# Is Computational Creativity Domain General?

**Róisín Loughran** and **Michael O'Neill**
Natural Computing Research and Applications Group (NCRA)
University College Dublin
Ireland
roisin.loughran@ucd.ie

## Abstract

The question as to whether creativity is domain general or domain specific is one of the most enduring and controversial topics in the field. Yet the importance or relevance of the chosen application domain has not been considered in the related field of computational creativity (CC). A recent study at ICCC demonstrated that the range of applications considered in the study of CC has been diverse with more novel topics being considered as the field progresses (Loughran and O'Neill 2017a). As the field grows, we propose that we need to consider the relevance of the application domain and any potential role or effect the choice of the domain may have on the outcome of the designed system. In this paper, we review what it means for CC to be domain-general. We consider the domain-dependence of creativity in human studies and what implications, if any, may arise from the choice of application domain in CC studies. We conclude that this is a multi-faceted question and that a simple yes or no answer may not be possible to acquire or sensible to suggest.

## Introduction

Computational systems that attempt to simulate, portray or genuinely exhibit creativity typically do so in a given application domain. A generative creative system can create some novel artefact such as a melody, piece of artwork or joke to be evaluated in order to ascertain the level of creativity exhibited by the system. Despite the diversity in possible domains there have been trends in the topics considered at ICCC with systems based around music and Natural Language Processing (NLP) remaining popular over the years (Loughran and O'Neill 2017a). Even so, the same study showed that there is a steady increase in systems based in novel application domains not considered before at ICCC. In another study that year on 'How to Build a CC System' it is proposed that the first step in the process is: choose a domain *D* (Ventura 2017). If this model is followed then all subsequent steps in the building of the system are dependent on this first step. If the application domain is such a fundamental choice and yet is so diverse — and getting increasingly so — between studies, it raises the question as to whether or not the chosen application domain has an effect on the potential creativity that could be displayed by the system.

The question, as posed in our title, requires clarification. The term Computational Creativity (CC) is defined as (Colton, Wiggins, and others 2012):

> 'The philosophy, science and engineering of computational systems which, by taking on specific responsibilities, exhibit behaviours that unbiased observers would deem to be creative.'

Thus CC refers to a field of study. Asking whether or not a field is domain-general is a much broader and more complex question than asking if a computational system is domain-general. In this paper we will mostly consider the latter question: Is the creativity demonstrated by or aspired to in a given computational system dependent on the application domain within which the system is placed? Even this more specific question raises many issues to consider: Is creativity itself domain-general? Is the creativity of a system dependent on the choice of domain by the programmer? If these questions were found to be true it could raise the issue as to whether or not the creative capacity of the system is dependent on the creativity of the programmer. Furthermore, we may need to consider if artefacts produced in certain domains *appear* more creative than others and if so does this make such domains more suitable for study than others? Do impressive results in traditionally 'creative' domains such as music or art give the impression of higher creative levels than those in more simple domains? A focus on impressive results without considering the underlying cognition of the creativity that produced said results can lead to a misunderstanding between weak CC — that which merely simulates human creativity. and strong CC — systems that exhibit autonomous creativity (al Rifaie and Bishop 2015). There is a place in the field for systems that focus on both weak and strong CC, but it is vitally important to be clear as to which is under consideration.

This paper examines the question of domain generality of systems developed within the field of CC. We review the domain-generality of creativity in human studies and consider what differences may lie in studying creativity from a computational standpoint. We review the types of creativity proposed in the literature and consider how they are approached by CC studies in various application domains.

## Human Creativity Studies

The current accepted definition of CC, given above, is in terms of creativity itself. The reason for this circularity lies in the fact that creativity remans to be such a difficult concept to define. It has been argued that whether creativity is domain dependent or domain general is one of the most controversial issues in (human) creativity research (Plucker 2004). For this reason, we consider general creativity — the manner in which it has been defined and evaluated and the question as to whether it is domain specific or domain general, before considering its relationship to computation.

### Definition of Creativity

The standard definition of creativity is succinct: 'Creativity requires both originality and effectiveness' (Runco and Jaeger 2012). But unfortunately, this is far from the only definition that is used. One of the main problems in the scientific study of creativity is that there have been so many different definitions proposed over the years. It has been stated that there exist over a hundred definitions for creativity within the relative literature (Meusburger 2009). However, considering the subject of creativity has been studied in many subject fields including philosophy, psychology, education, sociology plus all application and technical fields, it is likely that the number is significantly higher than this. Even different dictionary definitions of the word contain discrepancies. The Oxford English dictionary currently defines creativity as 'The use of imagination or original ideas to create something; inventiveness'[1], whereas the Cambridge English dictionary define it as 'The ability to produce original and unusual ideas, or to make something new or imaginative'[2]. In a study on the history of creativity for AI research it has been stated that 'Creativity needs creativity to explain itself'(Still and d'Inverno 2016). Hence even in general creativity, as in the definition of CC, the term is self-referential.

Despite the number of definitions that have been proposed, the common elements that have been present in any accepted definition are based in novelty (or originality) and value (or effectiveness). While the roots of the study of creativity began to emerge in the 1930s-1950s (Runco and Jaeger 2012) the first definition to include these elements was given by Stein: 'The creative work is a novel work that is accepted as tenable or useful or satisfying by a group in some point in time' (Stein 1953). More recent works that attempt to define or evaluate creativity do so by focussing on the two aspects of novelty and value (Ritchie 2001; Boden 2004). Value can be attributed to a concrete artefact or to a more abstract concept theory or interpretation. Novelty can refer to ideas that are new to the individual, know as Psychological (P) Creative, or those that are novel to the world — Historical (H) Creative. By this reasoning H-Creativity is a special case of P-Creativity (Boden 2004).

Csikszentmihalyi similarly separated the idea of creativity to the individual (P-creative) and to a culture (H-creative) in proposing the idea of 'Big C' Creativity (Csikszentmihalyi 2013). He posits that such a version of creativity cannot be experienced purely by an individual, but must have an influence on some aspect of culture. He considers that this Creativity can only be found in the interrelations between three parts of a system:

- Domain: a set of symbolic rules and procedures;
- Field: those who decide what is novel within the domain;
- Person: those who undertake the creative act or idea.

Only within the interplay of these parts can Creativity be exhibited. Thus he sees Creativity as an act that changes an existing domain or transforms a given domain into another domain. Boden has also proposed three distinct types of creativity: combinational, explorational and transformational (Boden 1998). She proposed that transformational creativity, which transforms the space within which one is searching, offers the most opportunity for discovery. Hence, transformational H-creativity is akin to Csikszentmihalyi's Creativity ('Big-C'), but due to the transformation of domains and requirement of historical domain knowledge, such creativity will be difficult to evaluate.

### Evaluating Creativity

The evaluation or measurement of creativity is often reported in relation to assessing creative ability, often in children, through a variety of 'paper-and-pencil' tests (Cropley 2000). Although a large variety of such tests exist, Cropley proposed concentrating on those developed during the 'modern creativity era' as first introduced by Guilford (Guilford 1950). He reviewed and organised such tests to reveal four dimensions relating to elements of creativity: product, process, motivation and personality/ability. From an analysis of a large number of tests he observed that creativity tests did not have as high a predictive value of success as traditional IQ tests. Cropley also stated a preference for models that encompass both thinking and personality such as the Test of Creative Thinking: Divergent Production (TCP-DP) (Urban and Jellen 2005) over those centred purely on divergent thinking (DT). DT tests, however have been used extensively as an indicator for creative potential (Runco and Acar 2012). While DT is acknowledged as different to creative thinking, psychometric tests have suggested that these tests can provide predictors for the potential of creative thinking. While Guilford tied divergent production to creative potential, Runco explicitly states that DT is not synonymous with creativity, but that tests based on this theory can indicate a potential for creative problem solving.

Although there have been many tests for identifying creative thinking, no one test has been agreed on as the most general or best. Furthermore, these DT style tests are not discussed in terms of application domain but rather in terms of general creative thinking potential. Nevertheless, the importance of the given application domain on creativity and creative ability has been discussed at length in the field.

### Domain Specificity of Creativity

Whether creativity is domain specific or domain general has been a hotly debated topic in creativity research for a number of years. In general creativity, this amounts

[1]https://en.oxforddictionaries.com/definition/creativity

[2]https://dictionary.cambridge.org/dictionary/english/creativity

to asking whether a person, process or product is considered creative only within one domain or across multiple domains and furthermore whether training in one domain can increase ability in another creative domain. There are strong opinions on both sides with some arguing for domain generality e.g. (Runco 1987) with others maintaining the domain specificity of creativity e.g. (Baer 1998; 2010). While the discussion continues, the dominant perspective appears to be leaning towards domain specificity. Even if a consensus cannot be reached, Baer still argues that it is better to assume specificity over generality (Baer 1998). He proposes that in assuming specificity, nothing is lost by training a subject in their specific domain even if it turns out that that creativity is domain general. Conversly if domain generality is assumed but domain specificity is the reality, much effort could be wasted in teaching and learning based on general domain creative-thinking tasks. Plucker and Beghetto further argue that too much focus on either position may hinder creativity. They conceptualise a model indicating a focus on generality can lead to superficiality whereby one may never fully engage in the creative task, while being too specific could lead to a fixedness whereby the goal is never satisfied (Plucker and Beghetto 2004). They conclude that creativity is most likely domain general, but that it can appear domain specific and that ultimately it is not beneficial to dwell on the concept

Whether or not creativity is domain-specific as defined above, it is evident that creative domains exist and are important to us. As Csikszentmihalyi has stated: 'The existence of domains is probably the best evidence of human creativity' (Csikszentmihalyi 2013). In contrast to many other living organisms, we as humans can choose from a number of responses when presented with a given stimulus; a flower will turn towards the light of the sun, but when we encounter such a flower we may choose to focus on visual art and paint it, we may be inspired to write a poem about it, or we may pass it by with barely a glance while focussing on other matters. An initial choice can lead to a decision to work in more specified sub-domains; if we want to make a visual artistic rendition of the flower do we paint, use charcoal or printing? Our individual responses to the world may be different, there is no standard domain in which me must work. We choose which specialist domain to focus on and the more we work on a domain, the more specialised we become. Animals other than humans are capable of actions other than the predefined responses of the above flower, however. Much creative and playful interactions have been observed in the social and exploratory behaviours or animals (O'Hara and Auersperg 2017). While these may not result in behaviour considered as traditionally creative as that of humans, creativity in nature has been noted to be heavily influenced by similar such social interactions (Saunders and Bown 2015).

## The Computational Comparison

A person must first find their passion to consider working in any given domain. Only through years of study can they then choose their own sub-domain and find their footing to build on what has been done in order to make any creative contribution. Can an autonomous creative system make such decisions before it undertakes a work? The decision as to what domain to work in is invariably decided on *a priori* by the programmer: we plan to write music generation programs or art generation programs. Is it possible to write a creative program without first specifying the domain —- and by specifying the domain have we paradoxically removed the potential for the system to intentionally display creativity?

While we may accept that the lack of clarity of a definition for creativity leads to the circularity in the CC definition, it does mean that this CC definition is reliant on an ill-defined concept. Computational measures, by definition, are based on the idea of enumeration — an exact process. Thus what we are trying to achieve in CC is to enumerate that which we cannot define. To some, this may be the best argument against applying computation to creativity; if we cannot define the concept and therefore cannot measure the concept, how can we expect computers to enumerate, generate or imitate the concept? Such an argument is overly defeatist however. Creativity is not a divine ability afforded to only a lucky talented few, it is merely a feature of human intelligence (Boden 1998). Much of the original theory on creativity was based on determining if it was distinct from intelligence (Runco 2014). AI systems are becoming increasingly important in our modern day society. With the push towards a general AI, it is imperative we recognise creativity as an important aspect of intelligence and do not merely dismiss the idea of computing it as too much of a challenge.

We have established that creativity requires novelty and value, but there is one extra important aspect that must be considered, particularly when considering autonomous creative systems: the aspect of intent. The requirement of intent is rarely addressed in humn creativity; presumably a person who is creating an artefact is doing so intentionally. Many recent CC studies however have stated that for creativity to be present, an agent must exhibit novelty, value and intentionality e.g. (Ventura 2017). Novelty and value are important but once we consider creation by a computative system, this idea of intent or ownership becomes equally important. If a system generates a joke — does it need to have intended to do so in order to display creativity? What would such an intent even mean — was the system attempting to make us laugh? This level of invoking an emotional response from a computative system is not possible yet, but for creativity to be displayed the system should provide evidence on some level of intending to produce its output. This issue of intent has raised discussion in recent years. (Guckelsberger, Salge, and Colton 2017) considered a non-anthropocentric model by adopting an enactive AI framework finding that CC systems that focussed on human creativity typically cannot provide a reason for intent as they lack intrinsic goal-ownership.

**CC Evaluation** A lack of evaluation has been noted many times throughout the development of the field of CC (Boden 1998; Cardoso, Veale, and Wiggins 2009; Jordanous 2011). This lack of evaluation could result in undermining any scientific progress of the field, resulting in a stricter focus towards evaluation within papers and the development of a number of frameworks according to which CC systems should be evaluated. Over the past decade a number of evaluation frameworks have been proposed including a set of empirical criteria (Ritchie 2007), the Creative Tripod (Colton 2008), numerous Turing-style tests (Ariza 2009) and the Standardised Procedure for Evaluating Creative Systems (SPECS) (Jordanous 2012). Ventura gave a series of milestones that a system must surpass in order to be in the 'realm' of creativity (Ventura 2016). He posits that systems should exhibit more than just randomisation, plagiarisation, memorisation, generalisation, filtration and inception to avoid the 'mere generation' trap, but concludes that the location of actual creativity is still somewhat out of reach from the computative community.

One test for creativity that is focussed on autonomy is the Lovelace Test (LT) (Bringsjord, Bello, and Ferrucci 2003). This involves an artificial agent $A$, its output $o$ and its human architect $H$. Simply put, the test is passed if $H$ cannot explain how $A$ produced $o$. It is important to note that this test is not about predicting results but explaining how they came about. Any programmer who can explain their written code can in theory explain how the their agent produced its output. The only situation in which this LT can be passed is one whereby the programmer cannot explain what they have written. In reality, if a system consists of multiple modules that interact to produce a final output, one could argue that not one individual programmer could explain the whole system, but theoretically their combined knowledge should be able to, and as such this scenario is merely an increase in complexity rather than a true solution to the test. As it stands, the LT does not appear to be passable, regardless of what application domain it is applied to.

By necessity, evaluation of the output or artefacts produced by any creative system is domain-dependent; if evaluation is performed on an artefact, then the domain of said artefact is imperative to the judgment. The field of CC is based on systems rather than artefacts however. While there still remains a tendency to evaluate purely on the final produced artefacts, such assumptions can lead to limitations within evaluations and hence in the growth of the field in general (Loughran and O'Neill 2017b). It would be beneficial instead to be able to make an evaluation of the running system, where this is more appropriate. Such a judgement should not be dependent on the application domain, but could in theory be dependent on the type of algorithm used in the development of the system. The LT is an example of such a metric as it is based on how the artefact is produced — not merely what is produced. Similarly the SPECS method could incorporate such a method as the definition of creativity and the standards used to evaluate it are both specified as part of the method. Hence, the question as to whether or not the evaluation of creativity is dependent on application domain is more dependent on the definition of the evaluation rather than the definition of creativity.

## Application Domains in a CC System

The design and implementation of any computational system involves a number of steps. Such a design may start with a choice of algorithm, analysis of data, a focus on inner workings or any high level consideration of the system. A recent proposal for the first step in creating a CC system is in choosing the application domain (Ventura 2017). Using this framework, once that is chosen, all representations both at a genotypic (internal) and phenotypic (external) level can be determined. Therefore, in the development or planning of a system, the application cannot be disregarded; what the system creates is quite often the initial purpose of making the system for many people. People naturally tend to work with systems that operate within a field that they themselves have domain knowledge, and they often stay within one application domain; musicians tend to make musical systems and artists work with visual systems. This results from personal interest but should not have any direct bearing on the potential creativity within the system if, that is, we can consider creativity to be domain general. This does ensure, however, that the initial intent in creating the system lies solely with the programmer rather than the system. Even if we adopt this framework and assume that it is acceptable to choose the application domain as the first step of building a system, the system should subsequently display some level of intentionality in creating its output.

### Concept

The analysis of application domains considered throughout ICCC listed 'Concept' as one singular domain (Loughran and O'Neill 2017a). In essence, however, many CC systems could be reduced to Concept, regardless of the given application domain. A story-telling system, such as Mexica (y Pérez 2015), may make use of NLP but this involves more than mere syntactic analysis of words. Yes, the text must make grammatical sense, but the meaning behind the words and the arch that the story follows — interaction between characters, emergence of themes, building and subsequent release of tensions — are what will engage the reader. These higher level features are less about the domain (NLP) as much as they are about the concept behind them. Similarly, in systems that deal with other domains such as music, games etc. the creativity within the system could be contained in the underlying concept, but this concept is wrapped up in layers of increasingly complicated representation. The application domain is the public front to such systems, and it can catch someone's eye, but is not necessarily where the creativity lies. While systems that generate very impressive aesthetic outputs may rely heavily on domain knowledge, others can be reduced to an underlying concept for more insightful understanding. The importance of the underlying concept within any CC system can be related to the given application domain.

Another particular topic of of interest often considered in CC studies is that of analogy. Studies in analogy generally use the written word and so could be broadly put into the domain of NLP. But again it is not the semantic understanding of the words that is under consideration, rather the interplay of the underlying meaning behind those words between two specific concepts. As such all studies in analogy consider two conceptual domains and the transition between them. Similarly, studies that consider conceptual blending (Fauconnier and Turner 2003) draw from more than one original domain. Conceptual blending integrates two or more mental spaces in order to create something in a new blended space.

CC studies that consider concept, analogy or blending do not always have distinct boundaries in regards to application domain. In such cases, the domain may not be simple to define and hence can appear to be more general than those studies focussed on specific aesthetic artefacts.

### CC Systems in Multiple Domains

Throughout the development of CC there have been a number of systems proposed and studies described that deal with more that one application domain. Some studies may use a specific domain to illustrate a point that is in fact domain-independent such as the theoretical model of creative inspiration proposed in (Wiggins 2012). A number of studies propose the examination of a new general principle and subsequently illustrate the point using a variety of examples from different domains such as considering intrinsic measures of fitness (Cook and Colton 2015), antagonistic and supportive behaviours (Guckelsberger et al. 2016) or multiple facets in preference functions (Bhattacharjya 2016). These studies typically propose a new method of considering or measuring CC, either formally or empirically, and then illustrate these concepts in concrete examples. Such studies may consider multiple domains, but are still limited to those under consideration, rather than generalising across all domains.

### Adaptive Systems

The above discussion focusses on generative systems. If a system is created with the purpose of generating 'something' then boundaries must be implemented within which this something will be generated. The limits of these boundaries constitute the domain and thus this domain must be specified by the programmer at the beginning of development. If instead, a system is developed whereby the ability or creativity of the system lies in the modification of an existing artefact or behaviour this may be considered *adaptive creativity* (Bown 2012). Adaptive creativity can be exhibited by flexible systems that adapt not just to their internal composition but respond to external perturbations within dynamic environments. As the domain or function within the system is already established, the explanation, evaluation or fitness of the adaptation should be possible; an adaptive system should be able to evaluate any changes to stay within viability boundaries (Guckelsberger, Salge, and Colton 2017).

If a system displays adaptive creativity, this may be exhibited within a certain domain but such creativity may not be *dependent* on this domain. The creativity emerges through the traversal of the behaviour or artefact being adapted within specified boundaries or limitations. In such a case it would be very difficult to argue that the creativity exhibited is not domain-general; if this creativity is exhibited within one given domain without being developed because of the domain, arguably this could be transferred to another domain to exhibit similar results.

### Big and Little Domains

We have established that creativity is not limited to extremely impressive or artistic feats. Creativity is a general aspect of intelligence, as described in Boden's 'P' and Csikszentmihalyi's 'small c' creativity. Yet, often by focussing on systems in specific domains that are associated with talent or passion, such as music, the purpose of the system can subconsciously focus on big-C Creative results. When someone hears that a computer system has written music, the expectations of the quality of the produced is automatically high — why would a system that creates mediocre music be of interest? It becomes difficult to focus on or even appreciate small-c creative achievements when one is effectively working within a big-C domain. In such a big-C domain, much *a priori* domain knowledge is required for a system to generate anything of value, hence the domain becomes important and it can be difficult to disentangle the creativity from the application domain.

If instead, systems were developed in less traditionally creative, aesthetic or artistic domains such as logic or problem solving, it is possible that small-c creativity would be more accessible to identify or study. Besold posited that CC systems belong to one of two families of 'artistic creativity' or 'problem-solving creativity' noting that the latter has stronger links to transformational creativity and is closer to strong creativity than more aesthetic based studies (Besold 2016). He acknowledged the lack of research focussed on computational cognitive systems with general creative capabilities that are mostly independent of a concrete domain, as similarly noted in (Loughran and O'Neill 2017a). If we wish to cut through the domain dependence of CC systems to consider a more general understanding, it is certainly worth considering a stronger focus on the cognitive aspects of computational agents as they undertake problem-solving tasks.

## 'Humanity' in CC

The current CC definition in (Colton, Wiggins, and others 2012), makes no reference to human opinion. Earlier definitions and discussions on the topic have made reference to human ability however. In Marsden's discussion on Intelligence, Music and Artificiality he discusses the 'intention to perform in a human-like fashion' as one of the two major topics of the paper (Marsden 2013). Ritchie justifies alluding to human-creativity when considering more general (non-human or machine) creativity for two reasons: that this is the established usage and secondly, that doing otherwise would risk circularity in claims about the process (Ritchie 2006). The definition of computational creativity offered in 2006 by Wiggins referred to behaviour of systems which would be 'deemed creative if exhibited by humans' (Wiggins 2006). As late as 2012 Jordanous' definition [3] referred to behaviour 'if observed in humans'.

Despite the lack of the term 'human' in the definition, there remains a lingering tendency to consider human opinion when evaluating or discussing CC systems. (Loughran and O'Neill 2016) have argued against a consistent human-comparison when it comes to evaluating generative musical systems. They proposed that evaluations which focus purely on human-based measurements would automatically be subjected to bias and therefore result in limitations in the development of the field, particularly in the area of au-

[3] quoted from the `computationalcreativity.net` website at that time

tonomous creativity. This is in mind with Guckelsberger's non-anthropocentric method of considering intent in CC and Amabile's earlier warnings of considering attributions of creativity purely as those of the individual (Amabile 1995). If the explicit mention of human opinion is not part of what defines CC, we should be mindful not to let a bias based on human opinion to creep back into evaluations and discussions on what it is to be creative.

### Machine vs. Human Capabilities

In considering CC, or any form of AI, we effectively attempt to model some aspects of brain function through computation. Many techniques developed in the field of ML are based on brain function such as connectionism and artificial neural nets. With the dramatic increase in computing power, the capabilities of such systems are likewise increasing. Thus there are two interrelating questions: are we modelling the brain correctly and if so, do we have enough computing power to emulate general human intelligence? While the theory of neural networks has been available for over half a century, it is only in the recent advancement of deep neural nets consisting of millions of neurons, that the computing power of such systems has become clear. These networks still offer a black-box approach to problem-solving however, rendering successfully trained networks difficult to analyse. The training of such networks is furthermore specific to a given task, there does not yet exist an AI of general intelligence. Nevertheless, with the advancement of such technologies it is becoming more feasible that an artificial mind as powerful as our own in on the horizon.

General AI or indeed strong (domain-general) creativity may not be possible until an artificial mind as powerful as that of a human is a reality. Even if the computational power is reached, it is not an absolute that it will model the world in the same manner as we experience it. For the moment, we must be content to work within the current bounded rationality of computer capabilities; we develop the models we can within the current limits. In developing weak creative systems or domain specific tasks we can consider creativity and intelligence from different aspects, thus moving towards a multi-faceted model or understanding of strong general creativity.

## Discussion

From an academic standpoint, should we consider creativity by computers in the same manner in which we consider creativity by humans? In defining what creativity actually means it appears that we must, yet in discussing creativity across domains it appears we cannot. In examining what constitutes creativity, the discussion revolves around notions of novelty and value regardless as to who or what displayed the creativity in question; if the term creativity is to refer to a clearly defined concept, then the entity involved should not matter. Discussing domain generality as it is proposed in human studies does not necessarily translate to the computational realm. Many of the human studies on domain generality discuss increasing the ability of a child in one domain from practice or learning in another domain. A similar transfer of knowledge between computational system from one domain to another could only be possible in the presence of general intelligence. While transformational creativity is a well known concept in CC, and knowledge transfer may be possible in certain circumstances, the transfer of knowledge from one domain to another requires a general level of adaptation or intelligence that is not presently available in an AI.

If a system is created for the purpose of generating an aesthetic artefact, as we have established that many — but not all — are, then the intentionality of working within the given application domain is wholly on the programmer rather than within the system. Furthermore, regardless as to how prestigious or accomplished a human creator is at their craft, they suffer from self-doubt and sometimes crippling self-criticism which can ultimately lead to a difficulty in ever finishing a piece of work (Nebel 1988):

> 'A work, finished or not, produces in many artists an aftershock ('choc en routour') of one type or another. As a result, artists are led to reflect upon the work, to change, modify, and even on occasion to destroy it.'

Such self-criticism is considered to be a built-in feature of the artistic personality, and it can lead to feeling of anguish at the thoughts of finishing a piece of work; to a perfectionist a piece may never be truly finished, rather they merely feel they must stop as they can do no more to it. Constant corrections, and an unattainable goal could be programmed into a computational system, but they cannot feel this self-doubt or anguish at their own self-imposed notion of mediocrity. Is such self-criticism necessary for true creativity or in contrast, are such feelings emphasised by authors describing the plight of the tortured artist?

The artists of the Renaissance period may have worked in many domains, but as the years progressed the art world has become more fragmented, more specialised with more subdomains not just in medium but in style and purpose within to place oneself. As expertise within a given application domain has become more specialised, it has been speculated that the emergence of 'Renaissance era' artist — one that exhibited such creativity in many areas as Da Vinci once did — are becoming less likely (Plucker and Beghetto 2004; Csikszentmihalyi 2013). If we no longer expect even the most talented contemporary artists to display such general creativity across multiple domains, is it fair to expect such achievements from computational systems that are under development? Again such arguments are based on the idea of working in Big-C domains, it is important to remember that such opinions propose much less of an issue in CC systems that focus on small-c logical problems.

At times the discussion around creativity has been framed as either creative ability or creative activity. It is important to consider which of these the term 'creativity' is actually referring to. Creative activity is surely domain-dependent, the activity in question must inherently take place within a given domain. But creative ability may be more ambiguous. The definition of creativity given in (Plucker and Beghetto 2004) stresses the interplay between ability and process:

> 'Creativity is the interplay between ability and pro-

cess by which an individual or group produces an outcome or product that is both novel and useful as defined within some social context.'

While they acknowledge that both aspects have been discussed within numerous definitions, they consider that it is in this interplay that creativity lies; regardless of ability, creativity must be enacted through a process.

It has been noted throughout this paper that a critical difference that arises between the discussion of creativity by machines and by humans is that of intentionality. CC is defined in terms of being creative, therefore if the term *creativity* is to be used in the same manner regardless of who or what is exhibiting it, we propose the concept of intentionality should be explicitly stated in the CC definition:

> 'The philosophy, science and engineering of computational systems which, by taking on specific responsibilities, exhibit behaviours that unbiased observers would deem to be intentionally creative.'

## Conclusion

We started this article by posing the question 'Is CC domain general?'. Asking whether or not human creativity is domain-general cannot necessarily be treated in the same manner as asking whether or not computational creativity is domain-general. When discussing the former, studies tend to be based on teaching people in one domain and ascertaining the effects such teachings would have on their abilities in other creative domains. This process does not make sense in computational systems unless we consider these systems to have enough general intelligence to understand both domains. Such a system would need both general intelligence and strong creativity. What we should consider is whether or not creativity itself is more suited, emergent or likely in one domain over another — as conducted in computational experiments. A comparison in this manner is not always logistically possible, however. Systems that operate in differing domains will have different representations, goals, methods of evaluations, plus further differences in relation to the algorithmic methods applied to the given problem. While it may be within the underlying concept that the creativity exists, this concept is generally wrapped up in layers of representation and computation specific to the given domain.

In examining the discussion of domains in relation to general and computational creativity, we find it difficult to prize creativity and application domain apart. According to (Plucker and Beghetto 2004), domain-dependency may not matter, yet when we consider a specific CC system it appears that it might. We have argued that these systems do not (and arguably can not until they acquire free-will) select the domain in which to work, this is generally chosen by the programmer. If the choice of domain is critically important or if CC is not independent in relation to domain, is this an argument against the idea that computational systems can be autonomously creative? If the application domain is chosen by the programmer we cannot separate the evaluation or impact of the CC system from its domain. This implies that reciprocally, we cannot separate the impact of a CC system from its programmer.

So can we answer the question 'Is CC domain general?'? In reality, a yes or no response to such a question trivialises that which is being asked. If we consider the field of CC, then yes all domains may be investigated but when we consider an individual system, more often than not we place all experimentation and evaluation into a given domain. The closer one looks at a given computational system, the more tied in and restricted to the boundaries of the application domain one becomes. This should not impede the development of any CC research, however. As long as the trend of considering new and diverse applications continues, then the scope of studies within the field and range of knowledge obtained through the field can only expand in the coming years. A strongly creative AI capable of general intelligence would arguably be domain-general, but until such a system exists, it appears that any implemented CC system is limited to the domain its user designed it to work in — which, of course, could be any domain at all.

## Acknowledgments

We would like to thank the reviewers for their detailed and insightful comments. This work is part of the App'Ed (Applications of Evolutionary Design) project funded by Science Foundation Ireland under grant 13/IA/1850.

## References

al Rifaie, M. M., and Bishop, M. 2015. Weak and strong computational creativity. In *Computational creativity research: Towards creative machines*. Springer. 37–49.

Amabile, T. M. 1995. Attributions of creativity: What are the consequences? *Creativity research journal* 8(4):423–426.

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

Baer, J. 1998. The case for domain specificity of creativity. *Creativity research journal* 11(2):173–177.

Baer, J. 2010. Is creativity domain specific. *The Cambridge handbook of creativity* 321–341.

Besold, T. R. 2016. The unnoticed creativity revolutions: Bringing problem-solving back into computational creativity. In *Proceedings of AISB Third International Symposium on Computatational Creativity*.

Bhattacharjya, D. 2016. Preference models for creative artifacts and systems. In *Proceedings of the Seventh International Conference on Computational Creativity*.

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

Boden, M. A. 2004. *The creative mind: Myths and mechanisms*. Psychology Press.

Bown, O. 2012. Generative and adaptive creativity: A unified approach to creativity in nature, humans and machines. In *Computers and creativity*. Springer. 361–381.

Bringsjord, S.; Bello, P.; and Ferrucci, D. 2003. Creativity, the Turing test, and the (better) Lovelace test. In *The Turing Test*. Springer. 215–239.

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.; Wiggins, G. A.; et al. 2012. Computational creativity: the final frontier? In *ECAI*, volume 12, 21–26.

Colton, S. 2008. Creativity versus the perception of creativity in computational systems. In *AAAI Spring Symposium: Creative Intelligent Systems*, 14–20.

Cook, M., and Colton, S. 2015. Generating code for expressing simple preferences: Moving on from hardcoding and randomness. In *Proceedings of the Sixth International Conference on Computational Creativity June*, 8.

Cropley, A. J. 2000. Defining and measuring creativity: Are creativity tests worth using? *Roeper review* 23(2):72–79.

Csikszentmihalyi, M. 2013. *Creativity: The Psychology of Discovery and Invention, reprint ed*. Harper Perennial.

Fauconnier, G., and Turner, M. 2003. Conceptual blending, form and meaning. *Recherches en communication* 19(19):57–86.

Guckelsberger, C.; Salge, C.; Saunders, R.; and Colton, S. 2016. Supportive and antagonistic behaviour in distributed computational creativity via coupled empowerment maximisation. In *Proceedings of the Seventh International Conference on Computational Creativity*.

Guckelsberger, C.; Salge, C.; and Colton, S. 2017. Addressing the "why?" in computational creativity: A non-anthropocentric, minimal model of intentional creative agency. In *Proceedings of the 8th International Conference on Computational Creativity, Atlanta*.

Guilford, J. 1950. Creativity. *American Psychologist* 5:444–454.

Jordanous, A. 2011. Evaluating evaluation: Assessing progress in computational creativity research.

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.

Loughran, R., and O'Neill, M. 2016. Generative music evaluation: Why do we limit to 'human' ? In *Computer Simulation of Musical Creativity (CSMC)*.

Loughran, R., and O'Neill, M. 2017a. Application domains considered in computational creativity. In *Proceedings of the 8th International Conference on Computational Creativity, Atlanta*.

Loughran, R., and O'Neill, M. 2017b. Limitations from assumptions in generative music evaluation. *Journal of Creative Music Systems* 2(1).

Marsden, A. 2013. *Music, intelligence and artificiality.* Routledge. 15–29.

Meusburger, P. 2009. Milieus of creativity: The role of places, environments, and spatial contexts. In *Milieus of Creativity*. Springer. 97–153.

Nebel, C. 1988. *The dark side of creativity: Blocks, unfinished works and the urge to destroy*. Whitston Publishing Company.

O'Hara, M., and Auersperg, A. M. 2017. Object play in parrots and corvids. *Current Opinion in Behavioral Sciences* 16:119–125.

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. *Creativity: From potential to realization.*

Plucker, J. A. 2004. Generalization of creativity across domains: Examination of the method effect hypothesis. *The Journal of Creative Behavior* 38(1):1–12.

Ritchie, G. 2001. Assessing creativity. In *Proc. of AISB'01 Symposium*. Citeseer.

Ritchie, G. 2006. The transformational creativity hypothesis. *New Generation Computing* 24(3):241–266.

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

Runco, M. A., and Acar, S. 2012. Divergent thinking as an indicator of creative potential. *Creativity Research Journal* 24(1):66–75.

Runco, M. A., and Jaeger, G. J. 2012. The standard definition of creativity. *Creativity Research Journal* 24(1):92–96.

Runco, M. A. 1987. The generality of creative performance in gifted and nongifted children. *Gifted Child Quarterly* 31(3):121–125.

Runco, M. A. 2014. *Creativity: Theories and themes: Research, development, and practice*. Elsevier.

Saunders, R., and Bown, O. 2015. Computational social creativity. *Artificial life* 21(3):366–378.

Stein, M. I. 1953. Creativity and culture. *The journal of psychology* 36(2):311–322.

Still, A., and d'Inverno, M. 2016. A history of creativity for future ai research. In *Title: Proceedings of the 7th Computational Creativity Conference (ICCC 2016). Universite Pierre et Marie Curie*.

Urban, K. K., and Jellen, H. G. 2005. *Test of Creative Thinking: Drawing Production: TCP-DP*. Harcourt Assessment, Incorporated.

Ventura, D. 2016. Mere generation: Essential barometer or dated concept. In *Proceedings of the Seventh International Conference on Computational Creativity, ICCC*.

Ventura, D. 2017. How to build a CC system. In *Proceedings of the 8th International Conference on Computational Creativity, Atlanta*.

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

Wiggins, G. A. 2012. The mind's chorus: creativity before consciousness. *Cognitive Computation* 4(3):306–319.

y Pérez, R. P. 2015. From mexica to mexica-impro: The evolution of a computer model for plot generation. In *Computational Creativity Research: Towards Creative Machines*. Springer. 267–284.