

A closer look at creativity as search

Graeme Ritchie
Computing Science
University of Aberdeen
Aberdeen AB24 3UE
g.ritchie@abdn.ac.uk

Abstract

Several papers by Wiggins (building on ideas by Boden) have outlined a view of creative concept generation as a very general search process, but that formalisation has not been developed much in the past few years. Also, there are some aspects where clarification or spelling out of details would be useful. We present a re-formulation of the central ideas in Wiggins's framework, with slightly more rigorous statements of the definitions and a number of minor extensions. We also explain how this framework relates to some hitherto completely separate proposals by Ritchie.

Introduction

In recent years, there have been various formalisations of aspects of the computational creative process ((Pease, Winterstein, and Colton 2001), (Colton, Pease, and Ritchie 2001), (García et al. 2006), (Thornton 2007), (Colton, Charnley, and Pease 2011)). Hence there is a consensus among at least some established researchers that it is methodologically beneficial to have fully precise, detailed and formal accounts of any mechanisms being considered as 'creative'.

A prominent example is Wiggins's *creative systems framework* (CSF), presented in a number of papers (particularly Wiggins(2006a; 2006b)). That framework emphasises the notion of *search* as the central mechanism for simulating creativity, and outlines how a metalevel search could represent some phenomena sometimes discussed as 'transformational' creativity. Although these ideas are very helpful in clarifying the nature of creative computation, the published versions of the CSF are at best a preliminary sketch: some details are unspecified, some natural extensions are undeveloped, and there are some formal errors or infelicities. The current paper starts from the central ideas of the CSF, but re-defines the formal mechanisms in a way which leaves fewer gaps, aspires to have fewer formal inconsistencies, and includes the description of more aspects of computationally creative processes. The central motivation for this is that, if we subscribe to a belief in the benefits of formal models (as noted above), then these models should not be left undeveloped, but should continue to be maintained, repaired, and extended as necessary.

It is important to realise that the underlying intuitive ideas – creation as the exploration of a 'conceptual space', and

possible 'transformation' of that space – have been set out in numerous articles by Boden, with many illustrative examples from human creativity. Wiggins's contribution was to take those informal, broadbrush ideas and outline a formal framework which both captured the core notions and made sense computationally. The reader is referred to publications by Boden, Wiggins, and many others for more about the intuitive motivation; our aim here is to refine and extend Wiggins's proposals.

A summary of Wiggins's CSF

Although this paper is centrally concerned with formalisation, we start with a very brief informal overview of the ideas in the existing version of the CSF.

The framework posits a universal set of all *concepts*, a term which covers both abstract ideas (e.g. a mathematical theorem, a design for a better political system) and concrete artefacts (e.g. a painting, a poem). Within this wide-ranging set, there are particular types of idea/artefact (e.g. stories, paintings, poems), and what counts as a recognisable example of a story/painting/poem/etc. may be highly dependent on socio-cultural norms. For many such creative genres, it is not realistic for there to be a firm definition of acceptability, as the specific concept may be *vague* in the sense of (van Deemter 2010). That is, the extent to which a text is or is not a well-formed story (or other artistic category) is a matter of degree, rather than an all-or-nothing decision. Hence, within the CSF, the criterion for acceptability/recognisability is represented as a rating between 0 and 1; in effect, the set of examples of a genre is treated as a *fuzzy set*. As well as whether something falls within the definition of some artistic genre, there is the separate question of whether it is a *good* (high quality) instance (e.g. a profound poem, a beautiful painting). This is similarly a 'vague' notion, and again is represented within the CSF as a score between 0 and 1. This means that an artefact can be an acceptable instance (e.g. recognisably a story) without being of high quality (e.g. it may be a poor story); hence the need for two different ratings (mappings from concepts to values).

The inspiration for the CSF was the work of Boden(1998; 2003), in which creativity was described as occurring within a *conceptual space*, which could be *explored*, or – in more radical creativity – *transformed* into a new space. This view has some resemblance to the traditional ideas of search

within AI (Nilsson 1971), and Wiggins set out both to clarify the relationship between conventional AI search and creative computation, and to provide a formal framework for describing the latter. The idea is that a creative system starts from some set of concepts, and by a series of steps creates further concepts one after another, thus ‘exploring the space’. Although the term ‘creative’ has connotations of ‘construction’, the CSF, following the practice in formalising AI search, regards ‘new’ concepts as not so much ‘constructed’ as ‘reached’. That is, notionally all the possible concepts are elements of some universal set, but the creative system computes a route through that set to particular concepts, and those which have been thus reached represent ‘discoveries’ or ‘inventions’.

The exploration of the space of concepts (the search method) is modelled by an operator which acts on a sequence of concepts (a list of the concepts the system has already processed), and yields as its result a new (presumably longer) list of concepts, which can then be processed in turn as the next cycle of search. Sequences are used because the search method, like an agenda-based AI search system, has to maintain some record of where it has reached within the space of available concepts, and what to work on next.

Wiggins points out that by separating the acceptability rating from the search method, we can describe the situation where different composers, each with a personal way of finding new artefacts (different search procedures) are working within the same style of music (a shared notion of acceptability). A creative agent should be able to recognise something as being a recognisable artefact, or a high quality artefact, without necessarily having a search method that would allow the agent to reach (create) that artefact.

More concretely, this is an intuitively plausible account of certain potentially creative programs. The MCGONAGALL poetry generator (Manurung, Ritchie, and Thompson 2012) uses an explicit search mechanism (a genetic algorithm) to find suitable candidate texts. Each text must at least be syntactically well-formed according to the system’s linguistic grammar; this could be regarded as the acceptability mapping. During the search, items are scored for rhythmic suitability and proximity to an initial semantic message; this would be the quality mapping. At each stage, the system keeps an ordered list of the current candidates, from which each cycle in the search starts.

A small formal detail is that the CSF search operator takes as arguments the two fuzzy criteria (for acceptability and quality), and from there computes a way of going from an existing concept-sequence to a new concept-sequence. This means that the search method can be sensitive to these two criteria if necessary, or that we can describe two systems as having the same search strategy relative to their own different definitions of validity and quality.

For Wiggins, the mappings (the two fuzzy sets and the search mechanism) are to be represented as expressions in some symbolic language, translatable to actual mappings.

Hence, in the CSF, an *exploratory creative system* consists of the following seven components:

- (i) the universal set of concepts
- (ii) the language for expressing the relevant mappings

- (iii) a symbolic representation of the acceptability mapping
- (iv) a symbolic representation of the quality mappings
- (v) a symbolic representation of the search mechanism
- (vi) an interpreter for expressions like (iii) and (iv)
- (vii) an interpreter for expressions like (v)

That constitutes the *object level* of the creative system, which searches through concepts in the domain (e.g. melodies). Wiggins also proposes that there can be a *metalevel*, which searches through possible object level systems to find an interestingly different ‘conceptual space’, thus modelling Boden’s idea of a ‘transforming’ the space. The metalevel in CSF is structured in the same way as the object level (i.e. the seven parts as set out above), except that its set for exploration (i.e. its universal set) contains expressions in the symbolic language used at the object level. In this way, the metalevel searches through expressions describing object-level systems, assessing these descriptions for acceptability and for quality (using the metalevel’s criteria for these two measures).

Relative to the published accounts of the CSF, the revisions or extensions made here are:

- The symbolic language for expressing the various mappings is given a much less central role.
- The way in which the metalevel defines the object level is explicitly stated. In particular, the notion of ‘transformation’ of an (object-level) space is defined.
- Some minor inconsistencies in definitions are removed.
- We outline how search methods within the CSF can be compared at the metalevel.
- The CSF is related to a proposal for formal assessment of creative systems (Ritchie 2001; 2007).

The object level

The structure of an object level system

Wiggins posits the existence of one universal set, \mathcal{U} , the set of all concepts, but then defines a creative system as a tuple, one component of which is the universal set. If the set is truly universal across all systems, it should not need to be mentioned as a defining component of a specific system. On the other hand, it would be useful to be able to allow different creative systems to consider only specific subsets of this hugely general set. The compromise here is to accept the existence of the wholly universal set, but for the definition of each creative system to specify a *subset* of this universe; this could, in principle, be a non-proper subset. We will use \mathcal{P} (mnemonic for ‘possibilities’) for these subsets in our definitions, below. The idea is that \mathcal{U} is universal enough to contain concepts for every type of artefact that might ever be conceived: it includes poems, stories, sculptures, jokes, paintings, theorems, architectural plans, designs for food mixers, etc. On the other hand \mathcal{P} represents the whole range of concepts within some narrower sphere, such as two-dimensional arrays of coloured pixels (which could act as a ‘universal’ set for the creation of visual art), or finite sequences of words and punctuation (a possible ‘universal’ set for various textual artistic forms).

Notation: For any sets A, B , B^A denotes the set of mappings from A to B . In particular, for any set X , $[0, 1]^X$ denotes the set of mappings from X to values between 0 and 1 inclusive. Since a fuzzy set is defined by a mapping from possible elements to values between 0 and 1, our fuzzy sets of ‘acceptable’ elements and of ‘valuable’ elements will be stated in this way; that is, as members of $[0, 1]^{\mathcal{P}}$. We also take $tuples(X)$ to denote the set of finite tuples (of any length) of elements of the set X .

Definition 1: An *exploratory creative system* comprises:

- (i) a subset \mathcal{P} of \mathcal{U} (possible concepts within this type or genre)
- (ii) $\mathcal{N} \in [0, 1]^{\mathcal{P}}$, the *acceptability mapping* (mnemonically, this describes *norms*)
- (iii) $\mathcal{V} \in [0, 1]^{\mathcal{P}}$, the *value mapping* (mnemonically, this describes *value*)
- (iv) a mapping \mathcal{Q} (the *exploration scheme*) from $[0, 1]^{\mathcal{P}} \times [0, 1]^{\mathcal{P}}$ to the set of mappings from $tuples(\mathcal{P})$ to $tuples(\mathcal{P})$ (mnemonically, this describes a *quest* for creations – ‘s’ for ‘search’ is used elsewhere).

The four components in our definition are direct counterparts of those in Wiggins’s CSF, but we have chosen different symbols for the components of a system, to avoid confusion; the relationships to Wiggins’s notation are: $\mathcal{N} \cong \llbracket \mathcal{R} \rrbracket$, $\mathcal{V} \cong \llbracket \mathcal{E} \rrbracket$, $\mathcal{Q}(\mathcal{N}, \mathcal{V}) \cong \langle \langle \mathcal{R}, \mathcal{T}, \mathcal{E} \rangle \rangle$. The intuitive meanings of the components are the same: \mathcal{P} is the set of possible concepts (e.g. arrays of pixels, sequences of words), \mathcal{N} defines the fuzzy set of acceptable instances of whatever domain/genre is being explored, \mathcal{V} indicates how ‘good’ an instance is, and \mathcal{Q} defines how to explore the space.

The component \mathcal{Q} , the search method, may need some explanation. Directly following Wiggins’s proposals, \mathcal{Q} is applied to a particular \mathcal{N} and \mathcal{V} , and from that produces a mapping which takes sequences of concepts into sequences of concepts; hence, \mathcal{N} and \mathcal{V} could in principle influence \mathcal{Q} , or could be ignored. It might seem odd to describe \mathcal{Q} as taking these two parameters, when the only possible values for the parameters seem to be fixed as \mathcal{N} and \mathcal{V} – why not just ‘compile in’ these two values, as they are specified in the same 4-tuple package as \mathcal{Q} ? At present, this level of parameterisation has no real advantage, but it leaves open the possibility, as the framework is elaborated, of considering a ‘transformed’ version of an exploratory creative system in which \mathcal{Q} is unchanged, but one or both of \mathcal{N} , \mathcal{V} are altered, with automatic consequences for the operation of \mathcal{Q} .

As in the original Wiggins formulation, $\mathcal{Q}(\mathcal{N}, \mathcal{V})$ maps from *sequences* of concepts to *sequences* of concepts, providing an agenda-like exploration of the set of possibilities, with the sequence representing the current search state

As noted earlier, the CSF includes a symbolic language in which mappings are expressed as rules, which are then interpreted into mappings. Here, we abstract away from the use of a language, and define a creative system using mappings. The advantage of this is that it states the essential relations within a creative system without regard for representational issues. In a later section, we show how the symbolic language can be incorporated, directly reflecting the Wiggins approach.

Characterising the conceptual space

As noted above, the basic definition of an exploratory creative system contains a fuzzy set, \mathcal{N} , of concepts, which are – intuitively – those concepts which conform to the current norms of the domain. In Wiggins’s formulation, this fuzzy set is *not* regarded as modelling Boden’s ‘conceptual space’. Instead, Wiggins stipulates that conceptual spaces are ordinary subsets (*not* fuzzy) of the universal set \mathcal{U} , and that the conceptual space \mathcal{C} (of a given creative system) consists of all concepts mapped by his $\llbracket \mathcal{R} \rrbracket$ (our counterpart is \mathcal{N}) to values greater than or equal to 0.5. Similarly, Wiggins defines the *valued set* as those concepts which his $\llbracket \mathcal{E} \rrbracket$ (our equivalent is \mathcal{V}) maps to values greater than or equal to 0.5. That is, although the CSF allows both these sets to have graded membership (0 to 1), Wiggins immediately simplifies them to non-graded sets by imposing a threshold.

Within our formulation, the counterpart to Wiggins’s non-fuzzy definitions would be as follows:

Definition 2: Given an exploratory creative system $S = (\mathcal{P}, \mathcal{N}, \mathcal{V}, \mathcal{Q})$, we define, for any $\alpha \in [0, 1]$ and $X \subseteq \mathcal{P}$:

- (i) $\mathcal{N}_\alpha(X) = \{c \in X \mid \mathcal{N}(c) > \alpha\}$ (the set of concepts which reach the threshold α in their ‘normality’).
- (ii) $\mathcal{V}_\alpha(X) = \{c \in X \mid \mathcal{V}(c) > \alpha\}$ (the set of concepts which reach the threshold α in their ‘quality’).
- (iii) the *fuzzy conceptual space* of S is \mathcal{N} (this just confirms the status of \mathcal{N} as outlined earlier).
- (iv) the *conceptual space* of S is $\mathcal{N}_{0.5}(\mathcal{P})$ (this is for backwards compatibility with Wiggins’s 0.5 threshold).
- (v) the *fuzzy valued set* of S is \mathcal{V} (this just confirms the status of \mathcal{V} as outlined earlier)..
- (vi) the *valued set* of S is $\mathcal{V}_{0.5}(\mathcal{P})$ (this is for backwards compatibility with Wiggins’s 0.5 threshold).

Searching

In the CSF, the searching process begins from an initial set of concepts. This is to allow for the situation where the creative system starts from some given concept set, representing the status quo. It is also useful later, when considering metalevel search. In Wiggins’s definitions, exploration always starts from the single totally unspecified concept, \top , representing a situation in which the system has no known concepts already. Here, we generalise this slightly.

If the sequences on which \mathcal{Q} operates are to be like an agenda in conventional AI search, then those sequences should contain only items which have been produced from previous steps of the search (i.e. applications of \mathcal{Q}). This means that the initial agenda has to include every item (concept) which could ever participate in discoveries but is not produced by an application of \mathcal{Q} .

Wiggins defines ‘reachable’ concepts using indefinitely many applications of the search operator. There is a minor slip in his definition, in that repeated applications of $\langle \langle \mathcal{R}, \mathcal{T}, \mathcal{E} \rangle \rangle$ (which corresponds to our \mathcal{Q}) will compute *sequences* (tuples) of concepts, not individual concepts. This is easily remedied (and we can add an intermediate version, for limited search). First we need a minor definition of all the items appearing within a set of tuples:

Definition 3: Given any set of tuples A , we define $elements(A) \equiv \{x \mid \exists \langle y_1, \dots, y_n \rangle \in A, \exists i 1 \leq i \leq n : x = y_i\}$

The Wiggins formalisation defines all the concepts which can be reached with any amount of search, i.e. without limit. Although this case is of theoretical interest, in practice any system will search only to some finite depth, and so we also define the notion of ‘reachable in a fixed number of steps’, relying on the fact that a single application of the search mapping Q corresponds to one step in the search process.

Definition 4: Given an exploratory creative system $(\mathcal{P}, \mathcal{N}, \mathcal{V}, \mathcal{Q})$, and a set $B \subset \mathcal{P}$ of concepts (B is the starting set of concepts for the search):

- (i) the set *reachable from B in m steps* is $\bigcup_{n=0}^m elements(\mathcal{Q}(\mathcal{N}, \mathcal{V})^n(B))$
(i.e. any number of repeated applications of \mathcal{Q} , up to m ; this describes search up to some depth.)
- (ii) the set *reachable from B* is $\bigcup_{n=0}^{\infty} elements(\mathcal{Q}(\mathcal{N}, \mathcal{V})^n(B))$
(i.e. any number of repeated applications of \mathcal{Q} ; this allows any depth of search.)
- (iii) the set of *reachable concepts* is the set reachable from $\{\top\}$. (This matches Wiggins’s notion, where all search starts from a single unspecified concept).
- (iv) the set of *concepts reachable in m steps* is the set reachable from $\{\top\}$ in m steps. (This is a bounded search variant of Wiggins’s ‘start from nothing’ definition.)

In considering the behaviour of a creative system, it is important to know which of its final output (i.e. creations) were provided to it initially and which were computed by the system itself. We can define these thus:

Definition 5: Given an exploratory creative system $(\mathcal{P}, \mathcal{N}, \mathcal{V}, \mathcal{Q})$, a subset B of \mathcal{P} , and a set of concepts K reachable from B , the *discoveries in K* is the set of concepts in $K - B$.

Wiggins defines the set of valued concepts as being those concepts reachable from the undefined concept, \top , which exceed a particular threshold value (0.5) when his ‘value’ mapping (our \mathcal{V}) is applied. That definition can be restated in the terminology here.

Definition 6: Given an exploratory creative system $(\mathcal{P}, \mathcal{N}, \mathcal{V}, \mathcal{Q})$, where RC is the set of reachable concepts, and a value $\alpha \in [0, 1]$:

- (i) the α -valued set of reachable concepts is $\mathcal{V}_\alpha(RC)$.
- (ii) the valued set of reachable concepts is the 0.5-valued set of reachable concepts (i.e. $\mathcal{V}_{0.5}(RC)$); this mirrors Wiggins’s definition.

The metalevel

Structure of the metalevel

For the metalevel, the first matter to be clarified concerns the set of items used for exploration. In the Wiggins papers, this

is the set of possible expressions in a symbolic language \mathcal{L} . An expression in \mathcal{L} is defined earlier as defining a rule-set representing either \mathcal{R} (acceptability rules), \mathcal{E} (value rules) or \mathcal{T} (search rules), with different interpreters applying depending on which of these is intended. If the metalevel is considering single \mathcal{L} -expressions, how does one such expression represent an entire object level system, which contains all three of \mathcal{R} , \mathcal{E} , and \mathcal{T} ? Within the language-based formalisation, a possible response would be to say that \mathcal{L} must contain notation which allows one expression to represent *three* rule sets. If following this path, Wiggins’s definitions of the language interpreters would also have to be patched. As we are separating out the language aspect, we have a different solution.

Here, the object level space is defined by the mappings \mathcal{N} , \mathcal{V} , \mathcal{Q} , so it seems reasonable to have the metalevel searching through triples (N, V, Q) where N, V, Q , are possible values for $\mathcal{N}, \mathcal{V}, \mathcal{Q}$ respectively (and hence are elements of the appropriate sets such as $[0, 1]^{\mathcal{P}}$). The exploration set at the metalevel will be the set of such triples; for brevity here, we will call this set of triples $ECS(\mathcal{P})$ (as it is the set of possible *exploratory creative systems* for \mathcal{P}).

Given a subset \mathcal{P} of \mathcal{U} , a *metalevel creative system* for \mathcal{P} is an exploratory creative system made up of:

- (i) $ECS(\mathcal{P})$ (i.e. this is the metalevel’s set to explore)
- (ii) an element \mathcal{N}^{meta} of $[0, 1]^{ECS(\mathcal{P})}$; this rates potential object-level systems as to how ‘normal’ they are, thus providing a (fuzzy) set of ‘acceptable’ triples (N, V, Q) .
- (iii) an element \mathcal{V}^{meta} of $[0, 1]^{ECS(\mathcal{P})}$; this rates potential object-level systems as to their ‘quality’, thus providing a (fuzzy) set of ‘valuable’ triples (N, V, Q) .
- (iv) a mapping \mathcal{Q}^{meta} from $[0, 1]^{ECS(\mathcal{P})} \times [0, 1]^{ECS(\mathcal{P})}$ to the set of mappings from $tuples(ECS(\mathcal{P}))$ to $tuples(ECS(\mathcal{P}))$; this is structured like the search device \mathcal{Q} in an object-level creative system, but operates on elements of $ECS(\mathcal{P})$ instead of \mathcal{P} .

That is, the structure at the metalevel is exactly parallel to the structure at the object level, as in the Wiggins version. A metalevel has information about what an object level creative system should look like (\mathcal{N}^{meta}) and what would count as a ‘good’ object level system (\mathcal{V}^{meta}). It also contains a way of searching through potential object level systems (\mathcal{Q}^{meta}).

Characterising an object level system

Given a definition of the components of a metalevel, it is essential to then define exactly how the parts of the metalevel characterise an object level system. This is not discussed in detail by Wiggins, but he indicates that the metalevel is to operate (in terms of search, etc.) exactly as an object level creative system.

An object level creative system will ascribe various characteristics to a set of concepts. Each concept will be: rated (by \mathcal{N}) as to how acceptable it is as a member of the conceptual space in question, rated (by \mathcal{V}) as to its value/quality, and defined (by \mathcal{Q}) as either reachable or not. As noted earlier, Wiggins proposes that the ratings by (his equivalents of) \mathcal{N} and \mathcal{V} are turned into non-fuzzy sets using a threshold. However, even then the object level does not characterise a

single object, or a unique set of systems: it defines three independent sets, via \mathcal{N} , \mathcal{V} and \mathcal{Q} . Since the metalevel has exactly the same structure as an object level system, the items which it explores (for Wiggins, expressions in a symbolic language \mathcal{L}) are presumably similarly allocated to 3 sets: the recognisable, the valued and the reachable (and for Wiggins, reachability is always relative to \top , not some specified starting set of items). Hence, the metalevel is assigning (potential) object level systems to these three categories. What the metalevel does not do is characterise a single object level system, or even a unique set of systems. This means that we do not, from the published papers, have a definition of how one object level system is a transformation of another, or how a computation at the metalevel will yield a new object level system – all that the metalevel provides is this tripartite classification. We will remedy this by defining how a metalevel can define (or transform) a specific object level system.

In the next few definitions, we assume two exploratory creative systems $S_{obj} = (\mathcal{P}, \mathcal{N}, \mathcal{V}, \mathcal{Q})$ and $S'_{obj} = (\mathcal{P}, \mathcal{N}', \mathcal{V}', \mathcal{Q}')$, and a metalevel system $\mathcal{S}^{meta} = (ECS(\mathcal{P}), \mathcal{N}^{meta}, \mathcal{V}^{meta}, \mathcal{Q}^{meta})$ for \mathcal{P} . Where the relation ' \neq ' is used here, this allows for the two items in question to have elements in common.

Definition 7:

- (i) S'_{obj} is a revision of S_{obj} using \mathcal{S}^{meta} if S'_{obj} is in the set reachable from $\{S_{obj}\}$ within \mathcal{S}^{meta} , and $S'_{obj} \neq S_{obj}$.
- (ii) for any $\alpha \in [0, 1]$, S'_{obj} is α -valued with respect to \mathcal{S}^{meta} if $\mathcal{V}^{meta}(\mathcal{N}', \mathcal{V}', \mathcal{Q}') \geq \alpha$.
- (iii) S'_{obj} is a transformation of S_{obj} using \mathcal{S}^{meta} if S'_{obj} is a revision of S_{obj} using \mathcal{S}^{meta} and also $\mathcal{N} \neq \mathcal{N}'$.

Here we have taken a 'transformation' to be a revision in which the definition of the conceptual space (acceptable set) changes, as indicated by the condition ' $\mathcal{N} \neq \mathcal{N}'$ '. As this could be true even if \mathcal{N} and \mathcal{N}' differ only on one element, proponents of transformation as a form of *radical* change might wish to enhance this definition.

Relationship to the original CSF

To clarify the amendments we have made to the formalisation, we can compare it with the original version in the papers by Wiggins. As mentioned earlier, the original CSF includes a symbolic language in which the components (the counterparts of our \mathcal{N} , \mathcal{V} , \mathcal{Q}) are expressed. We can add this to our framework by defining a symbolic-language version of an exploratory creative system, with appropriate links to the definitions given above. In order to mimic Wiggins's definitions, we first have to clarify certain aspects which are unclear in the published papers. Sometimes the mapping \mathcal{N} (or what corresponds to this in Wiggins's framework) is represented as a *single* expression in a symbolic language, and sometimes it is said to be a *set of* expressions. Either of these accounts could be made to work, if applied consistently. Here, we have opted for the single expression version, with the observation that the symbolic language could contain connective symbols such as 'conjunction', 'disjunction'

or other logical operators, thereby getting the effect of a *set* of rules in one syntactic expression.

Wiggins's version does not separate clearly the definition of the language from the specific language expressions used in a particular creative system. We have tried to draw this distinction in the next two definitions.

Definition 8: Given a set of concepts \mathcal{P} , a *creative systems language* for \mathcal{P} is a tuple $(\mathcal{A}, \mathcal{L}_{\mathcal{R}}, \mathcal{L}_{\mathcal{T}}, \llbracket \cdot \rrbracket, \langle\langle \cdot \rangle\rangle)$ where:

- (i) \mathcal{A} is a set of symbols, the alphabet.
- (ii) $\mathcal{L}_{\mathcal{R}}$ and $\mathcal{L}_{\mathcal{T}}$ are languages over \mathcal{A} (only 2 are needed because the language $\mathcal{L}_{\mathcal{R}}$ can be used for both the 'acceptability' rules and the 'value' rules, since these both describe fuzzy sets of concepts).
- (iii) $\llbracket \cdot \rrbracket$ is a mapping from $\mathcal{L}_{\mathcal{R}}$ to $[0, 1]^{\mathcal{P}}$ (this is the interpreter which takes an expression in the symbolic language and returns a mapping; that mapping is then a fuzzy set of concepts).
- (iv) $\langle\langle \cdot \rangle\rangle$ is a mapping from $\mathcal{L}_{\mathcal{R}} \times \mathcal{L}_{\mathcal{T}} \times \mathcal{L}_{\mathcal{R}}$ to $tuples(\mathcal{P})^{tuples(\mathcal{P})}$ (this is the interpreter which turns symbolic expressions specifying a search method into an actual mapping to carry out the search).

The above definition (which is closely modelled on Wiggins's proposals) provides the symbolic mechanisms, separately from any particular creative system which might use these representations. The next two definitions state how these mechanisms can be used to define a specific system.

Definition 9: Given a set of concepts \mathcal{P} and a creative systems language $(\mathcal{A}, \mathcal{L}_{\mathcal{R}}, \mathcal{L}_{\mathcal{T}}, \llbracket \cdot \rrbracket, \langle\langle \cdot \rangle\rangle)$ for \mathcal{P} , then a *symbolically represented exploratory creative system* for \mathcal{P} consists of a tuple $(\mathcal{W}_{\mathcal{R}}, \mathcal{W}_{\mathcal{E}}, \mathcal{W}_{\mathcal{T}})$ where:

- (i) $\mathcal{W}_{\mathcal{R}} \in \mathcal{L}_{\mathcal{R}}$; the norms or acceptability rules.
- (ii) $\mathcal{W}_{\mathcal{E}} \in \mathcal{L}_{\mathcal{R}}$; the rules assigning value to items.
- (iii) $\mathcal{W}_{\mathcal{T}} \in \mathcal{L}_{\mathcal{T}}$; rules which define the search method.

Definition 10: Given a set of concepts \mathcal{P} , a creative systems language $(\mathcal{A}, \mathcal{L}_{\mathcal{R}}, \mathcal{L}_{\mathcal{T}}, \llbracket \cdot \rrbracket, \langle\langle \cdot \rangle\rangle)$, and a symbolically represented exploratory creative system $SE = (\mathcal{W}_{\mathcal{R}}, \mathcal{W}_{\mathcal{E}}, \mathcal{W}_{\mathcal{T}})$, then the *exploratory creative system associated with SE* is the tuple $S = (\mathcal{P}, \mathcal{N}, \mathcal{V}, \mathcal{Q})$ where

- (i) $\mathcal{N} = \llbracket \mathcal{W}_{\mathcal{R}} \rrbracket$; i.e. the meaning of this rule expression is the normality mapping.
 - (ii) $\mathcal{V} = \llbracket \mathcal{W}_{\mathcal{E}} \rrbracket$; i.e. the meaning of this rule expression is the value mapping.
 - (iii) $\mathcal{Q}(\mathcal{N}, \mathcal{V}) = \langle\langle \mathcal{W}_{\mathcal{R}}, \mathcal{W}_{\mathcal{T}}, \mathcal{W}_{\mathcal{E}} \rangle\rangle$; i.e. the meanings of these rule expressions give the search mapping.
- (This directly mirrors the arrangement of Wiggins's CSF.)

Using his formalisation, (Wiggins 2006a) provides a number of definitions of specific behaviours that a creative system could display, in terms of what concepts are valued, which can be reached, etc. These terms can all be defined within the formalisation given here, as follows, assuming an exploratory creative system $S = (\mathcal{P}, \mathcal{N}, \mathcal{V}, \mathcal{Q})$, and using some of the terminology already defined above:

Hopeless uninspiration: The valued set of concepts is empty. That is, there are no concepts anywhere within the universal set that meet the 'quality' threshold.

Conceptual uninspiration: $\mathcal{V}_{0.5}(\mathcal{N}_{0.5}(\mathcal{P})) = \emptyset$. That is, there are no concepts within the acceptable ('normal') set of concepts that meet the 'quality' threshold.

Generative uninspiration: The valued set of reachable concepts is empty. That is, there are no concepts which the search mechanism can reach which meet the quality threshold.

Aberration: Where \mathcal{B} consists of exactly those elements in the reachable set which are not in $\mathcal{N}_{0.5}(\mathcal{P})$, *aberration* occurs if $\mathcal{B} \neq \emptyset$. That is, aberration is when the search mechanism goes outside the 'normal' set of concepts. *Perfect aberration* is where $\mathcal{V}_{0.5}(\mathcal{B}) = \mathcal{B}$ (i.e. all the non-normal concepts meet the quality threshold); *productive aberration* is where $\mathcal{V}_{0.5}(\mathcal{B}) \neq \emptyset$ and $\mathcal{V}_{0.5}(\mathcal{B}) \neq \mathcal{B}$ (i.e. just some of the non-normal concepts meet the quality threshold); *pointless aberration* is where $\mathcal{V}_{0.5}(\mathcal{B}) = \emptyset$ (i.e. no non-normal concepts meet the quality threshold).

Evaluating search methods

Ventura's analysis

Ventura(2011) gives an analysis of the limitations of uninformed search strategies in a creative context. His definitions and results are general enough that they should be applicable to the framework here, although there is one small formal point that needs to be stipulated first. Ventura (implicitly) makes the following assumption:

One concept per step: Each formal 'step' in the search process corresponds to the generation of exactly one concept/artefact.

That is, Ventura's analysis does not allow for intermediate computational steps behind the scenes which do not directly correspond to the generation of an artefact. The Wiggins definitions (and our reformulations) do not demand this restriction, but it is a plausible constraint, and could be formalised thus:

Given an exploratory creative system $(\mathcal{P}, \mathcal{N}, \mathcal{V}, \mathcal{Q})$, \mathcal{Q} is a *one-concept-per-step* scheme if, whenever $z' = \mathcal{Q}(\mathcal{N}, \mathcal{V})(z)$, $\exists c' \in \mathcal{P}$ such that:

- (i) c' is an element of z' ;
- (ii) c' is not an element of z ;
- (iii) for every element c of z' where $c' \neq c$, c is an element of z .

This perspective could be taken even further, by revising our definition of an exploratory creative system to include a set \mathcal{OP} of *operators*, which are mappings from tuples of concepts to concepts; that is, each operator is a member of $\mathcal{P}^{\mathcal{P}^k}$ for some integer k . Then we would stipulate that each step in a search meets the constraint that it corresponds to the invocation of one operator:

If two concept-sequences $\langle c_1, \dots, c_n \rangle, \langle d_1, \dots, d_m \rangle$ are such that $\mathcal{Q}(\mathcal{N}, \mathcal{V})(\langle c_1, \dots, c_n \rangle) = \langle d_1, \dots, d_m \rangle$, then:

there must be an operator $p \in \mathcal{OP}$, and concepts $\langle e_1, \dots, e_k \rangle$ (where $k \leq n$) such that for every $1 \leq i \leq k$, $e_i = c_j$ for some j , and $p(e_1, \dots, e_k) = d_r$ for

some $d_r \in \{d_1, \dots, d_m\}$, and $d_r \notin \{c_1, \dots, c_n\}$, and $\forall d_i \in \{d_1, \dots, d_m\}$, either $i = r$ or $d_i \in \{c_1, \dots, c_n\}$.

Ventura's analysis provides one possible formalisation of the intuitive notion of a search strategy being 'better' (or 'best'). It posits a set of target elements (concepts, in our terminology) and considers the probability of the search reaching one of these elements. In a footnote, Ventura also offers a definition where the desirability of elements is a function to $[0, 1]$ (cf. our \mathcal{V}), and computes the probability of reaching an element with a maximal value.

It is arguable that in the area of creative systems, there is less emphasis on finding specific target items (or even finding a maximal-scoring item) and more on generally reaching acceptable or highly valued concepts (a direction in which Ventura's footnote moves). Our formalisation allows an alternative perspective on the assessment or comparison of search methods; see the next subsection.

Comparing search

Our formulation of the CSF allows the comparison of two search methods according to how the concepts they reach are rated by the related mappings \mathcal{N} and \mathcal{V} , taking into account the depth of search involved. As we will want to apply certain definitions to various mappings (including \mathcal{N} and \mathcal{V}), we will start with a general schematic definition in which the mapping \mathcal{F} can be anything of the appropriate type (so \mathcal{F} is not mnemonic for anything, being just a placeholder for now). Also, *AGG* will stand for either *AVG* (arithmetic mean) or *MAX* (maximum) of a function applied to a set.

Definition 11: Suppose there are two exploratory creative systems $(\mathcal{P}, \mathcal{N}, \mathcal{V}, \mathcal{Q}_1)$ and $(\mathcal{P}, \mathcal{N}, \mathcal{V}, \mathcal{Q}_2)$, with $S_i(k, k')$ = the set of concepts reachable in no fewer than k and no more than k' steps in these two systems ($i = 1, 2$). Also, $\mathcal{F} \in [0, 1]^{\mathcal{P}}$ (i.e. a rating of concepts, of some sort). Then

- (i) \mathcal{Q}_1 has higher *AGG* \mathcal{F} -values up to k' steps than \mathcal{Q}_2 if $AGG(\mathcal{F}, S_1(0, k')) > AGG(\mathcal{F}, S_2(0, k'))$.
- (ii) \mathcal{Q}_1 has higher *AGG* \mathcal{F} -values beyond k steps than \mathcal{Q}_2 if $AGG(\mathcal{F}, S_1(k, k')) > AGG(\mathcal{F}, S_2(k, k'))$ for any $k' > k$.

This compares two variants of a system in which only the search method \mathcal{Q} is different. The above definitions will be applied, below, to specific values for \mathcal{F} .

Definition 12: Given a two exploratory creative systems $(\mathcal{P}, \mathcal{N}, \mathcal{V}, \mathcal{Q}_1)$ and $(\mathcal{P}, \mathcal{N}, \mathcal{V}, \mathcal{Q}_2)$:

- (i) \mathcal{Q}_1 is higher valued on average up to k steps than \mathcal{Q}_2 if \mathcal{Q}_1 has higher *AVG* \mathcal{V} -values up to k steps than \mathcal{Q}_2 .
- (ii) \mathcal{Q}_1 is more normal on average up to k steps than \mathcal{Q}_2 if \mathcal{Q}_1 has higher *AVG* \mathcal{N} -values up to k steps than \mathcal{Q}_2 .
- (iii) \mathcal{Q}_1 achieves higher value up to k steps than \mathcal{Q}_2 if \mathcal{Q}_1 has higher *MAX* \mathcal{V} -values up to k steps than \mathcal{Q}_2 .
- (iv) \mathcal{Q}_1 achieves greater conformity up to k steps than \mathcal{Q}_2 if \mathcal{Q}_1 has higher *MAX* \mathcal{N} -values up to k steps than \mathcal{Q}_2 .

For each of the 4 subparts of the above definition, we can frame a corresponding definition which says that there exists some depth k' after which one of the search methods \mathcal{Q}_i gives a higher value than the other; e.g.:

Q_1 is *higher valued eventually* than Q_2 if there is some integer $k' > 0$ such that Q_1 has higher AVG \mathcal{V} -values beyond k' steps than Q_2 .

Similar substitutions can be made in the other definitions.

In this way, we have several ways of describing one search method Q_1 as being ‘better’ than another, Q_2 . Next, we consider how comparisons of search methods can be more detailed.

Descriptive criteria

Ritchie(2001; 2007) defines a set of formal criteria which can be used to describe aspects of a potentially creative system’s behaviour. Central to these formal criteria are two *rating schemes* for assigning values in $[0, 1]$ to elements of the set of *basic items* (i.e. the set of possible artefacts). One rating scheme (*typ*) represents *typicality*, indicating the extent to which an item lies within the norm for this type of artefact. The other rating (*val*) is for *value*, and indicates the ‘quality’ of an item. Ritchie’s *typ* and *val*, Wiggins’s \mathcal{R} and \mathcal{E} and our \mathcal{N} and \mathcal{V} all appear to capture the same intuitive notions: that we can rate possible creations as to their membership of a concept set, and in terms of the quality of such creations.

There are some differences of nuance between Ritchie’s constructs and those in the CSF, to which we will return later, but for the moment let us consider how the central ideas in some of Ritchie’s criteria could be used within the CSF as stated here.

The first eight of Ritchie’s criteria are stated in terms of a *result set*, R , which is the set of artefacts produced by the computer program, and the two ratings *typ* and *val*. There is not space here to reproduce them all, but Criterion 7 illustrates the general idea:

$$ratio(V_{\gamma,1}(R) \cap T_{0,\beta}(R), T_{0,\beta}(R)) > \theta, \\ \text{for suitable } \beta, \gamma, \theta.$$

where $V_{\gamma,1}(R)$ is the set of elements of set R which are rated above γ by *val*, $T_{0,\beta}(R)$ is the set of elements of R which are rated below β by *typ*, and *ratio* computes the ratio of the sizes of two sets. That is, this computes the proportion of the untypical items which are of good quality.

Criteria like this could be applied to a creative system $(\mathcal{P}, \mathcal{N}, \mathcal{V}, \mathcal{Q})$, using \mathcal{N} to define $T_{i,j}$ and \mathcal{V} to define $V_{i,j}$. There are various ways in which the result set R could be defined in terms of reachable concepts: all reachable concepts? concepts reachable from a starting set B ? concepts reachable after some number of steps k ? All of these are plausible models of a ‘result set’. Hence there would be a few families of very similar formula, parameterised according to starting set or number of steps.

Ritchie(2007) emphasises that these criteria are not all measures of creative success, but can be used to ‘profile’ a (potentially) creative program by describing its behaviour in more detail. In the same way, they could give a more detailed picture of a creative system, in the CSF sense.

Ritchie also postulates an *inspiring set*, I , which are the existing artefacts upon which the design of the creative program was based. The remaining criteria (9 - 18 in Ritchie(2007)) make comparisons of different sorts between

I and the result set R . There is no exact counterpart within the CSF, as there is no ‘design’ stage in the formalisation. However, the formal structure of Ritchie’s criteria which involve I could be coerced into service within the CSF, by replacing I with some initial set of concepts B , from which search starts. That is, where Ritchie has a criterion such as:

$$ratio(V_{\gamma,1}(R - I) \cap T_{\alpha,1}(R - I), (R - I)) > \theta, \\ \text{for suitable } \alpha, \gamma, \theta$$

(informally, ‘a high proportion of the novel results are both highly valued and very typical of the genre’) we would have:

$$ratio(V_{\gamma,1}(R - B) \cap T_{\alpha,1}(R - B), (R - B)) > \theta, \\ \text{for suitable } \alpha, \gamma, \theta.$$

where R is the set of concepts reachable from B . (Again, there is a possible variant where a number of steps k is stipulated.)

Although this seems to indicate that we can simply port the Ritchie criteria into the CSF, there is one further issue to consider: there is a difference in the overall perspective of the two formal accounts. There are various viewpoints one could assume in devising a formal model in this area. For example, it would be possible to have an abstract declarative formalisation of the nature of the creative task, without including any details of how this task might be executed. Or a model might be proposed as describing (at some suitably abstract level) *how* a creative system operates. Ritchie’s criteria arguably take a third viewpoint, in which one treats the program as a conveyor of input/output data, and attempts, from an external viewpoint, to say more precisely how it has performed. The *typ* and *val* mappings are certainly not proposed as components of the program or system. Instead, these are measures which might be applied by, for example, having humans judge the program’s output.

In Wiggins’s CSF, the formal definitions of \mathcal{R} and \mathcal{E} (the symbolic counterparts of our \mathcal{N} , \mathcal{V}) would be compatible either with a characterisation of the abstract nature of creativity, or with a model of a creative system. However, the terminology used, and the inclusion of the \mathcal{T} (‘agenda’) mapping (our \mathcal{Q}) determine that this is a model (at a very abstract level) of the working of a creative system. Hence, the whole intent of the CSF is radically different from that of Ritchie’s definitions, even though there is a clear parallel between the intuitive meanings of \mathcal{N} and *typ*, \mathcal{V} and *val*.

This means that what we have sketched above is *not* the direction application of Ritchie’s criteria within the framework, but the definition of *counterparts* within the CSF, where the $T_{i,j}$ and $V_{i,j}$ mappings are defined using internal components (\mathcal{N} , \mathcal{V}) of the system, not external judgements. However, this means that if we are scrutinising the behaviour of a creative system, we can apply the conditions outlined above (i.e. the counterparts of Ritchie’s criteria) in distinct ways:

Internal: How is the system performing, in its own terms?

For this, we use \mathcal{N} and \mathcal{V} to define $T_{i,j}$ and $V_{i,j}$.

External: How is the system performing, in terms of independent measures such as human judgements? For this,

we use the independent measures to define $T_{i,j}$ and $V_{i,j}$.

In both of these, we retain the notions of B (initial set) and R (reachable concepts) discussed above. Given that the set of reachable concepts is in effect the ‘result set’, if the inspiring set I is known, then using I instead of B , with an External perspective, is effectively the scenario in the Ritchie papers.

These various adaptations of the criteria to the CSF can be viewed alongside the definitions in our section **Comparing search** above, and provide a slightly finer grained and more detailed vocabulary for comparing search strategies.

Guiding the metalevel

We have already shown how the metalevel of a creative system can start from an existing object level system and search for a variant system, using a metalevel value function \mathcal{V}^{meta} . What should the content of \mathcal{V}^{meta} be? Since the metalevel has access to the object level mappings \mathcal{N} and \mathcal{V} , it would be possible for \mathcal{V}^{meta} to be defined using composite criteria such as those we have outlined above, the counterparts of Ritchie’s criteria. That is, the metalevel search could be guided by how candidate object-level systems performed according to these criteria. For this, the distinction between internally-parameterised and externally-parameterised versions of the criteria is important. Whereas the externally-parameterised version (using human judgements or other measures) is exactly appropriate for profiling or assessing the success of the object-level system, the internally-parameterised version (using \mathcal{N} and \mathcal{V}) are the only ones that make sense within the creative (metalevel) system itself.

This glosses over the significant question of whether a real creative program would be implemented with the meta/object strata of the CSF, or whether the formal framework is only a way of describing, at some fairly abstract level, what creative systems do (or could do). It is possible that the actual use of structured criteria which compare ratings of initial sets and of reachable concepts would not be realistically applicable in implemented systems.

Summing up

We have presented a reformalisation of Wiggins’s CSF, which:

- makes the use of a symbolic language an optional extra
- indicates how search strategies can be formally compared within the CSF
- clarifies the metalevel, defining some metalevel constructs in more detail and making explicit some formal comparisons.
- shows how Ritchie’s criteria can be adapted, in a number of ways, to the CSF formalisation, thereby clarifying the relationship between these frameworks

In this way, we have extended the development of formal descriptive frameworks for creative systems.

Acknowledgments

The author is very grateful to Geraint Wiggins for detailed discussions of this material, and for comments on a draft of this paper.

References

- Boden, M. A. 1998. Creativity and Artificial Intelligence. *Artificial Intelligence* 103:347–356.
- Boden, M. A. 2003. *The Creative Mind*. London: Routledge, 2nd edition. First edition 1990.
- Colton, S.; Charnley, J.; and Pease, A. 2011. Computational creativity theory: The FACE and IDEA descriptive models. In Ventura et al. (2011), 90–95.
- Colton, S.; Pease, A.; and Ritchie, G. 2001. The effect of input knowledge on creativity. In Weber, R., and von Wangenheim, C. G., eds., *Case-Based Reasoning: Papers from the Workshop Programme at ICCBR 01*.
- García, R.; Gervás, P.; Hervás, R.; and Pérez y Pérez, R. 2006. A framework for the E-R computational creativity model. In *5th Mexican International Conference on Artificial Intelligence (MICAI-06)*, volume 4293 of *LNAI*, 70–80. Springer Verlag.
- Manurung, R.; Ritchie, G.; and Thompson, H. 2012. Using genetic algorithms to create meaningful poetic text. *Journal of Experimental & Theoretical Artificial Intelligence* 24(1):43–64.
- Nilsson, N. J. 1971. *Problem-solving methods in artificial intelligence*. New York: McGraw-Hill.
- Pease, A.; Winterstein, D.; and Colton, S. 2001. Evaluating machine creativity. In Weber, R., and von Wangenheim, C. G., eds., *Case-Based Reasoning: Papers from the Workshop Programme at ICCBR 01*, 129–137.
- Ritchie, G. 2001. Assessing creativity. In *Proceedings of the AISB Symposium on Artificial Intelligence and Creativity in Arts and Science*, 3–11.
- Ritchie, G. 2007. Some empirical criteria for attributing creativity to a computer program. *Minds and Machines* 17(1):67–99.
- Thornton, C. J. 2007. How thinking inside the box can become thinking outside the box. In *Proceedings of the 4th International Joint Workshop on Computational Creativity*.
- van Deemter, K. 2010. *Not Exactly: in praise of vagueness*. Oxford, UK: Oxford University Press.
- Ventura, D.; Gervás, P.; D. Fox Harrell; Maher, M. L.; Pease, A.; and Wiggins, G., eds. 2011. *Proceedings of the 2nd International Conference on Computational Creativity*.
- Ventura, D. 2011. No free lunch in the search for creativity. In Ventura et al. (2011), 108–110.
- Wiggins, G. 2006a. A preliminary framework for description, analysis and comparison of creative systems. *Knowledge-Based Systems* 19:449–458.
- Wiggins, G. A. 2006b. Searching for computational creativity. *New Generation Computing* 24(3):209–222.