques for improvised tool use by mapping the initial conceptual space of tool point clouds to a new conceptual space consisting of their geometric representation using superquadrics. Within the new representation, similarities between tools are utilized to identify how to adapt tool-use trajectories from one tool to another. Sinapov et al. (2007) learn novel tool affordances by transforming the sensory input of the robot (i.e., a 2D image), into five perceptual functions computed over the perceived sensory input. Each of the perceptual functions represent a geometric transform, such as computing image center or gripper location.

GRAPH TRANSFORMATIONS

Graph transformations re-represent the initial concepts as a set of features within a graph. Each initial concept is mapped to a feature node in a new graph space. The initial conceptual space may or may not be a graph, but the transformed conceptual space is always a graph. Prior work by Zhu et al. (2015) transform tools into three new graphs: spatial, temporal, and causal parse graphs. Together, they highlight 13 concepts associated with each tool and its use, e.g., material, density etc. Within this new conceptual space, a robot is able to identify appropriate substitute tools for a task. Closely related, prior work by Schoeler et al. (2015) transform tool point clouds into a set of part “signatures” indicating their shape and pose, mapped into a graphical representation of the tool. This graph serves as the new conceptual space for identifying tool substitutes. Prior work by Freedman et al. (2020) represent the initial conceptual space as a graph that captures various properties of objects such as height, weight, rigidity etc. The initial graph is then transformed to a new graphical representation by computing a *maximal common edge subgraph* (MCES) over the initial conceptual space. Within the new graph, the robot performs analogical reasoning to identify appropriate object substitutes for replacing missing objects in novel contexts.

OTHER TRANSFORMATIONS

Several approaches in CPS have also introduced other transformation functions. Kralik et al. (2016) use complexity reduction to restructure an internal belief representation during the incubation stage of CPS, providing a basis for spontaneous generalization to novelty in problem solving. In navigation, Saboia et al. (2019) introduce an approach for enabling robots to construct ramps for navigation by transforming the initial conceptual space representing the environment via a custom “height function”, defined as a mapping from a construction area Q to \mathbb{R}^+ , as $h : Q \rightarrow \mathbb{R}^+$. In similar work, Tosun et al. (2018) map the initial conceptual space representing the environment to a set of “templates” that characterize ledges within the environment. The templates are then used by a set of modular robots for the construction of ramps. Similar transformations have also been applied to the action space. Multiple works have shown transformation of actions through abstracting low-level trajectories into higher level behavior models, used for generalization to novel tasks (Kroemer et al., 2015). Xu et al. (2018) transform low-level trajectory-based tasks into high level actions by decomposing the tasks to exploit modular substructure. Similarly, prior work by Gizzi et al. (2019) transform the actions (defined at a symbolic level) into a representation of their low-level trajectories, in order to discover new actions via segmentation of the low-level representation. Thus, transforms are applied to the initial conceptual space consisting of high-level symbolic actions, to create a new conceptual space containing their low-level trajectory representation. Prior work has also focused on transforming actions at a strictly symbolic level in order




Figure 11: Evaluation modes in CPS: Existing physics-based simulators such as OpenAI (Schaff et al., 2019) (left), and physical evaluation on real robots, including custom platforms (Saboia da Silva et al., 2019) (right).

to discover new actions without any change to the underlying low-level mechanics of the actions (Sarathy et al., 2020; Silver et al., 2021). Murooka et al. (2019) re-represent motion trajectories using inverse kinematic models that account for the physics of screw tightening, such as force, momentum and friction. Within the new conceptual space, the robot was able to repair itself, as well as augment its capabilities by attaching hooks to its body for carrying bags. Lastly, Qin et al. (2020) re-represent tools using “Keypoints” that identify specific points of interest on a given tool point cloud. The tool keypoints include, *grasp point*, *function point*, and *effect point*. The keypoint representation of the tool is then used to derive novel and unconventional tool manipulation trajectories.

With the three modes of knowledge manipulation discussed here, we highlight an important connection between CC and AI, in the context of CPS. Note that existing work in CPS falls into one of the three modes of knowledge manipulation described above, although in the future, CPS frameworks could combine multiple modes of knowledge manipulation.

4.4 Evaluation

How is the novel conceptual space evaluated? This section discusses the different modes of evaluation adopted by existing CPS approaches. The mode of evaluation helps distinguish theoretical models from models that have been tested in real-world settings. The classes of evaluation here include, a) using a simulated environment (*Simulation*), b) using a real robot (*Real robot*), and c) developing a standard benchmark test to evaluate the approach.

4.4.1 SIMULATION-BASED EVALUATION

In this mode of evaluation, the conceptual space is evaluated in a simulated environment. This mode of evaluation is beneficial because it does not require physical access to agents, and can often generate a large number of trials in a small amount of time. Additionally, simulation-based evaluation allows for greater customization capability. Evaluation in simulation happens either as a holistic, end-to-end proof-of-concept evaluation, or as a way to assess a particular feature of a CPS framework.

PHYSICS-BASED SIMULATORS

In proof-of-concept cases, it is typical to use a 2D or 3D physics-based environments, so the

end-to-end behavior of the agent can be visually and algorithmically validated. Common examples of such environments include “Gazebo” (Gizzi et al., 2019; Levihn & Christensen, 2015; Toussaint et al., 2018; Gizzi et al., 2021; Oddi et al., 2020), “Unity” (Bapst et al., 2019; Pathak et al., 2019), “OpenAIGym” (Wang, Lehman, Clune, & Stanley, 2019; Ha, 2019; Schaff et al., 2019), “Mujoco” (Baker et al., 2019), “DeepMind Lab” (Leibo, Hughes, Lanctot, & Graepel, 2019), “PyBullet” (Qin et al., 2020; Chitnis et al., 2021; Silver et al., 2021), and “Bullet Physics Library” (Zhao et al., 2020). While all of these environments provide similar support, each cater well to different CPS cases. Gazebo is known for its realistic 3D graphics whereas Unity is known for its speed, especially for data-heavy CPS applications that use RL. OpenAIGym is a platform that is typically known for its broad range of RL problem domains, provided off-the-shelf. These domains range from text-based problems, to simple 2D problems, to 3D robotics problems (however these more complex domains are somewhat limited compared to Gazebo and Unity). The open-sourced DeepMind Lab platform provides support for game-based environments as a method for CPS evaluation. Both OpenAIGym and DeepMind Lab provide game-based evaluation environments. Other examples of less common simulation environments include “DART” (Levihh & Stilman, 2014), “CREATE” (Jain et al., 2020), and “BREVE” (Sinapov & Stoytchev, 2007).

CUSTOM SIMULATORS

In the case of validating specific aspects of a CPS approach, it may be more favorable to use a specialized or customized simulation environment for evaluation. In simple cases, validation may happen using basic RGB imaging in a 2-D world (Colin & Belpaeme, 2019; Sarathy et al., 2020). In more complex cases, computer vision algorithms built into the simulators are used to identify entities in a physical environment, which are then intelligently processed by a CPS method (Zhu et al., 2015; Abelha et al., 2016; Schoeler & Wörgötter, 2015). Here, CPS algorithms reason about the world without actually interacting with the physical environment. In a more elaborate case, Wang et al. (2019) focus on automatic or AI-generated testing environments, using the resultant performance to automatically improve the system with a feedback loop. Additional approaches in CPS have also evaluated their methods in simulation, not using a specific simulator, but rather evaluating their approach in terms of simulated metrics (Suárez-Hernández et al., 2020; Freedman et al., 2020).

While simulation-based evaluation is less precise than real-world dynamic environments, many approaches in CPS have been evaluated in this manner, owing to the ease of generating and testing in a variety of environments.

4.4.2 REAL WORLD EVALUATION

In this mode of evaluation, the approaches are evaluated on a real, physical robot. This mode of evaluation can be advantageous in that they can evaluate performance of the approaches against real-world dynamics and noise models.

ROBOTS

Several existing works in CPS have focused on evaluating the models on various physical platforms such as Sawyer (Xie et al., 2019; Xu et al., 2018), Baxter (Wicaksono & Sheh, 2017; Wicaksono & Sammut, 2020; Yang et al., 2020), Kinova (Fitzgerald et al., 2017, 2019; Nair et al., 2019, 2019; Nair & Chernova, 2020), PR2 (Gajewski et al., 2019; Murooka et al., 2019), KUKA (Hangl et al., 2017), Darias (Kroemer et al., 2015), and custom platforms

that are specifically built for the tasks that the robot is required to perform (Saboia da Silva et al., 2019; Boteanu et al., 2015; Tosun et al., 2018; Hangl et al., 2017). The Sawyer and Kinova platforms each have a single 7-DOF robot arm with a stationary base. Hence, they are not well suited for mobile manipulation applications. The PR2 and ATLAS platforms each consist of two 7-DOF robot arms for bi-manual manipulation, while also supporting navigation. Both the Kinova and Sawyer platforms are quite commonly used for approaches that involve learning from demonstration (LfD) since the robot arms are easily manipulated by humans, particularly when mobility is not a requirement. The custom platforms used in (Saboia da Silva et al., 2019; Boteanu et al., 2015) consist of robot arms that are attached to a mobile platform in order to enable mobile manipulation, in addition to specific design aspects that make the platform well suited for the task at hand, such as navigation. For example, the custom platform used in (Saboia da Silva et al., 2019) is built to handle highly uneven terrains. The custom platform used in (Hangl et al., 2017) consists of two KUKA robot arms that are attached to Schunk SDH grippers in order to perform bi-manual manipulation of objects. The custom platforms developed in (Tosun et al., 2018) consist of populations of small modular robots that can attach to one another, with navigation capabilities.

SENSORS

In terms of sensing capabilities of the platforms, the vast majority of evaluations predominantly test the approaches in unimodal settings (i.e., using a single sensor). These approaches often use only RGB-D sensors that capture partial and noisy point clouds. A small subset of CPS papers evaluate the approaches in multi-modal settings as well, such as using spectrometers in addition to visual inputs (Nair et al., 2019, 2019). Multi-modal sensing capabilities can be especially useful for CPS, since the robot may have to reason about multiple modalities of the objects, including weight, materials, forces etc., rather than solely relying on visual properties, in order to derive creative solutions.

Despite the benefits of evaluating CPS approaches in real-world settings, evaluation on physical robots tend to be slower than simulators, due to hardware related shortcomings.

4.4.3 BENCHMARK

In this section, we discuss papers that introduce benchmarks for evaluating CPS algorithms. Benchmarks allow a standardized comparison of different methods. Designing systems to evaluate creativity often depends on the subjective definitions of creativity proposed by the designer, and more general benchmarks of CPS are yet to be developed. A universal benchmark of CPS would allow a uniform and fair comparison of different approaches. We discuss this limitation in greater detail in Section 7. Note that a subset of the papers in this section only introduce benchmarks, but not a specific CPS approach itself. However, we include these papers here for completeness, discussing them in this section only.

Existing benchmarks in CPS typically evaluate specific types of creative tasks, rather than a test of general creative intelligence, e.g., benchmarks that focus specifically on creative tool use (Allen et al., 2019). In many cases, the CPS system is evaluated using an output-based approach which compares the output to that of a human participant. In the case of the “Alternative Uses Test,” a CPS system is evaluated in its ability to generate alternative uses of an object by bench-marking its output against the set of corresponding alternative uses generated by a human participant (Guilford, 1967b). This benchmark has been used

specifically in evaluating CPS systems which employ combinational creativity (Olteşeanu & Falomir, 2016; Lieto et al., 2019). In some cases, the percentage of CPS problems successfully solved in a given set is used in bench-marking, either in comparison to human performance, or in a standalone manner. Bisk et al. (2019), propose a benchmark to evaluate common-sense physical reasoning capabilities for CPS by comparing the percentage of accurate reasoning cases to human performance. Kralik et al. (2016), compare the performance of their Hierarchical RL system to empirical data generated from trials of CPS in Rhesus monkeys. Guzdial et al. (2018), propose a cross domain metric for evaluating the ability of a CPS system to obtain a goal state, where given a specific domain, an “Uncreative Max” (UM) baseline is developed to represent a solution most similar to the desired goal state, given CPS was not employed. Creativity is then measured as any positive score differential between a creative agent and its corresponding UM score. The challenge with this method is in the domain specific crafting of the UM baseline, which is shown to have variance across different domains. This metric is also limited in that it is specific to combinational creativity.

Benchmarks that extend existing simulation software have also been proposed. The “POET” (Paired Open-Ended Trailblazer) benchmark (Wang et al., 2019), modifies the “Bipedal walker hardcore” in OpenAI Gym to create challenging traversal environments for a bipedal walker, requiring the generation of creative traversal policies. Similarly, “NovelGridWorlds” (Goel, Tatiya, Scheutz, & Sinapov, 2021) present an OpenAI Gym environment framework for evaluating agents that can adapt to sudden novelties in their environments. NovelGridWorlds includes crafting tasks for a bow and pogostick, that requires collection of different resources and presents the AI agents with different classes of novelties, such as object, attributes and action novelties.

Current CPS benchmarks exist solely in simulation, limiting the practical applicability of the methods. Designing benchmarks that evaluate different domains of creative problem-solving, and can be tested on real-robots, remains an open problem (see Section 7).

5. Types of Conceptual Spaces

In the previous sections, we discussed the overall CPS framework and presented three approaches for discovering a new conceptual space. However, *what specific information does the conceptual space contain?* As described previously, the conceptual space can be associated with states or actions. While all problem formulations consist of states and actions, either the action space or the state space is manipulated to discover new actions or states. In this section, we categorize existing CPS literature on the basis of the information that is manipulated (also shown in Figure 12).

5.1 States

Based on our review, we categorize the existing literature in CPS based on the types of states that are manipulated. These include, a) states of objects in the environment (object modification); b) states of the environment terrain, e.g., holes or ledges in the environment (terrain modification); and c) states of the agent itself (agent modification). In these cases, creativity arises from modifying concepts related to objects, concepts related to the terrain (i.e., modification of the terrain), and concepts related to the agent’s design or morphology.

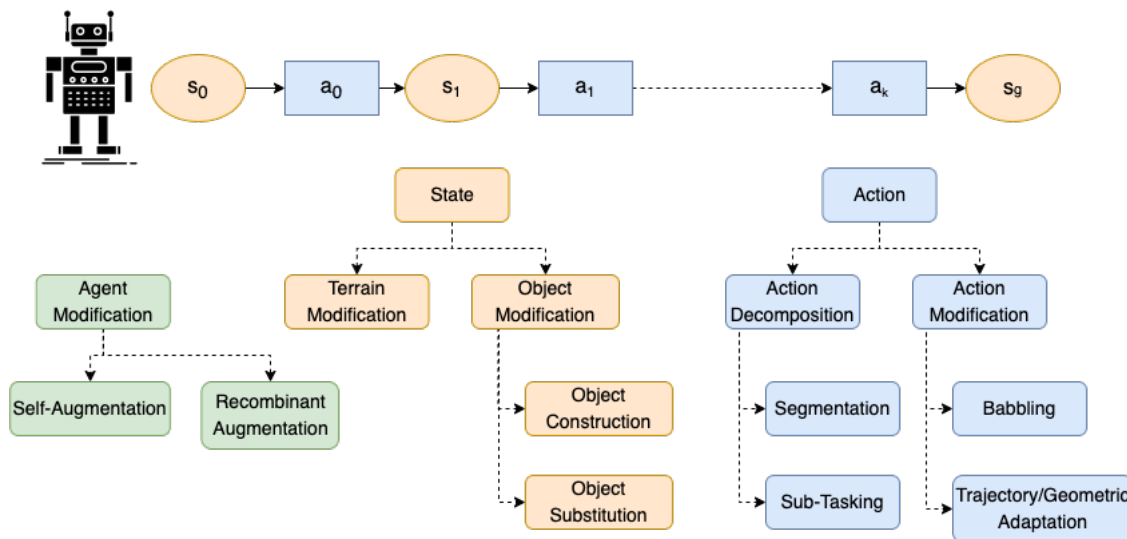


Figure 12: The conceptual spaces of the agent that are manipulated for CPS can relate to states or actions. The state space can involve concepts regarding the environment (terrain modification and object modification), or concepts regarding the agent (agent modification). Similarly, manipulating the action space can involve babbling in the existing action space, or modifying existing actions.

5.1.1 OBJECT MODIFICATION

The CPS literature presented in this subsection involves *modeling and reasoning about the objects themselves*, such as their visual and material properties. Hence, the conceptual space manipulated in these approaches is the states of objects in the environment. Commonly, these problems are referred to as “Macgyvering”, defined as “solving problems creatively using whatever objects are available at hand” (Dictionary, 1989). Macgyvering as a sub-class of CPS has been researched in both AI (Olteteanu & Falomir, 2016; Sarathy & Scheutz, 2018) and Robotics (Erdogan & Stilman, 2016; Nair et al., 2019, 2019). Existing CPS research in object modification can include a) object substitution that involves adapting objects for non-prototypical uses, e.g., using a pan as a hammer, or b) object construction that involves creating new objects, e.g., constructing a hammer from a rock and stick.

OBJECT SUBSTITUTION

Existing research in object substitution often involves comparing the physical attributes of the available objects to those of the missing object in order to identify potential substitutes. Prior work has focused on modeling objects in terms of geometric shapes (Abelha et al., 2016; Gajewski et al., 2019) or shape histograms (Schoeler & Wörgötter, 2015), and reasoning about the objects by computing their similarities to the missing object that is being substituted. The similarity is computed based on the difference between the geometric parameters or a score computed from the shape histograms. In contrast to using geometric similarities alone, Shrivatsav et al. (2019) compare both material and shape properties of objects to identify good substitutes. Prior work has also focused on modeling and reasoning about objects via semantic networks representing object properties such as affordances and visual properties




Figure 13: Examples of object modification in CPS: Substitution where robot uses non-prototypical objects to chop wood (Zhu et al., 2015) (left), and construction where robot constructs tools by combining individual objects, such as a screwdriver and foam block, to make a squeegee (Nair et al., 2019) (right). In both cases, the agents reason about object properties, e.g., shape and materials.

(Boteanu et al., 2015; Olteteanu & Falomir, 2016; Lieto et al., 2019; Freedman et al., 2020). Zhu et al. (2015) present a “hybrid” approach that reasons about visual properties of objects as well as their physical attributes like mass and density to identify substitute tools, and to suggest appropriate tool-use trajectories for the substitutes. In similar work, Levihn et al. (2015, 2014) reason about weight and mass in order to identify objects that can supply the required amount of force, or support the desired amount of weight in order to accomplish navigational tasks.

OBJECT CONSTRUCTION

In contrast to object substitution, object construction involves combining available objects to construct new objects with desired capabilities. These approaches often involve reasoning about individual objects that can be joined to create a new object that has a combination of their individual properties. There is currently very limited research in the area of object construction. The OROC framework (Olteteanu & Falomir, 2016) proposed by Olteteanu and Falomir reason about “object composition” using semantic knowledge. In this case, the semantic concepts encapsulate affordances of objects, enabling the agent to reason about individual capabilities of objects, e.g., matchbox as a container, and tacks as an attachment method. These can be combined to create a composite object with a new capability, e.g., combining the matchbox and tacks to create a container that can be attached to the wall. In contrast to reasoning about semantic attributes, prior work has also focused on combining objects through visual reasoning alone (Yang et al., 2020) or visual and material reasoning (Nair et al., 2019, 2019; Nair & Chernova, 2020). These approaches decompose a reference tool into sub-parts, reasoning about objects that are similar to each sub-part and can be combined to create the reference tool. Similar work by Erdogan and Stilman (2016) focus on the autonomous construction of simple machines, further incorporating reasoning about physical concepts such as mass and weight. Wicaksono and Sammut (2017, 2020) focus on creating novel tools from polymers through 3D printing by encoding visual attributes of the tool such as length of the handle or angle of the hook (at the end of the tool). In all of the cases described above, the manipulated conceptual space specifically encodes properties of the objects.



Figure 14: Examples of terrain modification in CPS: Example of ADFS where the agent constructs a stable structure to traverse a gap in the terrain (Erdogan & Stilman, 2013) (left), and terrain modification in hide-and-seek where the blue agents build a fort to hide from red agents (Baker et al., 2019) (right). In contrast to object modification, both these cases involve reasoning about environment properties, such as gaps and walls, as opposed to reasoning about the objects themselves.

5.1.2 TERRAIN MODIFICATION

The CPS literature discussed in this section involves *modeling and reasoning about the terrain*, e.g., modeling and reasoning about gaps or elevations in the terrain that the agent is unable to cross. Hence, the manipulated conceptual space involved in these approaches is the state of the terrain. Terrain modification often involves modifying unstructured environments to facilitate navigation. Note that terrain modification *can* involve using objects in the environment to modify the terrain. However, the key distinction between terrain and object modification is that the newly discovered concepts relate to the state of the terrain rather than some object in the environment. Further, terrain modification includes approaches that explicitly reason about terrain attributes such as elevation.

NAVIGATIONAL TASKS

In the context of terrain modification for navigational tasks, “Automated Design of Functional Structures” (ADFS) (Erdogan & Stilman, 2013) deals with construction of structures such as bridges to improve terrain navigability. Prior work in ADFS has focused on visual reasoning about the height of elevation points on the terrain (Erdogan & Stilman, 2013; Tosun et al., 2018) and symbolic reasoning of terrain properties such as the widths of a gap (Choi et al., 2018) for traversing the environment in novel ways. More recently, Saboia et al. (2019) focus on the construction of makeshift ramps using compliant bags, by modeling and reasoning about the reachability or navigability between any two given points on the terrain, in addition to other factors such as terrain elevation.

NON-NAVIGATIONAL TASKS

Beyond ADFS, terrain modification has also been applied to other domains such as games (Baker et al., 2019), and construction of structures for purposes that are non-navigational (Bapst et al., 2019; Colin & Belpaeme, 2019). Prior work by Baker et al. (2019) demonstrated terrain modification using available objects within a multi-agent system, where the agents attempt to succeed in a game of hide-and-seek. These agents reason about the environment, e.g., walls, to devise strategies such as fort building to *create* unreachable locations in the terrain (as hidiers), or to *traverse* unreachable locations to find other agents (as seekers).




Figure 15: Examples of agent modification in CPS: Self-augmentation where a single agent augments its own capabilities, e.g., through self-repair and self-extension (Murooka et al., 2019) (left), and recombinant augmentation where individual agents (“limbs”) combine to form a composite agent, in which case both policy and agent design are optimized (Pathak et al., 2019) (right).

Prior work by Bapst et al. (2019) construct physical structures that reason about, and modify the terrain, for achieving different goals such as connecting separate structures or covering existing structures. In the CPS task of Colin et al. (2019), the agent learns to move a box to a certain location in order to obtain a previously unreachable reward. Lastly, in Kralik et al. (2016), the agent restructures its representational belief about its environment as a means for CPS. Note that in the mentioned non-navigational cases, agents *implicitly* reason about the terrain in the context of reward functions, in contrast to the approaches in ADFS that explicitly model the terrain attributes.

5.1.3 AGENT MODIFICATION

The CPS literature presented in this subsection involves approaches wherein the *agent models and reasons about itself*, e.g., modeling and reasoning about the joints of the agent’s body. Hence, the conceptual space involved in these approaches is the state of the agents’ body. Depending on whether the approach reasons about a single agent or multiple agents, we divide the existing literature into two classes: Self-augmentation (for single agents), and recombinant augmentation (for multiple agents). Note that in most cases, creativity in agent modification arises from discovering novel agent designs.

SELF-AUGMENTATION

Self-augmentation involves cases where a single agent augments or modifies its own body, in order to accomplish the task. Murooka et al. (2019) describe a novel algorithm for self-repair and self-extension of robots. In this work, they model the physical body of the robot using CAD to enable the robot to reason about itself, further utilizing inverse kinematic models of the joints of the robot. By applying these models, the robot is able to tighten screws located on its own body, either for self-repair or to augment its capabilities by attaching hooks to enable them to carry more bags. Prior work by Ha (2019) jointly optimize the agent’s design and policy for navigation tasks. They model various properties of the agent’s morphology such as mass, and the orientation of the agent’s body parts and joints. Each of these properties are parameterized and incorporated into the policy network. In similar work, Schaff et al. (2019) parameterize the lengths and radii of the links within the robot’s configuration, wrapped into an optimization function that is used to guide the policy learning.



Figure 16: Examples of action modification in CPS: Action babbling through repeated interactions with the tool and object (Xie et al., 2019) (left), and action modification that adapts known tool use trajectories to new target objects, e.g., from a knife to a spatula (Gajewski et al., 2019) (right).

In contrast to modeling agents within an RL framework, prior work by Zhao et al. (2020) introduce “RoboGrammar”, that models robot designs as a graph. These graphs represent various joints, body parts, and connectors, that make up the configuration space of the robot.

RECOMBINANT AUGMENTATION

Recombinant augmentation involves populations of agents that combine in order to augment the capabilities of the individual agents by reasoning about their individual as well as collective capabilities. Existing research has looked at modular agents that combine to create a new *single agent* with capabilities that the modular agents do not possess individually. Specifically, prior work by Pathak et al. (2019) have looked at self-assembling morphologies, where individual agents combine to create a composite agent that is capable of efficiently navigating its environment. Each primitive agent in the population can be construed as a “limb” wherein the limbs may choose to link up to form a single agent. Note that there is currently significantly limited work in the area of recombinant augmentation, thus making it an open question for future research in the area.

5.2 Actions

In some cases of CPS, the agent may modify the conceptual space associated with actions, i.e., discovering new actions that enables the agent to accomplish its goal. Thus, in these cases the agent *models and reasons about actions*. As opposed to learning from human demonstration or instruction, we review methods for autonomous, fully and/or partially unassisted action discovery. In autonomous action learning, it is important that the agent learns a representation of the task on a motor/trajectory level for its execution, *and* on a semantic level for enabling the agent to determine when the action should be used. Overarching methods of action discovery include action decomposition, and action modification.

5.2.1 ACTION DECOMPOSITION

In action decomposition, the agent evaluates known or demonstrated actions, to extract useful and reusable substructures. These substructure element(s), in turn, are utilized as standalone actions. CPS methods for decomposition include segmentation, and sub-tasking.

SEGMENTATION

In the segmentation method of action decomposition, actions are broken into smaller sub-actions through an evaluation of the action’s trajectories to identify change points. While several segmentation techniques have been proposed in this context, specifically CPS approaches have used state-based transition autoregressive hidden markov model (STARHMM) (Kroemer et al., 2015) and Bayesian Change-point Detection (BCP) (Gizzi et al., 2019). In Kroemer et al. (2015), following segmentation using STARHMM, a combination of RL and Dynamic Movement Primitives (DMPs) are used to sequence new actions and execute them in problem solving. Gizzi et al. (2019), use BCP to segment recorded trajectories of known actions to generate the controllers of candidate action primitives. They ground each controller by evaluating their end effects at a logical descriptor level (using PDDL, see section 4.2.1), and consider only those actions which symbolically change the environment to be new action primitives. For example, the action of pushing a button is segmented into a ‘press’ and ‘release’ action, which can then be used in a generalized manner on any object in the environment.

SUB-TASKING

In sub-tasking methods for action decomposition, a high-level task (sequence of actions) is broken down into sub-tasks (smaller sub-sequences of actions). In contrast to action segmentation which deals with individual actions, sub-tasking methods decompose larger tasks into smaller sub-groups of actions (sub-tasks). Xu et al.(2018), use neural task programming (NTP) for learning hierarchical decomposition of a task into sub-tasks, which are then generalized to other novel tasks with different task length, topology, and semantics. They demonstrate their method for successful completion of an object sorting, block stacking, and table clean-up task, all of which were initially unseen tasks. In similar work, Hangl et al. (2017) first break down larger tasks into smaller sub-behaviors, and later “compose” the sub-behaviors to enable a robot to learn new behaviors. Given a set of behaviors B , they define new behaviors as a composition of behaviors $b_i \in B$, as $b_l \circ \dots \circ b_2 \circ b_1 \circ b^\sigma$. Compound behaviors that are successful at accomplishing the task are added to list of known behaviors for the robot.

5.2.2 ACTION MODIFICATION

Action modification techniques *adapt* pre-existing capabilities to discover novel actions. Modifications can happen at a high level representational level (i.e. operator learning), at a low-level trajectory level, or both. Reviewed methods include behavior babbling adaptations and trajectory/geometry adaptation.

BABBLING-BASED ADAPTATION

A common method for autonomous action discovery is ‘babbling,’ where the agent interacts with its environment with the goal of learning about the environment. Action discovery happens through such interactions, in conjunction with validation of new behaviors. Behavior

validation can happen through the use of perceptual information, or through intrinsic motivation. In the case of perceptual validation, Gizzi et al. (2021) demonstrate babbling through systematic variations in continuous parameters of known actions, in order to discover new actions to be used in CPS tasks. New actions are validated through novelty/usefulness in predicate end effects. For example, when a “push” action results in an object falling off of a surface, the speed parameter can be varied such that a “nudge” action is discovered. In the case of tool use in sequential manipulation tasks, predictive models of consequences of actions are developed through repeated interactions of using tools on objects in the environment (Toussaint et al., 2018; Allen et al., 2019; Sinapov & Stoytchev, 2007; Xie et al., 2019). The known actions are then adapted based on the outcomes of the interactions. The work of Silver et al. (2021) show how action operators can be learned from data in a task and motion planning (TAMP) domain. Their approach is to learn symbolic actions as probabilistic transition operators, and to learn their controllers through a combination of clustering, predicate search, and parameter estimation. Chitnis et al. (2021) use a technique called “goal literal babbling” (GLIB) to learn object-relational transition models to enable generalizable (lifted) planning. Intrinsically motivated goal-based exploration is driven by a novelty measure, encouraging babbling in unvisited state space as a way to learn a complete model for general action policies. The agent is able to use GLIB to update its previously flawed transition model in order to handle novelty, where formerly learned actions are used on novel objects. Qin et al. (2020) learn novel ways to use tools through repeated, self-supervised interactions with the tools in a simulated environment. Oddi et al. (2020) develop a framework for learning skills through intrinsically motivated reinforcement learning (via self-generated goals), and then abstracting those newly acquired skills to high level, set theoretic action representations. They demonstrate CPS in a grasp task, where the robot is able to learn how to pick up a newly encountered object (vase shaped rock) which exceeds its gripper span. The newly acquired action is encoded in PDDL for generalized use.

TRAJECTORY/GEOMETRIC ADAPTATION

In some cases of action modification, the agent adapts known symbolic/non-symbolic action representations to discover new actions. Prior work by Suarez-Hernandes (2020) discover new STRIPS actions from execution traces, by introducing cost functions to effectively search the action space. The discovery process consists of four phases: “Initialization” that involves generating a new planning problem; “Searching” for new actions and validating whether they satisfy the planning problem; “Expanding” which involves adding new compilations to an existing open list; and “Induction” where the finalized set of actions are induced based on the best planning solution. Prior work has also looked at adapting non-symbolic representations of tool-use trajectories to new tool use scenarios, either adapting to new tool specifications such as a different handle size (Fitzgerald et al., 2019, 2017) or to a completely new tool (Gajewski et al., 2019).

6. Examples of Existing Creative Problem Solving Architectures

In this section, we describe four existing architectures capable of creative problem solving, namely, CreaCogs, Robogyver, ICARUS and DIARC. We note that CreaCogs and Robogyver are architectures targeted specifically for CPS, whereas ICARUS and DIARC are more general problem-solving architectures that support CPS within them.

	Formulation	Representation	Manipulation	Evaluation	Concepts
CreaCogs	N/a	Symbolic	Combinational	N/a	Objects
Robogyver	Planning	Non-symbolic	Combinational	Physical	Objects
ICARUS	Planning	Symbolic	Exploratory	Simulation	Terrain
DIARC	Planning	Symbolic	Exploratory	Simulation	Objects

Table 2: Table highlighting the classifications of the frameworks discussed in Section 6

6.1 CreaCogs

CreaCogs is an architecture that enables agents to solve creative problems through Object Replacement and Object Construction (OROC) (Olteteanu & Falomir, 2016). CreaCogs seeks to accomplish three creative tasks: i) replacement of a missing object for a task; ii) composing objects to construct a final object for a task; and iii) decomposing objects. The architecture reasons about object affordances and capabilities in order to accomplish each task. CreaCogs consists of a symbolic large-scale knowledge base that organizes concepts at three levels: i) a feature space that encodes features or attributes of objects; ii) a concept level that represents all the relationships between the known features of an object; and iii) a problem template where the relations between different concepts are encoded. In their work, the information regarding the three levels is symbolically encoded a-priori. Given the knowledge base, replacement objects are identified on the basis of the similarity of their features to the missing object. Similarly, composite objects are represented as a conjunction of the features (denoted as Y) of the individual objects (denoted as O_k) that make up the composite object, e.g., $relation(O_1, Y) \wedge relation(O_2, Y)$. Hence, CreaCogs uses combinational manipulation for identifying composite objects. Additionally, topological relationships are used to define the relative positioning of the objects when constructing the composite object. Lastly, for decomposition, a larger object is represented as a conjunction of its individual components, e.g., $object(fishing\ rod) \rightarrow component(hook, fishing\ rod) \wedge component(line, fishing\ rod) \wedge component(rod, fishing\ rod)$ etc. Thus, new concepts (i.e., objects) can be identified through various combinations of relevant concepts.

In Olteteanu (2015), the authors discuss re-representation of the problem in order to highlight concepts that are relevant for object replacement and composition, e.g., “containability” is a relevant concept when creating a candle holder. For this problem, a candle holder is re-represented to highlight the containability feature that then helps identify substitute objects for it. They discuss the notion of “seeing as” referring to: i) the ability to represent a group of features as a meaningful object, and ii) ability to represent a group of objects as a meaningful structure that can solve the problem on hand. Thus, the initial problem description is *transformed* into an alternate representation that highlights the relevant concepts or features of object(s) that can enable the agent to solve the problem at hand. While the authors discuss these ideas theoretically in (Olteteanu, 2015), the CreaCogs implementation in (Olteteanu & Falomir, 2016) does not include the re-representation, hence we categorize it as combinational manipulation within our survey.

6.2 Robogyver

Robogyver is an architecture that enables robots to accomplish tasks where required tools are missing, either through tool construction or tool substitution (Nair, Shrivatsav, & Chernova, 2020; Nair, 2020). While substitution and construction is closely related to object replacement and composition in CreaCogs, Robogyver focuses specifically on tool-based problems through non-symbolic rather than symbolic reasoning (Nair & Chernova, 2020). Robogyver formulates tool macgyvering as a planning problem (Nair & Chernova, 2020), introducing the Feature Guided Search (FGS) approach that extends classical planning using supervised learning to guide the search through the planning space. The initial conceptual space comprises of multi-modal sensory inputs (object point clouds and spectral readings). The authors perform multi-objective optimization in order to identify viable objects for replacing the missing tool, based on their point clouds (indicating shape) and spectral measurements (indicating material). The objectives within the multi-objective function are learned through supervised learning techniques. Robogyver uses a hybrid symbolic and non-symbolic representation wherein the objects are represented in a non-parametric and non-symbolic manner, whereas, the task planning itself uses symbols represented in PDDL. To perform tool construction, the multi-objective functions are used to evaluate object combinations to identify viable constructions. Specifically, the multi-objective function, denoted as f , acts as a combinational function that takes object point clouds and spectral readings as inputs, i.e., $f(o_1, o_2, \dots, o_m)$ (i.e., combinational methods). Within the new conceptual space, the multi-objective functions are used to guide the search for a valid task plan. Specifically, the object-based symbols within the planner are assigned the output scores of the multi-objective functions. The scores are then incorporated into planning heuristics in order to guide the search (using A^*) towards valid object combinations. Hence, given a task and a set of objects, the agent is able to adapt the task plan to construct an appropriate tool for the task. The approach is evaluated on a physical robot (Kinova) for the construction of tools for six different tasks.

6.3 ICARUS

ICARUS is a cognitive architecture that provides an infrastructure for building intelligent AI systems (Langley, Choi, & Rogers, 2009). Choi et al. (Choi et al., 2018) extend the cognitive architecture ICARUS to incorporate the capability to create and use structures, specifically navigational structures, such as bridges and ramps. Hence, their work is an example of Automated Design of Functional Structures (ADFS). ICARUS formulates ADFS as a planning problem, and introduce a Unified Planning System (UPS) that incorporates primitive skills, hierarchical skills, and their combination to generate plans. ICARUS consists of two types of long-term knowledge, namely, *concepts* and *skills*. In their work, concepts refers to percepts and their corresponding attribute values. For instance, ICARUS represents objects as concepts, defined by various attributes. On the other hand, skills build upon conceptual knowledge, by including conditions that must hold for their application as well as the effects that their application produces. For example, ICARUS represents building or construction actions as skills. Both concepts and skills are hard-coded within the system’s knowledge base. Further, they are symbolically encoded, e.g., blocks are defined as `block ?b`, or left and right positions are denoted by `?left` and `?right`. Given the initial symbolic conceptual space, the agent searches for a task plan that involves creating navigational

structures such as bridges and staircases to traverse the environment. As described in Section 4.3.1, they apply numeric heuristics to guide a best-first search (exploratory manipulation), where the heuristics favor states that satisfy more elements in the goal formulae. In order to construct the structures, the planner specifically reasons about the terrain, such as the width of gaps, and hence is an example of terrain modification. Note that the properties of the objects are encoded a-priori, and the planner specifically reasons about the terrain in order to identify a suitable construction of objects that makes the terrain traversable. Lastly, the approach is evaluated in simulation for the construction of ramps.

6.4 Distributed Integrated Cognition Affect and Reflection (DIARC)

DIARC is a cognitive architecture comprising of action, vision, and natural language sub-systems (Scheutz, Williams, Krause, Oosterveld, Sarathy, & Frasca, 2019). More recently, DIARC was extended to incorporate creative capabilities for dealing with novelties in the environment (Muhammad, Sarathy, Tatiya, Goel, Gyawali, Guaman, Sinapov, & Scheutz, 2021). The authors propose an architecture that uses a knowledge base in conjunction with a PDDL-based planning framework. The authors further define novelty as an encounter between the agent and an item that a) the agent has not experienced before, and b) cannot derive it from its knowledge base. The agent uses a symbolic knowledge representation using first-order language (FOL) and a knowledge base (KB) containing facts associated with the agent’s environment/task specification. Given the FOL description of the world-state σ_t at time t , the agent computes $KB \cup \sigma_t$ to generate a task plan $\pi = \langle \alpha_t, \alpha_{t+1} \dots \alpha_n \rangle$ to achieve a task goal G . In the event that an action fails or leads to an unexpected state, the agent updates its KB and attempts a re-plan. Hence, DIARC is an example of exploratory manipulation of the initial KB (i.e., conceptual space) for deriving a new conceptual space that enables the agent to deal with unexpected novelties. In particular, the agent reasons about objects, and seeks to characterize object properties to resolve the novelties. Hence, the conceptual space focuses on object properties for generating creative solutions (object modification). The overall architecture proceeds as follows: the agent interacts with the environment, and if novelties are encountered, the corresponding world states are incorporated into the KB via a novelty handling mechanism that deals with different categories of novelties e.g., class and attribute novelties. In the case of class novelty, a new class of object is introduced and in the case of attribute novelty, the property of an existing object changes. Upon handling the novelty, the agent then attempts to re-plan for the task. This proceeds until the task goal is achieved. Most recently, the architecture has been extended to allow the agent to formulate a reinforcement learning problem on-the-fly in cases where manipulating the existing KB is insufficient to adapt to the environmental change (Goel, Shukla, Sarathy, Scheutz, & Sinapov, 2022; Lorang, Goel, Zips, Sinapov, & Scheutz, 2022). The authors evaluate their approach in simulation using a “Polycraft” domain, and an internal grid-world variant that mimics Polycraft (Goel et al., 2021). The task in the domain involves crafting a pogo-stick using available materials. Here, the agent must reason about objects in the environment and novelties specifically associated with objects, to successfully construct the desired object (Object construction). The authors demonstrate that the agent is able to successfully reason about novelties at different levels of complexity.

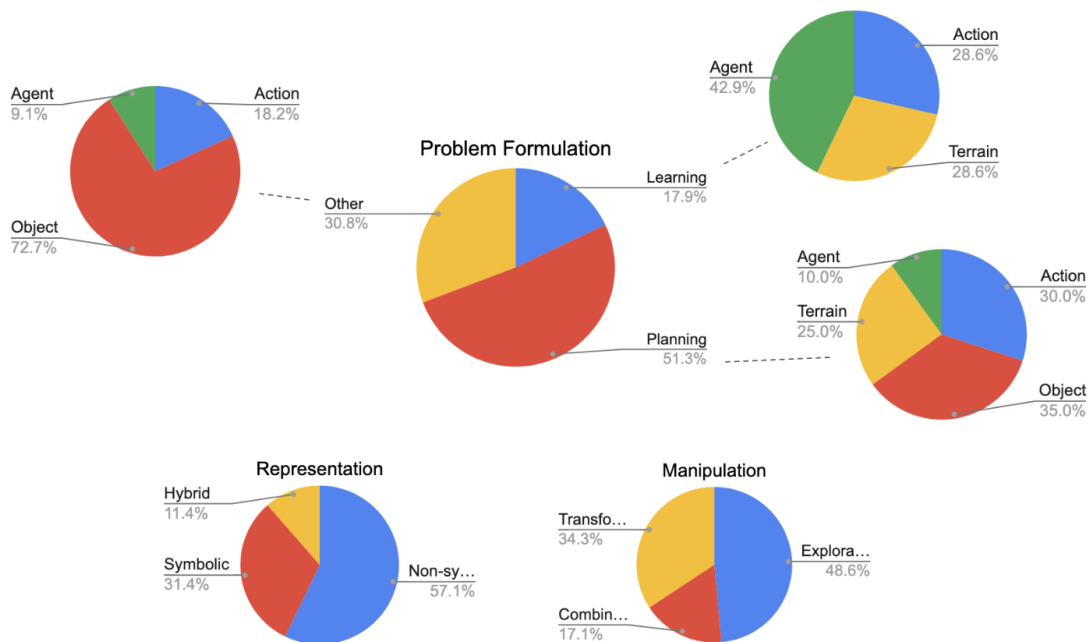


Figure 17: Distribution of the 51 papers in our survey. To note limited work in CPS, other paradigms of knowledge manipulation have not focused on terrain modification, and learning approaches have not looked at object modification indicating clear gaps in existing CPS research. A tabulated list of papers in each category is provided in the Appendix.

7. Open Research Questions

In our survey of creative problem solving, we identified open research questions that have not been adequately addressed in the field of CPS. In highlighting these problems, we hope to stimulate further work and suggest future research directions. We divide this section into two subsections. The first subsection discusses open research questions in the context of our CPS framework. The second subsection discusses broader CPS questions that are not specifically related to our framework or taxonomy, but need to be addressed, and potentially incorporated into future iterations of the CPS framework.

7.1 Open Questions Relating to CPS Framework

In this section, we discuss open research questions in the context of our CPS taxonomy.

7.1.1 HYBRID SYMBOLIC AND NON-SYMBOLIC REPRESENTATIONS

There has been limited work utilizing hybrid symbolic and non-symbolic representations in CPS. Such hybrid representations can be useful in order to effectively represent different types of information required by the CPS method. In particular, some types of information such as visual properties are more easily represented through non-symbolic means, whereas symbolic representations can be useful to encode hierarchies such as in semantic networks.

There are several challenges to further explore in this space. In particular, what specific information should be represented as symbolic vs. non-symbolic (or vice-versa)? Moreover,

how can we design hybrid systems that can effectively leverage the strengths of the hybrid representation? As noted in our survey, planning approaches are strongly correlated to symbolic representations, whereas learning approaches are correlated to non-symbolic representations. Hence, hybrid systems that combine planning and learning, can be useful (and perhaps necessary) for effectively utilizing hybrid representations. To this end, existing approaches such as hybrid planning and RL (Strens & Windelinckx, 2004) can be beneficial. Additionally, recent advances in Neuro-symbolic AI (d’Avila Garcez & Lamb, 2020), can also be useful in this context. Neuro-symbolic AI seeks to combine neural networks with symbolic representations of the problem, and can greatly enhance the reasoning capability, as well as explainability of CPS systems.

7.1.2 RETHINKING AGENT-REASONING IN CPS

Among the conceptual spaces discussed in our survey, the papers relating to conceptual spaces of agent states, focus on agent modification alone. However, a key point to note is that modeling and reasoning about the agent can have wider implications for CPS, beyond just agent modification. Modeling the capabilities of the agent/robot, can be highly beneficial for identifying practically viable creative solutions for performing the task, which is especially important when operating in real-world domains rather than simulations. A key question here is, how can CPS systems account for the agent when deriving solutions? In this case, the solutions derived would depend on the agent’s capabilities, e.g., it could vary from the two-armed Baxter robot to a single-armed Sawyer robot. Consider the example of opening a tightly closed jar. The Baxter robot can potentially grasp the jar with one hand and use a knife with the other to pry it open. In contrast, this would not be a viable solution for Sawyer, that may instead run it under hot water to loosen the cap, or break it open to access its contents. However, existing CPS research does not account for the agent’s capabilities in this manner, making it an open question for future work.

7.2 Broader CPS Research Questions

In this section, we discuss broader CPS questions relating to, a) extension of additional computational creativity methods to CPS, b) universal metric for evaluating CPS solutions, c) generalizability and CPS, and d) lifelong CPS.

7.2.1 EXTENSION OF ADDITIONAL COMPUTATIONAL CREATIVITY METHODS TO CPS

Conceptual blending, introduced by Fauconnier and Turner (Fauconnier & Turner, 2008) is a widely studied approach in computational creativity. However, its applicability to creative problem solving in AI has not been well explored. In Section 4.3.2, Multi-Concept Combination, we presented work by Lieto et al. (2019) that uses conceptual blending, a mode of combinational knowledge manipulation, to discover new concepts within a semantic network. However, apart from the presented research, approaches to CPS have not focused on the potential application of conceptual blending techniques. Recent approaches have developed computational models for conceptual blending (Falomir & Plaza, 2020; Schorlemmer & Plaza, 2021). An interesting question is the extension of similar methods to knowledge representations observed in the context of CPS in AI. Approaches that use convolutional neural networks for style transfer, referred to as Neural Style Transfer (Jing, Yang, Feng, Ye,

Yu, & Song, 2019), demonstrate that it is possible to learn conceptual blending. However, the application of such techniques to the problem of task learning and planning in AI remains an open question. An additional area of CC research which should be considered in its application to CPS is the concept of “framing”, where an intelligent agent attempts to “explain” its own creativity (Colton, Charnley, & Pease, 2011; Guckelsberger et al., 2017; Cook et al., 2019). Examples include explanation of creative intentions, inward reflection of intrinsic goals, and justification of actions in creativity. Integrating framing into CPS could provide a computational method for isolating the novelty/salience responsible for converting a general problem solving scenario into one which requires CPS. This in turn could expedite the time required for CPS resolution.

7.2.2 UNIVERSAL METRIC FOR CPS SOLUTIONS

As described in Section 3, there does not yet exist a universally agreed upon definition of computational creativity. Because of this lack of a concise definition of CC, it has been challenging for the CPS research community to develop a universal measure of a CPS solution, with consideration of *both* the traditional problem solving aspects of CPS, and the creativity portion of CPS. Most evaluation metrics of CPS which have been developed in past research are domain and/or problem specific (See Section 4.4.3). While some degree of universality has been considered, most work has focused on developing tests for CPS success, rather than a specific metric for objectively quantifying how *creative* a solution is. Other works have developed *theoretical* measures of CPS solutions, without evidentiary implementation. Developing such a metric for CPS would be highly beneficial, as it would help streamline and benchmark existing and future CPS methods. For instance, an objective metric for categorizing the subset of papers in RL that constitutes CPS needs to be defined, as opposed to the subjective novelty-based criterion used in this survey.

7.2.3 GENERALIZABILITY AND CPS

A key consideration and theme in CPS research is the flexibility of systems, which allows agents to handle the inherent novelty of CPS tasks. While current research has made this consideration of generalizability within individual problem domains, there has yet to be extensive research and testing of cross-domain generalizability of CPS systems. Additionally, despite the domain-agnostic theoretical formulations of CPS methods which exist in current research, testing specific methods for cross-domain generalizability at the implementation level remains a largely unexplored challenge.

7.2.4 LIFELONG CPS

Techniques in CPS have focused on flexibility in handling unforeseen and ill-structured problems. Despite extensive research into creating flexible CPS methods for “problem at hand” solutions, 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

that the agent may need in a different, future CPS task. Existing research in lifelong learning (Parisi, Kemker, Part, Kanan, & Wermter, 2019) and transfer learning (Pan & Yang, 2010) could be useful venues to explore in this context. While there does exist preliminary work in the area of life-long CPS (Gizzi, Lin, Castro, Harvey, & Sinapov, 2022), 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.

8. Conclusions

This survey discussed creative problem solving in artificial intelligence, by formally defining CPS, and introducing a framework that encompasses existing CPS approaches. In contrast to existing CPS formalizations, our definition 1) introduces a broader notion of concepts, and 2) connects Boden’s three levels of creativity from computational creativity literature, to problem solving in AI, thus leveraging theoretical aspects from both CC and AI. Our presented CPS framework consists of four key steps including, 1) problem formulation, 2) knowledge representation, 3) knowledge manipulation, and 4) evaluation. Additionally, we categorized existing CPS research within our framework. We further expanded on our CPS framework to discuss the two types of conceptual spaces that are modified, i.e., states (including objects, terrain and agent) and actions. We further categorized existing approaches along this dimension. Next, we presented existing CPS architectures, also organizing them within our CPS taxonomy. Our survey of the literature concluded with a list of open research questions, which we believe will serve as a useful guide for future work in CPS. We further hope that this survey will encourage research into this relatively unexplored field, bridging the gap between CC and problem solving in AI.

Paper	Formuln.	Repn.	Manipn.	Eval.	Concepts
(Xie et al., 2019)	Learning	Non-symbolic	Expl.	Phys	Action
(Fitzgerald et al., 2014, 2017)	Planning	Non-symbolic	Trans	Phys	Action
(Wicaksono & Sheh, 2017)	Planning	Symbolic	Expl.	Phys	Object
(Wicaksono & Sammut, 2020)	Planning	Symbolic	Expl.	Phys	Object
(Saboia da Silva et al., 2019)	Planning	Non-symbolic	Transf.	Phys	Terrain
(Toussaint et al., 2018)	Planning	Hybrid	Expl.	Sim	Action
(Zhu et al., 2015)	Other	Hybrid	Transf.	Sim	Object
(Boteanu et al., 2015)	Planning	Symbolic	Expl.	Phys	Object
(Olteteanu & Falomir, 2016)	Other	Symbolic	Comb.	n/a	Object
(Erdogan & Stilman, 2013)	Planning	Symbolic	Expl.	Sim	Terrain
(Erdogan & Stilman, 2016)	Planning	Symbolic	Expl.	Sim	Object
(Choi et al., 2018)	Planning	Non-symbolic	Expl.	Sim	Terrain
(Levihn & Christensen, 2015)	Planning	Symbolic	Expl.	Sim	Object

(Tosun et al., 2018)	Planning	Non-symbolic	Transf.	Phys	Terrain
(Abelha et al., 2016)	Other	non-symbolic	Transf.	Sim	Object
(Schoeler & Wörgötter, 2015)	Other	Non-symbolic	Transf.	Sim	Object
(Gajewski et al., 2019)	Planning	Non-symbolic	Transf.	Phys	Action
(Baker et al., 2019)	Learning	Non-symbolic	Expl.	Sim	Terrain
(Allen et al., 2019)	Learning	Non-symbolic	Expl.	Sim	Action
(Sinapov & Stoytchev, 2007)	Other	Non-symbolic	Transf.	Sim	Action
(Hangl et al., 2017)	Other	Non-symbolic	Comb.	Phys	Action
(Bapst et al., 2019)	Learning	Non-symbolic	Expl.	Sim	Terrain
(Pathak et al., 2019)	Learning	Non-symbolic	Expl.	Sim	Agent
(Nair et al., 2019, 2019, 2020)	Other	Symbolic	Comb.	Phys	Object
(Nair & Chernova, 2020)	Planning	Hybrid	Comb.	Phys	Object
(Shrivatsav et al., 2019)	Other	Non-symbolic	Comb.	Phys	Object
(Gizzi et al., 2019)	Planning	Hybrid	Transf.	Sim	Action
(Gizzi et al., 2021)	Planning	Hybrid	Expl.	Sim	Action
(Murooka et al., 2019)	Planning	Non-symbolic	Transf.	Phys	Agent
(Suárez-Hernández et al., 2020)	Planning	Symbolic	Expl.	Sim	Action
(Freedman et al., 2020)	Planning	Symbolic	Transf.	Sim	Object
(Ha, 2019)	Learning	Non-symbolic	Expl.	Sim	Agent
(Zhao et al., 2020)	Other	Symbolic	Expl.	Sim	Agent
(Yang et al., 2020)	Other	Non-symbolic	Comb.	Phys	Object
(Qin et al., 2020)	Planning	Non-symbolic	Transf.	Sim	Action
(Colin et al., 2016)	Learning	Non-symbolic	Comb.	Sim	Terrain
(Colin & Belpaeme, 2019)	Learning	Non-symbolic	Comb.	Sim	Terrain
(Lieto et al., 2019)	Planning	Symbolic	Comb.	B.mark	Object
(Silver et al., 2021)	Planning	Hybrid	Transf.	Sim	Action
(Chitnis et al., 2021)	Learning	Symbolic	Expl.	Sim	Action
(Oddi et al., 2020)	Learning	Hybrid	Expl.	Sim	Action
(Kralik et al., 2016)	Learning	Non-symbolic	Transf.	B.mark	Terrain
(Xu et al., 2018)	Other	Hybrid	Transf.	Sim	Object
(Fitzgerald et al., 2019)	Learning	Non-symbolic	Transf.	Phys	Action
(Kroemer et al., 2015)	Learning	Non-symbolic	Transf.	Phys	Action

(Ha, 2019)	Learning	Non-symbolic	Expl.	Sim	Agent
(Schaff et al., 2019)	Learning	Non-symbolic	Expl.	Sim	Agent

Table 3: Table highlighting the classifications of the papers discussed in this survey.

9. Acknowledgements

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

- Abelha, P., Guerin, F., & Schoeler, M. (2016). A model-based approach to finding substitute tools in 3d vision data. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 2471–2478. IEEE.
- Allen, K. R., Smith, K. A., & Tenenbaum, J. B. (2019). The tools challenge: Rapid trial-and-error learning in physical problem solving. *Arxiv preprint, arXiv:1907.09620*.
- Atkeson, C. G., Benzun, P. B., Banerjee, N., Berenson, D., Bove, C. P., Cui, X., DeDonato, M., Du, R., Feng, S., Franklin, P., et al. (2018). What happened at the darpa robotics challenge finals. In *The DARPA Robotics Challenge Finals: Humanoid Robots to the Rescue*, pp. 667–684. Springer.
- Baker, B., Kanitscheider, I., Markov, T., Wu, Y., Powell, G., McGrew, B., & Mor-datch, I. (2019). Emergent tool use from multi-agent autotutorials. *Arxiv preprint, arXiv:1909.07528*.
- Bapst, V., Sanchez-Gonzalez, A., Doersch, C., Stachenfeld, K. L., Kohli, P., Battaglia, P. W., & Hamrick, J. B. (2019). Structured agents for physical construction. *Arxiv preprint, arXiv:1904.03177*.
- Bengio, Y., Courville, A., & Vincent, P. (2013). Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8), 1798–1828.
- Bird, C. D., & Emery, N. J. (2009). Insightful problem solving and creative tool modification by captive nontool-using rooks. *Proceedings of the National Academy of Sciences*, 106(25), 10370–10375.
- Bishop, M., & Boden, M. A. (2010). The turing test and artistic creativity. *Kybernetes*, 3(39), 409–413.
- Bisk, Y., Zellers, R., Bras, R. L., Gao, J., & Choi, Y. (2019). Piqa: Reasoning about physical commonsense in natural language. *Arxiv preprint, arXiv:1911.11641*.
- Boden, M. A. (1998). Creativity and artificial intelligence. *Artificial Intelligence*, 1-2, 347–356.
- Boesch, C., & Boesch, H. (1990). Tool use and tool making in wild chimpanzees. *Folia primatologica*, 54(1-2), 86–99.
- Boteanu, A., Kent, D., Mohseni-Kabir, A., Rich, C., & Chernova, S. (2015). Towards robot adaptability in new situations. In *2015 AAAI Fall Symposium Series*.

- Cass, S. (2005). Apollo 13, we have a solution. *IEEE Spectrum On-line*, 04, 1.
- Chitnis, R., Silver, T., Tenenbaum, J., Kaelbling, L. P., & Lozano-Pérez, T. (2021). Glib: Efficient exploration for relational model-based reinforcement learning via goal-literal babbling. In *Proc. AAAI*.
- Choi, D., Langley, P., & To, S. T. (2018). Creating and using tools in a hybrid cognitive architecture. In *2018 AAAI Spring Symposium Series*.
- Colin, T. R., & Belpaeme, T. (2019). Reinforcement learning and insight in the artificial pigeon. In *CogSci*, pp. 1533–1539.
- Colton, S., Charnley, J. W., & Pease, A. (2011). Computational creativity theory: The face and idea descriptive models. In *The International Conference on Computational Creativity (ICCC)*, pp. 90–95.
- Colton, S., Wiggins, G. A., et al. (2012). Computational creativity: The final frontier?. In *Ecai*, Vol. 12, pp. 21–26. Montpellier.
- Cook, M., Colton, S., Pease, A., & Llano, M. T. (2019). Framing in computational creativity—a survey and taxonomy. In *The 10th International Conference on Computational Creativity (ICCC)*, pp. 156–163.
- d’Avila Garcez, A., & Lamb, L. C. (2020). Neurosymbolic ai: The 3rd wave. *Arxiv preprint, arXiv:2012.05876*.
- Dictionary, O. E. (1989). Oxford english dictionary..
- Ellamil, M., Dobson, C., Beeman, M., & Christoff, K. (2012). Evaluative and generative modes of thought during the creative process. *Neuroimage*, 59(2), 1783–1794.
- Erdogan, C., & Stilman, M. (2013). Planning in constraint space: Automated design of functional structures. In *2013 IEEE International Conference on Robotics and Automation*, pp. 1807–1812. IEEE.
- Erdogan, C., & Stilman, M. (2016). Autonomous realization of simple machines. In *Experimental Robotics*, pp. 471–486. Springer.
- Erol, K., Hendler, J. A., & Nau, D. S. (1994). Umcp: A sound and complete procedure for hierarchical task-network planning.. In *Aips*, Vol. 94, pp. 249–254.
- Falomir, Z., & Olteţeanu, A.-M. (2019). Special issue on problem-solving , creativity and spatial reasoning. *Cognitive Systems Research*, 58, 31–34.
- Falomir, Z., & Plaza, E. (2020). 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.
- Fauconnier, G., & Turner, M. (2008). *The way we think: Conceptual blending and the mind’s hidden complexities*. Basic books.
- Fikes, R. E., & Nilsson, N. J. (1971). Strips: A new approach to the application of theorem proving to problem solving. *Artificial intelligence*, 2(3-4), 189–208.
- Fitzgerald, T., Goel, A., & Thomaz, A. (2015). A similarity-based approach to skill transfer. In *Women in Robotics Workshop at Robotics: Science and Systems Conference (RSS)*.

- Fitzgerald, T., Goel, A., & Thomaz, A. (2017). Human-robot co-creativity: Task transfer on a spectrum of similarity. In *2017 International Conference on Computational Creativity (ICCC)*, pp. 104–111.
- Fitzgerald, T., Goel, A. K., & Thomaz, A. L. (2014). Representing skill demonstrations for adaptation and transfer. In *2014 AAAI Fall Symposium Series*.
- Fitzgerald, T., Short, E., Goel, A., & Thomaz, A. (2019). Human-guided trajectory adaptation for tool transfer. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, pp. 1350–1358. International Foundation for Autonomous Agents and Multiagent Systems.
- Fox, M., & Long, D. (2003). Pddl2. 1: An extension to pddl for expressing temporal planning domains. *Journal of artificial intelligence research*, 20, 61–124.
- Freedman, R., Friedman, S., Musliner, D., & Pelican, M. (2020). Creative problem solving through automated planning and analogy. *AAAI Workshop on Generalization in Planning*, -.
- Gajewski, P., Ferreira, P., Bartels, G., Wang, C., Guerin, F., Indurkha, B., Beetz, M., & Śnieżyński, B. (2019). Adapting everyday manipulation skills to varied scenarios. In *2019 International Conference on Robotics and Automation (ICRA)*, pp. 1345–1351. IEEE.
- Gerevini, A. E., Haslum, P., Long, D., Saetti, A., & Dimopoulos, Y. (2009). Deterministic planning in the fifth international planning competition: Pddl3 and experimental evaluation of the planners. *Artificial Intelligence*, 173(5-6), 619–668.
- Gibson, J. (1977). The concept of affordances. *Perceiving, acting, and knowing*, 1.
- Gizzi, E., Castro, M. G., & 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 (ICDL-EpiRob)*, pp. 228–233. IEEE.
- Gizzi, E., Hassan, A., Lin, W.-W., Rhea, K., & Sinapov, J. (2021). Toward creative problem solving agents: Action discovery through behavior babbling. In *The International Conference on Development and Learning (ICDL)*. IEEE.
- Gizzi, E., Lin, W. W., Castro, M. G., Harvey, E., & Sinapov, J. (2022). Toward life-long creative problem solving: Using world models for increased performance in novelty resolution. In *13th International Conference on Computational Creativity*.
- Goel, S., Shukla, Y., Sarathy, V., Scheutz, M., & Sinapov, J. (2022). Rapid-learn: A framework for learning to recover for handling novelties in open-world environments. In *IEEE International Conference on Development and Learning (ICDL), London, UK, September 12-15, 2022*, pp. 1–8. IEEE.
- Goel, S., Tatiya, G., Scheutz, M., & Sinapov, J. (2021). Novelgridworlds: A benchmark environment for detecting and adapting to novelties in open worlds. *Adaptive and Learning Agents (ALA) Workshop*, -.
- Guckelsberger, C., Salge, C., & 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 (ICCC)*.

- Guilford, J. P. (1967a). Creativity: Yesterday, today and tomorrow. *The Journal of Creative Behavior*, 1(1), 3–14.
- Guilford, J. P. (1967b). *The nature of human intelligence*. McGraw-Hill.
- Guzdial, M., Liao, N., Shah, V., & Riedl, M. O. (2018). Creative invention benchmark. *Arxiv preprint, arXiv:1805.03720*.
- Ha, D. (2019). Reinforcement learning for improving agent design. *Artificial life*, 25(4), 352–365.
- Hangl, S., Dunjko, V., Briegel, H. J., & Piater, J. (2017). Skill learning by autonomous robotic playing using active learning and creativity. *Arxiv preprint, arXiv:1706.08560*.
- Hélie, S., & Sun, R. (2010). Incubation, insight, and creative problem solving: a unified theory and a connectionist model. *Psychological review*, 117(3), 994.
- ICCC (2022). Association of computational creativity: What is computational creativity?. <https://computationalcreativity.net/home/about/computational-creativity/>. Online, accessed in 2022.
- Isaksen, S. G., & Treffinger, D. J. (2004). Celebrating 50 years of reflective practice: Versions of creative problem solving. *The Journal of Creative Behavior*, 38(2), 75–101.
- Jain, A., Szot, A., & Lim, J. (2020). Generalization to new actions in reinforcement learning. In *International Conference on Machine Learning*, pp. 4661–4672. PMLR.
- Jing, Y., Yang, Y., Feng, Z., Ye, J., Yu, Y., & Song, M. (2019). Neural style transfer: A review. *IEEE transactions on visualization and computer graphics*, 26(11), 3365–3385.
- Jordanous, A. K. (2013). *Evaluating computational creativity: a standardised procedure for evaluating creative systems and its application*. Ph.D. thesis, University of Sussex.
- Kitchen, O. E. D. (2020). Apollobvm - emergency use ventilator.. <http://oedk.rice.edu/apollobvm/>.
- Knoblich, G., Ohlsson, S., Haider, H., & Rhenius, D. (1999). Constraint relaxation and chunk decomposition in insight problem solving. *Journal of Experimental Psychology: Learning, memory, and cognition*, 25(6), 1534–1555.
- Kralik, J. D., Mao, T., Cheng, Z., & Ray, L. E. (2016). Modeling incubation and restructuring for creative problem solving in robots. *Robotics and Autonomous Systems*, 86, 162–173.
- Kroemer, O., Daniel, C., Neumann, G., Van Hoof, H., & Peters, J. (2015). Towards learning hierarchical skills for multi-phase manipulation tasks. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1503–1510. IEEE.
- Kroemer, O., Van Hoof, H., Neumann, G., & Peters, J. (2014). Learning to predict phases of manipulation tasks as hidden states. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4009–4014. IEEE.
- Lamb, C., Brown, D. G., & Clarke, C. L. (2018). Evaluating computational creativity: An interdisciplinary tutorial. *ACM Computing Surveys (CSUR)*, 51(2), 1–34.
- Langley, P., Choi, D., & Rogers, S. (2009). Acquisition of hierarchical reactive skills in a unified cognitive architecture. *Cognitive Systems Research*, 10(4), 316–332.

- Leibo, J. Z., Hughes, E., Lanctot, M., & Graepel, T. (2019). Autocurricula and the emergence of innovation from social interaction: A manifesto for multi-agent intelligence research..
- Levihn, M., & Christensen, H. (2015). Using environment objects as tools in unknown environments. In *2015 IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids)*, pp. 160–165. IEEE.
- Levihn, M., & Stilman, M. (2014). Using environment objects as tools: Unconventional door opening. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 2502–2508. IEEE.
- Lieto, A., Perrone, F., Pozzato, G. L., & Chiodino, E. (2019). Beyond subgoalng: a dynamic knowledge generation framework for creative problem solving in cognitive architectures. *Cognitive Systems Research*, *58*, 305–316.
- Lorang, P., Goel, S., Zips, P., Sinapov, J., & Scheutz, M. (2022). Speeding-up continual learning through information gains in novel experiences. In *4th Planning and Reinforcement Learning (PRL) Workshop at IJCAI-2022*.
- McDermott, D., Ghallab, M., Howe, A., Knoblock, C., Ram, A., Veloso, M., Weld, D., & Wilkins, D. (1998). Pddl-the planning domain definition language. *Technical Report, CVC TR-98-003/DCS TR-1165*.
- McDermott, D. M. (2000). The 1998 ai planning systems competition. *AI magazine*, *21*(2), 35–35.
- Muhammad, F., Sarathy, V., Tatiya, G., Goel, S., Gyawali, S., Guaman, M., Sinapov, J., & Scheutz, M. (2021). A novelty-centric agent architecture for changing worlds. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*, pp. 925–933.
- Mumford, M. D., Mobley, M. I., Reiter-Palmon, R., Uhlman, C. E., & 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., & Threlfall, K. V. (1997). Process-based measures of creative problem-solving skills: V. overall prediction. *Creativity Research Journal*, *10*(1), 73–85.
- Murooka, T., Okada, K., & Inaba, M. (2019). Self-repair and self-extension by tightening screws based on precise calculation of screw pose of self-body with cad data and graph search with regrasping a driver. In *2019 IEEE-RAS 19th International Conference on Humanoid Robots (Humanoids)*, pp. 79–84. IEEE.
- Nair, L., Balloch, J., & Chernova, S. (2019). Tool macgyvering: Tool construction using geometric reasoning. In *2019 International Conference on Robotics and Automation (ICRA)*, pp. 5837–5843. IEEE.
- Nair, L., & Chernova, S. (2020). Feature guided search for creative problem solving through tool construction. *Frontiers in Robotics and AI*, *7*, 205.
- Nair, L., Shrivatsav, N., & Chernova, S. (2020). Tool macgyvering: A novel framework for combining tool substitution and construction. *Arxiv preprint, arXiv:2008.10638*.

- Nair, L., Srikanth, N. S., Erickson, Z. M., & Chernova, S. (2019). Autonomous tool construction using part shape and attachment prediction.. In *Robotics: Science and Systems*.
- Nair, L. V. (2020). *Robogyver: Autonomous Tool Macgyvering for Inventive Problem Solving*. Ph.D. thesis, Georgia Institute of Technology.
- Oddi, A., Rasconi, R., Santucci, V. G., Sartor, G., Cartoni, E., Mannella, F., & Baldassarre, G. (2020). Integrating open-ended learning in the sense-plan-act robot control paradigm. In *24th European Conference on Artificial Intelligence (ECAI)*.
- Olteteanu, 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 (ICCC)*.
- Olteteanu, 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*, p. 105.
- Olteteanu, A.-M., & Falomir, Z. (2016). 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.
- Osborn, A. F. (1952). *Wake up your mind: 101 ways to develop creativeness*. Scribner.
- Pan, S. J., & Yang, Q. (2010). A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10), 1345–1359.
- Parisi, G. I., Kemker, R., Part, J. L., Kanan, C., & Wermter, S. (2019). Continual lifelong learning with neural networks: A review. *Neural Networks*, 113, 54–71.
- Pathak, D., Lu, C., Darrell, T., Isola, P., & Efros, A. A. (2019). Learning to control self-assembling morphologies: a study of generalization via modularity. In *Advances in Neural Information Processing Systems*, pp. 2292–2302.
- Pednault, E. P. (1987). Formulating multiagent, dynamic-world problems in the classical planning framework. In *Reasoning about actions & plans*, pp. 47–82. Elsevier.
- Qi, C. R., Su, H., Mo, K., & Guibas, L. J. (2017). Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 652–660.
- Qin, Z., Fang, K., Zhu, Y., Fei-Fei, L., & Savarese, S. (2020). Keto: Learning keypoint representations for tool manipulation. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 7278–7285. IEEE.
- Rowe, J., & Partridge, D. (1993). Creativity: A survey of ai approaches. *Artificial Intelligence Review*, 7(3), 43–70.
- Saboia da Silva, M., et al. (2019). *Autonomous adaptive modification of unstructured environments*. Ph.D. thesis, State University of New York at Buffalo.
- Şahin, E., Çakmak, M., Doğar, M. R., Uğur, E., & Üçoluk, G. (2007). To afford or not to afford: A new formalization of affordances toward affordance-based robot control. *Adaptive Behavior*, 15(4), 447–472.

- Sarathy, V. (2018). Real world problem-solving. *Frontiers in Human Neuroscience*, 12.
- Sarathy, V., Kasenberg, D., Goel, S., Sinapov, J., & Scheutz, M. (2020). Spotter: Extending symbolic planning operators through targeted reinforcement learning..
- Sarathy, V., & Scheutz, M. (2018). Macgyver problems: Ai challenges for testing resourcefulness and creativity. *Advances in Cognitive Systems*, 6, 31–44.
- Schaff, C., Yunis, D., Chakrabarti, A., & Walter, M. R. (2019). Jointly learning to construct and control agents using deep reinforcement learning. In *2019 International Conference on Robotics and Automation (ICRA)*, pp. 9798–9805. IEEE.
- Scheutz, M., Williams, T., Krause, E., Oosterveld, B., Sarathy, V., & Frasca, T. (2019). An overview of the distributed integrated cognition affect and reflection diarc architecture. *Cognitive architectures*, -, 165–193.
- Schoeler, M., & Wörgötter, F. (2015). Bootstrapping the semantics of tools: Affordance analysis of real world objects on a per-part basis. *IEEE Transactions on Cognitive and Developmental Systems*, 8(2), 84–98.
- Schorlemmer, M., & Plaza, E. (2021). A uniform model of computational conceptual blending. *Cognitive Systems Research*, 65, 118–137.
- Shrivatsav, N., Nair, L., & Chernova, S. (2019). Tool substitution with shape and material reasoning using dual neural networks. *Arxiv preprint, arXiv:1911.04521*.
- Silver, T., Chitnis, R., Tenenbaum, J., Kaelbling, L. P., & Lozano-Pérez, T. (2021). Learning symbolic operators for task and motion planning. *Arxiv preprint, arXiv:2103.00589*.
- Sinapov, J., & Stoytchev, A. (2008). Detecting the functional similarities between tools using a hierarchical representation of outcomes. In *2008 7th IEEE International Conference on Development and Learning*, pp. 91–96. IEEE.
- Sinapov, J., & Stoytchev, A. (2007). Learning and generalization of behavior-grounded tool affordances. In *2007 IEEE 6th International Conference on Development and Learning*, pp. 19–24. IEEE.
- Sosa, R., & Gero, J. S. (2016). Multi-dimensional creativity: a computational perspective. *The International Journal of Design Creativity and Innovation*, 4(1), 26–50.
- Speer, R., & Havasi, C. (2013). Conceptnet 5: A large semantic network for relational knowledge. In *The People’s Web Meets NLP*, pp. 161–176. Springer.
- Strens, M., & Windelinckx, N. (2004). Combining planning with reinforcement learning for multi-robot task allocation. In *Adaptive Agents and Multi-Agent Systems II*, pp. 260–274. Springer.
- Suárez-Hernández, A., Segovia-Aguas, J., Torras, C., & Alenyà, G. (2020). Strips action discovery. *Arxiv preprint, arXiv:2001.11457*.
- Tosun, T., Daudelin, J., Jing, G., Kress-Gazit, H., Campbell, M., & Yim, M. (2018). Perception-informed autonomous environment augmentation with modular robots. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6818–6824. IEEE.

- Toussaint, M., Allen, K., Smith, K. A., & Tenenbaum, J. B. (2018). Differentiable physics and stable modes for tool-use and manipulation planning.. In *Robotics: Science and Systems*.
- Turner, M., Duggan, L., Glezerson, B., & Marshall, S. (2020). Thinking outside the (acrylic) box: a framework for the local use of custom-made medical devices. *Anaesthesia*, -.
- Varshney, L. R., Wang, J., & Varshney, K. R. (2016). Associative algorithms for computational creativity. *The Journal of Creative Behavior*, *50*(3), 211–223.
- Wang, R., Lehman, J., Clune, J., & Stanley, K. O. (2019). Paired open-ended trailblazer (poet): Endlessly generating increasingly complex and diverse learning environments and their solutions. *Arxiv preprint, arXiv:1901.01753*.
- Wicaksono, H., & Sammut, C. (2020). A cognitive robot equipped with autonomous tool innovation expertise.. *International Journal of Electrical & Computer Engineering (2088-8708)*, *10*(2).
- Wicaksono, H., & Sheh, C. S. R. (2017). Towards explainable tool creation by a robot. In *IJCAI-17 Workshop on Explainable AI (XAI)*, p. 63.
- Xie, A., Ebert, F., Levine, S., & Finn, C. (2019). Improvisation through physical understanding: Using novel objects as tools with visual foresight. *Robotics Science and Systems. Freiburg, Germany*, -.
- Xu, D., Nair, S., Zhu, Y., Gao, J., Garg, A., Fei-Fei, L., & Savarese, S. (2018). Neural task programming: Learning to generalize across hierarchical tasks. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3795–3802. IEEE.
- Yang, C., Lan, X., Zhang, H., & Zheng, N. (2020). Autonomous tool construction with gated graph neural network. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 9708–9714. IEEE.
- Younes, H. L., & Littman, M. L. (2004). Ppddl1. 0: An extension to pddl for expressing planning domains with probabilistic effects. *Techn. Rep. CMU-CS-04-162*, *2*, 99.
- Zhang, W., Sjoerds, Z., & Hommel, B. (2020). Metacontrol of human creativity: The neurocognitive mechanisms of convergent and divergent thinking. *NeuroImage*, *210*, 116572.
- Zhao, A., Xu, J., Konaković-Luković, M., Hughes, J., Spielberg, A., Rus, D., & Matusik, W. (2020). Robogrammar: graph grammar for terrain-optimized robot design. *ACM Transactions on Graphics (TOG)*, *39*(6), 1–16.
- Zhu, Y., Zhao, Y., & Chun Zhu, S. (2015). Understanding tools: Task-oriented object modeling, learning and recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2855–2864.