Can a Computer be Lucky? And Other Ridiculous Questions Posed by Computational Creativity

Abstract. Given the fragility of today's intelligent systems, we consider the necessity of creativity in systems designed for artificial general intelligence. We examine an archetypical creativity "algorithm" suggested by Czikzentmihalyi in the context of computational systems, and, in particular consider the computability of such an algorithm. We argue that it is likely not computable, in the Turing sense, but that this need not necessarily preclude the building of computationally creative systems, and, by extension, (potentially) systems with a level of artificial general intelligence.

Keywords: computational creativity, computability, inspiration

1 Introduction

It is not difficult to argue that the promise of artificial intelligence is beginning to be fulfilled by a variety of modern intelligent systems, from chess programs to autopilots to loan underwriters to search engines. It is also not difficult to argue that no extant intelligent system is yet in danger of exhibiting artificial general intelligence (AGI). There are likely many reasons for this, with one certainly being the fragility of today's systems—it is clear that if not sufficient, robustness is certainly a necessary attribute for any system to claim general intelligence. And, while there may be multiple approaches to endowing today's fragile systems with the requisite robustness, one promising approach is that of *computational creativity*—imbuing computational systems with the ability to do things that an unbiased observer would deem creative. While perhaps not yet a common talking point in AGI discussions, creativity in computational systems has begun to be mentioned in this context [2, 21]. Here, we continue that discussion, considering ideas from computational creativity in the context of the theory of computability.

While the field of computational creativity is still relatively nascent, with Boden generally credited for beginning the discussion [3, 4], there have been a number of recent attempts at building systems that exhibit creativity in a variety of non-trivial domains, including visual art [6, 16], music [9, 10], language [25, 11], poetry [12, 23], humor [14, 22], narrative [18, 19], mathematics [5] and even cooking [15, 24]. In addition, there have been some attempts at the beginning of a generalization from these domains to an abstract theory of computational creativity [20, 26, 7], though much work remains to be done. One obvious tack is to attempt to understand creativity in humans and then translate that to a computational stratum, and this is how we will approach the problem here. It should

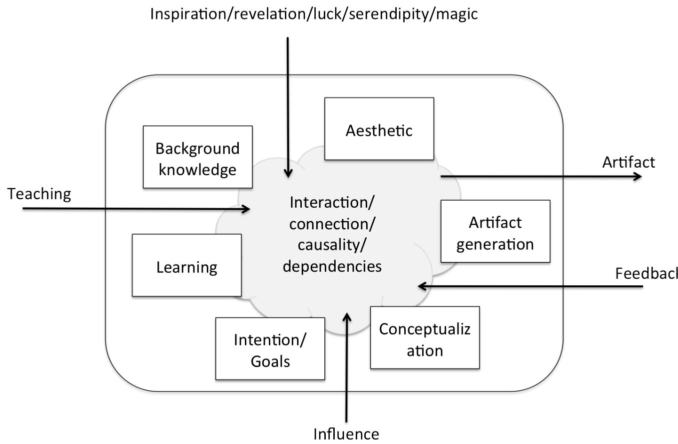

Fig. 1. *Possible logical overview of a creative agent.* The component internal mechanisms are meant to be representative rather than exhaustive. In the same spirit, no attempt here is made to accurately visualize the dependencies and communication between these mechanisms. The agent (potentially) communicates with the environment in several ways, represented by labeled arrows entering or leaving the agent.

be noted that there are difficulties with this, including the common airplanes-don't-flap-their-wings-but-they-still-fly argument and the fact that creativity is at best an ill-defined concept; indeed, most practitioners of computational creativity eschew any direct attempts at such an ambitious analogical transfer, but we will pursue the topic here nonetheless.

Many investigators have attempted to elicit the "creativity algorithm" used by people when they are being creative. There are many variations on this "algorithm", but it most often looks something like the general steps distilled by Czikzentmihalyi [8]:

1. preparation
2. incubation
3. insight
4. evaluation
5. elaboration

In what follows we will treat this "algorithm" as a surrogate for all of these proposals and discuss each step in the context of computability. Note that, as Czikzentmihalyi and others have observed, these steps should not be taken as a single-iteration process but rather as parts of a multiple-iteration, possibly recursive process in which the various steps are revisited multiple times, in varying order as necessary. Here, we will ignore this obviously important issue of flow control and focus only on the five steps. In what follows we will consider an archetype agent (see Figure 1) whose ambition is creativity, and we will consider how that agent might follow Czikzentmhalyi's archetypal "algorithm" for being so.

Such an agent is composed of many internal mechanisms/processes that interact with each other in some unspecified way, and these internal mechanisms and their interactions are the subject of much ongoing research, with both human and computational subjects. However, they are not of specific interest here. Because the agent exists in an environment, the agent interacts with the environment in multiple ways, some of which are shown as labeled arrows that enter or leave the agent abstraction. Both on the human and computational fronts, there have been significant advances in understanding many of the individual mechanisms shown in the figure. What is still less understood is how these mechanisms interact to realize the creativity "algorithm", and it is this question that we will try to say something about here.

2 Computability of Creativity

We will treat each of the steps of the "algorithm" in turn, positing something about the salient agent mechanisms and their interactions and what the prospects are for its implementation in a computational setting.

2.1 Preparation

Preparation is the initial process of learning about the domain in which an agent will attempt creativity. It entails significant interaction with the environment for the acquisition of background knowledge and understanding accepted practices and open problems. In addition, an agent must acquire or develop some aesthetic sense of the domain, where we use aesthetic here in the sense of some abstract notion of quality. Initially this sense could be taught to the agent by the environment in just the same way that the background knowledge is. Of course, agents that develop new aesthetic sensibilities (a meta-level creative act?) are likely to be considered more creative in their output. Eventually, an agent may use its acquired background information to learn/develop such novel aesthetics. It is sometimes argued that too much preparation can result in the repression of creativity as old, set ideas are assimilated too thoroughly. However, it is certainly the case that a good deal of preparation is necessary to facilitate downstream processes, particularly those of evaluation and elaboration.

Computational challenges inherent in this step include the acquiring, encoding, and understanding of knowledge, ontologies, formalization, etc. as well as methods for learning/developing evaluation strategies. These are nontrivial tasks, to be sure, but many proof-of-concept structured, semi-structured and unstructured projects (cf., Wikipedia[1], WordNet [1], ContextNet [13], the semantic web[2] and even the World-Wide-Web itself) put the knowledge acquisition aspects squarely in the category of difficult-but-manageable engineering tasks. As for learning/developing an aesthetic, general purpose machine learning

[1] http://www.wikipedia.org
[2] http://www.w3.org/2013/data/

techniques exist for inferring structural relations from data. In many respects, this preparation step is not unlike developing pedagogy for human students, and many AI approaches to the problem, from ontologies to machine learning would be recognized to some extent by educational practitioners.

2.2 Incubation

Incubation is the process of "putting it on the back burner"—allowing ideas to simmer in a possibly unconscious way, the exploration of unusual connections, brainstorming, etc. This is often described as an open-ended process without a clear time line or quantifiable goals, other than "finding something interesting". The agent conceptualizes and generates ideas using its knowledge base and additional outside environmental influences. These concepts and ideas are judged against the agent's aesthetic sense and very often discarded immediately. While this step can be performed consciously and intentionally, as in the aforementioned brainstorming session, it is often described as best happening when the conscious mind is otherwise engaged (with another task, while exercising, while in the shower or even while sleeping). It is unclear whether this unconscious aspect is necessary or simply catalyzing and whether intentionality may be uncoupled from consciousness.

Given an effective organization and acquisition of knowledge, it is not difficult to argue that computational systems will actually (eventually) enjoy a significant advantage over human intelligence in this step—speed, lack of bias, nonsusceptibility to fatigue, distraction, boredom, etc. all favor computational approaches to the exploration of potentially interesting connections and the generation of ideas and conceptualizations at scale. Of course, any "intelligent" biases should be formalized and leveraged by computational systems for obvious reasons; however, determining whether a bias is useful or potentially detrimental is likely classifiable as a creative task itself (another meta-level concern?)

2.3 Insight

Insight is most often described as having nothing explicitly to do with any action or intention of the agent; indeed, many people will describe it as originating from outside themselves. Depending on a person's bent, this might be called inspiration or revelation or luck or serendipity or magic or something else. It is often associated with an "Aha" moment, when things just fall into place, the answer suddenly becomes clear, etc. This presents us, apparently, with something of a Gödelian quandary, which may (or may not be) resolvable in one of several ways.

One possibility is that insight is an agent fabrication that is not really necessary for creativity; a second possibility is that insight, though a necessary part of the "algorithm", does not, in fact, originate outside the agent at all[3]; a third

[3] The agent's belief that it does may be explainable by appeal to the unconscious, insufficient understanding of neuropsychological and cognitive processes, etc.

possibility is that insight is somehow necessary for human creativity but may not be for a computational variant[4], and it is therefore unimportant for the current discussion; a fourth possibility is that, in fact, insight is necessary for creativity in any medium and does also, in fact, represent a Gödelian process over which the agent can never have any control.

The computational challenge faced in realizing this step of the "algorithm" depends upon which, if any, of the possibilities above best explains insight. In the first three cases, the simplest solution must involve some variation on a brute force search (in what space? Is identification/construction of the search space another meta-level problem?) Such an approach will (eventually) produce artifacts that satisfy the agent's aesthetic and are potentially considered creative. Of course in any interesting domain, the search space is very likely to be infinite and so the first real computability concern raises it's head. Such a search will not be computable in the strong sense of decidability (see more on this in Section 2.4); however, it will be in the weaker sense of recognizability, and this could be argued to be no more egregious than is the case for human creativity—we can't define or guarantee it, but we know it when we see it. Of course, the next obvious solution is to introduce search heuristics to circumvent the complexity/computability issues associated with the brute force approach. These may be learned from the environment[5] or invented by the agent (meta-level process, again) and there will be a tradeoff between computational guarantees and likelihood of success.

In the fourth case, we have the possibility that creativity has an analog to Gödel's incompleteness theorem in that something from outside the agent is necessary. This would, of course, preclude any general (strictly) computational creative system and will perhaps seem appealing to some who may see creativity as a last bastion of humanity or as something *ex vi termini* impossible computationally. And yet, if the premise were indeed true, the same would have to be said about the human variety as well. Even if this is the case, we still see creative acts, both personal and historical, occurring with regularity, and we might yet simulate this productivity computationally by acting ourselves as the requisite extra-agent component of insight. That is, computational creativity would be effective at least some of the time only with the aid of human intervention, suggesting something of co-creativity and, at the same time, allowing some to maintain a small toe-hold on the precipice of human superiority.

As a last comment, we note that in at least one theory insight has been equated with re-representation [17]. That is, creativity is difficult (or impossible) when the agent's representation of the problem is not amenable to it—the agent can (figuratively) wander around forever and not discover anything useful until—Aha—it "lucks into" the right representation (this appears like yet another potential meta-level issue, with at least the outside possibility that there may be no access to the meta-level by the agent).

[4] Again, the airplane vs. bird analogy.

[5] And may in fact simulate some unconscious cognitive or sub-cognitive process.

2.4 Evaluation

Evaluation is the application of the aesthetic measurement process to the product of the generation process. Both of these processes may be learned during preparation or they may be the product of a (meta)creative process themselves. This is an internal evaluation, not to be confused with the external appraisal and feedback from the environment to which all potentially creative acts must be subject[6]. A result that passes the aesthetic test will be elaborated and eventually presented to the environment for that external assessment.

Though the high-level process description is deceptively simple, the computational challenges posed at this step are non-trivial. Assume that evaluation is computable in principle, so we have an algorithm E that computes it. What we want is another algorithm F that can tell us whether an artifact a is accepted by E; that is, we are interested in the language $L(F) = \{a|E \text{ accepts } a\}$. Initially, let's optimistically assume that E is even computable in the strong Turing sense, that is, it is *decidable*. Then, we have an easy algorithm for F (run E on input a), and, thus the rudimentary makings of an algorithm C for using F to solve whatever the problem is (that is, to be creative):

```
C()
  do
    choose a
  until a in L(F)
  return a
```

Of course, in any interesting case, the space to explore is infinite and E may be very selective, so this algorithm may be worthless, but at least it is computable—it has been reduced to a "simple" question of complexity. To make it useful, we need a good exploration strategy. It is possible that this might be learned from the environment during the preparation step, but if such a search strategy is already known, then the problem to which the search strategy is to be applied is likely already (nearly) solved. So, for non-trivial problems, it is likely that the agent must discover a search strategy. This is yet again a meta-level problem, and one we'll examine in a bit more detail.

We now have a new (meta)space to explore (this one containing exploration strategies for the original space containing the artifacts a) in which we are looking for a new (meta)artifact (the exploration strategy), so we have to reconsider the

[6] In a very real sense, creativity is a social phenomenon. It is impossible to attribute creativity in a vacuum—both a creator and one or more receivers are necessary for creative attribution. The creator produces an artifact, or engages in a process, and the receiver(s) experience the result and attribute creativity based upon their perception of the artifact's, the process' and/or the creator's characteristics.

five steps in that context. Of course this meta-problem immediately suggests yet another meta-meta-problem—how do we search the space of search strategies?[7]

(meta)preparation —Is this a new domain with its own background knowledge, etc? How should knowledge be represented at this abstract level? Does (already) knowing the base domain suggest the strategy? Is there some level of abstract exploration strategy domain that an agent must master before it can reasonably expect any success at this level? Or, perhaps there is often not much to consider here, and one just hops between meta- and base-level steps 2-4...

(meta)incubation —How does the agent make connections at this level? How are abstract connections related to base-level connections? Another way to think about this is that the agent is looking for ways to structure the base space so that it is easy to explore. So, the dual problem is one of representation rather than exploration strategy—if the agent can re-represent the base domain so that, for example it is (approximately) convex, the exploration problem becomes trivial.

(meta)insight —This is still an "Aha" moment. Or not. The same arguments apply as were given for the base level.

(meta)evaluation —The agent must now have some (meta)aesthetic for recognizing a good search strategy/representation, which suggests the following interesting philosophical question: Can an agent elaborate this (meta)aesthetic without recognizing where it points in the base search (and thus already solving the base problem)? A more concrete version of this question is whether it is possible to recognize a good fitness function without knowing what objects score well under that function.

(meta)elaboration —In many cases, this likely degenerates to simply applying the exploration strategy (or the re-representation) back in the base domain. There may be situations in which the search strategy or re-representation itself is of general interest and perhaps even supersedes anything discovered in the base domain using it. In such cases, consideration must be given to communicating the (meta)discovery and its import.

Returning to our base-level discussion, we first note the potential difficulty this apparent recursion introduces—it is not clear that there is a base case for terminating the recursion. Perhaps there exists a level of abstraction sufficient so that no further meta-level issues can arise. Or perhaps there will always be a point at which an "Aha" moment must be provided (by a human) that will serve the purpose of tipping the process out of the recursion.

---

[7] It is possible that this third-level question is still related to the base domain in a non-trivial way, so that perhaps we don't have a really complete abstraction.

Finally, we will mention that it is very probably unrealistic to suppose that the evaluation function E is decidable; rather, it is likely more realistic to suggest that E is at best *semi-decidable*—a quality artifact can be recognized, but it is not possible to recognize an artifact that does *not* measure up to the aesthetic.[8]

Now, the algorithm for F cannot simply consist of running E on a[9] because E may not halt. In this case, we need $F(E, a)$ to be decidable in some other way. Unfortunately, the obvious trivial reduction from the classical Halting Problem[10] means that this is not possible. So, in the absence of a decidable aesthetic, the problem of computational creativity is not computable in the strong sense, independent of whether the insight problem is real and independent of any difficulties (or lack thereof) due to meta-level recursion issues.

2.5 Elaboration

The elaboration step is often described as the "99% perspiration" that complements the "1% inspiration" of insight. The process is deliberate and intentional—it is Edison trying 2000 different materials while looking for the perfect filament—the artifact is situated relative to the background knowledge, additional variations and details are generated and evaluated against the aesthetic, feedback from the environment may drive additional iterations and local refinement (or, even potentially major revisions). Herein lies all the hard work of development and polishing ideas, framing results and selling the finished product, and these processes may themselves require additional creativity, both large and small—iterating or recursing on some or all of the five "algorithmic" steps.

The computational challenges here are in many ways similar to those at the preparation stage, only in the reverse. Now, the system, rather than needing to acquire knowledge must dispense it, communicating both results and their import. The hard work of filling in details, exploring possible processes, etc. may again be argued to be advantage computer for the same reasons cited above. The difficulty of framing or marketing the result is a more complex consideration, and may be regarded as a creative act itself—what story to tell, how to write the research paper, how to market a product, how to explain a piece of art.

3 Final Thoughts

It is unlikely that the creativity "algorithm" is computable in the strong Turing sense of decidability. If this is the case, and if creativity is necessary for artificial

[8] Perhaps the environment itself accepts those artifacts that everyone appreciates and rejects those that no one appreciates but isn't sure about those with mixed reception. Any aesthetic that accurately models such a scenario will not be decidable given the existence of all three types of artifact.

[9] Unless it is acceptable to have programs that may not terminate. If the insight issue resolves to the sticky fourth case, this will be unavoidable, in which case F may remain a simple simulation of E without incurring any additional computational penalty for the overall "algorithm".

[10] Actually, the most obvious reduction is from the related Acceptance Problem.

general intelligence (as we've suggested without substantiation), it follows that AGI would also not be Turing computable. It is somewhat more likely that creativity is weakly Turing computable in the sense of recognizability (semi-decidability), though this is not yet proven. And, even given this result, the weak computability of AGI would of course not immediately follow unless we can argue the sufficiency of computational creativity (and we do not suggest this here even without substantiation).

Still, Turing computability is a very strong result, and it is not surprising that a creativity "algorithm" might resist this level of constraint; indeed, most of human intelligence, if held to the strict standards of the current theory of computability, is a failure. That is not to say that efforts at computationally simulating it are failures but that humans themselves fail to "compute" by such strict standards. Also, it is certainly true that other uncomputable problems of interest, in many instances, do yield themselves to computational attacks of varying efficacy, so it is not unreasonable to expect that computational creativity may yield significant advances toward a theory of AGI.

Of course, there is also the (remote) possibility that in fact all the assumptions necessary to render creativity strongly computable will prove true, and we will discover that we can, simply, brute force an "Aha" moment. Wouldn't that be lucky?

References

1. Wordnet. http://wordnet.princeton.edu/wordnet/
2. Abdel-Fattah, A.M.H., Besold, T., Kühnberger, K.U.: Creativity, cognitive mechanisms, and logic. In: Proceedings of the 5th International Conference on Artificial General Intelligence. pp. 1–10. Springer-Verlag, Berlin, Heidelberg (2012)
3. Boden, M.: The Creative Mind. Abacus, London (1992)
4. Boden, M.: Creativity and artificial intelligence. Artificial Intelligence 103, 347–356 (1998)
5. Colton, S.: Computational discovery in pure mathematics. In: Deroski, S., Todorovski, L. (eds.) Computational Discovery of Scientific Knowledge, Lecture Notes in Computer Science, vol. 4660, pp. 175–201. Springer Berlin Heidelberg (2007)
6. Colton, S.: The painting fool: Stories from building an automated painter. In: McCormack, J., dInverno, M. (eds.) Computers and Creativity, chap. 1, pp. 3–38. Springer, Berlin; Heidelberg (2012)
7. Colton, S., Charnley, J., Pease, A.: Computational creativity theory: The FACE and IDEA descriptive models. In: Proceedings of the 2nd International Conference on Computational Creativity. pp. 90 – 95 (2011)
8. Csikszentmihalyi, M.: Creativity: Flow and the Psychology of Discovery and Invention. Harper Perennial (1996)
9. Dannenberg, R.B.: A vision of creative computation in music performance. In: Proceedings of the 2nd International Conference on Computational Creativity. pp. 84 – 89 (2011)
10. Eigenfeldt, A., Pasquier, P.: Negotiated content: Generative soundscape composition by autonomous musical agents in *Coming Together: Freesound*. In: Proceedings of the 2nd International Conference on Computational Creativity. pp. 27 – 32 (2011)

11. Gatti, L., Guerini, M., Callaway, C., Stock, O., Strapparava, C.: Creatively subverting messages in posters. In: Proceedings of the 3rd International Conference on Computational Creativity. pp. 175 – 179 (2012)
12. Gervás, P.: Exploring quantitative evaluations of the creativity of automatic poets. In: Proceedings of the 2nd Workshop on Creative Systems (2002)
13. Liu, H., Singh, P.: ConceptNet—a practical commonsense reasoning tool-kit. BT Technology Journal 22, 211–226 (2004)
14. Manurung, R., Ritchie, G., Pain, H., Waller, A., O'Mara, D., Black, R.: The construction of a pun generator for language skills development. Applied Artificial Intelligence 22, 841 – 869 (2008)
15. Morris, R., Burton, S., Bodily, P., Ventura, D.: Soup over bean of pure joy: Culinary ruminations of an artificial chef. In: Proceedings of the 3rd International Conference on Computational Creativity. pp. 119 – 125 (2012)
16. Norton, D., Heath, D., Ventura, D.: Finding creativity in an artificial artist. Journal of Creative Behavior 47, 106– 124 (2013)
17. Ohlsson, S.: Information-processing explanations of insight and related phenomena. In: Keane, M., Gilhooly, K. (eds.) Advances in the Psychology of Thinking, vol. 1, pp. 1 – 44. Harvester–Wheatsheaf (1992)
18. Pérez y Pérez, R., Sharples, M.: Three computer-based models of storytelling: BRUTUS, MINSTREL and MEXICA. Knowledge-Based Systems 17, 15 – 29 (2004)
19. Riedl, M., Young, M.: Story planning as exploratory creativity: Techniques for expanding the narrative search space. New Generation Computing 24, 303 – 323 (2006)
20. Ritchie, G.: Some empirical criteria for attributing creativity to a computer program. Minds and Machines 17, 76–99 (2007)
21. Schmidhuber, J.: Artificial scientists and artists based on the formal theory of creativity. In: Proceedings of the 3rd Conference on Artificial General Intelligence. pp. 145 –150 (2010)
22. Stock, O., Strapparava, C.: The act of creating humorous acronyms. Applied Artificial Intelligence 19, 137 – 151 (2005)
23. Toivanen, J.M., Toivonen, H., Valitutti, A., Gross, O.: Corpus-based generation of content and form in poetry. In: Proceedings of the 3rd International Conference on Computational Creativity. pp. 211 – 215 (2012)
24. Varshney, L.R., Pinel, F., Schörgendorfer, A., Chee, Y.M.: Cognition as a part of computational creativity. In: Proceedings of the 12th IEEE International Conference on Cognitive Informatics and Cognitive Computing. pp. 36 – 43 (2013)
25. Veale, T., Hao, Y.: Comprehending and generating apt metaphors: A web-driven, case-based approach to figurative language. In: Proceedings of the Association for the Advancement of Artificial Intelligence. pp. 1471–1476. AAAI Press (2007)
26. Wiggins, G.A.: A preliminary framework for description, analysis and comparison of creative systems. Knowledge-Based Systems 19, 449 – 458 (2006)