Robotics and Autonomous Systems 86 (2016) 196–206

Contents lists available at ScienceDirect

Robotics and Autonomous Systems

journal homepage: www.elsevier.com/locate/robot

Hierarchical reinforcement learning as creative problem solving

Thomas R. Colin [a,*], Tony Belpaeme [a], Angelo Cangelosi [a], Nikolas Hemion [b]

[a] Plymouth University, Drake Circus, Plymouth, Devon, United Kingdom
[b] AI-Lab, Aldebaran Robotics, 48 Rue Guynemer, 92130 Issy-les-Moulineaux, France

HIGHLIGHTS

- Reinforcement learning's option switches are analogous to psychological insight.
- Insight and options reveal comparable capabilities for transformational creativity.
- Open problems remain: lifelong learning, switching when exploring, option discovery.

ARTICLE INFO

Article history:
Received 28 July 2015
Received in revised form 17 May 2016
Accepted 23 August 2016
Available online 11 September 2016

Keywords:
Creativity
Insight
Hierarchical reinforcement learning
Robotics

ABSTRACT

Although creativity is studied from philosophy to cognitive robotics, a definition has proven elusive. We argue for emphasizing the creative process (the cognition of the creative agent), rather than the creative product (the artifact or behavior). Owing to developments in experimental psychology, the process approach has become an increasingly attractive way of characterizing creative problem solving. In particular, the phenomenon of insight, in which an individual arrives at a solution through a sudden change in perspective, is a crucial component of the process of creativity.

These developments resonate with advances in machine learning, in particular hierarchical and modular approaches, as the field of artificial intelligence aims for general solutions to problems that typically rely on creativity in humans or other animals. We draw a parallel between the properties of insight according to psychology and the properties of Hierarchical Reinforcement Learning (HRL) systems for embodied agents. Using the Creative Systems Framework developed by Wiggins and Ritchie, we analyze both insight and HRL, establishing that they are creative in similar ways. We highlight the key challenges to be met in order to call an artificial system "insightful".

1. Introduction

People achieving extraordinary creative breakthroughs are not born creative; they require extensive hands-on experience to become capable of brilliance in a particular domain. This is not a straightforward result. Indeed, if creativity consists in the production of novelty, one might expect that habits inherited from past experience are a hindrance, and indeed past experience can cause mismatched transfer to new tasks [1]. Nonetheless, there is wide agreement that considerable domain experience is required for people to discover solutions to so-called "insight problems" [1,2] — problems considered difficult precisely because of negative transfer. Animal insight seems no different: macaques with significant laboratory experience are excellent at seeing right through novel experiments in a moment of insight, including for problems in which past experiments seem to discourage the successful behavior [3].

Consider embodied agents, such as robots, situated in an unknown environment, gathering experience from repeatedly interacting with their environment. How can such agents cope with a novel situation, one for which the learned response fails? Their situation resembles that of naive human beings or animals confronted with a previously unseen problem or environment. While some animals, and specifically humans, show remarkable adaptability by creatively developing novel and useful behaviors, artificial agents typically fall short when faced with change. We believe that discoveries from the cognitive sciences offer a way forward, despite the methodological difficulties associated with integrating contributions from multiple disciplines. Below, we integrate results from psychology and machine learning, and suggest a research program for giving artificial agents the capacity to be creative in problem-solving.

Essential to our approach is the notion that creativity is a process — that is, we assume that creativity is a specific manner in which individuals reason and make decisions. We leave aside the domain of artistic creativity which involves socially and culturally

* Corresponding author.
E-mail address: thomas.colin@plymouth.ac.uk (T.R. Colin).

http://dx.doi.org/10.1016/j.robot.2016.08.021
0921-8890/

construed value; and we restrict our contribution to the domain of creative problem-solving. In the last 30 years, significant developments have been made in our understanding of human creativity in problem-solving, leading to gradual convergence towards a single integrated theory combining analytic search and insight [2,4].

When tying creativity in with machine learning, it seems that many properties of Hierarchical Reinforcement Learning (HRL) techniques match the description of insight in human beings — including both analytic progress and the ability to restructure the search space. When we analyze HRL and insight as creative systems, using the Creative Systems Framework (CSF) [5,6], these similarities become more striking. This suggests how HRL might be used to produce insightful behavior in artificial agents.

Such an approach is especially relevant to robotic creativity because it is based on control techniques: it manipulates policies in a sensorimotor space, rather than features or parameters in a conceptual space. This distinguishes it from methods used for some of the more abstract domains in computational creativity research, such as e.g. joke invention or musical composition.

Our contribution is two-fold. First, we propose a process-focused theoretical analysis of creativity in problem-solving. We develop this analysis in two disciplines: psychology (insight) and machine learning (HRL). Second, we use the CSF to unveil connections between psychological theories of insight and HRL methods. This sheds light on a novel way to build an agent whose behavior exhibits parallels with human creativity.

2. Assessing creativity

2.1. Concerns for computational creativity

Researchers involved in computational creativity often face two concerns.

The first one is the difficulty of defining creativity. There is widespread agreement on the following working definition: creativity is "the ability to generate novel, and valuable, ideas" [7]. This definition states that something merely novel[1] (alternatively, surprising, original, unusual) is not necessarily creative; otherwise random behavior would be considered highly creative [8]. This justifies the introduction of value or usefulness in the definition. But creativity researchers are aware of the limitations of this working definition; this is especially the case in computational creativity, where precise criteria are needed to assess proposed algorithms. What exactly constitutes sufficient novelty, and in relation to what do we measure value or usefulness?

The second concern is directed at the notion that a machine could be creative. The skeptic's arguments are similar to those disputed by Turing as he ponders the question "Can machines think?" [9]. Turing dismisses the initial question as uninteresting (because it is dependent on the conventional usage of the terms "machine" and "to think" in English), and replaces it with the eponymous test: can a digital computer do well in the imitation game? The substitution clarifies the debate and grounds it in experience. In the same spirit, Boden proposed a version of the Turing Test for computational artwork, which would be passed by creative products "indistinguishable from one produced by a human being; and/or, [seen] as having as much aesthetic value as one produced by a human being" [10].

In an attempt to answer both concerns, Colton and Wiggins [11] shift the burden of finding criteria onto an unbiased observer. They call computational creativity "the philosophy, science and engineering of computational systems which, by taking on particular responsibilities, exhibit behaviors that unbiased observers would deem to be creative". Below, we discuss this concept of creativity and challenge it in the context of creative problem-solving.

2.2. Products and processes

Boden [8,10] and Colton and Wiggins [11] focus on the end product (behavior or artifact) of creative agents or software. They ask whether the *output* of the algorithm is creative — implicitly excluding any reference to the inner workings of the agent. But their Turing test of creativity, by concentrating on appearances, rewards front-end improvements and variations on a given style over genuine novelty [12]. Recognizing these limitations, theorists of computational creativity have proposed increasingly sophisticated assessment methods for creative outputs [13,14], while also acknowledging the specificity of creative processes [5,15].

We take one step further in that direction. Assessment methods for creativity that focus on end-products, such as those inspired by the Turing test, imply the following:

- There is something special about creative products (i.e. novelty and value).
- By extension, processes are creative when they result in creative products.

We propose to turn this view around:

- There is something special about creative processes.
- By extension, products are creative when they result from the successful use of creative processes.

This resolves some issues with the product view: rather than assessing whether a product is novel, we can check whether it is the result of a copying process; rather than determining the value of a product, we can identify how it was produced, what it was produced for, and how well it fulfills that function. This approach also correctly classifies instances of mere chance as non-creative, even when the end product is indistinguishable from "the real thing" (e.g. when made by the proverbial monkeys with typewriters). And it implies that, should a computer achieve a result that would be called creative if achieved by a human being, but by using a non-creative process (such as exploiting its speed advantage to compute every possibility), we still should not call that computer creative. But this raises a question: what is special about creative processes?

2.3. Creativity as search

Fortunately, the computational creativity community has worked towards characterizing creative processes in a general manner [5,6,16]. Wiggins' Creative Systems Framework (CSF) [5] proposes an analysis of creativity as search, focusing especially on performing search in at least two levels: (1) the search *in* a problem space, and (2) the (meta-)search *of* a problem space.

The first level of search achieves exploratory creativity, such as the analytic discovery of new theorems from axioms, or of a control policy to achieve a task. The second level achieves the more elusive transformational creativity, which consists of a radical change of the domain being investigated; this appears to occur when solving insight problems.[2]

We will consider the simplified version of the framework proposed by Richie [6]. Because the framework aims at clarifying the nature of creative computation, it can be considered a definition of creativity, and that will be our interpretation. That is, any cognitive process that can be accurately described as performing the CSF's search and meta-search, without assistance from a human programmer, can be considered creative. If the framework is indeed

[1] Novel to the creative agent, in what Boden calls "Psychological-creativity" (as opposed to "Historical-creativity") [8].

[2] See [8] for a description of exploratory, combinatorial and transformational creativity.