Advances in Cognitive Systems X (20XX) 1-6 Submitted X/20XX; published X/20XX

An observer-oriented theory of creativity for the evaluation of creative systems and processes

Chris Miller CWMILLE6@NCSU.EDU
Arnav Jhala AHJHALA@NCSU.EDU
Department of Computer Science, North Carolina State University, Raleigh, NC, 27695, USA

Abstract

As a step towards formalizing the cognitive process and interpretation of creativity, we present an extended theory of creative systems that aims to better define the evaluation of creative systems and processes. We extend previous research defining a formal framework for describing and comparing creative systems by taking into account the judgment of artefacts that they generate with respect to an observer. Prior work on formalizing creativity has mainly focused on the generating system with respect to the evaluation of created artefacts. This paper extends this formalism by modeling the judgment of creativity through models of observer systems. We show how our extended theory can support the interpretation of MacGyver problems—problems defined in the cognitive systems research community as classical planning problems designed to elucidate the cognitive process of creative problem solving. Finally, we expand an existing definition of evaluation in creative systems by putting forth general criteria of subjective evaluation based on initial empirical work regarding aesthetic measure across artistic domains. This was done by using statistical methods to model individual and group preferences over creative artefacts. These pilot experimental results model abstract metrics of balance, symmetry, and readability across these artistic domains.

1. Introduction

The motivation for our work on understanding creativity and the creative process begins with Boden's seminal work on understanding creative behavior described in her book, *The Creative Mind* (2004). In this book, Boden sets out to describe human creativity, creative behavior, and how computers can help us understand creativity. However, Boden's approach is purely qualitative and does not provide a formal framework to explain the ideas in a way that is amenable to computational analysis and representation.

To address the need to formalize Boden's ideas, Wiggins (2006) put forth a detailed framework for describing, analyzing, and comparing creative systems directly motivated by Boden's hierarchy of creativity. Wiggins' work represents one of the few attempts to formalize Boden's approach to creativity with the explicit intent of moving towards a better understanding of systems that exhibit creative behavior. Inspired by Boden's description of *psychological creativity*, we propose an extension to Wiggins' framework to account for the subjective evaluation of potentially creative systems and artefacts. In this extended framework, we present the perception of creativity as the subjective judgment of an observer with varying amounts of domain knowledge. We then expand Wiggins' initial definition of evaluation in creative systems by putting forth general criteria of subjective evaluation based on empirical work regarding aesthetic measure across artistic domains.

In this paper we also highlight interesting parallels between our work on formalizing creative systems with work in cognitive systems research. In particular, we reinterpret MacGyver problems described by Sarathy and Scheutz (2018) in our extended theoretical framework. The formulation of MacGyver problems represents a step towards designing creative cognitive systems while also introducing a fruitful class of problems for the cognitive systems research community to tackle. Better understanding creativity, both in terms of generative systems and evaluation systems, is a vitally important part of modeling human-like cognitive systems. Accordingly, we believe that this work, and continued research on formalizing creativity, can greatly benefit the cognitive systems research community as researchers explore the cognitive process of creativity.

2. Background: related work

2.1 Boden and creativity

Boden (2004) defines creativity as the ability to "come up with ideas or artefacts that are *new, surprising and valuable*". Although researchers have debated the extent to which surprise is an essential criteria for creativity or just a kind of novelty, researchers generally agree that *novelty* and *value* are two essential criteria for creativity (Ritchie, 2007; Oman & Tumer, 2009). We will do the same in our framework. Boden's analysis of creativity starts by defining two particular cases of creativity: *psychological creativity* and *historical creativity* (also referred to as P-creativity and H-creativity respectively). Psychological creativity describes a creative process through which an idea is created that is new to the person who came up with the idea. Historical creativity is a special case of P-creativity wherein an idea is created that no one else has ever had before (i.e. completely new in a societal context). P-creativity and H-creativity (and frankly creativity as a whole) are *very* context-specific. If we were to break human society into subgroups, an idea that might be considered H-creative in one subgroup might not be considered H-creative in another subgroup and vice versa. Similarly, an individual's idea that is considered H-creative in their specific subgroup might only be considered P-creative in another subgroup that had already conceived of that individual's idea. For the majority of this paper however, we will not distinguish between P-creativity and H-creativity. Our analysis of creativity will focus primarily on the fundamental idea of P-creativity which frames any evaluation of creativity as an evaluation of creativity from an individual's perspective.

In her work, Boden also puts forth the idea of the creative process involving a search through a conceptual space of ideas and artefacts. If one were to consider a content generation system as a potentially creative system, its conceptual space could be loosely interpreted as a state space of artefacts. In order to better explain how a creative process involves exploring these conceptual spaces, Boden defines two particular types of creativity: *exploratory creativity* and *transformational creativity*. Exploratory creativity is defined as discovering new concepts in the conceptual space via an exploration of the conceptual space whereas transformational creativity is defined as transforming the conceptual space via rules which "sculpt" the conceptual space. As Wiggins (2006) points out, Boden does not provide any explicit details of how conceptual space constraints may be defined or what the difference between exploring and transforming the space is.

To address these concerns and to formalize Boden's ideas, we present Wiggins' descriptive framework of creative systems. However, before attempting to explain creativity—a term which is already ill-defined—we will adopt a useful definition of creativity as mentioned by Wiggins analogous to early definitions of intelligence. This definition describes creativity as "the performance of tasks which, if performed by a human, would be deemed creative."

2.2 Wiggins' creative systems framework

Wiggins' (2006) creative systems framework begins by defining a *universe*, $\mathscr{U}$, a multidimensional space that can represent anything. All possible concepts exist as distinct points in the multidimensional space defined by $\mathscr{U}$. Following the state-space analogy of Boden's initial conceptual space definition, $\mathscr{U}$ is capable of representing all abstract and complete concepts and contains both complete and incomplete artefacts; this includes the empty concept $\top$. Hence, $\mathscr{U}$ contains all possible concepts, including the empty concept $\top$. Additionally, all concepts in $\mathscr{U}$ are distinct (i.e. no two concepts in $\mathscr{U}$ are identical to one another).

To define the conceptual spaces of $\mathscr{U}$, two rule sets, $\mathscr{R}$ and $\mathscr{T}$, are introduced. $\mathscr{R}$ represents the rules that constrain (i.e. sculpt) the conceptual space and $\mathscr{T}$ represents the rules that determine how to traverse the conceptual space. In the context of answer set programming, $\mathscr{R}$ can be likened to constraints which omit potential solutions exhibiting properties that the constraints specify as invalid—thereby constraining the space of potentially valid solutions (Smith & Mateas, 2011). In the context of state-space search, $\mathscr{T}$ can be likened to search strategies (e.g. heuristics) for informed search through states. $\mathscr{R}$ and $\mathscr{T}$ are further defined as belonging to a language $\mathscr{L}$ composed from an alphabet $\mathscr{A}$. The last step required for formally defining a conceptual space is defining an interpretation function. [[.]] is defined as a partial function that maps from $\mathscr{L}$ to real number functions that return a real number in [0,1] which is then mapped to either `true` or `false` depending on its value. With the above definitions of $\mathscr{U}$, $\mathscr{R}$, and the interpretation function [[.]], a conceptual space, $\mathscr{C}$, is defined as:

$$\mathscr{C} = [[\mathscr{R}]](\mathscr{U}) \tag{1}$$

By this definition, a conceptual space, $\mathscr{C}$, is defined by the interpretation, [[.]], of rules defining how to constrain *and* define a conceptual space, $\mathscr{R}$, applied to a larger multidimensional space, $\mathscr{U}$. Additionally, now the distinction between $\mathscr{R}$ and $\mathscr{T}$ is more apparent; $\mathscr{R}$ defines rules that apply to $\mathscr{U}$ to describe a whole domain of acceptable artifacts whereas $\mathscr{T}$ defines traversal rules (e.g. heuristics) for exploring possible artifacts within the domain defined by $\mathscr{R}$.

In order to account for the traversal (i.e. exploration) of conceptual spaces via $\mathscr{T}$, another interpreter $\langle\langle .,.,. \rangle\rangle$ is defined. $\langle\langle .,.,. \rangle\rangle$ takes as input three subsets of $\mathscr{L}$, namely $\mathscr{R}$, $\mathscr{T}$, and $\mathscr{E}$ which will be briefly described soon. $\langle\langle .,.,. \rangle\rangle$ operates on a totally ordered subset of $\mathscr{U}$, c_{in}, and maps it to another totally ordered subset, c_{out}. Namely,

$$c_{out} = \langle\langle \mathscr{R}, \mathscr{T}, \mathscr{E} \rangle\rangle(c_{in}) \tag{2}$$

These formalizations more rigorously define Boden's idea of conceptual spaces for creativity and can already describe Boden's (2004) exploratory and transformational creativity. Transformational creativity can be achieved by either transforming $\mathscr{R}$ or $\mathscr{T}$.

$\mathscr{E}$ is described as a set of rules that allows for the evaluation of concepts in $\mathscr{C}$ and determination of their value. Formally defining $\mathscr{E}$ is complicated by the fact that the criteria used for $\mathscr{E}$ can vary depending on desired attributes of artefacts and the application domain. Accordingly, Wiggins' framework consciously chooses to avoid discussing how to evaluate these concepts and leaves it as an important issue to be discussed at a later time. In summation, Wiggins' framework rigorously formalizes the process of creativity in systems and creative behavior in agents but leaves the evaluative aspect open for future discussion. We will revisit $\mathscr{E}$ when we describe our analysis of empirical work on defining a cross-domain metric of aesthetic measure for the comparison of created artefacts later in the paper.

It should be noted that previous research has explored how to evaluate artefacts in specific domains and how to provide preliminary frameworks for evaluating creativity but, to our knowledge, little research has been done to define the evaluation of creativity in terms of an explicit observer of varying knowledge and expertise (Boden, 1998; Pearce & Wiggins, 2001; Ritchie, 2001). We believe that an observer's subjective evaluation is fundamental to theories of creativity analysis given the strong connection between Boden's P-creativity and the personal context of the individual(s) evaluating some creative process or artefact. We will make this observer-oriented subjective evaluation explicit in our extended framework of creative systems.

3. Extended theory of creative systems

3.1 Defining an observer

We start by defining an observer, $\mathscr{O}$, of a creative process:

> **Definition 1** (*Observer*). An observer, $\mathscr{O}$, is an entity capable of representing concepts in the multidimensional space $\mathscr{U}$.

Given that $\mathscr{U}$ defines the universe of possibilities of a creative process' conceptual space, $\mathscr{O}$ describes an observing entity with the *potential* to re-represent concepts in the universe of possibilities $\mathscr{U}$. By our definition an observer can be a human or even a computational system. In fact, observers can be represented as generative systems that produce artefacts in their respective conceptual spaces. This idea of *potential re-representation* is particularly important because not every observer, $\mathscr{O}$, can be assumed to be capable of representing any *and* all concepts in $\mathscr{U}$. It is for this reason that we define a special type of observer which we shall call the *omniscient oracle* $\mathscr{O}_o$. The omniscient oracle is an observer capable of representing any *and* all concepts in $\mathscr{U}$. As we will later show, the omniscient oracle as described here is an observer for which nothing is creative due to its complete representation of $\mathscr{U}$. This distinction between observers and the omniscient oracle leads to the following axiom:

> **Axiom 1** (*Restricted observer representation*). Any observer, $\mathscr{O}$, that is not the omniscient oracle can only represent conceptual spaces, $\mathscr{C}$, that are proper subsets of $\mathscr{U}$. $\mathscr{C} \subsetneq \mathscr{U}$.

The above axiom restates the conclusion that any observer that is not the omniscient observer can only re-represent parts of $\mathscr{U}$ but never the entirety of $\mathscr{U}$. We would also like to explicitly state that *multiple observers can exist simultaneously*. Given the definition stating that an observer, $\mathscr{O}$, can potentially re-represent concepts in the universe $\mathscr{U}$, we now define *how* observers represent these concepts.

3.2 Defining an observer's conceptual space

A fundamental property of observers is that all observers possess their own, individual rule sets $\mathscr{R}_o$ and $\mathscr{T}_o$. These rule sets are analogous to $\mathscr{R}$ and $\mathscr{T}$ defined earlier except that they are particular to an observer, $\mathscr{O}$, rather than being generally defined for a creative process. Hence, $\mathscr{R}_o$ represents an observer's specific rules for constraining a conceptual space and $\mathscr{T}_o$ represents their specific rules for traversing a conceptual space. In practice, these rule sets represent the extent to which an observer is able to define domains within $\mathscr{U}$ as well as explore concepts within it. As such, these rule sets effectively represent an observer's knowledge of domains in $\mathscr{U}$.

Using the interpretation function $[[.]]$ from earlier which works on well-formed rule sets described in a language $\mathscr{L}$, we can define the conceptual space of an observer $\mathscr{O}$ as follows:

$$\mathscr{C}_o = [[\mathscr{R}_o]](\mathscr{U}) \tag{3}$$

There are a few things which should be stated:

1. The interpretation function $[[.]]$ is defined such that as long as a constraining rule set, $\mathscr{R}$, is a well-formed rule set, it can be successfully interpreted and this interpretation works for any well-formed rule set in $\mathscr{U}$.

2. Given that the conceptual spaces of any observer, $\mathscr{C}_o$, must exist within $\mathscr{U}$, interpreting $\mathscr{R}_o$ within $\mathscr{U}$ yields $\mathscr{C}_o$.

3. The same universe, $\mathscr{U}$, can be used to describe the conceptual spaces of different observers with different rule sets.

This definition formalizes the idea that observers' conceptual spaces are subsets of the universe of possibilities $\mathscr{U}$. It also becomes more apparent that an observer and its conceptual spaces always exist within the context of $\mathscr{U}$.

We would also like to revisit how conceptual space restraining rules may be interpreted by a system. Let r be a rule defined in the rule set $\mathscr{R}$. We refrain from stating that concepts that violate r are immediately excluded from the constrained conceptual space given that numerous rules can be defined in $\mathscr{R}$ and they need not be mutually exclusive. For example, consider the following rules in the context of music composition:

1. Produce musical scores using a tempo of 120 beats per minute (bpm).

2. Produce musical scores using at least two different instruments.

If the rules are not explicitly defined as being mutually exclusive, acceptable concepts can include artefacts that only use a tempo of 120 bpm, only use at least two different instruments, or both. If these are the only two rules applied to $\mathscr{U}$, then $\mathscr{C}$ will include concepts that satisfy at least one of the rules in $\mathscr{R}$ while excluding the rest. Regardless, the issue of rule exclusiveness is a design choice which can be delegated to the rule interpretation function [[.]] which is shared by the creative process and observers. We envision that a key challenge in an implemented system will be being able to recognize and handle rules that conflict and or contradict each other. Whether it is by excluding one of the contradictory rules or by finding another set of compatible rules, this decision can greatly impact how well a potentially creative system can combine and analyze rules to generate creative artifacts.

3.3 Defining the observer in relation to the creator

We begin our theory's components on subjective evaluation by first highlighting the relationship between the observer and the general creative process. The general creative process refers to the potentially creative process defined by Wiggins' (2006) preliminary creative systems framework. This process is a formalization of how some entity (henceforth called the *creator*) defines its conceptual space and explores and transforms conceptual spaces to exhibit potentially creative behavior in the course of creating new and valuable artefacts in the context of a universe of possibilities $\mathscr{U}$. We use the term "creator" to emphasize the fact that the entity engaging in exploring and redefining conceptual spaces in $\mathscr{U}$ seeks to generate creative artefacts and can be a human creator or another generative system. We can now phrase an observation from earlier as follows:

> **Observation 1** (*Shared universal context*). The creator and observer(s) share the same universal context $\mathscr{U}$.

Another way to interpret the first observation is that the conceptual spaces of the creator and the observer(s) all exist within the same universe. Given that our viewpoint of the observer defines an observer as an independent onlooker to the creator's creative process, the creator's conceptual space defining rule sets, $\mathscr{R}$ and $\mathscr{T}$, are independent from those of the observer(s), $\mathscr{R}_o$ and $\mathscr{T}_o$. This leads to the follow observation which is illustrated in Figure 1:

> **Observation 2** (*Independence of creator/observer concepts*). The conceptual space of the creator, $\mathscr{C}$, and the conceptual space of the observer, $\mathscr{C}_o$, are defined independently of each other but are both subsets of $\mathscr{U}$. (Note: If there are multiple observers, each respective $\mathscr{C}_o$ is defined independently of the other observers' conceptual spaces.)

The idea of independence between these conceptual spaces only extends to how they are independently defined by their individual originators. If $\mathscr{R}$ and $\mathscr{R}_o$ contain similar rules, $\mathscr{C}$ and $\mathscr{C}_o$ will overlap which in turn means that the observer is able to conceive of *some* concepts in the conceptual space of the creator (as determined by the amount of intersection) and vice versa. This is evident in the image on the right in Figure 1. With this knowledge we are now able to state how and when the observer might consider the creator (and its corresponding creative process) as being creative.

3.4 An initial definition of the subjective evaluation of creativity

Our definition of an observer's subjective evaluation of creativity begins with the following postulate:

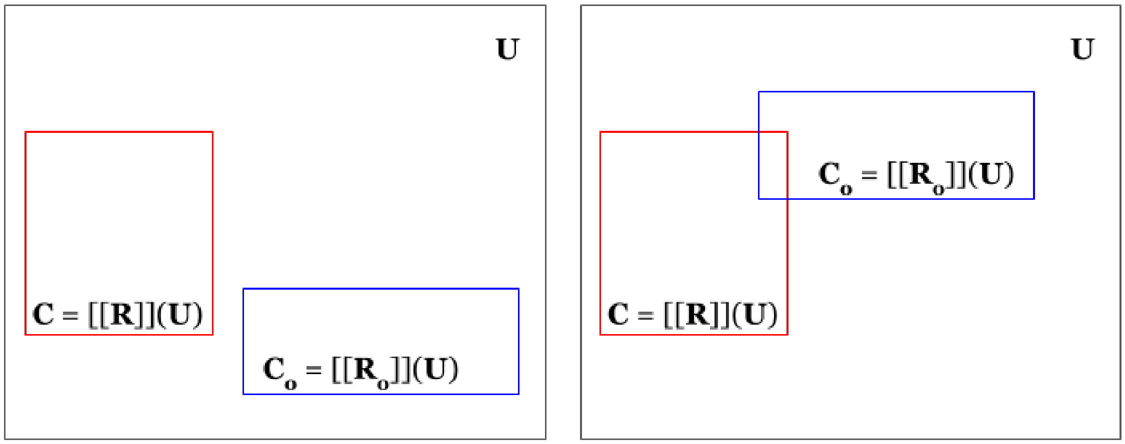

Figure 1. Illustration of how two example conceptual spaces, $\mathscr{C}$ and $\mathscr{C}_o$, exist within $\mathscr{U}$. **Left**: No intersection of concepts belonging to the creator and the observer. **Right**: Intersection of concepts belonging to the creator and the observer.

> **Postulate 1**: an observer is likely to consider concepts that are not explicitly defined in their conceptual space, $\mathscr{C}_o$, but appear in the conceptual space of a creator, $\mathscr{C}$, as being creative.

These concepts represent artefacts and ideas that lie outside of the observer's explicitly known artefact range and represent a class of concepts that are new to the observer. We have based this postulate on the essential criteria of creativity which suggests that creative ideas and artefacts are novel.

We can also represent concepts that are new to the creator by framing the creator as its own observer. Therefore, even if an observer may not consider a creator's concept as creative, if the creator came up with a novel idea within their own perspective, the creator can still think that their newly formed idea is creative. This is directly related to Boden's initial formulation of P-creativity in that concepts in $\mathscr{C}$ but not in $\mathscr{C}_o$ represent new concepts that an observer is unable to represent (or at least initially conceive of). If the creator is also its own observer, then P-creativity results from the creator's own initial observer state (i.e. prior to the creative process) identifying a new result as creative. As an aside, if we represent a society as consisting of the creator and any number of observers, then we can describe ideas that are novel to all of the observers and the creator as being historically creative within that societal context.

We can use rule sets to define how these conceptual spaces may overlap or be distinct from each other. Recall that conceptual spaces are defined by the interpretation of space constraining rules $\mathscr{R}$ in the context of a universe of possibilities $\mathscr{U}$. Based on the rule interpretation applied to $\mathscr{R}$, each rule r in $\mathscr{R}$ constrains $\mathscr{U}$ to *include* concepts that satisfy r. Hence, for each r that is both in $\mathscr{R}$ and $\mathscr{R}_o$, the conceptual spaces $\mathscr{C}$ and $\mathscr{C}_o$ overlap each other over concepts that satisfy r. More concisely:

$$\exists r : r \in \mathscr{R} \wedge r \in \mathscr{R}_o \iff \exists c : c \in \mathscr{C} \wedge c \in \mathscr{C}_o. \tag{4}$$

Similarly, if neither of the rule sets share a common rule then there is no overlap between the two conceptual spaces.

It should now be more intuitive that the rule sets of the creator and the observer affect the observer's perceived creativity of the creator and its concepts. When the creator's conceptual space includes concepts that are outside of the observer's conceptual space, these concepts are more likely to be deemed creative by the observer. These "external" concepts result from rules that are exclusive to the creator's rule set which allows them to represent these concepts. In other words, these concepts exhibit one of the essential criteria of creativity: *novelty*. This leads to the following claim:

> **Claim 1**: an observer's perceived creativity of a creator results from rules that the creator possesses but the observer has not yet incorporated into their own rule set with respect to some application domain.

In consideration of the claim of perceived creativity mentioned above, one might also come to the conclusion that if that an observer adopts previously unknown rules into their rule set, then the observer may no longer consider concepts relying on those rules to appear creative as creative. This observation describes an observer who expands their domain knowledge as they encounter new concept rules. In this case, the "learned observer" expands their threshold of creativity thereby requiring a creative process to generate even newer concept rules that the learned observer will now consider creative. Hence, the creative process can also be framed as a struggle to continue being perceived as creative by observers who learn from their exposure to new ideas.

Understanding the perception of creativity through this framework can also aid in the development of systems that aim to produce interesting content for users and designers. One such application is the area of procedural content generation via machine learning (PCGML) as defined by Summerville et al (2018). By using machine learning methods to learn an abstracted form of an individual's knowledge of a domain, such systems can better focus on what properties of content may appear creative to specific individuals. In the context of video game level design, this could extend to a level generation system generating levels with concepts that are newer and newer to a player. Of course, balancing the expressiveness of the generator as it adjusts to an individual while making sure its content is functionally valuable is a challenging task.

3.5 Connection to MacGyver problems in cognitive systems research

Our formulation of the creative systems framework—in particular its emphasis on rule comparison-based creativity evaluation—shares an interesting connection to MacGyver Problems in cognitive systems research (Sarathy & Scheutz, 2018). MacGyver problems are defined as a class of classical planning problems that are initially unsolvable for an agent given their current knowledge. However, MacGyver problems are also defined in such a way that if an agent sufficiently expands its domain representation or understanding of its world within the context of a larger universe, the problem can ultimately be solved. Similar to Wiggins' initial framework and subsequently our extended framework, MacGyver problems define a universe consisting of numerous worlds whose *conceptual spaces* within the universe are shaped by the abilities of individual agents of which these agents are either aware or unaware. A key component of the world definition is that worlds only represent

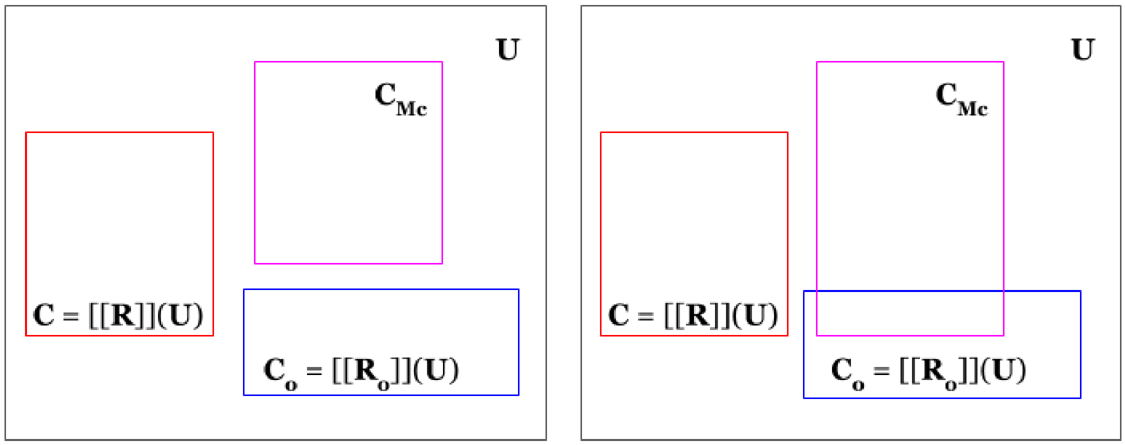

Figure 2. Illustration of the initial states of sample MacGyver problems informally described in the creative systems framework. **Left**: No intersection between the space of valid solutions and the observer's known world conceptual space at the initial state. **Right**: Intersection between the space of valid solutions and the observer's known world conceptual space at the initial state.

portions of the universe that are both *perceivable* and *actionable* by their respective agents. Hence, worlds can be likened to conceptual spaces in $\mathscr{U}$.

MacGyver problems are designed to showcase how humans exhibit flexible and creative behavior to solve problems that are initially unsolvable given their current knowledge of a problem. They are specifically presented as an initial step towards designing creative cognitive systems. The core component of MacGyver problems that relates to the process of creative problem solving is how an agent expands its understanding of its world to solve these problems. This particular trait of MacGyver problems is comparable to how our described observers consider a creator to be creative.

We can informally represent a MacGyver problem in our framework by letting a conceptual space $\mathscr{C}_{Mc}$ describe a class of solutions to a MacGyver problem and by letting the agent trying to solve the MacGyver problem be the creator. By the definition of MacGyver problems, we know that the agent's conceptual space will not initially include any solutions in $\mathscr{C}_{Mc}$. In order to solve the problem, the agent (i.e. creator) is free to explore and transform their conceptual space by defining and redefining their rule sets to hopefully achieve a valid solution defined in $\mathscr{C}_{Mc}$. This is analogous to an agent in a MacGyver problem expanding its domain representation and world understanding. We also define an observer of the agent which represents an onlooker of the agent's attempts at discovering a solution.

Once the agent discovers a solution, the onlooking observer can then determine whether or not the agent's solution is creative. For instance, if the agent adopted a rule to solve the problem that the observer had not previously considered or defined within their own world's conceptual space, then the observer may very well consider the agent's overall problem solving process as creative. The subjective aspect of the observer's evaluation of the agent's creativity defined by our framework can also describe a situation wherein the observer does not consider a particular solution to be creative. An example of this would be when the creator's solution (i.e. expanded world knowledge to reach a solution) already exists within the observer's world. In other words, the observer's conceptual space intersects with some part of the class of solution $\mathscr{C}_{Mc}$. This would mean that the agent's solution happened to be a solution of which the observer was already aware. Hence, even if the agent's solution was creative based off of their initial knowledge when solving the problem, an observer might not consider it to be creative in their own context. Figure 2 illustrates these two scenarios via their respective images.

4. Finding a general aesthetic measure metric for $\mathscr{E}$

We now return our attention to the evaluation rule set $\mathscr{E}$ from Wiggins' (2006) creative systems framework. Once again, $\mathscr{E}$ is defined as a set of rules that allows for the evaluation of concepts in $\mathscr{C}$. In terms of the essential criteria of creativity, $\mathscr{E}$ represents *value*. One of the most difficult aspects of formalizing $\mathscr{E}$ results from the fact that the criteria used for $\mathscr{E}$ can vary depending on the desired attributes of artefacts. Furthermore, the criteria for $\mathscr{E}$ can vary significantly across different domains describing the general types of artefacts that should be produced by the creator. If we consider a creator pursuing a creative process as a generative system, we find that this difficulty in characterizing $\mathscr{E}$ mirrors one of the key challenges in the design, development, and evaluation of generative systems: defining a metric for the *comparison* of created artefacts. Numerous researchers have also noted that a further challenge in characterizing the *creativity* of a generative system is to utilize the metric of comparison with respect to an *observer* to evaluate the perceived creativity of generated artefacts of the generative system (Karimi et al., 2018; Grace & Maher, 2016).

In order to better define $\mathscr{E}$ in the context of a generalizable, formal framework with an observer, we present empirical work in two artistic domains that illustrates abstract, general properties of aesthetic measure for the evaluation of artefacts. These abstract metrics can be applied across numerous artistic domains and suggest general attributes for $\mathscr{E}$ as an evaluation metric.

Our inspiration for this analysis is the quest for a precise mathematical definition of a cross-domain aesthetic measure as pursued by Birkhoff (1933) and a more recent reformulation of it by Moles (1968) , Bense (1969), and Rigau, Feixas, and Sbert (2007) in terms of information theory and computational complexity. This is our starting point because any generative system needs self- and external-validation in terms of being creative by means of *comparison* of the artefacts it generates. Such a comparison, particularly in terms of creativity, can be seen as a relative measurement of its aesthetic (broadly speaking) appeal to its audience (e.g. an observer).

Birkhoff, in his detailed monograph titled *Aesthetic Measure*, defined aesthetic measure as the ratio of *Order* over *Complexity*. This can be understood as the complexity of interpreting the number of features an artifact includes and the amount of regularity observed in the construction of the artifact. For example, in a generative system that constructs shapes consisting of 4 axis-aligned line segments that form closed figures, its range consists of axis-aligned quadrilaterals of all sizes. For any given quadrilateral in its range, the complexity of shapes can be considered equal but squares and rectangles have a certain orderly arrangement with right angles between edges. Furthermore, among rectangles, certain aspect ratios are perceived to be more pleasing due to the perceived regularity in their shape. Birkhoff provides extensive calculations to justify the aesthetic scores that are calculated with his measure for this domain. Some shapes are reproduced in Figure 3.

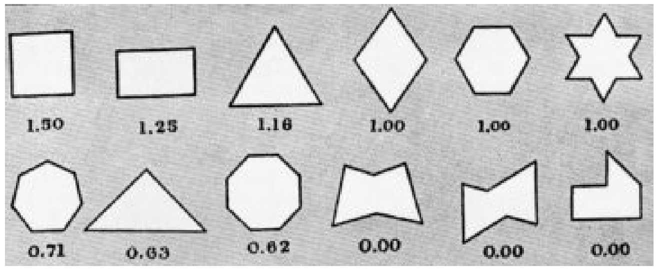

Figure 3. Applied to polygons, Birkhoff's formula for aesthetic value gives a square the highest rating. A five-pointed star (not shown) similar to the one that appears on the flag of the United States has a rating of 0.90. (Ivas Peterson, Science News, May 2004)

Another notable thread of research that directly connects with Birkhoff is that by Rigau et al (2007). They limit their quantification of aesthetic measure to paintings and use paintings by Mondrian, Pollock, and Van Gogh to illustrate the applicability of their metric. In this work, the measure of order is the Shannon Entropy of the palette of the image which corresponds to the selection of palette colors in relation to a uniform palette. The measure of complexity is defined as the Kolmogorov complexity which is estimated as the compression ratio of a compression algorithm applied to the image. Taking these two metrics, the aesthetic measure of an image is then defined as the ratio of *initial information content* in the image represented by Shannon Entropy to *reduction in uncertainty* represented by the compression ratio. One key distinction between the original approach by Birkhoff and latter work is the choice of artefacts. The examples Birkhoff evaluates range from abstract forms like the shapes of vases to musical arrangement. At first glance, it appears that the choice of vase form is primarily motivated by the shapes of ancient vases and have a historical connection. But this approach of evaluation of forms is quite relevant in modern design as can be seen by the story of the design of the Coke bottle[1]. There is a difference between a generally pleasing form (e.g. photographic composition that adheres to the golden ratio) that becomes commonplace in society and is generally appreciated in terms of features of design and a specific and unique artifact.

5. Practical measurement of aesthetics

Computational models of aesthetic measure have many applications ranging from creativity support tools, evaluation of algorithmic creativity, and personalization with respect to user preferences. Such measures are challenging to develop due to the lack of functional and interpretable inputs, noise in naturally occurring data sets, and unreliable self-reported data due to different interpretations of embedded content. Early work in psychology on art and visual perception identified concepts like symmetry, rhythm, contrast, etc. that correspond to cross-domain features of visual aesthetics (Arnheim, 1974). As a significant component of the evaluation rule set $\mathscr{E}$, aesthetic evaluation of generated artefacts is an important piece in characterizing creativity. For this, we present empirical work in two domains, namely photographic composition and full-body gestural performance (dance), and

1. `https://www.coca-colacompany.com/stories/the-story-of-the-coca-cola-bottle`

Figure 4. Collage of in-game photographs taken from the game Panorama, developed at UC Santa Cruz.

present experimental results on modeling abstract metrics of Balance, Symmetry, and Readability across these artistic domains. In these experiments, researchers respectively created corpora with automatically annotated features, and collected human preference ratings through crowdsourcing platforms to get a judgment of *goodness* of artefacts. We discuss issues related to the design of experiments, feature modeling and selection, and applications of machine learning algorithms for learning preferences.

5.1 Composition preferences for synthetic photographs

For understanding visual composition preferences, Swanson, Escoffery, and Jhala (2012) created a corpus of synthetic photographs through a game, *Panorama*, that procedurally generated panoramas with limited palettes of shapes, colors, and objects (Morgens & Jhala, 2013). Their goal was to create models to compute the overall composition quality of an image with respect to composition features as well as capturing the difference in preferences of individual viewers (see Figure 4). Top-down design of composition metrics, based on photography rules, were combined with bottom-up statistical analysis to correlate aesthetic quality with viewer preferences. Color coded badges in the game functioned as an aesthetic meter, encouraging players to take pictures that rated highly for Rule of Thirds, Balance, and Symmetry. Generating their own corpus of images reduced the feature set in a way that would be difficult using photos from image repositories on the Web. This gave them additional control over feature dimensions through design abstraction, such as removing color or minimizing cultural references in the game's representation of landscape photography. Using their corpus to conduct a preference study, viewer perception of image quality was rated through crowd sourcing on Mechanical Turk.

Within this simplified vocabulary of greyscale panoramic images with a small number of shapes for buildings, trees, meadows, and windmills, the game was still able to get a large number of photographs that varied in terms of the number of objects, size of objects, location in frame, angle of the camera, horizon line in the frame, and other features totaling to over 250 low level features. They conducted a study with pairwise comparison with 4 alternative forced choices. After pairwise preference data collection, they used SVN based models to check for accuracy of prediction. The

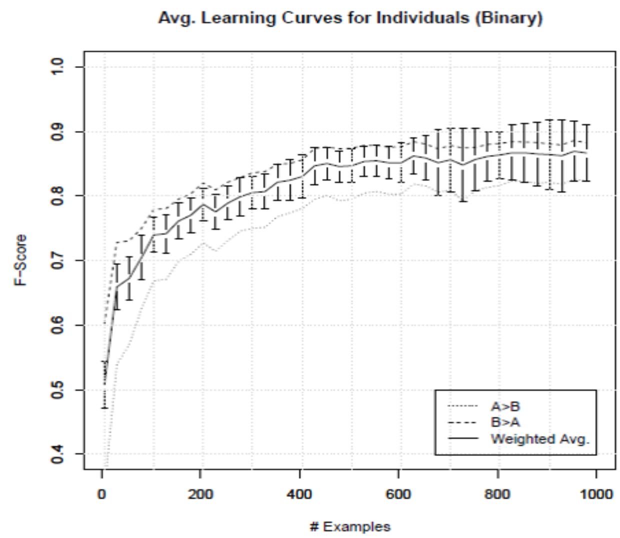

Figure 5. Learning rate for model that uses high level Balance, Spacing, Symmetry features on individual preferences. Higher f-score indicates more consistency in higher ratings of images with significant features.

key insight was that human raters, for whom the algorithm was able to predict with high accuracy ratings for unseen photographs, were consistent in their choice of features for aesthetic evaluation. Users for whom the algorithm was unable to get high prediction accuracy were not consistent with the presence of features in their high rated images.

Figure 5 summarizes results for individual learning rate over user ratings.

5.2 Evaluating gestural aesthetics

Maraffi, Ishikawa, and Jhala (2013) took the notions of Balance, Symmetry, and Readability in full body gestural performances and defined these high level features in terms of low level joint positions and angles that were captured through a motion-capture camera. Similar to the setup for the Panorama project, performers played the Michael Jackson Experience dance game with the Microsoft Kinect but were captured in high resolution motion capture setup as shown in Figure 6.

An example performance with high level features is shown in Figure 7. Through an initial survey, performers were divided into two clusters, gamers with experience playing Kinect games but not good dancers, and skilled dancers without much experience playing games. The hypothesis was that if the definitions of high level aesthetic features were correct then the algorithm would be able to classify good and bad dancers from observation of their gameplay regardless of their score on the game. They also looked at how much the game's evaluation of poses matched the aesthetics of dance. For the experiment, data was recorded on simplified animation rigs as shown in Figure 8. Results of this work indicate that the high level features of balance, symmetry, and readability are good indicators of performer skill as judged by independent observers outside the context of gameplay.

Figure 6. Setup for gesture capture. Players played the dance game Michael Jackson Experience inside a high resolution Motion Capture system.

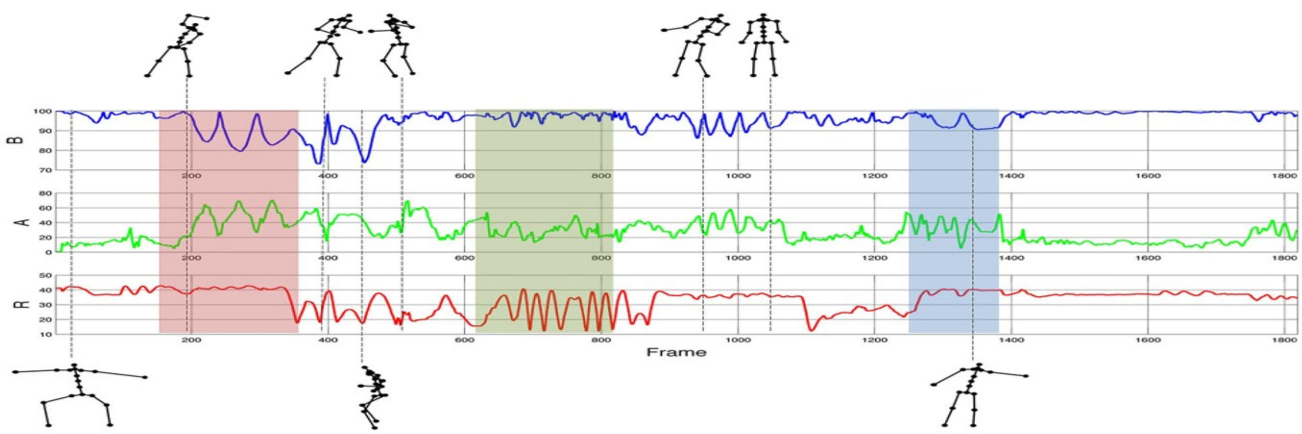

Figure 7. Aesthetic evaluation of dance performance using Balance, Readability, and Asymmetry as higher level features and joint positions and angles as low-level features.

Figure 8. Setup for gesture capture. Players played the dance game Michael Jackson Experience inside a high resolution Motion Capture system.

Along the lines of Birkhoff's idea of Order and Complexity, computational frameworks for the area of study of visual aesthetics have been influenced by psychology methods and visual art theory related to aesthetics. Composition rules have been studied in the popular arts of painting, photography, and cinema. For instance, Rule of Thirds has been implemented as a quality inferring feature used to distinguish professional from amateur photography. The Gestalt psychology concept of goodness configuration, where perception is organized according to properties like symmetry and simplicity, has also influenced feature representation. Processing fluency theory suggests readability features for image appeal. Challenges include an "aesthetics gap" that stems from the inherent semantics gap between low-level computable visual features and high-level subjective semantics. Core problems are predicting aesthetic and emotional responses for cliques in the general population, and understanding individual preferences that make some images more appealing than others are key concerns in that research area.

6. Conclusion

In this paper, we extended Wiggins' (2006) creative systems framework by adding a formalism describing an observer. This extension was inspired by Boden's (2004) description of psychological creativity in which creativity is judged from the perspective of an individual and their specific domain context. This observer exists within the same universe of possibilities described by Wiggins and, as such, is capable of representing concepts in the universe. By reframing the general cognitive process defined by Wiggins as an entity pursuing some creative process to generate artefacts via conceptual spaces, we defined the observer in relation to the creator as being an observer to a potentially creative process. The creativity of a potentially creative process is then formalized as the subjective judgment of an observer with varying amounts of domain knowledge. Our main claim from this formalization is that perceived creativity results from a creative process generating concepts by utilizing descriptive domain rules that an observer (or oneself) is initially unaware of within the context of an application domain. Our framework also supports situations in which the perceived creativity of some process or artefact differs amongst multiple agents.

We reinterpreted MacGyver problems within our framework as an agent's creative search process of expanding its domain descriptive rules well enough to reach any number of MacGyver problem solutions defined in a conceptual space within a universe of possibilities (Sarathy & Scheutz, 2018). The judgment of whether or not an agent's process of solving a MacGyver problem is creative is then determined by an observer of that agent's creative process.

We then concluded by putting forth general criteria for the subjective evaluation of artefacts based on aesthetic measure across artistic domains. These general criteria were proposed as an initial expansion of Wiggins' (2006) definition of evaluation in creative systems. These pilot experimental results modeled abstract metrics of balance, symmetry, and readability across these artistic domains.

We have described our work on formally defining creativity and cross-domain aesthetic measure as a step towards better understanding creativity. Our approach focuses on both the generative aspect of creativity as well as the evaluative aspect of potentially creative artefacts and ideas. We have presented this work with the belief that research on understanding creativity from the perspective of generative systems and evaluation systems can greatly benefit the cognitive systems research community.

7. Acknowledgements

This work was partly supported through a start-up grant from the NC State University Chancellor's faculty excellence cluster program. Authors would like to thank Christopher Maraffi, Dustin Escoffery, Reid Swanson, and Sasha Ishikawa for their support with the Panorama game and Dance Game Gesture corpus collection, as well as members of the Computational Cinematics Studio at UC Santa Cruz.

References

Arnheim, R. (1974). *Entropy and art: An essay on disorder and order*. Univ of California Press.

Bense, M. (1969). Introduction to the information-theoretical aesthetics, foundation and application to the text theory.

Birkhoff, G. D. (1933). *Aesthetic measure*. Cambridge, Massachusetts: Harvard University Press.

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

Boden, M. A. (2004). *The creative mind: Myths and mechanisms*. Routledge.

Grace, K., & Maher, M. L. (2016). Surprise-triggered reformulation of design goals. *Thirtieth AAAI Conference on Artificial Intelligence*.

Karimi, P., Grace, K., Maher, M. L., & Davis, N. (2018). Evaluating creativity in computational co-creative systems. *arXiv preprint arXiv:1807.09886*.

Maraffi, C., Ishikawa, S., & Jhala, A. (2013). Inferring performer skill from aesthetic quality features in a dance game gesture corpus. *Ninth Artificial Intelligence and Interactive Digital Entertainment Conference*.

Moles, A. (1968). Information theory and esthetic perception.

Morgens, S.-M. C., & Jhala, A. (2013). Synthetic photographs for learning aesthetic preferences. *Workshops at the Twenty-Seventh AAAI Conference on Artificial Intelligence*.

Oman, S., & Tumer, I. Y. (2009). The potential of creativity metrics for mechanical engineering concept design. *DS 58-2: Proceedings of ICED 09, the 17th International Conference on*

Engineering Design, Vol. 2, Design Theory and Research Methodology, Palo Alto, CA, USA, 24.-27.08. 2009.

Pearce, M., & Wiggins, G. (2001). Towards a framework for the evaluation of machine compositions. *Proceedings of the AISB'01 Symposium on Artificial Intelligence and Creativity in the Arts and Sciences* (pp. 22–32).

Rigau, J., Feixas, M., & Sbert, M. (2007). Conceptualizing birkhoff's aesthetic measure using shannon entropy and kolmogorov complexity. *Computational Aesthetics* (pp. 105–112).

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

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

Sarathy, V., & Scheutz, M. (2018). Macgyver problems: Ai challenges for testing resourcefulness and creativity. *Advances in Cognitive Systems*, *6*, 31–44.

Smith, A. M., & Mateas, M. (2011). Answer set programming for procedural content generation: A design space approach. *IEEE Transactions on Computational Intelligence and AI in Games*, *3*, 187–200.

Summerville, A., Snodgrass, S., Guzdial, M., Holmgård, C., Hoover, A. K., Isaksen, A., Nealen, A., & Togelius, J. (2018). Procedural content generation via machine learning (pcgml). *IEEE Transactions on Games*, *10*, 257–270.

Swanson, R., Escoffery, D., & Jhala, A. (2012). Learning visual composition preferences from an annotated corpus generated through gameplay. *2012 IEEE Conference on Computational Intelligence and Games (CIG)* (pp. 363–370). IEEE.

Wiggins, G. A. (2006). A preliminary framework for description, analysis and comparison of creative systems. *Knowledge-Based Systems*, *19*, 449–458.