# The Creative Machine

**James Hodson***

Jožef Stefan International Postgraduate School
Jamova 39, 1000 Ljubljana, Slovenia
james.hodson@ijs.si

**Abstract**

This paper offers a critical review of the underlying assumptions in the field of *Computational Creativity*. We present and integrate the state of the art in the search for machines that could be considered creative by human standards. Through the lens of existing literature, philosophical thought, and empirical experimentation, we propose ways to better understand the roots of creativity, and a new approach for its investigation within the field of Artificial Intelligence.

## 1 Introduction

What is creativity, and what would make a machine *creative*? This is the question that beguiles the literature in the field of Computational Creativity, oft repeated and cited as too difficult to nail down (see, for example, (**?**), (**?**), (**?**)). The field of Computational Creativity is tasked with both defining the philosophical foundations of the search for creativity, and transferring this tentative understanding into real machines that are convincing and valuable to society. A rather sizeable literature has emerged in recent years that attempts to provide grounding for the field by defining it, in general terms, as "building software that exhibits behavior that would be deemed creative in humans." (**?**). This is a functional definition that provides little guidance as to the areas that should be examined in our search for knowledge, and begs the philosophical question of the roots of creative behaviour more generally.

Computational Creativity is not a new field. It is considered a sub-field of Artificial Intelligence, and the thirst to understand how a machine might appear to undertake creative acts akin to humans has lasted at least as long as AI itself. Among the first attempts at defining the requirements for creative acts, Newell, Shaw, and Simon (**?**) provided the following four heuristics:

1. The answer has novelty and usefulness (either for the individual or society).
2. The answer demands that we reject ideas we had previously accepted.
3. The answer results from intense motivation and persistence.
4. The answer comes from clarifying a problem that was originally vague.

In what is considered one of the foundational works in the field, Margaret Boden (**?**) states that creative acts represent (1) *novel combinations* of familiar ideas, (2) *explorations* of the potential of conceptual spaces, and (3) *transformations* that enable the generation of novel ideas. In short, Boden believes creative acts need to be novel, surprising, and valuable.

More recently, the literature appears to tacitly accept that a notion of creativity is difficult to pin down. In their *apology* for the field of Computational Creativity, Colton and Wiggins state:

> "[P]erhaps creativity is, for some proponents of AI, the place that one cannot go, as intelligence is for AIs opponents. After all, creativity is one of the things that makes us human; we value it greatly, and we guard it jealously." (**?**)

Among further refinements to these definitions is the idea that {em creative processes} in varying disciplines should be viewed as "family resemblances" rather than a rigid set of requirements (for an overview of family-resemblances in concept formation from a category-theory perspective, see (**?**)). We make no claim regarding how specific concepts more or less suggest {em creativity} to observers across disciplines, and fully allow for the contextualisation of these markers.

However, none of these suggestions provides a satisfying definition of what it is to be creative, though each in turn captures certain qualities that may be associated with particular instances of creative processes in one domain or another. In fact, the elephant in the room appears to be that the notion of "creative" is no clearer or more meaningful in our minds than would be "imaginative", or "playful". These concepts that we use to personally categorize behaviour ex-post do not have a stable grounding in the processes that generate them. The search for an a priori understanding of an ex-post phenomenon may be a red herring, which would go someway to explaining the case-based example driven approach that researchers in the field have exhibited.

To date, experimentation in the area of Computational Creativity has tended in two primary directions. The first

*The research project leading to these results received funding from the EU Research and Innovation programme Horizon 2020 under grant agreement No. 675044 (BigDataFinance).

camp aims to discover the set of domain-independent creative processes, able to generate creative artefacts when applied in any domain. Such processes include *conceptual blending* and *bisociative discovery* of new concepts. The second camp looks to creative processes that are active in specific domains, such as music production, painting, pun generation, and poetry. Rather than find governing principles, the second approach emulates and extends forms specific to each discipline. (**?**)

In this paper, we aim to provide a more rigid foundation for the exploration of creative behaviour in a computational framework. We will proceed by offering a detailed philosophical discussion of the roots of creativity as perceived and examined by human beings. This discussion will provide a backbone of definitions for further exploration of the kinds of questions that are both useful and interesting to those engaged in scientific research in this area. Our philosophical discussion will be augmented with an empirical exploration of human creativity (and the expectations thereof), through a comprehensive dataset of web searches for creative artefacts and processes. Although not definitive, our empirical exploration offers the first attempt to tie a functional definition to a concrete understanding of this area of study. Finally, we will draw from a natural experiment of creativity by comparing the occurrence of *creative* assertions in human and machine-based reasoning. This test will serve to highlight the areas of investigation noted in our philosophical discussion, allowing us to tie together the structural foundations and future aims.

## 2 Philosophical Foundations

As with most discussions of a philosophical nature, it is often necessary to start with definitions. In our case, definitions are also the object of our search. As we set out on this path, it is important to make clear that efforts to date under the flag of *Computational Creativity* have value independent of its taxonomical roots. No part of our discussion is aimed at undermining the great progress that has been made in a variety of disciplines, from automating the process of mathematical discovery (**?**) to visual arts (**?**), slogan generation (**?**), musical composition (**?**), and many more. Our aim in this discussion is to better understand what we mean by *computational creativity*, and how this relates to creativity and creative processes more broadly. Through a renewed understanding of the concepts we are trying to model, and how they fit together with other concepts nearby, we hope to provide a unifying theme for Computational Creativity and Artificial Intelligence that can help guide future research efforts.

### 2.1 Definitions of Creativity

Our modern understanding of the term *creativity* has been affected by several tides over the course of many hundreds of years. Its essence is the literal "process of making", stemming from Ancient Greek culture, which had no concept of *creation*, since all arts were seen as following pre-defined rules that strived for the ultimate ideal forms of each concept. In particular, in *The Republic* (see (**?**), Republic X, 597a-e), Plato asks whether a painter and craftsman should be considered *makers* of things, or simply imitators of the true forms and ultimate ideals. In fact, he goes on to elucidate and agree that only the god could be the *natural maker*, and *all* other makers are in one way or another *imitators* of the forms in nature.

Later, in Roman culture, there were two concepts, "facere" and "creare". However, the term "creare", which is the root of "create" and "creativity" in modern usage, referred strictly to the creation of something out of nothing (*ex-nihilo*) (**?**). Although poets were also sometimes allowed to partake in acts of creation, the meaning was reserved for things and concepts that did not previously exist, and were not merely combinations of existing artefacts. In fact, the term was most usually used to mean *the divine act of creation*. For the manufacturing of objects, such as tables, chairs, houses, Romans would use the functional term *construct*, whereas the arts had specific words for *painting*, *sculpting* and so on. The concept of creation bore no connection to the activities of daily life, though these were no more or less *creative*, arguably, than today.

In matters of scientific progress, such as the invention of flying machines, computers, and wind turbines, the general belief would have been that these ideas were incrementally "discovered" in nature. Analogously, today we might claim that all scientific discoveries, even the most important of our time, are simply incremental and nuanced shifts in perception atop a mountain of communal technical and scientific knowledge. It is rare that discoveries happen in isolation, indicating some general state of uncovering from pre-existing knowledge and pre-existing mental processes, rather than creation out of nothing.

To this day, the distinction between "making" and "creating" is heavily blurred. Colloquially, English speakers often refer to creating *meetings*, *websites*, *companies*, etc., despite the fact that none of these acts would habitually be considered *creative* in their own right. On the other hand, the determination of the kinds of artefacts considered to be *creative* by society is highly subjective. There exist very few, if any, universally accepted examples of creative behaviours and artefacts, and the individual decision can vary according to several parameters, including:

- Familiarity with the subject matter;
- Cultural norms; and
- Local context.

In the literature on Computational Creativity, there is little resistance to the idea that judgement of creative artefacts is contextual, subjective, and knowledge-dependent (**?**). In fact, most human endeavours fall prey to these factors– *Intelligence* itself is no more stable in the minds of individuals. However, whereas intelligence has come to refer to the specific cognitive processes that allow an agent to learn from and reason about the information available to them (the manipulation of symbols towards a discriminative objective), creativity possesses no analogous claim. In the AI sub-field of Machine Learning, we can talk about *learnability* through the PAC framework, allowing us to make strong statements

about what can be learned, and how quickly (**?**). In Computational Creativity, although there have been numerous attempts to specify theoretical frameworks for the purpose of evaluation (see, e.g. (**?**)), none of them fundamentally captures a concept of creativity from the underlying processes.

**A fundamental misunderstanding** Earlier, we provided an overview of the most broadly cited characterizations of *creativity*. These, usually vague lists of property attributes that creative actions must conform to, do not capture a universal essence of creativity. However, work in the field of computational creativity often refers to creativity as if it were an absolute phenomenon (as a universally accepted fact), whether absolute globally, or with respect to a particular domain. We dispute such an implicit assumption, both from the perspective of its necessity to the study of computational creativity, and from the perspective of objective reality–there is no need for it, and it does not exist.

Creativity is an ex-post phenomenon,[1] whereas Computational Creativity has tacitly assumed it to be an ex-ante set of cognitive processes.[2] The same sets of cognitive process that allow agents to reason about the world around them, to survive, and evolve knowledge for the future, already envelope the properties necessary and sufficient for the generation of artefacts that fit our classifiers for "creativity". Attempting to reverse engineer the infrequent by-product of these processes is problematic, and will not lead to a stable-set characterization. That is, creative acts do not systematically appear through particular and differentiable processes of cognitive systems (including brains). They can be the result of any combinations of the available symbolic and statistical manipulations available. More often than not, the same processes that yield creative artefacts, yield useful-but-not-*creative* outcomes.

Does this mean that we should give up our search for creativity? Should we let the Artificial Intelligence community at large explore this arena as an indirect consequence of their work? Should we re-brand the field as *Computational Arts and Sciences*?

We should not give up so quickly. But, we must come to terms with the fact that creativity is not separate and special. It is a label applied subjectively to certain societal artefacts, which we might study in two ways. First, by way of analyzing specific classes of artefact considered to be creative, in order to identify the processes of *societal value formation*. This is useful on many levels, and achieving this goal computationally can help improve our understanding of how humans learn, the information used, and the dynamics of "expert" behaviour. Second, by way of harnessing computational models to explore the space of possibilities in the arts and sciences–extending and re-interpreting models of creation while making use of the advantages that Artificial Intelligence techniques can bring (scale, data-driven modeling, reasoning over diverse types of information, avoiding bias, etc.). The more we help computational models understand the arts and sciences, the more they will surprise us with creative artefacts, but this investigation does not start with the creative artefact as an objective, since this framing is not useful for making broad scientific progress.

In short, by adopting this alternative framing for work in the area of Computational Creativity, we can start to measure progress in a more encompassing manner. Through a better understanding of the processes of value formation, we can better inform future frameworks for evaluation. This reframing does not invalidate any of the work done to date, but re-poses it in slightly broader and clearer terms, potentially opening up new possibilities for investigation.

**On mixtures of cognitive processes** When we imagine the vast numbers and types of creative acts that intelligent agents are capable of, it seems logical that these do not stem from particular cognitive processes devoted only to producing creative artefacts. The cognitive abilities that we hold so dear emerge at the intersection of many sub-systems, the repurposing of cognitive pathways, and the particular circumstances/information available. This emergent architecture allows for a richness of generative behaviour that is practically unbounded. In the same way as language is infinitely complex, so too can be the mixtures of creative processes (the grammatical forms) of the mind.

**On the question of general vs domain-specific creativity** Of ongoing debate in the community is the question whether creativity and creative processes are domain-general or domain-specific (see (**?**),(**?**)). That is, are there types of creative processes specific to *musical composition* vs. *sculpting* vs. *mathematical discovery*? It should be fairly clear from the discussion above that the distinction is redundant–there are no creative processes. If every creative artefact could conceivably have originated from a different mixture of underlying cognitive processes, then the question of generality and specificity disappears entirely.

**On the question of convincing** Several authors in the Computational Creativity space bemoan the difficulty of convincing biased observers to attribute creativity to a machine (see, for example (**?**), (**?**), (**?**)), even when those algorithms produce artefacts indistinguishable from those produced by skilled humans. We should note that biases in evaluation performed by humans exist in all corners of human endeavours, from judges in court (**?**), to evaluations of creative work by different genders (**?**), and music generation systems (**?**). In particular, individuals' relative overconfidence in their own abilities tends to lead to an underestimation of the abilities of others (**?**), which bias increases as the "other" is more different. It stands to reason that the evaluation of a machine is about as "other" as one might perceive, even in this era of growing habituation to computational devices.

**Appropriate evaluation criteria** Although a full discussion of the selection of appropriate evaluation criteria for creative systems is beyond the scope of this article, it is clear that this is a problem that beguiles all branches of scientific discovery, and needs to be treated with care. In order

[1] It only exists after-the-fact, by virtue of an observer imbuing an action with such an attribute, independently of the manner of its creation.

[2] That is, that some cognitive processes are inherently *creative*, and others are not

to develop convincing theories, and empirical evidence for or against them, we must be capable of dealing with a variety of human biases, and be wary of the statistical pitfalls of insufficient evidence for the phenomena we claim. More work is necessary towards the understanding of optimal incentive structures, proxy measures, and hypothesis testing in the field of Computational Creativity.

Wiggins (**?**) notes that an overarching principle for the evaluation of creative acts in machines should be the following:

> "The performance of tasks by a computer, which, if performed by a human, would be deemed creative."

In (**?**), the authors argue that "creativity is in the eye of the beholder." Seemingly, the first proposal requires us to define what it means for a human act to be considered creative, and the second tells us that any creative act may only be judged subjectively, individual by individual. Within the framework we are proposing, we want to:

- promote the development of novel techniques in each relevant sub-field;
- reward interesting mixtures of "cognitive processes" used to generate artefacts; and
- encourage the development of systems which can introspect about the processes being implemented.

We believe that a focus on the underlying pillars of intelligence, learning, and reasoning, rather than "creativity" directly, is the best way to maximize the perception of creative value by external observers. As discussed earlier, a deeper understanding of the processes of societal value formation will be the most important precursor to the correct measurement of the ability of machines to imbue those values in the artefacts they produce. The FACE and IDEA frameworks proposed by Colton, Charnley, and Pease (**?**) are an attempt to pose the question of what a plausible account of computational creativity might look like. The claim that the Computational Creativity Theory (CCT) framework is equivalent to the Computational Learning Theory (CLT) framework for Machine Learning is not convincingly elaborated, but it is fair to say that the FACE model (Framing information, Aesthetic measures, Concepts, and Expressions of concepts) does provide some unifying themes for the discussion of relevant artistic artefacts, and the IDEA model ((I)terative (D)evelopment- (E)xecution-(A)ppreciation cycle) can be seen as covering a type of interaction between such artefacts and their social context/environment. However, by focusing on creativity as the first order goal of the generative act, such models may actually force narrower and narrower explorations of the space, which is generally an undesirable outcome.

## 2.2 Structure of Investigation

In (**?**) Pease and Colton provide three main reasons to pursue work in the area of Computational Creativity:

- to provide a computational perspective on human creativity, in order to help us to understand it (cognitive science);
- to enable machines to be creative, in order to enhance our lives in some way (engineering); and
- to produce tools which enhance human creativity (aids for creative individuals).

It is important to note that our clarifying discussion on the nature of creativity, and how this relates to the notion of *computational* creativity, does not require us to deviate from these objectives. In fact, these aims are perhaps more attainable now than ever.

**Virtue and Creativity** We provide an anecdote from a different area to guide the discussion and investigation of creative acts. Specifically, we will treat the idea of Virtue in the *Virtue Ethicist* normative framework (**?**). If one is virtuous, then one makes decisions because that is the only way one can imagine behaving–it is an inherent quality. If there is any level of hesitation, questioning, or ulterior motives that drive a decision not to deviate from the virtuous path, then it is not a virtuous act. For instance, a decision not to steal from the local store must be driven by the understanding that it is ethically wrong to do so, rather than because one fears the consequences, or understands the pain it might cause to others, or if the store does not stock the coveted goods. Similarly, the intention to act according to virtuous doctrine cannot be coming from a desire to be recognised as virtuous–one must be good for the sake of being good. Just like silence, when you talk about it, it disappears. Creativity, like virtue, is not taught (**?**).[3] If the aim is to be creative, then the act cannot be deemed creative. Unlike virtue, creativity can only be judged ex-post (virtue requires evaluation of both the process and the outcome), but for the rest, this provides a neat analogy to current philosophical thinking on the matter, from a Virtue Ethics perspective.

**Grounding the discussion** One may ask whether our motivation of creativity is not at odds with the manner in which it is *taught* in the design professions (such as architecture, marketing, etc.), where people strive directly and explicitly to be creative. Though this may seem like an alluring argument, there are two distinctions to be made. The fact that we strive to teach something does not necessarily mean that it is so. In fact, by teaching our students about creative actions and artefacts, we imbue them with templates that potentially allow them to imitate and develop these themes, but it is unclear which lesson exactly teaches creativity in a pure sense. In addition, although the design professions certainly appear creative, this may be a simple consequence of the design process being iterative and developmental–leading to many poor ideas being dropped, and better ideas validated. This would be well aligned with the argument we have presented thus far. In any case, we direct the interested reader to Plato's *Protagoras*, for a Socratic exposition of this debate.

We now move in a more practical direction, attempting to find some empirical grounding and suggestive evidence for the statements made so far. Primarily, we will present two empirical investigations of the perception of creativity from both a human and machine perspective.

[3] The veracity of this statement can be questioned from a variety of angles and philosophical schools of thought, none of which would be helpful to those interested in claiming the existence of a concrete concept: *creativity*

The intention is to motivate two claims:

*Firstly*, that the common concept of creativity is not as it is often characterized in the literature, that it borrows from and overlaps heavily with other difficult concepts, and that much of the time it is used in ways that have been unexplored by the Computational Creativity field (for good reason). Our claim is that the ambiguity in usage is a consequence of the post-hoc nature of the assignment of *creativity* as a label.

*Secondly*, that even with fairly simple reasoning infrastructure, machines and humans generate linguistic artefacts of some level of creativity at a non-negligible background rate, independent of the process used. This claim directly supports our view that the same processes that sometimes generate creative artefacts, often (if not usually) generate artefacts not deemed to be creative, and are therefore not relevant to defining creativity.

# 3 Human Expectations of Creativity

In order to ground our discussion of creative acts, we wish to explore the occurrence of such acts in the wild, to better understand the general expectations, related concepts, and domain specificity of the concept's usage. Each day, English speakers the world over query the web billions of times, and some portion of these queries naturally encodes the relation between creativity and other concepts. We will leverage these publicly available data sets to infer the types of things associated with creativity in the minds of everyday users of language, and expect to receive an unbiased picture of modern usage.

## 3.1 Data

We make use of two different sources of web data in order to explore the trends in co-occurrence patterns and concept formation over time. Specifically, we use the Google Trends API (**?**) and *The Google Web 1T 5-Gram Database* from the Corpus Linguistics Group at FAU Erlangen-Nrnberg. [4]

**Google Trends**

> **What is Trends data?**
>
> "Trends data is an unbiased sample of our Google search data. Its anonymized (no one is personally identified), categorized (determining the topic for a search query) and aggregated (grouped together). This allows us to measure interest in a particular topic across search, from around the globe, right down to city-level geography." (Simon Rogers, Data Editor at Google)

The Google Trends API is used to query the frequency of occurrence of concepts across a large cross-section of web searches, allowing the exploration of interest patterns in a particular topic. Topics represent keywords and groups of keywords that are taxonomically disambiguated into classes (e.g. location, organization, person, field of study).

**The Google Web 1T 5-Gram Database** Google's 5-gram web language corpus from 2006 provides counts over n-grams (from 1 to 5) of keywords from the public web. The files contain over 1 Trillion tokens, over 95 Billion sentences and from 13 Million unique uni-grams to 1.1 Billion unique 5-grams.[5] The corpus provides a way to observe keyword co-occurences across one of the largest cross-domain samples of the English language ever made available. We make use of a web API for exploring the web corpus that has been provided by the Corpus Linguistics Group at FAU Erlangen-Nrnberg. The API allows researchers to choose different statistical measures of importance to rank co-occurrence frequencies among terms of interest.

## 3.2 Methodology

**Web Trends around Creativity** In order to explore and compare web search trends relating to creativity and creative processes, we leverage the powerful topic structure available to define a set of structured search terms. "Creative (process)", "Creative (thing)", "Creativity(concept)" are combined to provide a baseline. We then search for terms that represent similar concepts to gauge their level of exchangeability through co-occurrence. The additional terms relate to "Imagination", "Innovation", "Intelligence", and "Novelty". We compare trends for these terms between the beginning of 2008 and the end of 2015. For each concept, we visualize the relative frequency (see Figure 1). For every pair of concepts, we show the correlation between the relevant time series (see Figure 3, Results).

Figure 1: The relative frequency of searches across the 5 concepts of interest: "Creativity", "Imagination", "Innovation", "Intelligence", and "Novelty".

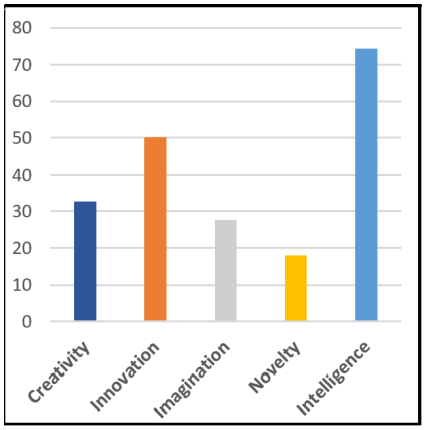


**Describing Creativity** The second of our explorations into the usage of "creativity" in the wild revolves around the large corpus of web text provided by Google, as described above. We want to understand better the sorts of concepts that people associate with creativity, on average. The literature on Computational Creativity is heavily tilted towards the arts, with painting, musical composition, and literature relatively over-represented. In order to make best use

[4] Accessible at http://corpora.linguistik.uni-erlangen.de/demos/cgi-bin/Web1T5/ as of March 3rd, 2017.

[5] The Google Web 1T 5-gram Database is available from the Linguistic Data Consortium: https://catalog.ldc.upenn.edu/LDC2006T13 as of March 3rd, 2017.

of the phrasal statistics available in the corpus, we construct a set of search terms that seem to capture the essence of how people might talk about creativity and creative artefacts. Our constructed queries include "creativity", "creative", "creatively", among others. We aggregate the returned co-occurrence frequencies, and rank results following the modified dice coefficient (**?**). Finally, we strip out obvious noise terms such as "Adobe Creative Cloud" and other proper names. We are left with a set of the most popular terms co-occurring with text about creativity on the public web.

### 3.3 Results

Our analysis of the trends (Figure 2) shows that "Creativity" is very highly correlated (Figure 3) with "Imagination", "Innovation", and "Intelligence", and that "Novelty" (or "Newness" and "Surprise", among other concepts) seems not to be correlated with these terms.

Figure 2: The trends of searches across the 5 concepts of interest: "Creativity", "Imagination", "Innovation", "Intelligence", and "Novelty".

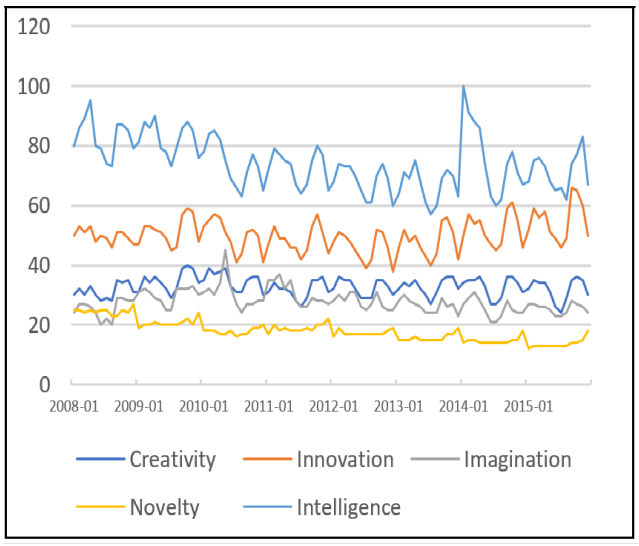


It is difficult to make strong claims about the absence of correlation in this particular setting, so we focus on the positive associations. In particular, we argue that since these concepts trend together, they likely share many of the same properties. One could certainly imagine similar terminological debates around the foundations of "Computational Imagination" or "Computational Innovation", were these to become areas of interest to the research community. The results of this particular experiment are merely suggestive, and should not be over-interpreted.

Our analysis of web text data relating to creativity offers interesting insight into the kinds of concepts that people are most interested in when discussing creativity. The most informative terms, ordered by the Modified Dice Coefficient, can be seen in Table 1. The first thing to note is the overlap in terminology from our experiment on search trends among

Figure 3: The correlations between 5 concepts of interest: "Creativity", "Imagination", "Innovation", "Intelligence", and "Novelty".

| | Innovation | Imagination | Intelligence | Novelty |
|---|---|---|---|---|
| **Creativity** | 0.70 | 0.58 | 0.58 | 0.03 |
| **Innovation** | | 0.25 | 0.59 | -0.12 |
| **Imagination** | | | 0.38 | 0.08 |
| **Intelligence** | | | | 0.43 |

terms. "Innovation and "imagination" are the top two concepts related to creativity, but further down we see "artistic", "originality", "inspiration", "talent". It certainly seems that these terms correlate with what we have seen in the literature (see for example (**?**), (**?**), (**?**)). It is reassuring that many of the attributes that the various evaluation frameworks look for, are actually everyday relevant concepts for those that speak about creativity. The question of how we measure the contribution of each of these factors effectively still remains, but we leave this discussion open for future elaboration.

Table 1: Collocated keywords with the concept of "Creativity", by Modified Dice Co-efficient

| Collocate | Modified Dice | Frequency | Expected |
|---|---|---|---|
| innovation | 26.9433 | 349035 | 1673.04 |
| imagination | 15.4379 | 113050 | 642.4 |
| originality | 10.4318 | 46434 | 116.84 |
| unleash | 7.7493 | 32807 | 77.01 |
| ideaflow | 6.8198 | 26233 | 6.19 |
| artistic | 6.3052 | 48396 | 706.94 |
| stimulate | 5.2166 | 28843 | 314.11 |
| ingenuity | 4.8916 | 20775 | 79.49 |
| bipolar | 4.8203 | 24871 | 246.49 |
| inspire | 3.7247 | 20865 | 327.41 |
| latitude | 3.6512 | 27493 | 680.29 |
| flexibility | 3.4415 | 34976 | 1162.15 |
| encourages | 3.1393 | 20010 | 468.76 |
| talent | 3.1184 | 31711 | 1163.24 |
| enthusiasm | 2.9654 | 17810 | 401.37 |

## 4 Reasoning Creatively

In this section we continue our empirical exploration of creativity in the wild, providing supporting evidence from a very different source, namely one of the largest manually curated ontology projects, Cyc (**?**). Our aim is to measure the background occurrence of interesting "creative" assertions from a knowledge base that was not designed with any creative output in mind. Once again, this is merely supposed

to add suggestive evidence to our arguments so far, and by no means should it be seen as proof positive for any of the arguments put forward.

### 4.1 Data

Our set of assertions from Cyc consists of 2,200 sentences generated via the Cyc infrastructure. For each sentence, a random starting node in the ontology is chosen, and an assertion is extracted at random rooted at that node. Since the ontology is the product of individual concepts and rules added manually, along with inference over the graph, many of the assertions that are retrieved follow links that were automatically generated due to logical constraints and reasoning by the Cyc engine. Fortunately, Cyc has a Natural Language Generation model (Logic2NL (**?**)) that provides human-readable sentences, rather than abstract logical compositions.

Some sample sentences that were generated by this procedure:

- A feeling of anger is likely to be accompanied by a feeling of hatred.
- Every paper contains some weather report.
- In dropping, the ”falling” step ends later than the ”release of support” step.

### 4.2 Methodology

For each of the generated 2,200 sentences of natural language, we reviewed them manually according to the novelty, surprise, and value (each on a scale of 1-5) presented by the ideas and concepts evoked. This evaluation loosely follows the criteria suggested by Boden in her seminal work on the subject of Computational Creativity. We chose a cut-off of 12 points to mark sentences as “creative”. This choice was based on a ranked review of the quality of output, and can be seen as fairly arbitrary, although sufficient for the point we wish to make.

### 4.3 Results

We present results of our exploration, again, as providing ad-hoc suggestive evidence for the arguments that we have presented. We find that the vast majority of the sentences reviewed have little to no trace of a creative foundation. However, in just over 0.5% of instances we do find some sort of noticeable creativity, which does not seem so far away from the background rate of creative artefact generation in our every-day life. Of course, we would need a much more thorough and independent evaluation in order to make any more conclusive statements on the matter. Some of the more noticeable structures include:

- Every resurrection destroys some dead person.
- Haircuts generally affect mammal hair.
- Holding one’s breath requires the use of one or more lungs.
- Many actors are charismatic people.
- Middle-class American adults accomplish brushing of teeth easily.
- Monetary values are only monetary values.
- Most auto bodies are taller than most flags.
- No euthanasia is suicide.
- No time lasts longer than forever.
- Sauron The Enemy heads the government of Mordor.

One will notice immediately that novelty and surprise are playing a larger role in the above examples than value, but as an illustrative point some mixture of processes over the Cyc ontology generates somewhat creative artefacts at some rate above zero. See figure 4 for the distribution of scores over the sentences in our corpus.

Figure 4: The frequency of occurrence of each evaluation score for 2,200 Cyc logic2NL generated sentences.

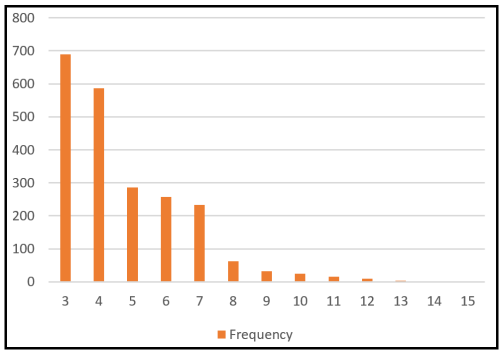


## 5 Conclusion

In this paper we have proposed a new way in which creative acts might be understood, whether produced by humans, machines, or other intelligent agents. At its core is the argument that creativity is an ex-post phenomenon, and as such the search for the roots of creativity might be a red herring. Instead, we should look for value creation and novelty that emerges from mixtures of all available cognitive processes. None of these will fully characterize what we judge to be creative after the fact, but by broadening our understanding of the driving forces behind creative artefacts we can perhaps start to better understand and model what is going on. Relaxing the assumption that creative processes somehow exist as identifiable entities in the brain needn’t change the frameworks for evaluation that are under active development (although it may offer interesting new avenues to explore), nor will it undermine the work that has happened to date in the field. However, by understanding this core distinction we can hopefully contribute even more to our scientific understanding of the cognitive processes involved, and enhance creative output across the board. In addition to a philosophical and theoretical discussion of the nature of creativity, we have presented some suggestive empirical evidence for the concept of creativity in the wild, and for the natural occurrence of creative artefacts even from a stylized setting such as the Cyc ontology. We encourage future work to elaborate these concepts to augment existing evaluation frameworks, and to find new avenues to analyze and expand our understanding of creativity across disciplines.

## References


Ariza, C. 2009. The interrogator as critic: The turing test and the evaluation of generative music systems. *Computer Music Journal* 33(2):48–70.

Biles, J. A. 1994. Genjam: A genetic algorithm for generating jazz solos. In *ICMC*, volume 94, 131–137.

Bluck, R. S. 1961. Plato's" meno". *Phronesis* 94–101.

Boden, M. A. 1998. Creativity and artificial intelligence. *Artificial Intelligence* 103(1):347–356.

Bourigault, D.; Jacquemin, C.; and L'Homme, M.-C. 2001. *Recent advances in computational terminology*, volume 2. John Benjamins Publishing.

Bradeko, L.; Starc, J.; Mladenic, D.; Grobelnik, M.; and Witbrock, M. 2016. Curious cat conversational crowd based and context aware knowledge acquisition chat bot. In *2016 IEEE 8th International Conference on Intelligent Systems (IS)*, 239–252.

Cardoso, A.; Veale, T.; and Wiggins, G. A. 2009. Converging on the divergent: The history (and future) of the international joint workshops in computational creativity. *AI Magazine* 30(3):15.

Colton, S., and Wiggins, G. A. 2012. Computational creativity: The final frontier? In *Proceedings of the 20th European conference on artificial intelligence*, 21–26. IOS Press.

Colton, S.; de Mántaras, R. L.; Stock, O.; et al. 2009. Computational creativity: Coming of age. *AI Magazine* 30(3):11–14.

Colton, S.; Bundy, A.; and Walsh, T. 2000. On the notion of interestingness in automated mathematical discovery. *International Journal of Human-Computer Studies* 53(3):351–375.

Colton, S.; Charnley, J.; and Pease, A. 2011. Computational creativity theory: The face and idea descriptive models. In *Proceedings of the Second International Conference on Computational Creativity*, 90–95.

Colton, S. 2008. Creativity versus the perception of creativity in computational systems. In *AAAI spring symposium: creative intelligent systems*, volume 8.

Colton, S. 2012. The painting fool: Stories from building an automated painter. In *Computers and creativity*. Springer. 3–38.

Cooper, J. M., and Hutchinson, D. 1997. Plato complete works.

David, A.; Marianne, B.; and Sendhil, M. 2011. Do judges vary in their treatment of race? *Journal of Legal Studies*.

Fedyk, A. 2015. Overcoming overconfidence: Teamwork and self-control.

Google. 2010. *How does Google Trends Work*.

Jordanous, A. 2012. A standardised procedure for evaluating creative systems: Computational creativity evaluation based on what it is to be creative. *Cognitive Computation* 4(3):246–279.

Lenat, D. B. 1995. Cyc: A large-scale investment in knowledge infrastructure. *Communications of the ACM* 38(11):33–38.

Lenney, E.; Mitchell, L.; and Browning, C. 1983. The effect of clear evaluation criteria on sex bias in judgments of performance. *Psychology of Women Quarterly* 7(4):313–328.

McMullin, E. 2010. Creation ex nihilo: Early history.

Neuman, Y., and Nave, O. 2008. A mathematical theory of sign-mediated concept formation. *Applied Mathematics and Computation* 201(1):72–81.

Newell, A.; Shaw, J. C.; and Simon, H. A. 1959. *The processes of creative thinking*. Rand Corporation Santa Monica, CA.

Pease, A., and Colton, S. 2011a. Computational creativity theory: Inspirations behind the face and the idea models. In *Proceedings of the Second International Conference on Computational Creativity*. Citeseer.

Pease, A., and Colton, S. 2011b. On impact and evaluation in computational creativity: A discussion of the turing test and an alternative proposal. In *Proceedings of the AISB symposium on AI and Philosophy*.

Plucker, J. A., and Beghetto, R. A. 2004. Why creativity is domain general, why it looks domain specific, and why the distinction does not matter.

Plucker, J. A. 1998. Beware of simple conclusions: The case for content generality of creativity. *Creativity Research Journal* 11(2):179–182.

Ritchie, G. 2007. Some empirical criteria for attributing creativity to a computer program. *Minds and Machines* 17(1):67–99.

Sreenivasan, G. 2002. Errors about errors: Virtue theory and trait attribution. *Mind* 111(441):47–68.

Tomašic, P.; Znidaršic, M.; and Papa, G. 2014. Implementation of a slogan generator. In *Proceedings of 5th International Conference on Computational Creativity, Ljubljana, Slovenia*, volume 301, 340–343.

Valiant, L. G. 1984. A theory of the learnable. *Communications of the ACM* 27(11):1134–1142.

Ventura, D. A. 2008. A reductio ad absurdum experiment in sufficiency for evaluating (computational) creative systems.

Wiggins, G. A. 2006. Searching for computational creativity. *New Generation Computing* 24(3):209–222.