

Dr. Chris Thornton

Word counts: abstract 120 words, main text 8553, references 600, entire text
9500

A Quantitative Reconstruction of Boden's Creativity Theory

Chris Thornton
Department of Informatics
University of Sussex
Brighton
BN1 9QH
UK

Email: c.thornton@sussex.ac.uk
Tel: (44)1273 678856

October 31, 2005

1 Short abstract

Boden's theory of creativity explains the creative process in terms of exploration in, and transformation of generatively-represented conceptual spaces (Boden, 1990, 1998, 2003). An advantage of the theory is that its terms of reference are ordinary computational processes, readily understood in terms of conventional models. Her core definition however, which attempts to identify creativity specifically with transformation (generation of new spaces), has been questioned, since both transformation and exploration are ultimately applications of search. This paper attempts to follow the suggestion of one reviewer (Bundy, 1994) to re-cast the definition in quantitative terms, using concepts of computational complexity. The approach involves giving the complexity of Boden's search processes a formal evaluation and using this as the foundation for a revised definition.

2 Long abstract (240 words)

Boden's theory of creativity explains the creative process in terms of exploration in, and transformation of generatively-represented conceptual spaces. An advantage of the theory is that its terms of reference are ordinary computational processes, readily understood in terms of conventional models. Her core definition however, which attempts to identify creativity specifically with transformation (generation of new spaces), has been questioned since both transformation and

exploration are ultimately applications of search. Boden's aim was to present a mechanistically precise theory of creativity offering greater explanatory power than Koestler's 'matrix bisociation' framework. While her own exposition was no more formal than Koestler's, she hoped to achieve explanatory precision by mapping out the ways in which existing artificial intelligence systems exemplified the processes she was describing. In practice, however, the complexity and opacity of her example systems meant that they were not able to play the envisaged role. The debate surrounding Boden's theory focussed on her own, qualitative descriptions of process and revolved particularly around the difficulties arising from its lack of formal precision. This paper attempts to move the debate forward by providing a *quantitative* version of Boden's account. Following a suggestion from Bundy, a reviewer of the initial target article, the approach involves the establishment of a method for measuring the computational complexity of conceptual development. This is then used for purposes of grounding a revised version of Boden's definition.

Keywords: complexity, conceptual space, creativity, representation

3 Introduction

While there is a long tradition of introspective analysis of creativity (e.g. Wallas, 1926, Perkins, 1981), recent years have seen the emergence of more principled and objective accounts. Of these, Arthur Koestler's 'The Act of Creation' (1964) was an early example. This presented a theory that was both general and mechanistic. More recently, Margaret Boden published 'The Creative Mind: Myths and Mechanisms' (1990). This presented a theory of creativity that was not only general and mechanistic, but also computational.

Unlike Koestler's theory, where the envisaged mechanism was idiosyncratic, Boden's theory dealt in terms of conventional computational structures and processes. Moreover, her theory described mechanisms which could, in principle, be constructed using standard computational devices. This gave Boden's theory a clear advantage. As Boden notes 'The thought processes [Koestler] described do happen, and they do seem to be involved in creativity. But because *how* they happen was not detailed, he did not fully explain how creativity is possible.' [ibid, p. 24]

Boden's starting point for the development of her account is the observation that the concept of creativity contains a paradox. By definition, creativity *creates*, i.e., it produces something new. But if we are committed to a mechanistic account of the world — no miracles allowed — we believe that everything that occurs is predictable in principle. We also believe that any new thing must be constructed from existing components. This implies that nothing can ever be *intrinsically* new. How, then, should we reconcile the definitional requirement that creativity produces novelty with the assumption that there is no such thing as novelty? Boden's aim is to deal scientifically with this question.

She takes it as given that creativity is always mediated by conceptual development of some form. Any creative act is thus always founded on conceptuali-

sation or the realisation of a ‘point’ within a particular ‘conceptual space’. But, she notes, there are two ways in which this might happen. If the conceptual space has an existing mental representation, realisation of a new point is simply a matter of identifying a new location within that space. If no such representation exists, then realisation of a new point necessarily involves construction of the representation as a preliminary step.

For Boden, this offers the means of distinguishing two forms of conceptualisation: a straightforward form, involving the identification of a new point in an existing space, and a more complex form which involves as a preliminary step the construction of the relevant conceptual space. The process of identifying a new point she terms *exploration*; the process of generating a new space she calls *transformation*. Her key idea is then to use the distinction between exploration and transformation to define what is to count as true creativity.

‘We can now distinguish first-time novelty from radical originality’ she writes. ‘A merely novel idea is one which can be described and/or produced by the same set of generative rules¹ as are other, familiar ideas. A genuinely original, or creative, idea is one which cannot.’ (Boden, 1990, p. 40) A new concept, then, is to be considered genuinely creative just in case its construction involved some element of transformation.

For Boden, this way of framing the definition effectively resolves the creativity paradox. A concept whose construction involves some element of transformation cannot be generated on the basis of existing mental representations. In this sense, it is mentally ‘impossible’. Boden then proposes we can justify saying a concept is ‘new’ on the grounds that it was previously ‘impossible’. Or, as she puts it ‘To justify calling an idea creative ... one must identify the generative principles with respect to which it is impossible.’ [ibid., p. 40]

With this transformation-based definition, Boden establishes an intuitively clear characterisation of the process of creativity. However, consideration of further evidence compels her to take steps to generalise the characterisation in several ways. First, she observes that ordinary exploration of existing conceptual spaces may also be creative. With regard to mathematics for example, she notes how the ‘creative mathematician explores a given generative system, or set of rules, to see what it can and cannot do.’ [ibid., p. 45]. The recognition that exploration may be creative also surfaces in her comment that ‘creativity is a matter of using one’s computational resources to explore, and sometimes to break out of, familiar conceptual spaces.’ [ibid., p. 108] However, it is not Boden’s intention to broaden the definition to include brute-force (uninformed) exploration. Exploratory processes are only to be considered creative if they are guided in some way by ‘heuristics’ or ‘maps’.

Boden also makes a point of allowing that transformation may constitute creativity in its own right, regardless of any ensuing concept development. This is evidenced in her comments on creative mathematics: ‘By creative mathematics, I do not mean adding 837,921 to 736,017 to get 1,573,938 ... I mean

¹For reasons that will be explained, it can be assumed that mental representation of a conceptual space must be generative in character.

producing new generative systems, new styles of doing mathematics.’ [ibid. p. 45] The construction of new generative systems, then, would appear to fall into the category of creative processes, regardless of whether any concepts or ideas within that new space are ever actuated.

The indications are, then, that Boden wishes to generalise her original definition so as to allow that creativity may involve either

- (1) *guided exploration*, the use of heuristics and maps to identify valuable concepts within an existing conceptual space, or
- (2) *transformation alone*, the development of new conceptual spaces (i.e., new generative systems) in which useful exploration may take place.

Indeed, the fact that she subsequently refers to the original formulation as the ‘strong definition’ [ibid, p. 49] would seem to confirm that we should treat the broadened formulation as a ‘weak’ or generalised alternative.

However, caution is in order. In neither edition of her book does Boden go so far as to explicitly differentiate a ‘weak’ from a ‘strong’ definition. In fact, in the first edition, she offers no final count of the number of different types of creativity she has identified. It seems to be her intention to distinguish the two types noted, and this is certainly a common interpretation. Yet in the ‘nutshell’ summary of her theory, added as a prologue to the second edition (Boden, 2003), and in (Boden, 1998), she states that her account distinguishes *three* main forms of creativity, these being exploration, transformation and *combination*. Adding to the uncertainty is the observation that only the strong definition has the power to resolve the creativity paradox, arguably forcing us to recognise not two forms of creativity, or three, but one: transformation.²

There is a choice to be made, then, regarding the definition. There is the original formulation, which seeks to identify creativity specifically with transformation. Then there is the more general characterisation which allows that exploration may figure. Finally, there is the view from the post-1998 publications which adds ‘combination’ as a possibility.

The original formulation is clearly stated, but as Boden herself has shown, not completely general. Her broadened characterisation on the other hand may be more general, but it is too all-encompassing to be useful as a definition. As for the inclusion of ‘combination’ as a separate type of creative process (as suggested in the second edition) this would seem to be a matter of explanatory expedience. As Richie notes, the distinction between *combinational* and *exploratory* is ‘hard to pin down’ (Richie, 2001). Indeed, any combinational process must presumably operate through exploration of a space of possible combinations. So it seems reasonable to view combinational processes as subsumed within the exploratory category.

²Boden’s own remarks about the definition tend to heighten the element of doubt. The characterisation is ‘vague’ she writes, and no more than ‘intuitive talk’. For her, it cannot stand on its own: ‘Anyone hoping for a scientific explanation of creativity must be able to discuss mental spaces, and their exploration more precisely.’ [ibid. p. 73]

Insofar as there *is* a definition underpinning the theory, then, it has to be the ‘strong definition’. But this raises a difficulty. Even taking on board the withdrawal of any generality claim, the definition cannot be regarded as a firm foundation for the theory. It contains an ambiguity which undermines its application in practice. Although the two processes Boden describes appear to be distinct and separate, in computational terms they are not. They are both applications of search. (Wiggins, 2001)

Boden herself notes (Boden, 1995, p. 161) that many of the reviewers of her account have pointed to problems with the definition, and specifically to the problem of ambiguity. Samples are to be found in Fetzer (1994) O’Rourke (1994), Bundy (1994), Ram et al (1994) and (Ram *et al.* 1995) Treisman (1994) and Weisberg (1994). Several of these writers have noted the fact that transformation is really just a type of search, so would seem to logically belong in the same category of process as exploration (e.g. Ram *et al.* 1995, p. 114). Wiggins has used a rationalisation of Boden’s theory to demonstrate the same point more formally (Wiggins, 2001). Some have noted that transformation is a form of ‘meta-search’. Schank and Forster describe application of the definition as a ‘mare’s nest’ (Schank and Forster, 1995, p. 138) while Turner (Turner, 1995, p. 152) notes that while Boden wants to view computers as potentially creative, by her ‘definition of creativity, it is impossible for a computer program to be creative.’ (Turner, 1995, p. 153)

In her responses to this criticism, Boden pleads guilty to the charge that the definition is insufficiently clear. ‘Most of the reviewers point out, quite rightly,’ she says, ‘that my definition of creativity ... was vague.’ (Boden, Behavioral Reviewers, 1994, p. 559) Her defence 1994, p. 559) Her defence is that the necessities of writing for a general audience dictated the limiting of technical content. But, to be fair to the reviewers, the problem they highlight is not so much that the definition is vague, but rather that it has no computational foundation.

According to the proposed rule, any thought which is classified as creative on the grounds that its construction has involved the creation of a new conceptual space, may also be regarded as *uncreative*, on the grounds that the creation of that space itself involved ordinary search (or *meta-search* as we might call it) in another, over-arching conceptual space. Thus, classification done using Boden’s definition will generally produce two, contradictory results: one positive, the other negative.

In later writings, Boden concedes there is a problem but suggests that it results from mis-application of the criterion. ‘If we are to apply [my definition] to the entire resources of the person’s mind (or the computer’s program), of course we get an inconsistency’ she says. (Boden, 1995, p. 163) But the definition’s inconsistent results do not come about as a result of mis-application. They result from the fact that transformation, as a process, cannot easily be distinguished computationally from exploration, since both are forms of search. Search is just the generation of candidates in a space of possibilities and as Bundy has noted ‘Transformation is [simply] a kind of generation.’ (Bundy, 1994, p. 534)

There is a deep problem, then, with the definition. But it would be a mis-

take to let these difficulties undermine the theory, which has proved its value in many ways. The essence of Boden’s definition is that creativity involves some sort of special, representational effort. And there are other ways in which we might carve out what this ‘extra effort’ might involve. Wiggins has proposed the idea that it might take the form of *meta-traversal* of conceptual space representations (Wiggins, 2001), arguing that this might form the basis for a different interpretation of the meaning of ‘transformation’. An alternative approach involves re-working the definition in terms of some notion of complexity, i.e., a computational measure of ‘effort’. This line of attack was suggested by one of the initial reviewers (Bundy, 1994) but, to my knowledge, has never been followed up.

Doing so is the specific aim of the present paper. The plan will be as follows. The next section will review the core concepts in Boden’s theory and clarify their technical aspects. Following this, there will be an analysis of the complexity of conceptual-space search. Finally, the complexity analysis will be used for purposes of developing a quantitative reconstruction of Boden’s creativity definition.

4 The structure of generative conceptualisation

To develop a complexity-based reformulation of Boden’s model, the core ideas underlying the account need examination. To start, Boden’s key notion of ‘conceptual space’ needs to be reviewed.

A conceptual space in her terms is simply a set of associated concepts, assumed to be represented generatively. In computational terms, the conceptual space is a *virtual set* of objects representing concepts. But the generative nature of the representation is essential. Any form of explicit, ‘point-by-point’ representation for conceptual spaces would be impractical and, in the case of an infinite space, impossible. The generative approach overcomes these difficulties. As Boden comments ‘a generative system [has] ... the potential (in principle) to generate every location within the conceptual space. The number of these locations may be very large, even infinite.’ [ibid. p. 78]

In operation, any generative system follows a particular methodology. To a set of originating objects — the ‘prototypes’ or ‘primitives’ — it applies processes which have the effect of combining those objects to form new objects. If there is the potential for recursion, these objects may then become constituents in the construction of further objects and so on. Any generative system must therefore embody (a) a set of originating objects and (b) functionality for the construction of new objects from existing objects.

Programming languages provide a convenient framework for the implementation of such systems. For example, consider this sequence of instructions:

```
To each number between 0 and 9 apply the next 2 instructions
To each number between 0 and 9 apply the next instruction
Print the first number followed by the second number
```

This program has the effect, when ‘executed’, of printing out all 100 points in a 10x10 2-dimensional space; i.e.,

```
0 0, 0 1, 0 2, 0 3, 0 4, 0 5, 0 6, 0 7, 0 8, 0 9
1 0, 1 1, 1 2, 1 3, 1 4, 1 5, 1 6, 1 7, 1 8, 1 9
2 0, 2 1, 2 2, 2 3, 2 4, 2 5, 2 6, 2 7, 2 8, 2 9
.
. <60 pairs omitted>
.
9 0, 9 1, 9 2, 9 3, 9 4, 9 5, 9 6, 9 7, 9 8, 9 9
```

The program provides a generative representation for the 2-dimensional space based on the integer range 0..9. The primitive objects are the integers. The construction method is the procedure for printing pairs of digits.

Boden’s preferred example of a generative representation is the grammar. And indeed, grammar notation is a particularly convenient way of detailing generative systems. In writing out a grammar we only have to specify the bits of information which are essential for object construction — *what* object results from *what* combination. Everything else is implicit.

Consider this:

```
S -> NP V
NP -> Det N

Det -> the | a | one
N -> man | woman
V -> ran | walked | sang | laughed
```

The first two lines of this grammar are rewrite rules showing how the object named on the left of the arrow can be constructed by combining objects named on the right. The remaining three lines are lexical rules which show how each of the words on the right is an example of the object named on the left. We may think of the whole as a generative representation for all the sequences of words which result from all the different ways of constructing an ‘S’ object. These include ‘the man ran’, ‘the man walked’, ‘the man sang’, ‘the man laughed’, ‘the woman ran’, ‘the woman walked’ and so on.

For present purposes, we are interested in generative representation of conceptual spaces rather than languages. These have special features which go beyond what we expect to find in a grammar or program representation. As noted, any generative system must embody a set of originating objects and functionality for the formation of new objects from combinations of existing objects. In the case of the conceptual space the constructed objects must obviously be concepts. But what of the originating objects? Ultimately, these are simply ‘phenomena’; but the system’s representation of any phenomenon can only be itself a concept. Originating and constructed objects are thus all of the same type — they are all just concepts.

5 Shallow vs deep conceptual space

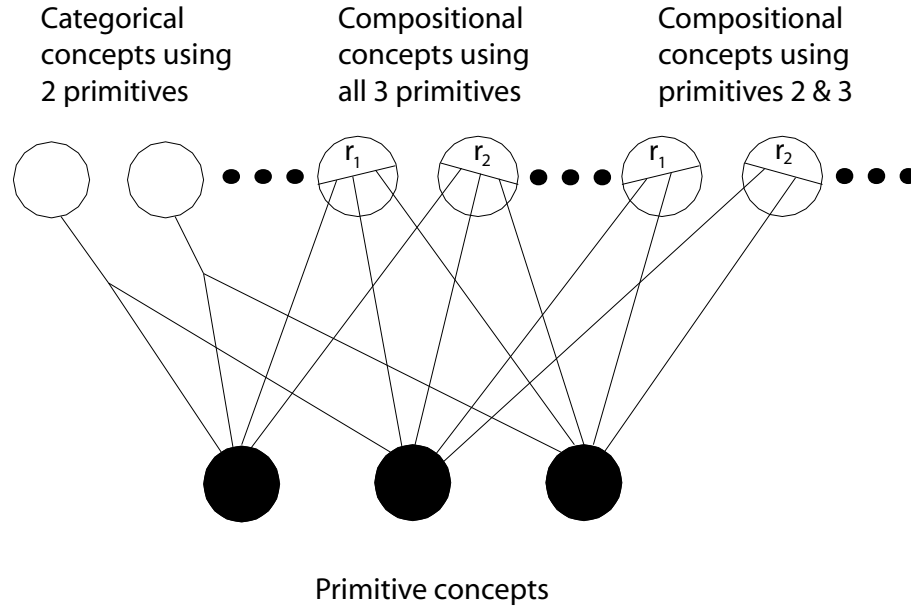


Figure 1: Possible forms of conceptualisation.

A generative representation of a conceptual space consists, then, of a set of originating concepts together with functionality for forming new concepts from combinations of existing concepts. But there are two quite different ways in which this may happen. The combination may be done *compositionally*, in which case the existing concepts become components of a new whole; or it may be done *categorically*, in which case they are brought together to form a class. Concept combination may therefore exploit composition or categorisation. In addition, there may be multiple ways of producing compositional combinations of concepts. (e.g., multiple ways to place the concepts into a relationship).

In the language of logic, categorisation is simply the arrangement of the concepts as a *disjunction*. Composition, on the other hand, is the arrangement of the concepts in a relationship or function such that each concept has a constituent role. In a simple case this might mean the formation of a *conjunction* (i.e., application of the AND relationship). But the total number of compositional possibilities depends on the number of relationships/functions which can be brought to bear.

Figure 1 illustrates the ways in which concept combination may occur. The filled circles at the bottom of the figure represent the primitive or originating concepts. Empty circles in the layer above represent categorical concepts. Circles containing internal structure represent compositional concepts. In this

scenario, the conceptualiser is assumed to be capable of forming two distinct forms of compositional concept, one based on the relationship r_1 and the other based on the relationship r_2 . The organisation of the intersections of the connecting lines in these concepts is intended to represent the way in which the concept imposes a particular relationship over the components. In the case of the categorical concepts the arcs combine, representing the fact that the components in this case are treated as members of a class of alternatives. In the case of compositional concepts, the arcs are independent, representing the fact that the components are *constituents*.




Figure 2: Possible avenues of concept development.

Using the distinction between composition and categorisation, we can identify three basic forms of conceptual space. The most elementary is the *undeveloped* conceptual space. This embodies an empty set of construction methods and therefore lacks the ability to generate any new concepts. The only concepts within the space are the originating, primitive concepts themselves.

More generally, we have the *deep* conceptual space. This results when concept-construction may be recursive. The result of any compositional process is a new concept which may be used in a further round of concept construc-

tion. Thus, with compositional construction, there is always the potential for incremental, hierarchical growth of the space, an effect illustrated in Figure 2 .

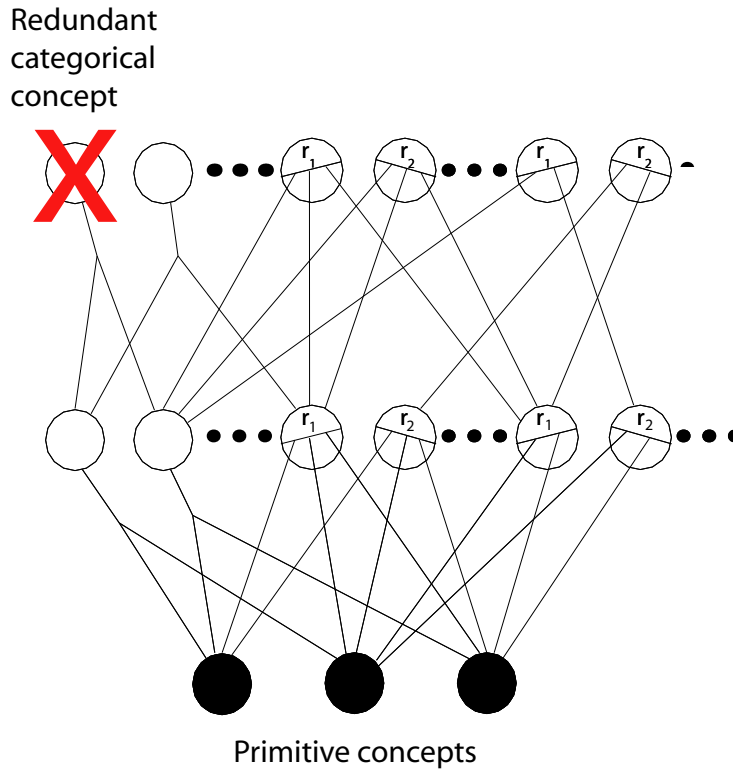


Figure 3: Possible avenues of recursive concept development.

In some cases, there may be constraints on the compositional methods which disallow recursive construction. Or it may be that only categorical construction is possible. In this situation, the conceptual space is *shallow*, since it lacks the capacity for hierarchical development. It is important to note that categorical construction *alone* cannot support recursive development. The construction of classes of concepts which themselves represent classes is a redundant operation: it can only produce a concept representing a larger class of originating objects, i.e., a concept which may be produced via a simple act of categorisation applied to the primitives. (The crossed-out concept in Figure 3 is a case in point.)

The distinction between deep and shallow conceptual spaces is qualitative rather than quantitative. As a rule, the deep conceptual space contains more concepts than a shallow one. Often it will be infinitely large. But there may be constraints which overturn the normal balance. Treating the sample grammar from above as a concept-generation system, one would have to say that it is

‘deep’ on the grounds that it makes use of recursive compositional rules. However, constraints are in place which limit the recursion. The total number of constructs is therefore finite. Were one to introduce a new rule so as to allow for infinite recursion, however (S → N for example), the situation would be altered.

The programmatic representation of 2-dimensional space, in contrast, is in the ‘shallow’ category due to the fact that it makes no use of recursive construction. Compositional construction is featured but only a tiny fraction of the possible combinations of primitives feature may appear in any generated construct — in particular the ones containing exactly two objects. The total number of these is related to the size of the represented space. But this may be arbitrarily large. Thus, this type of ‘shallow’ space has the potential to be larger than the ‘deep’ space defined by the grammar.

6 Conceptual space complexity

The contents of a particular conceptual space depend on (a) the originating concepts and (b) the available construction methods. But, as seen, the number of objects within the resulting conceptual space may vary. To identify the complexity of conceptualisation we need to know how many there are.

Assume we start with n originating concepts and make use of unconstrained construction methods, i.e., composition and categorisation applied to *any* combination of components. The undeveloped conceptual space, in which construction plays no role, contains just the originating n primitives. There are no derived concepts.

Introducing a categorical construction method produces a space with a distinct concept for every subset of the primitives. The maximum number of concepts is then

$$2^n$$

Adding a recursive construction method creates a situation in which the number of concepts may grow, level by level. The space is then a structured hierarchy rather than an unstructured set. At the lowest level are the originating concepts. At the first level are the concepts which can be constructed from the originating concepts. At the second level, there are the concepts which can be constructed from the concepts at the first level, and so on.

Let us say both categorical and compositional construction methods are available. At the bottom level of the hierarchy we have the n primitives. At the layer above, the number of concepts is some multiple of 2^n , since the number of possible combinations of n objects is 2^n , and our assumed system can generate a fixed number of concepts for each combination. Let’s say the system is capable of generating (for each combination) a categorical concept and r compositional concepts. In other words, it is capable of arranging the components as a class or in r different relationships. In this case the total number of concepts which can be generated is 2^n multiplied by $r + 1$. Given we are exponentiating to base 2, this can be expressed as

$$2^{(n+r)}$$

To identify the number of concepts which can be generated at *any* level, we might now try to derive a recursive version of this. However, there is a complication. The expression makes no distinction between concepts formed using categorisation and concepts formed using composition. The result is that it effectively ‘counts in’ redundant concepts formed using categorisation applied to concepts which were themselves constructed categorically. Some balancing discount is required.

With r methods of compositional construction, exactly $\frac{n}{r}$ of any new set of concepts will be class-based. Therefore in specifying the true number of potential concepts at the next level we should subtract

$$2^{\frac{n}{r}}$$

since this is the number of ways of constructing a categorical concept in terms of constituents that are also categorical concepts. The revised formula for the number of possible concepts then becomes

$$2^{(n+r)} - 2^{\left(\frac{n}{r}\right)}$$

Arguably, a further modification should be made to take account of the fact that exactly n of the 2^n possible combinations are singleton sets, i.e., they simply yield ‘stand-ins’ for nodes at the layer below. (We could discount these by subtracting $2n$.) However, if we wish to allow that construction may utilise concepts from different levels of the hierarchy — and at present we have no reason not to — it simplifies matters if we allow the singleton sets to remain. This way, every level of the tree contains a ‘stand-in’ for every node at every level below. The possibility of cross-level combination is automatically taken into account.

To render the formula in a general form is now straightforward. If n_0 is set equal to the number of originating concepts, the number of potential concepts n_i at the i 'th level of the hierarchy may be calculated using the recursive formula

$$n_{i+1} = 2^{(n_i+r)} - 2^{\left(\frac{n_i}{r}\right)} \tag{1}$$

This, then, is the *growth formula* for conceptualisation using unrestricted categorical and compositional construction. It enables calculation of the maximum number of concepts which may be developed at any level of construction. In fact, since each level contains a ‘stand-in’ for every concept contained below, the formula actually gives the maximum size of the entire space when developed *up to* this level. In effect, it bounds the size of the achievable conceptual space, when resource limitations are fixed to prevent development of concepts above this level.

As may be clear, the unconstrained, deep conceptual space grows at a prodigious rate, driven by two distinct exponential forces.³ After construction has proceeded above the initial few levels, the number of potential concepts becomes astronomically large. In any real context, only a small minority will have any value. But all are syntactically feasible.

We can illustrate the strength of the effect by considering an example. Were we to start out with just two primitive concepts and to apply just one method of (unconstrained) composition in addition to (unconstrained) categorisation, we could develop four concepts at the first level of construction and 16 concepts at the second level. At the third level, there are more than 65000 to take into consideration but at the fourth level there are nearly than 10^{20000} . Processing these at even the implausible rate of one quadrillion per second, it would still take far longer than the estimated age of the universe to get through them all. It can be inferred that any sort of brute-force exploration of conceptualisations is impractical.

7 Conceptual spaces as universal sub-spaces

The growth formula for unconstrained conceptualisation provides a mathematical formalisation for Boden's notion of conceptual space. Every possible way of processing every possible combination of concepts at every level of the expansion is taken into account. The space of possible conceptual developments derivable from certain originating concepts thus includes within it every possible conceptual space that might be defined in terms of any subset of those concepts. In other words, a conceptual space defined in terms of the growth formula contains within it every relevant conceptual *sub-space*.

This prompts the making of a distinction between conceptual spaces and conceptual universes. An agent's *conceptual universe* may be defined as the set of concepts that may be developed from the agent's original concepts and construction methods, i.e., whatever concepts it possesses at the start of its concept-using existence.⁴

This conceptual universe contains every possible concept that the agent might develop. It also subsumes every possible conceptual *space* that the agent might develop. In other words, the sub-spaces of the conceptual universe *are* the agent's potential conceptual spaces. (A schematic illustration of the situation appears in Figure 4.)

On this basis, we see that the construction of new conceptual spaces may be viewed, not as a separate and distinct process, but as a natural feature of ordinary concept development. In fact, the expansion of concepts based on a given set of primitives, using unconstrained composition and categorisation, must eventually construct *every* feasible conceptual space, although the process may take infinitely long. From the mathematical point of view, then, Boden's

³In addition to the explicit exponential value, there is a second multiplicative effect implemented through the recursion.

⁴Wiggins makes a similar distinction in setting up his formalisation (Wiggins, 2001).

Conceptual universe

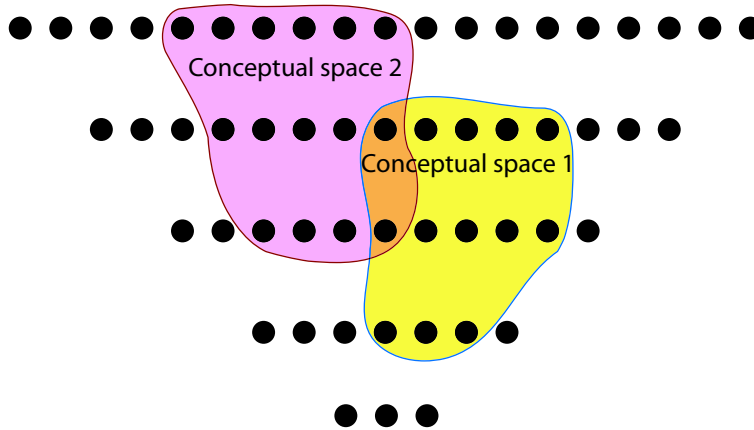


Figure 4: Conceptual spaces as sub-spaces within a conceptual universe.

two processes of exploration and transformation look to be reducible to a *single* mechanism and this appears to be ordinary concept construction.

Boden is committed to the idea that the two processes are distinct. Indeed she goes to some lengths to show how existing artificial intelligence models allocate functionality between the two. The examples she provides cover a wide territory. Working forwards from her Chapter 6 [ibid, pp. 112-216], we see that creative, conceptual-space operations may involve ‘pattern-completion’ and ‘contextual memory’ as featured in connectionist networks (Rumelhart *et al.* 1986); ‘reasoning’, and ‘problem solving’ as featured in traditional problem-solving systems; combining of ‘general and specific knowledge’, ‘constraint satisfaction’ and the exercise of ‘judgement’ as featured in Harold Cohen’s AARON system (Cohen, 1981).

It may involve operating in terms of complex, nested, hierarchical formulations, as in Johnson-Laird’s jazz improvisation program (Johnson-Laird, 1988); the application of ‘knowledge of planning’ as in TALE-SPIN (Schank and Riesbeck, 1981, pp. 197-258); ‘analogical thinking’ on the basis of ‘multiple constraint satisfaction’ and a recognition of ‘semantic similarity’ as featured in ACME (Holyoak and Thagard, 1989); the creation and mapping of analogies (using generate and test) as featured in ARCS (Holyoak *et al.* 1988); induction and learning, as featured in systems like AQ-11 (Michalski and Chilausky, 1980) and ID3 (Michie and Johnston, 1984, pp. 110-12); searching for new rules, as featured in meta-DENDRAL (Buchanan *et al.* 1976), or for quantitative relationships, as featured in ‘discovery’ systems such as BACON (Langley *et al.* 1987) or for qualitative laws, as featured in GLAUBER [ibid]; defining second-level theoretical terms as in later versions of BACON [ibid]; ‘hypothe-

sis generation’, ‘heuristic application’ and ‘concept transformation’ as featured in AM (Lenat, 1977). Last but not least it may involve rule-specialisation, as featured in EURISKO (Lenat, 1983).

This is certainly a broad catalogue. But the published descriptions for these systems make little or no use of Boden’s ‘conceptual space’ terminology as such. Indeed, in most cases, no reference is made to operations on any kind of concept *packaging* structure. Where these systems make use of concept representation, they do so on the basis of conceptual content rather than on the basis of conceptual-space membership. In general conceptual spaces are not significant as an implementational construct in these references.

There would appear to be nothing in the example set, then, which *requires* maintenance of the distinction between exploration and transformation. In fact, consideration of the likely representational requirements suggests, rather, that *enforcing* the distinction in a conceptual agent might impose prohibitive resource costs. The originating concepts for any new conceptual space must clearly be represented somewhere in the overall conceptual system. But they cannot be located within the conceptual space itself, for fear of infinite regress. The content of a primitive of a new conceptual space is thus necessarily located *outside* the space itself. Since every concept within a conceptual space ultimately references the content of its primitives, content is then always to some degree cross-referenced against other conceptual spaces.

Exacerbating the effect is the fact that the situation is ‘recursive’. Concept A may invoke concept B, which invokes concept C, and so on. We therefore have to expect the content of a given concept to be distributed across a series of spaces. In an extreme case, it might be ‘smeared’ across the entire conceptual system.

Achieving independent representation of conceptual spaces would necessitate some sort of ‘de-smearing’ operation, i.e., introduction of a mechanism which has the effect of encapsulating within a given space all the chains of reference underpinning the content of member concepts. But except in the most trivial scenarios, the combinatoric properties of any such operation would generate exponentially increasing resource requirements. We have to assume, then, that conceptual spaces *cannot* be represented ultimately as independent objects. Rather, they must be acquired from a representation in which constituent structure serves to represent content in the usual way. At its most basic level, then, it would seem that a concept-representing system *cannot* be structured as a set of independent conceptual spaces. It must be structured in the familiar way, as a hierarchy where between-concept connections serve to reference content.

Two arguments thus converge in support of the idea that the distinction between exploration and transformation is essentially subjective, and one which we should not expect to find consolidated at implementation level. Nothing in Boden’s example set of systems, or in our understanding of how those systems work, suggests a concrete necessity for it. And the observation that any transformational process must operate, ultimately, on a structure in which conceptual spaces are not *explicitly* represented suggests, rather, that transformation has to be reducible to the more mundane process of conceptual search. Finally, we

also have Wiggins formal demonstration that the two processes are effectively equivalent (Wiggins, 2001).

8 Steps towards a complexity-based definition

With no clear distinction being identifiable between transformation and exploration, the creativity definition stands in need of new foundation. Bundy's suggestion (1994) is that the definition could be re-constructed in terms of some notion of computational complexity. Given the analysis developed above, this idea can now be progressed.

As an initial venture, we might try substituting the complexity measure (Equation 1) *directly* into the definition, as a self-contained replacement for the all-or-nothing criterion.

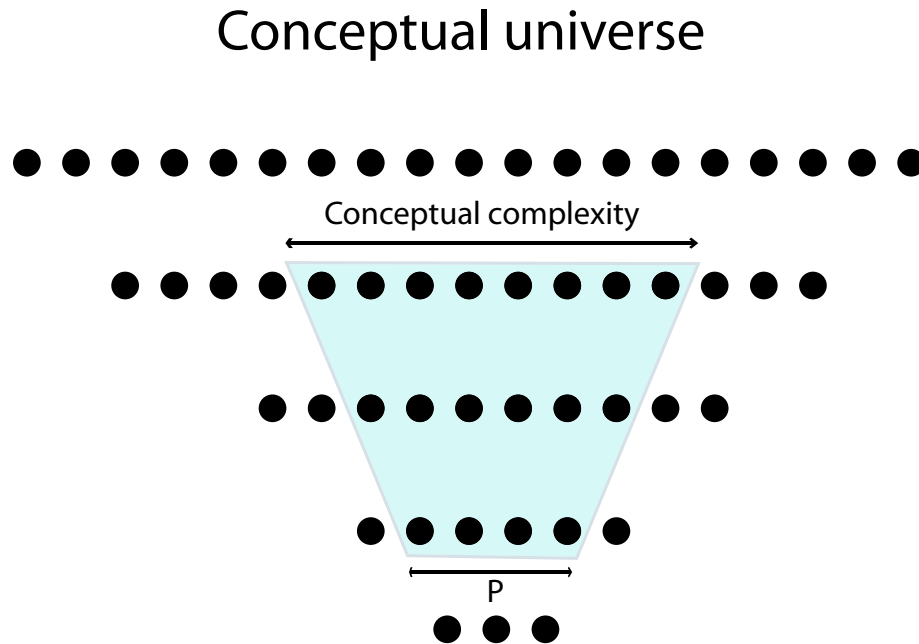


Figure 5: Visual illustration of conceptual complexity.

Where Boden's original definition states that the creativity of an idea depends on whether or not its conception involves the development of a novel conceptual space, our revised definition would state that the *degree* of creativity of an idea is to be measured in terms of the computational complexity of its construction. (A graphical illustration of the derivation of the complexity of concept c in terms of primitives P appears in Figure 5 .) To the extent that this captures the idea that creativity involves extra representational effort, it seems

a plausible replacement, while the fact that it eliminates the ‘all-or-nothing’ property of that definition might also be regarded as attractive. How well does it stand up in other respects?

Boden herself has commented on the general idea of using a complexity rule as a means of evaluating creativity, pointing out specific requirements that any satisfactory measure would have to satisfy. In discussing the general question of whether creativity can be *measured* she writes ‘Some form of complexity measurement, as used by computer scientists, would be useful. However, depth within the space must be recognised too. The appropriate method of assessment would have to take into account the fact that conceptual spaces are multidimensional structures, where some features are “deeper”, more influential, than others.’ (Boden, 1994, p. 113)

She goes on to say that a satisfactory measure should somehow return a ‘structured distance’ and be sensitive to ‘interesting differences’ between artefacts of different styles or genres.⁵ It should somehow reconcile different dimensions of creativity and deal properly with the tradeoff between conceptual ‘depth’ and ‘number’. It should also properly relate transformational creativity to exploratory creativity. [ibid., p. 114-115] How well does the proposed definition stand up to these demands?

Because of the way the growth formula generalises exploratory and transformational processes, it is capable of capturing the differences between these two modes of operation without modification. It also handles the tradeoff between conceptual ‘depth’ and ‘number’ in the sense it measures both things in the same way — in terms of the amount of work involved in the generation of the relevant concept.

The problem relating to incommensurability among concepts, is also solved but with a caveat. On the assumption that different genres embody concepts drawn from different conceptual spaces, no *space-specific* measure can deliver an unambiguous cross-genre comparison. However, measures produced this way are commensurable in the sense that they are calculated in the same units. Thus there is the potential for cross-genre comparison provided that the relative contribution embodied in the respective primitives is taken into account.

Overall, the replacement of the all-or-nothing criterion with the complexity measure produces a revised definition which implements Bundy’s suggestion reasonably well, and does so in a manner which satisfies most of Boden’s requirements (always remembering that these were not actually intended for this use). But it seems to me that, even so, it is not quite what is wanted. Something important from the original formulation is being lost.

Although the systems evidence Boden brings to bear may not quite compel us to believe in the mediational role of conceptual-space structures, the human evidence is more persuasive. The anecdotal evidence Boden assembles appears to suggest it is key facet of creativity that it is mediated within structures or frameworks of concepts, or theories. For a satisfactory re-working of the defini-

⁵She gives the example of a comparison of the degree of creativity in the *Mona Lisa* as opposed to the *Demoiselles d’Avignon*.

tion, this needs to be brought out. What is really needed is a complexity-based definition which somehow gives the conceptual space the prominent position it formerly occupied.

9 The net/gross distinction

The complexity of a concept (as measured by the growth formula) is a measure of the amount of work that needs to be applied to bring the concept into representation. Let this now be called the *gross conceptual complexity*. The *net conceptual complexity* can then be defined as the original amount of work less the amount of work that bringing the concept into existence *eliminates*.

To calculate how much this is, we only have to look at the extension or coverage of the relevant concept. The establishment of a concept covering all the entities in a particular set X means that they may subsequently be processed as a single entity. That is, after all, the key goal of conceptualisation. The representational gain may then be measured in terms of the difference between having to process all members of X individually and being able to process them as a single entity. In other words, the representation gain may be measured in terms of the size of X .

This leads directly to the following formulation for *net conceptual complexity*:

$$\text{NCC}(c, P) = n_i - |E(c)| \quad (2)$$

Here i is the level at which concept c appears in a hierarchy based on primitives P . n_i is the complexity of concept c , derived using Equation 1 with $n_0 = |P|$. $|E(c)|$ is the size of the extension of concept c . The formula asserts that the net conceptual complexity (NCC) of concept c derived in a hierarchy based on primitives P is the objective conceptual complexity less the size of c 's extension. A schematic illustration appears in Figure 6.

Net conceptual complexity is a measure of the true, representational gain which results from bringing a particular concept into representation. NCC values may be high or low; they can also be *negative*. This occurs in situations where the representational benefit of bringing a concept into representation outweighs the cost of doing so. Treating the NCC as a foundation for a creativity definition, there is the possibility of equating ‘radical creativity’ specifically with the production of a negative NCC value.

The NCC seems to be a better basis on which to ground the definition, then, since it allows us to have our cake and eat it too with respect to grain-size. Should we prefer the all-or-nothing approach, we can achieve this by the identification of ‘radical creativity’ with the production of negative NCC values. Should we prefer the graded view, we can achieve this by using NCC values directly.

In addition, utilisation of NCC as a foundation for the definition allows the incorporation of an evaluative element. The original revision, like Boden’s criterion itself, ignores the *value* of conceptualisation, focussing exclusively on the amount and nature of work involved. On this type of view, a useful concept




Figure 6: Visual illustration of net conceptual complexity.

may be rated as being just as ‘creative’ as a useless concept, provided only that the conceptualisation effort is equivalent.

Using the concept of NCC within the definition allows us to get beyond this, at least with respect to contexts in which representational power is the significant factor. It allows the definition to capture not only the idea that creativity involves representational work at a specific level, but also the idea that it should involve movement in the direction of greater representational coverage.

NCC does not offer a *general* solution to the problem of evaluation however. While it is sensitive to a key element of intellectual value (explanatory coverage), it ignores artistic value completely. Even with regard to the measurement of explanatory value, it presents a somewhat blinkered view, being responsive specifically to coverage rather than to problem-solving potential. As a relational construct, it also fails to distinguish high-cost/high-coverage concepts from low-cost/low-coverage concepts, treating both as embodying the same degree of creativity (Bundy, 2005).

10 Sponge zones

Any concept in a conceptual universe may be labelled with its NCC value but most of the syntactically feasible concepts in a particular conceptual universe are likely to be of no value or significance. The NCC values for such concepts will be positive values, and relatively high, indicating that the cost of conceptualisation considerably outweighs any representational benefit. Other NCC values will be lower, indicating that representational benefits exist, i.e., that the corresponding concept is of some use. Some values will be negative indicating that the representational benefits *outweigh* the conceptualisation cost.

With this labelling in mind, it is possible to view the conceptual universe as a landscape of peaks and troughs. Peaks occur in the vicinity of concepts with high NCC, small elevations where concepts show low NCC, and where concepts show negative NCC, there are marked depressions. The landscape thus shows regions of high, low and negative NCC.

In the case where a set of concepts form the primitive constituents of an explanatory theory, we would expect to find a marked depression ‘upstream’ in the hierarchy. This depression will be made up of all the concepts which utilise the theory and which share in its explanatory power. The depression is a zone in which conceptualisation has the effect of *reducing* conceptual complexity. Putting it another way, the depression acts like a ‘sponge’, soaking up and eliminating complexity from the concept hierarchy. See Figure 7 for an illustration.

A conceptual universe is thus a landscape in which we expect to find pronounced areas of low or negative complexity — sponge zones — located just upstream of any set of concepts which form the constituents of an explanatory theory or conceptual framework. By accessing concepts in these zones, the conceptual agent moves into representational ‘profit’ reaping a representational reward that exceeds the cost of development.

Since every conceptual space is based on primitives which, by definition, provide the basis of an explanatory framework or theory, there is a direct correlation between conceptual spaces and sponge zones. Every conceptual space *is* a sponge zone with the explanatory power of the space corresponding to the ‘absorbency’ of the sponge, i.e., its capacity to eliminate complexity.

Boden’s observations about the ways in which creative thinkers operate with respect to conceptual spaces can thus be understood in terms of sponge zones and complexity reduction. Exploration is the attempt to focus conceptualisation *within* zones of maximum dividend. Transformation can be understood as the attempt to open up new zones, in which exploration may then be profitably applied.

11 Evaluation

Instead of Boden’s all-or-nothing creativity criterion, the NCC formula may be brought to bear as a means of grounding a core definition of creativity. Using

Conceptual universe

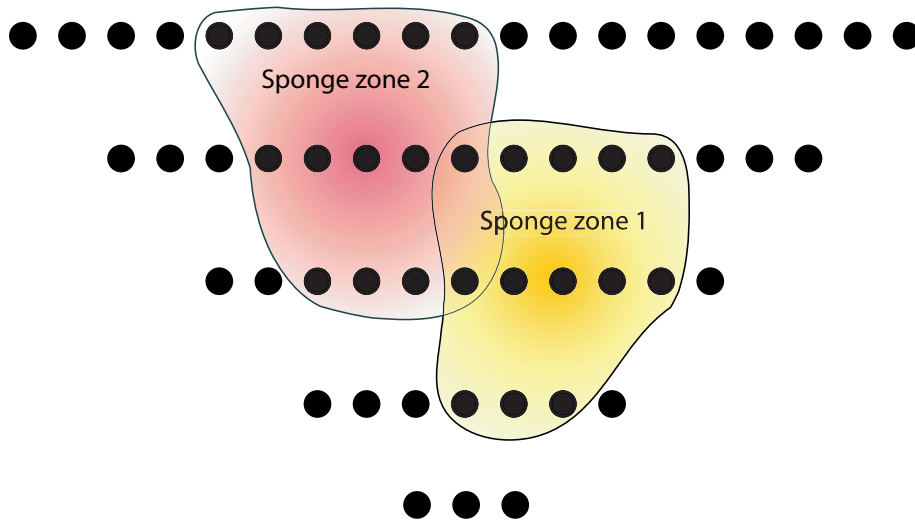


Figure 7: Sponge zones in a conceptual space.

NCC, we can have either an all-or-nothing or a graded definition, depending on preference. Degree of conceptual creativity may be measured directly in terms of the corresponding NCC value, while the identification of ‘radical creativity’ may be reserved for situations in which the NCC is actually negative.

Recognition of the role played by sponge zones allows us to see that this approach in no way undermines Boden’s usage of the conceptual space as an explanatory entity. Rather it supports and confirms it by illuminating the relevant efficiency implications. The original account, with the NCC-based definitions substituted for the all-or-nothing criterion, would therefore appear to form a valid reconstruction of her theory.

Boden herself might not be completely happy with it, however. She has registered some resistance to the *idea* of a complexity-based definition on the grounds that it could not truly reflect the multi-dimensional nature of creativity. As she comments in a response to a reviewer, ‘I prefer to avoid speaking of the “degree” of creativity, since to do so implies a continuous spectrum along which individual thoughts are to be ordered. To the contrary, a main theme of [The Creative Mind] is that creativity is multi-dimensional.’ (Boden, 1995, p. 169) To me this seems unnecessarily restrictive. We regularly use 1-dimensional measures of multi-dimensional phenomena. Indeed we do so every time we speak about distance in 3-dimensional space. So I feel we should not necessarily balk

at the idea of a definition based on a 1-dimensional measure.⁶

Others may find the reconstruction more appealing since it does appear to answer some of the criticisms which have been raised. Indeed, it would appear to answer Ram et al.'s requirements (Ram *et al.* 1995, p. 122) directly. Though quite happy with the overall form of Boden's account, these reviewers worry that it does not sufficiently clarify the way transformation is grounded in normal conceptual operations. As they note, 'The question is though how the search space comes to be expanded to facilitate creative thought using ordinary mechanisms.' In particular they want to know how 'ordinary operators and processes can take the reasoner out of the space that would usually be explored.' [ibid, p. 122]

Since it takes account of the reducibility of transformation to ordinary conceptual search, the reconstruction is able to answer this question while still staying within the general scope of the original account. Ram et al. should therefore find it acceptable. As they note 'we would prefer a model of long-term conceptual development in which the individual evolves a search space, that, when explored by normal thought processes, still includes many thoughts that would be considered *creative*. [ibid, p. 114] The reconstruction delivers precisely this enhancement.

12 The Universal Thought Year

Unfortunately, there remains at least one problem which is no more resolved in the revised theory than it was in the original. This is the problem of *subjectivity*, noted above. Values of NCC depend not only on the degree of work involved in conceptualisation and the degree of representational profit, they also depend on the set of available primitives. With a higher-level set of primitives, NCC values will be lower, and vice versa. Changing the set of primitives assumed to be available impacts NCC values directly. It might even have the effect of making a 'radical creativity' value positive, thus changing its qualitative classification.

An ideal definition of creativity would be fully *objective*, of course. But this could only be achieved by establishing a fixed starting point for conceptual development, i.e., a 'standard conceptual agent' equipped with a specific set of primitive concepts and constructive operations. On this basis, it would then be possible to measure the creativity of any concept against a fixed background, perhaps producing values expressed in *thought years* (analogous to light years) which would be the amount of work required from the standard agent in order to produce the relevant concept. (Figure 8 provides a schematic illustration of the way this might work with respect to some financial concepts.) However, agreeing a specification for the standard agent will be controversial. The innovation of the standard 'thought year' and emergence of a fully objective creativity definition is thus probably still some way off.

⁶Were we to do so, we should arguably have the same reaction to the original definition.




Figure 8: Imaginary thought-year distances for some financial concepts.

13 Integration and Equality

An additional benefit of the reconstruction is that it tends to reconcile Boden's theory with Koestler's earlier account. The key difference between the two is that Koestler envisaged one basic mechanism of creativity while Boden envisages two. But in the reconstruction we find that Boden's two processes may be reduced to a common core which, once analysed, emerges as a reasonable model for Koestler's analogy-like 'matrix bisociation'. This points towards a possible unification of the two accounts.

One issue on which the two accounts are already in complete agreement is the question of 'special powers.' In Boden's view, it was a key strength of Koestler's approach that it 'appealed to no special creative faculty, granted only to an elite.' (Boden, 1990, p. 24). Boden's own theory closely follows Koestler's line on this, stressing that 'creativity is not a single ability, or talent.' [ibid., p. 12] Both writers, then, see creativity as a process closely integrated within other processes of cognition.

The reconstruction does not in any way contradict this 'diffusional' view. Rather it consolidates it in the form of a still more radical proposition: that

as a process, creativity is actually *indistinguishable* from ordinary conceptual development. In other words, there is no substantive difference to be discerned in the underlying mechanisms of conceptual development, and those of creativity. Of course, as a result of having different experiences, and perhaps of being equipped with a different set of original concepts, different concept-using agents will inevitably produce different results. But, on the present account, the deed is done in the same way.

This idea encourages an egalitarian viewpoint in which ‘creative masters’ are understood to be using the same mental equipment as ‘common-or-garden thinkers’. While they may produce very different results, this would seem to be only because of different experiences and starting conditions. On this view, Kekule’s dream-inspired innovation of the benzene ring concept was achieved by exactly the same process as would have been involved in the imaginary crossing-sweeper’s innovation of the ‘nicely swept crossing’ concept. Indeed, if we believe that animals generate concepts in the same way we do, we might have to go further still, perhaps accepting that Kekule’s ring concept was derived by the same mechanism as that used by, say, a large labrador developing the concept of ‘small-furry-chasable-thing’. In the context of science, one innovation may be of greater magnificence than another. In the context of mechanism, there may, after all, be very little difference.

14 Final summary

The aim of the paper has been to progress Bundy’s suggestion of replacing Boden’s all-or-nothing creativity definition with one based on a quantitative notion of complexity. Key elements in the exposition were the clarification of the distinction between shallow and deep conceptual development, the derivation of a complexity analysis for various forms of conceptual development and the demonstration that both of Boden’s creative processes are implicitly carried out in deep conceptualisation.

The complexity formula for conceptualisation was utilised for purposes of constructing a revised definition of creativity. The final version of this defines creativity in terms of net conceptual complexity, with the category of ‘radical creativity’ being reserved for concepts carrying negative NCC. The reconstruction was found to answer requests for a theory which would clarify the basis of creativity in ‘ordinary thought’ and to deal more directly with the question of evaluation. Finally, it was found to add additional stress to some of the fundamental themes of Boden’s account, including the insistence that creativity should be regarded as an integral part of ordinary cognition. In sum, the reconstruction bears out much of what Boden has argued, but allows it to rest, definitionally-speaking, on firmer ground.

References

- Boden, M. (1990). *The Creative Mind: Myths and Mechanisms*. London: Weidenfeld and Nicolson.
- Boden, M. (Ed.) (1994). *Dimensions of Creativity*. MIT Press.
- Boden, M. (1995). Modelling creativity: reply to reviewers. *Artificial Intelligence*, No. 79 (pp. 161-182).
- Boden, M. (1998). Creativity and artificial intelligence. *Artificial Intelligence*, 103 (pp. 347-356). 1-2.
- Boden, M. (2003). *The Creative Mind: Myths and Mechanisms* (2nd edition). London: Weidenfeld and Nicolson.
- Buchanan, B., Smith, D., White, W., Gritter, R., Feigenbaum, E., Lederberg, J. and Djerassi, C. (1976). Applications of artificial intelligence for chemical inference: XXII automatic rule formation in mass spectrometry by means of the meta-dendral program. *Journal of the American Chemistry Society*, No. 98 (pp. 6168-78).
- Bundy, A. (1994). What is the difference between real creativity and mere novelty?. *Behavioral and Brain Sciences*, 17, No. 3 (pp. 533-534).
- Bundy, A. (2005). *Personal communication*.
- Cohen, H. (1981). *On the Modelling of Creative Behavior*. No. 6681, Santa Monica, Calif: Rand Corporation.
- Fetzer, J. (1994). Creative thinking presupposes the capacity for thought. *Behavioral and Brain Sciences*, 17, No. 3 (pp. 539-540).
- Holyoak, K. and Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science*, No. 13 (pp. 295-396).
- Holyoak, K., Thagard, P., Nelson, G. and Gochfield, D. (1988). *Analog Retrieval by Constraint Satisfaction*. Unpublished research-paper.
- Johnson-Laird, P. (1988). Freedom and constraint in creativity. In R.J. Steinberg (Ed.), *The Nature of Creativity: Contemporary Psychological Perspectives* (pp. 209-19). Cambridge.
- Koestler, A. (1964). *The Act of Creation*. London: Hutchinson.
- Kuscu, I. and Thornton, C. (1994). Design of artificial neural networks using genetic algorithms: review and prospects. In C. Bozsahin, U. Halici, K. Oflazer and N. Yalabik (Eds.), *Proceedings of the Third Turkish Symposium on Artificial Intelligence and Neural Networks* (pp. 411-420). Ankara.

- Langley, P., Simon, H., Bradshaw, G. and Zytkow, J. (1987). *Scientific Discovery: Computational Explorations of the Creative Processes*. Cambridge, Mass.: MIT Press.
- Lenat, D. (1977). The ubiquity of discovery. *Artificial Intelligence* (pp. 257-86).
- Lenat, D. (1983). EURISKO: a program that learns new heuristics and domain concepts: the nature of heuristics III: program design and results. *Artificial Intelligence*, 21 (pp. 61-98).
- Michalski, R. and Chilausky, R. (1980). Learning by being told and learning from examples: an experimental comparison of the two methods of knowledge acquisition in the context of developing an expert system for soybean disease diagnosis. *International Journal of Policy Analysis and Information Systems* (pp. 125-61).
- Michie, D. and Johnston, R. (1984). *The Creative Computer: Machine Intelligence and Human Knowledge*. London: Penguin.
- O'Rourke, J. (1994). The generative-rules definition of creativity. *Behavioral and Brain Sciences*, 17, No. 3 (p. 547).
- Perkins, D. (1981). *The Mind's Best Work*. Cambridge, Mass.: Harvard University Press.
- Ram, A., Domeshek, E., Wills, L., Neressian, N. and Kolodner, J. (1994). Creativity is in the mind of the creator. *Behavioral and Brain Sciences*, 17, No. 3 (p. 549).
- Ram, A., Wills, L., Domeshek, E., Neressian, N. and Kolodner, J. (1995). Understanding the creative mind: a review of margaret boden's 'creative mind'. *Artificial Intelligence*, No. 79 (pp. 111-128).
- Richie, G. (2001). Assessing creativity. In G. Wiggins (Ed.), *Proceedings of the AISB'01 Symposium on AI and Creativity in Arts and Sciences* (p. AISB).
- Rumelhart, D., McClelland, J. and the PDP Research Group, (Eds.) (1986). *Parallel Distributed Processing: Explorations in the Microstructures of Cognition. Vols I and II*. Cambridge, Mass.: MIT Press.
- Schank, R. and Forster, D. (1995). The engineering of creativity: a review of boden's 'the creative mind'. *Artificial Intelligence*, No. 79 (pp. 129-143).
- Schank, R. and Riesbeck, C. (1981). *Inside Computer Understanding: Five Programs Plus Miniatures*. Hillsdale, N.J.: Lawrence Erlbaum Associates.
- Treisman, M. (1994). Creativity: myths? mechanisms?. *Behavioral and Brain Sciences*, 17, No. 3 (p. 554).

- Turner, S. (1995). Margaret boden, 'the creative mind'. *Artificial Intelligence*, No. 79 (pp. 145-159).
- Wallas, G. (1926). *The Art of Thought*.
- Weisberg, R. (1994). The creative mind versus the creative computer. *Behavioral and Brain Sciences*, 17, No. 3 (pp. 555-557).
- Wiggins, G. (2001). Towards a more precise characterisation of creativity in AI. *Proceedings of the ICCBR 2001 Workshop on Creative Systems*. Vancouver, British Columbia.